Graphs, Gain Graphs, and Geometry

a.k.a.
Signed Graphs and their Friends

Course notes for

Math 581: Graphs and Geometry
Fall, 2008 — Spring–Summer, 2009 — Spring, 2010

Math 581: Signed Graphs
Fall, 2014

Binghamton University (SUNY)

Version of November 23, 2014

Lectures by

Thomas Zaslavsky

Notes by

Richard Behr, Peter Cohen, Melissa Fuentes, Simon Joyce, Jackie Kaminski, Yash Lodha, Nathan Reff, Ting Su, and Thomas Zaslavsky

Redaction by

Thomas Zaslavsky

Contents

Chapter II. Signed Graphs

Oct 6: Jackie Kaminski

Now at last we've arrived at the meat of the course.[1] Our purpose is to generalize graph theory to signed graphs. Not all of graph theory does so generalize, but an enormous amount of it does—or should, if the effort were made. Since that has not happened yet, there is plenty of room for a fertile imagination to create new graph theory about signed graphs.

A. Introduction to Signed Graphs

[[LABEL 2.basics]][[LABEL 2.sg]]

A signed graph is a graph with signed edges. But what, precisely, does that mean? In fact, not every edge has a sign; it is only ordinary edges—links and loops—that do.

A.1. **What a signed graph is.** [[LABEL 2.sgintro]]

We give two definitions. The first is the simpler: every edge gets a sign. The cost is that we cannot have loose or half edges; but as we shall see in the treatment of contraction (Section E.1) that is rather too constraining, whence the second, more general definition.

Definition A.1. [[LABEL D:1006 Ord. Signed Graph]] An *ordinary signed graph* is a signed ordinary graph, that is, $\Sigma = (\Gamma, \sigma) = (V, E, \sigma)$ where Γ is an ordinary graph (its edges are links and loops) and $\sigma : E \to \{+, -\}$ is any function.

Definition A.2. [[LABEL D:1006 Signed Graph]] For any graph Γ, which may have half or loose edges, we define $E^* = \{e : E(\Gamma) : e \text{ is a loop or link}\}$. A *signed graph* is $\Sigma = (\Gamma, \sigma) = (V, E, \sigma)$ where Γ is any graph and σ is any function $\sigma : E^* \to \{+, -\}$.

In either case, we call σ the *(edge) signature* or *sign function*. Not surprisingly, we refer to $\{+, -\}$ as the *sign group*. One may instead use the additive group $\mathbb{Z}_2^+ = \mathbb{F}_2^+ = \{0, 1\}$ as the sign group, or the group of signed units $\{+1, -1\}$. We prefer the strictly multiplicative point of view implied by $\{+, -\}$ for reasons that will become clear when we discuss equations (Section I); a hint appears when we define the sign of a walk.

A subgraph $\Gamma' = (V', E')$ of Γ is naturally signed by the signature $\sigma' = \sigma|_{E'}$. An edge subset $S \subseteq E$ makes the natural signed spanning subgraph $(V, S, \sigma|_S)$.

Definition A.3. [[LABEL D:1006 Isomorphism]] Signed graphs Σ_1 and Σ_2 are *isomorphic* if there is a graph isomorphism between Σ_1 and Σ_2 that preserves edge signs.

Many people write $+1$ and -1 instead of $+$ and $-$. This is harmless as long as we remember the symbols are not numbers to be added. (I will eat these words when it comes to defining the adjacency matrix.)

To some signed graph theorists (in particular, Slilaty), loose edges are positive, and half edges are negative. This is *not* a convention I will use.

For general culture, I point out that it is well known that graph theory has been invented independently by many people. Signed graph theory was also independently invented by multiple people for multiple purposes.

A curious sidelight is that, in knot theory, there is a similar-looking assignment of labels $\{+1, -1\}$ to edges. This is not a signed graph in our sense because the "signs" $+1$ and -1

[1]For benefit of vegetarians: the term "meat" is intended in its early sense of 'substantial food', not 'flesh'.

are interchangeable, so they do not form a group. The sign group $\{+, -\}$ is present implicitly through the action of swapping or not swapping the edge labels.

A.2. **Examples of signed graphs.** [[LABEL 2.examples]]

The *restriction* of Σ to an edge set S is $\Sigma|S := (V, S, \sigma|_S)$.

Oct 22: Jackie Kaminski

Several different signed graphs can be constructed from an ordinary graph Γ.

Example A.1. [[LABEL X:1022all+]] $+\Gamma = (\Gamma, +)$, where all edges are positive. We call it the all-positive Γ.

Example A.2. [[LABEL X:1022all-]] $-\Gamma = (\Gamma, -)$, where all edges are negative. Unsurprisingly, we call it the all-negative Γ.

Example A.3. [[LABEL X:1022homogeneous]] If Σ is all positive or all negative we call it *homogeneous*. Otherwise it is *heterogeneous*.[2]

Example A.4. [[LABEL X:1022pmG]] $\pm\Gamma = (+\Gamma) \cup (-\Gamma)$, where we differentiate the positive and negative edges, i.e., $e \in \Gamma \mapsto +e, -e \in \pm\Gamma$. Thus, $V(\pm\Gamma) = V(\Gamma)$ and $E(\pm\Gamma) = +(E(\Gamma)) \cupdot -(E(\Gamma))$. We call this the *signed expansion* of Γ. Loops and especially half edges can be problematic here, so we generally would assume Γ is a link graph.

Example A.5. [[LABEL X:1022full]] Σ is *full* if every vertex supports a half edge or a negative loop. In other words, referring ahead to Section A.5, every vertex supports an unbalanced edge. This is not the same as having every vertex supporting a negative edge. Note that the terms positive and balanced, or negative and unbalanced, are equivalent for circles, but not for edges.

Example A.6. [[LABEL X:1022filled]] Σ° is Σ with a negative loop adjoined to every vertex that doesn't have a negative loop or half edge. $\Sigma^\bullet$ is Σ with a negative loop or half edge adjoined to every vertex that doesn't have one.

We can also define $+\Gamma^\circ$, $-\Gamma^\circ$, and $\pm\Gamma^\circ$. We think of doing the $+$, $-$, or $\pm$ before adding the negative loops, so that the final result has one negative loop at each vertex, not a positive loop or two loops.

---

10 Nov. 2014: T.Z.

Example A.7. [[LABEL X:1022signedKn]] If Δ is a simple graph, K_Δ is the signed complete graph with vertex set $V(\Delta)$ and negative edge set $E^-(K_\Delta) := E(\Delta)$.

A.3. **Degrees.** [[LABEL 2.degrees]]

A vertex of Σ has several different degrees.

2 Nov. 2014: T.Z.

Definition A.4. [[LABEL D:20141102degrees]] Let's define the various signed degrees of a vertex of Σ. The degree $d(v)$ is the degree in $|\Sigma|$. The *positive degree* $d^+(v)$ and the *negative degree* $d^-(v)$ are the degrees in, respectively, the positive and negative subgraphs, Σ^+ and Σ^-. In all these degrees a loop counts twice and a half edge once—but a half edge, so far, has not entered into any of the signed degrees. Let the *half-edge degree* be $d^{1/2}(v) :=$ the number of half edges incident with v. Then the *net degree* of v is

$$d^\pm(v_i) := d^+(v) - d^-(v) + d^{1/2}(v).$$

A loose edge, of course, appears nowhere in the degrees.

[2] These impressive and also useful names are due to my late friend Dr. B.D. Acharya.

A.4. **Walk and circle signs.** [[LABEL 2.walksigns]]

Oct 6: Jackie Kaminski

The signs of walks and (especially) circles are fundamental in the subject of signed graphs.

Definition A.5. [[LABEL D:1006 Walk Sign]] For a walk $W = e_1 e_2 \cdots e_l$, the sign of W is

$$\sigma(W) := \sigma(e_1)\sigma(e_2)\cdots\sigma(e_l).$$

We call a walk positive or negative according to its sign; in particular each circle is positive or negative.

(Contrastingly, a walk is *all-positive* if it is homogeneously positive, *all-negative* if homogeneously negative.)

Note that this definition does not depend on the edge set of the walk, but on precisely how often each edge appears in the walk. As a circle is simply a closed walk, we can define the sign of a circle similarly; but since each edge of the circle appears exactly once, the sign of a circle (as a walk) is also its sign as an edge set.

Definition A.6. [[LABEL D:1006 Set Positive Circles]] In a signed graph Σ,

$$\mathcal{B}(\Sigma) := \{\text{positive circles of } \Sigma\}.$$

The complementary subset in $\mathcal{C}(|\Sigma|)$ is $\mathcal{B}^c(\Sigma) := \{\text{negative circles of } \Sigma\}$.

A.5. **Balance.** [[LABEL 2.balance]]

We are now ready to define the key concept of signed graph theory (as I interpret it): balance.

Definition A.7. [[LABEL D:1006 Balanced Graph]] We say Σ is *balanced* if it has no half edges and every circle is positive. A subgraph is balanced if it is a balanced signed graph. An edge set S is balanced if the restriction $\Sigma|S$ is balanced.

The negation of balance is *imbalance*. We say Σ is *unbalanced* if it is not balanced. But the truly opposite concept is being entirely unbalanced.

Definition A.8. [[LABEL D:1107 contrabalance]]A signed graph is *contrabalanced* if it contains no loose edges or balanced circles.

Exercise A.1. [[LABEL Ex:1107 contrabalance]] Which signed graphs are contrabalanced? Prove the following characterization. A *cactus* is a connected graph in which every nontrivial block is a circle or an isthmus. If every component of a graph is a cactus, we call it a *cactus forest*.

Proposition A.1. *[[LABEL P:1107 contrabalance]] A signed ordinary graph is contrabalanced if and only if it is a cactus forest in which every circle is negative.*

A hint is that a graph is a cactus forest if and only if it contains no theta subgraph. (Prove that as part of your solution.)

Two types of signed graph, which will be essential as subgraphs, are the two varieties of circuit.

Definition A.9. [[LABEL Df:1022fcircuit]] A *frame circuit* is a signed graph, or an edge set, that is either a positive circle, or a pair of negative circles that have in common just one vertex (a *contrabalanced tight handcuff*), or a pair of vertex-disjoint negative circles together with a minimal connecting path (a *contrabalanced loose handcuff*). Throughout this definition, a half edge may substitute for any negative circle.

Minimality of the connecting path means that it must intersect each circle in one vertex, which is an endpoint of the path; but the path need not have minimum length.

Definition A.10. [[LABEL Df:1022lcircuit]] A *lift circuit* is a signed graph, or an edge set, that is either a positive circle, or a pair of negative circles that have in common just one vertex (a *contrabalanced tight bracelet*), or a pair of vertex-disjoint negative circles (a *contrabalanced loose bracelet*). In this definition, a half edge may substitute for any negative circle.

Note that the presence of loose edges has no effect on balance.

We now state Harary's Balance Theorem (known in psychology as the "Structure Theorem").

Definition A.11. [[LABEL D:1006 bipartition]] A *bipartition* of a set X is an unordered pair $\{X_1, X_2\}$ of complementary subsets, that is, subsets such that $X_1 \cup X_2 = X$ and $X_1 \cap X_2 = \varnothing$. X_1 or X_2 could be empty.

A bipartition isn't simply a partition into two parts, since in a partition the parts are not allowed to be empty.

Theorem A.2 (Balance Theorem (Harary 1953a))**.** *[[LABEL T:1006 Harary]] Σ is balanced $\iff$ there is a bipartition $V = V_1 \cup V_2$ such that every negative edge has one endpoint in V_1 and the other in V_2 and every positive edge has both endpoints in V_1 or both in V_2, and Σ has no half edges.*

We call a bipartition as in the Balance Theorem a *Harary bipartition* of Σ. That is, a Harary bipartition is a bipartition of V into $\{X, X^c\}$ such that every positive edge is within X or within X^c, and every negative edge has one endpoint in each. Notice that we are ignoring half edges in this definition; thus, the statement of Harary's theorem is that Σ is balanced iff it has a Harary bipartition and it has no half edges. (Although, in fact, Harary's signed graphs had no half edges!)

The original proof is somewhat long. We'll have a shorter but more sophisticated proof soon (Section A.6). For now we make a few observations about balance. First, in a balanced graph (or balanced subgraph) all loops must be positive. Second is a lemma that can be useful in many proofs.

Lemma A.3. *[[LABEL L:1006 balanced blocks]] Σ is balanced if and only if every block of Σ is balanced.*

Recall that a graph is inseparable if every pair of edges is in a common circle. (Some people define a graph to be 2-connected if does not contain any cutpoints, where a cutpoint is a vertex whose deletion leaves more connected components than there were before. The two definitions disagree on whether or not a loop is its own connected component. The lemma is true in either case, but we shall prove it with the first definition.)

Proof. The forward direction is trivial, since Σ's being balanced means *every* circle in Σ is balanced, so any circles in a particular block are certainly balanced. Also, there can be no half edges.

For the reverse direction, assume every block of Σ is balanced. That rules out half edges. Let C be a circle in Σ. Each circle is contained within a single block, so C is balanced since it is a circle in a balanced subgraph. Therefore Σ is balanced. □

Finally, here is Harary's second theorem about balance.

Theorem A.4 (Path Balance (Harary 1953a))**.** *[[LABEL T:1006pathbalance]] A signed link graph is balanced if and only if every path with the same endpoints has the same sign.*

This proof, also shorter than the original proof, will also be postponed to Section A.6. The astute reader will have noticed the similarity to conservative vector fields, σ corresponding to the vector field and ζ to a potential function (on which, more later).

A.6. **Switching.** [[LABEL 2.switching]]

I will now introduce one of the most useful and powerful techniques in signed graph theory.

Definition A.12. A function $\zeta : V \to \{+, -\}$ is called a *switching function*, or sometimes a *selector*. The *switched signature* is $\sigma^\zeta(e) := \zeta(v)\sigma(e)\zeta(w)$, where $e{:}vw$, and the *switched signed graph* is $\Sigma^\zeta := (\Gamma, \sigma^\zeta)$.

Looking at examples we notice that switching a single vertex doesn't change the sign (or equivalently balance) of any circle. We formalize this with a proposition, in preparation for the Switching Theorem.

Proposition A.5. *[[LABEL P:1006 Switching Circles]] Switching leaves the signs of all circles unchanged.*

Proof. Let ζ be a switching function and $C = v_0e_0v_1e_1v_2 \cdots v_{n-1}e_{n-1}v_0$ be a circle. (So e_i has endpoints e_i and e_{i+1} with the indices understood modulo n.) Now

$$\sigma^\zeta(C) = \big(\zeta(v_0)\sigma(e_0)\zeta(v_1)\big)\big(\zeta(v_1)\sigma(e_1)\zeta(v_2)\big)\dots\big(\zeta(v_{n-1})\sigma(e_{n-1})\zeta(v_0)\big).$$

Since for each $v_i \in V(C)$, $\zeta(v_i)$ appears twice in the product above, and $\zeta(v_i)\cdot\zeta(v_i) = +$, the product above reduces to $\sigma^\zeta(C) = \sigma(e_0)\sigma(e_1)\cdots\sigma(e_{n-1}) = \sigma(C)$. □

In particular, switching never changes the sign of a loop.

For circles, the terms 'balanced' and 'positive' are equivalent, as are the terms 'unbalanced' and 'negative', although this certainly isn't the case for arbitrary edge sets.

An alternative (and equivalent) point of view on switching is that switching Σ by ζ means negating every edge with one endpoint in $\zeta^{-1}(+)$ and the other in $\zeta^{-1}(-)$. (This is immediate from the definition.) We call this *switching the vertex set* $\zeta^{-1}(-)$, or equivalently $\zeta^{-1}(+)$.

Definition A.13. [[LABEL D:1006 vertex set switching]] For $X \subseteq V$, the signed graph Σ^X is the result of negating every edge with one endpoint in X and the other not in X; that is, every edge of the cut (X, X^c). We call this operation *switching* X and we say Σ^X is Σ *switched by* X.

Vertex switching means switching a single vertex v, i.e., switching $\{v\}$. We write Σ^v for Σ switched by v.

Note that set switching is simply a change in perspective, from the switching function $\zeta : V(\Sigma) \to \{+, -\}$ to the vertex set $X = \zeta^{-1}(-)$, or conversely from X to ζ_X which is $-$ on X and $+$ on all other vertices. We will use whichever notation is more convenient.

Notice also that switching by X is equivalent to switching by X^c, and similarly $\Sigma^\zeta = \Sigma^{-\zeta}$ for any switching function ζ. Any switching is the product of vertex switchings (in any order). Specifically, $\Sigma^X = \big(\cdots\big(\big(\Sigma^{v_1}\big)^{v_2}\big)\cdots\big)^{v_n}$ where $X = \{v_1, v_2, \dots, v_n\}$

Switching and balance.

Theorem A.6 (Switching Theorem)**.** *[[LABEL T:1006 Switching]]*
(1) *Switching leaves $\mathcal{B}$ unchanged, i.e., $\mathcal{B}(\Sigma^\zeta) = \mathcal{B}(\Sigma)$.*
(2) *If $|\Sigma_1| = |\Sigma_2|$ and $\mathcal{B}(\Sigma_1) = \mathcal{B}(\Sigma_2)$, then there exists a switching function ζ such that $\Sigma_2 = \Sigma_1^\zeta$.*

Proof. We notice that (1) follows immediately from Proposition A.5, since switching doesn't create or destroy any circles, and it doesn't change the sign of any circles.

For part (2), notice that Σ_1 and Σ_2 have the same vertices and edges (since $|\Sigma_1| = |\Sigma_2|$); we will write $\Gamma := |\Sigma_1| = |\Sigma_2|$. Since switchings of different components are independent, we may assume Σ_1 is connected. Now pick a spanning tree T in the underlying graph, and list the vertices in such a way that v_i is always adjacent to a vertex in $\{v_0, \ldots, v_{i-1}\}$ (this is a fairly standard exercise in basic graph theory, and the list is not usually unique). Let t_i denote the unique tree edge connecting v_i to $\{v_0, \ldots, v_{i-1}\}$.

We take a brief pause to recall that every circle in Γ is the set sum (symmetric difference) of the fundamental circles of the non-tree edges of C. Precisely, $C = \bigoplus_{e\in C\setminus T} C_T(e)$. (This is Proposition **??**.)

We now define (recursively) a series of switching functions, ζ_i for $0 \leq i < n$, where $\zeta_0 \equiv +$ and

$$\zeta_i(v_j) = \begin{cases} \zeta_{i-1}(v_j) & \text{if } j < i, \\ \sigma_1^{\zeta_{i-1}}(t_i) \cdot \sigma_2(t_i) & \text{if } j = i, \\ + & \text{if } j > i. \end{cases}$$

(Here σ_k is the signature of Σ_k and $\sigma_1^{\zeta_{i-1}}$ denotes the signature of $\Sigma_1^{\zeta_{i-1}}$.) Notice that for each of the edges in T, $t_1, \ldots, t_{n-1}$, $\sigma_2(t_k) = \sigma_1^{\zeta_i}(t_k)$ for $i \geq k$, so in particular, $\sigma_2(t_k) = \sigma_1^{\zeta_{n-1}}(t_k)$ for all t_k tree edges.

We now consider a non-tree edge $f \in C \setminus T$. Since $\mathcal{B}(\Sigma_1) = \mathcal{B}(\Sigma_2)$, we conclude that $\sigma_1(C_T(f)) = \sigma_2(C_T(f))$, and by Proposition A.5, $\sigma_1^{\zeta_{n-1}}(C_T(f)) = \sigma_2(C_T(f))$, since ζ_{n-1} is a switching function. Finally, we notice that by construction $\sigma_1^{\zeta_{n-1}}$ and σ_2 agree on each edge in $C_T(f)$ except f, and on the product (the entire fundamental circle), they must agree on f. Therefore, $\sigma_1^{\zeta_{n-1}}$ and σ_2 agree on every edge in Γ. Hence, ζ_{n-1} is the desired switching function. $\square$

This theorem can be regarded as the natural generalization of the standard characterization of bipartite graphs.

Corollary A.7. *[[LABEL C:1006 bipartite]] An ordinary graph Γ is bipartite $\iff$ it has no odd circles.*

Proof. All circles in Γ are even $\iff$ all circles in $-\Gamma$ are positive $\iff$ (by definition of balance) $-\Gamma$ is balanced $\iff$ (by Theorem A.6) $V = X_1 \mathbin{\dot\cup} X_2$ so that all negative edges (that is, all edges) have one endpoint in X_1 and the other in X_2 $\iff$ Γ is bipartite. $\square$

(Next time we will work on a proof that when T is a maximal forest, Σ can be switched to be any desired value on the edges of T, which is accidently included in the proof of Thm A.6.)

One thing to observe is that for a walk W from v to w in a signed graph Σ we have $\sigma^\zeta(W) = \zeta(v)\sigma(W)\zeta(w)$. In particular, the sign of a closed walk is fixed under switching. Oct 8: Simon Joyce

Lemma A.8. *[*[LABEL L:1008 tree signs]] *Given a signed graph Σ and a maximal forest T of Σ, there exists a switched graph Σ^ζ such that Σ^ζ has any desired signs on T. Furthermore, ζ is unique up to negation on each component.*

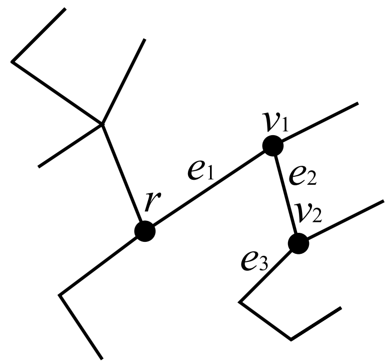

FIGURE A.1. F:1008 We want $\sigma^\zeta(e_i) = \tau(e_i)$.

Proof. We can treat each component of Σ separately so we'll assume Σ is connected. Then T is a spanning tree. Let $\tau : E(T) \to \{+, -\}$ be the desired edge sign function. Pick a root vertex r in $V(\Sigma)$. Then

$$\tau(e_1) = \sigma^\zeta(e_1) = \zeta(v_1)\sigma(e_1)\zeta(r),$$

so

$$\zeta(v_1) = \tau(e_1)\sigma(e_1)^{-1}\zeta(r)^{-1} = \tau(e_1)\sigma(e_1)\zeta(r).$$

For $v \in V(\Sigma)$, let P_{rv} be the unique path in T between r and v. Thus, $P_{rv} = re_1v_1e_2v_2 \dots e_lv$. Then $\sigma(P_{rv}) = \sigma(e_1)\sigma(e_2)\cdots\sigma(e_l)$. We want to show $\sigma^\zeta(e_i) = \tau(e_i)$. We know $\sigma^\zeta(e_i) = \zeta(v_{i-1})\sigma(e_i)\zeta(v_i)$, so we have

$$\begin{aligned}\sigma^\zeta(P_{rv}) &= [\zeta(r)\sigma(e_1)\zeta(v_1)][\zeta(v_1)\sigma(e_1)\zeta(v_2)]\cdots[\zeta(v_{l-1}\sigma(e_l)\zeta(v)]\\ &= \zeta(r)\sigma(P_{rv})\zeta(v).\end{aligned}$$

Therefore we must have $\zeta(r)\sigma(P_{rv})\zeta(v) = \tau(P_{rv})$, so $\zeta(v) = \tau(P_{rv})\sigma(P_{rv})\zeta(r)$. Choosing $\zeta(r)$ to be $+$ or $-$, the rest of ζ is completely determined. Switching by ζ,

$$\begin{aligned}\sigma^\zeta(e_i) &= \zeta(v_{i-1})\sigma(e_i)\zeta(v_i)\\ &= \tau(P_{rv_{i-1}})\sigma(P_{rv_{i-1}})\zeta(r)\sigma(e_i)\tau(P_{rv_i})\sigma(P_{rv_i})\zeta(r)\\ &= \sigma(e_i)\tau(e_i)\sigma(e_i)\\ &= \tau(e_i).\end{aligned}$$

□

The following immediate corollary is a very useful result.

Proposition A.9. *[*[LABEL C:1008 balanced positive]] *If Σ is a balanced signed graph, then there is a switching function ζ such that all ordinary edges of Σ^ζ are positive.*

Proof. Since Σ is balanced it has no half edges. Let T be a maximal forest of Σ. By the previous result there is a switching function ζ such that all the edges of T are positive. Consider an edge e not in T. Either e is a loose edge, it is a balanced loop, or it is a link. If e is a loose edge then it has no sign. If it is a balanced loop it is positive before and also after switching. If e is a link, its sign in Σ^ζ equals the sign in Σ of its fundamental circle, which is

+. Therefore, $\sigma^{\zeta}(e) = +$; consequently, switching by ζ does in fact make all ordinary edges positive. □

In particular, this result tells us that for any balanced component $\Sigma{:}B$ of a signed graph, there exists a switching function such that all the edges of $\Sigma{:}B$ are positive. More broadly,

Corollary A.10. [[LABEL C:1008 balanced positive subgraph]] *If S is a balanced edge set in Σ, then there is a switching function such that all ordinary edges of S are positive.* □

Switching equivalence and switching isomorphism.

Now we examine the relationships between signed graphs that are induced by switching.

Definition A.14. [[LABEL D:1008 switch class]] We say two signed graphs Σ_1 and Σ_2 are *switching equivalent* if $|\Sigma_1| = |\Sigma_2|$ and there is a switching function ζ such that $\Sigma_1^{\zeta} = \Sigma_2$. Switching equivalence is an equivalence relation; we call an equivalence class a *switching class* of signed graphs.

A related concept is that of *switching isomorphism*, which means that Σ_1 is isomorphic to some switching of Σ_2. We call an equivalence class a *switching isomorphism class.* (In the literature, the terms "switching equivalence" and "switching class" often refer to switching isomorphism; one has to pay close attention.)

[THE FOLLOWING REPEATS A THEOREM FROM THE PREVIOUS DAY:]

Theorem A.11. [[LABEL T:1008 switch equiv]] *Given two signed graphs Σ_1 and Σ_2, if $|\Sigma_1| = |\Sigma_2|$ and $\mathcal{B}(\Sigma_1) = \mathcal{B}(\Sigma_2)$ then Σ_1 and Σ_2 are switching equivalent.*

Proof. Let T be a maximum forest of $|\Sigma_1|$. First we take ζ_1, a switching of Σ_1 and ζ_2, a switching of Σ_2 such that $\Sigma_1^{\zeta_1}$ and $\Sigma_2^{\zeta_2}$ are all positive on T. We want to show $\Sigma_1^{\zeta_1} = \Sigma_2^{\zeta_2}$.

Take an edge $e{:}vw$ in the graph. If $e \in T$, then $\sigma_1^{\zeta_1}(e) = \sigma_2^{\zeta_2}(e) = +$. If $e \notin T$, then there is a unique path T_{vw} joining v and w in T. Let $C = T_{vw} \cup e$. So C is a circle. Since $\mathcal{B}(\Sigma_1) = \mathcal{B}(\Sigma_2)$, we know $\sigma_1(C) = \sigma_2(C)$. Furthermore, since the sign of a closed walk is fixed under switching we have,

$$\sigma_1^{\zeta_1}(C) = \sigma_1(C) \text{ and } \sigma_2^{\zeta_2}(C) = \sigma_2(C).$$

Therefore,

$$\sigma_1^{\zeta_1}(C) = \sigma_2^{\zeta_2}(C).$$

In particular we have,

$$\sigma_1^{\zeta_1}(C) = \sigma_1^{\zeta_1}(T_{vw})\sigma_1^{\zeta_1}(e) = \sigma_1^{\zeta_1}(e) \text{ and } \sigma_2^{\zeta_2}(C) = \sigma_2^{\zeta_2}(T_{vw})\sigma_2^{\zeta_2}(e) = \sigma_2^{\zeta_2}(e)$$

Therefore e has the same sign in both $\Sigma_1^{\zeta_1}$ and $\Sigma_2^{\zeta_2}$, ie $\Sigma_1^{\zeta_1} = \Sigma_2^{\zeta_2}$. □

This theorem means two signatures of a graph Γ are switching equivalent if and only if they have the same circle signs. With this result we can efficiently prove Harary's original theorem. **[THIS PROOF BELONGS IN THE PREVIOUS DAY'S NOTES.]**

Proof of Harary's Balance Theorem A.2. Suppose Σ has the stated form. Then Σ is obviously balanced; but we also note that Σ^{V_2} is all positive, hence balanced, hence Σ is balanced by Proposition A.5.

Conversely, suppose Σ is balanced. Switching by a suitable vertex set X so a maximal forest F is all positive (which is possible by Theorem A.11), every other edge must be positive because its fundamental circle $C(e)$ is positive and all edges in $C(e)$ other than e are in F.

Calling this all-positive graph Σ_1, $\Sigma = \Sigma_1^X$ has every edge within X or X^c positive and every edge between X and X^c negative. □

Proof of Harary's Path Balance Theorem A.4. Suppose first that Σ is balanced. Switch by ζ so it is all positive. Every path is now positive! Since the sign of a uv-path P in Σ^ζ is $\zeta(u)\sigma(P)\zeta(v) = +$, in Σ the sign of P is $\zeta(u)\zeta(v)$, independent of the particular path.

Conversely suppose that for each pair $u, v \in V$ the sign of every path P from u to v is the same, say $\bar\sigma(u, v)$. Choose a vertex $r \in V$ and define $\zeta(v) := \bar\sigma(r, v)$. Then in Σ^ζ every path is positive; in particular, every link is positive. Since Σ switches to all positive, it is balanced. □

Oct. 8 supplement: T.Z.

Exercise A.2. [[LABEL X:1008supp swKn]]

(a) Find all the switching isomorphism classes for signed K_n's, $n = 3, 4, 5$. Find their frustration indices and all their minimal members.
(b) What is the number of negative triangles in each switching isomorphism class? Make a histogram of these numbers by switching isomorphism class. Do you notice any patterns? Are there any generalizations?
(c) How many switching equivalence classes are there in each switching isomorphism class? What is the significance of that number?

Oct 13: Nate Reff

B. CHARACTERIZING SIGNED GRAPHS

[[LABEL 2.characterization]][[LABEL 2.basic]]

The next question is: Which circle sign patterns are possible for a signed graph? We give two kinds of answer: one algebraic and one combinatorial.

B.1. **Signature as a homomorphism.** [[LABEL 2.shomomorphism]]

A function $f : V_1 \to V_2$, where V_1 and V_2 are binary vector spaces, is a homomorphism if and only if it is additive (we can ignore the scalar multiplication axioms because for $\mathbb{F}_2$ they are satisfied automatically). So any function $\sigma : E \to \mathbb{F}_2$ gives a unique extension $\sigma : \mathcal{P}(E) \to \mathbb{F}_2$ that is a vector space homomorphism by the identification $\sigma(S) = \sum_{e\in S} \sigma(e)$.

Oct 15 (draft): Peter Cohen and T.Z.

In this discussion we take signs as elements of the field $\mathbb{F}_2 = \{0, 1\}$ and we write Z_1 as short notation for the binary cycle space $Z_1(\Gamma; \mathbb{F}_2)$.

Theorem B.1 (Signature as a linear functional)**.** [[LABEL T:1015signhomomorphism]] *Given any function $\bar\sigma : \mathcal{C} \to \mathbb{F}_2$, the following properties are equivalent:*

(1) $\bar\sigma = \sigma|_{\mathcal{C}}$ *for some signature $\sigma : E \to \mathbb{F}_2$ (extended to Z_1 by linearity).*
(2) $\bar\sigma$ *is the restriction to $\mathcal{C}$ of a homomorphism $\tau : Z_1 \to \mathbb{F}_2$.*
(3) $\bar\sigma^{-1}(0) = \mathcal{C} \cap U$ *for some subspace U of Z_1, with codimension 0 or 1.*

Proof. We prove a chain of implications:

1-2) If we let $\tau = \sigma : P(E) \to \mathbb{F}_2$, where $P(E)$ is essentially a subspace with the form $\mathbb{F}_2^E$, then σ is restricted to a homomorphism, so 1) implies 2).

2-3) Given τ, we can set $\bar\sigma = \tau|_{\mathcal{C}}$. In this case, $U = \operatorname{Ker}\tau$, so $\bar\sigma^{-1}(0) = C \cap \operatorname{Ker}\tau$. Since $\operatorname{Ker}\tau$ is a subspace of Z_1, 2) implies 3).

3-2) $\tau : Z_1 \to \mathbb{F}_2$ can be defined by $\tau^{-1}(0) = U$, or we could define $\tau : Z_1 \to Z_1/U \in \mathbb{F}_2$, in which case $U = \operatorname{Ker}\tau$, so therefore $\bar\sigma = \tau|_{\mathcal{C}}$, and 3) implies 2).

2-1) Given that $\bar\sigma = \tau|_{\mathcal{C}}$, we define $\sigma' : P(E) \to \mathbb{F}_2$ to be any function with the condition $\sigma'(C) = \bar\sigma(C)$. Given this σ', we obtain the desired $\sigma : E \to \mathbb{F}_2$, since σ' is defined on $P(E)$. □

[PLEASE WRITE UP the presentation I gave on Tuesday 21 about the different levels of sign functions: σ, σ', $\bar\sigma = \bar\sigma$; and why σ' (defined on Z_1) is determined by $\bar\sigma$ (defined on $\mathcal{C}$).] ? Oct 2014: Richard Behr

A signature $\sigma : E \to \mathbb{F}_2$ extends linearly to a functional $\sigma_1 : \mathbb{F}_2E \to \mathbb{F}_2$, since E is a basis for $\mathbb{F}_2E$. Conversely, a functional $\sigma_1 : \mathbb{F}_2E \to \mathbb{F}_2$ determines its values on E, so we get a signature $\sigma : E \to \mathbb{F}_2$. As the correspondence between σ and σ_1 is bijective, we'll now assume any function on E is automatically extended to $\mathbb{F}_2$ and we'll use the same name for both functions.

Given σ, we obtain a functional $\sigma' : Z_1 \to \mathbb{F}_2$ by restriction, since $Z_1 \subseteq \mathbb{F}_2E$ (recall that Z_1 is the kernel of the boundary mapping $\partial : \mathbb{F}_2E \to \mathbb{F}_2V$). However, given a mapping $\sigma' : Z_1 \to \mathbb{F}_2$, we can extend it back to a function $\sigma : \mathbb{F}_2E \to \mathbb{F}_2$, but this extension will not be unique in general.

We can define a third type of signature, a function $\bar\sigma : \mathcal{C} \to \mathbb{F}_2$. This function determines a function $\sigma' : Z_1 \to \mathbb{F}_2$, since $Z_1 = \langle \mathcal{C}(\Gamma) \rangle$ (by Theorem **??**). We will see a function of this form in the characterization of positive circles below.

Oct 15 (draft): Peter Cohen and T.Z.

B.2. **Balanced circles and theta oddity.** [[LABEL 2.oddity]]

Now we give a combinatorial condition characterizing the class of balanced circles of a signed graph. For a subclass $\mathcal{B}$ of all circles of a graph, *theta additivity* or *theta oddity* (called "circle additivity" in Zaslavsky (1981b)), is the property that every theta subgraph contains 1 or 3 members of $\mathcal{B}$.

Theorem B.2 (Characterization of Positive Circles)**.** *[[LABEL T:1015poscircles]] Let $\mathcal{B}$ be any subclass of $\mathcal{C}(\Gamma)$. Then $\mathcal{B}$ is the class of positive circles of some signature of Γ if and only if it has an odd number of circles in every theta subgraph.*

We need a lemma about expressing circles as theta sums, that will let us use induction in proving the theorem. A *theta sum* is a representation of a circle C as the set sum $C_1 \oplus C_2$ of two other circles such that $C_1 \cup C_2$ is a theta graph whose third circle is C. Given T, a maximal forest in Γ, define $\nu(C) := |E(C) \setminus E(T)|$.

Lemma B.3. *[[LABEL L:1015thetasum]] Each circle is either a fundamental circle with respect to T, or a theta sum $C_1 \oplus_\theta C_2$ of two circles with smaller values of ν.*

2014 Oct 22: Richard Behr

Proof. The smallest possible value of $\nu(C)$ is 1; in this case C is a fundamental circle by definition. Now suppose that $\nu(C) \geq 2$, and choose two of the edges of C that are not in T, and call them f and g. Thus, $C \setminus f \setminus g$ is disconnected and consists of two disjoint paths P_1 and P_2 (if f and g are adjacent, one of the paths will be a single vertex). Since all the vertices of P_1 and P_2 lie in T (it is spanning), there must be some path P_3 contained in T that connects a vertex of P_1 to a vertex of P_2. Thus, P_3 is a chord of C, and $C \cup P_3$ is a theta graph $C_1 \oplus_\theta C_2$, where C_1 contains f and C_2 contains g. Therefore, $\nu(C_1) < \nu(C)$ and $\nu(C_2) < \nu(C)$. $\square$

Oct 15 (draft): Peter Cohen and T.Z.

Proof of the theorem. A theta graph is made up of three internally disjoint paths, all with the same endpoints, which we will call P_1, P_2, P_3. We denote by C_{ij} the circle made by the two paths P_i and P_j.

First suppose we have a signature σ. The signs of the circles can be found by multiplying the signs of the paths:

$$\begin{aligned}\sigma(C_{12}) &= \sigma(P_1)\sigma(P_2),\\ \sigma(C_{23}) &= \sigma(P_2)\sigma(P_3),\\ \sigma(C_{13}) &= \sigma(P_1)\sigma(P_3).\end{aligned}$$

Therefore,

$$\begin{aligned}\sigma(C_{12})\sigma(C_{23})\sigma(C_{13}) &= \sigma(P_1)\sigma(P_2)\sigma(P_2)\sigma(P_3)\sigma(P_1)\sigma(P_3)\\ &= \sigma(P_1)\sigma(P_1)\sigma(P_2)\sigma(P_2)\sigma(P_3)\sigma(P_3)\\ &= +.\end{aligned}$$

The number of negative circles is even, so theta oddity is satisfied.

Now suppose a class $\mathcal{B}$ is given that satisfies theta oddity. Let $\bar\sigma : \mathcal{C} \to \mathbb{F}_2$ be the characteristic function of $\mathcal{B}^c$, that is, $\bar\sigma(C)$ equals 1 if C is not in $\mathcal{B}$, 0 if it is in $\mathcal{B}$.
[THIS IS $1_{\mathcal{B}^c}$ (in $\mathcal{C}$). CHANGE NOTATION?]
(In this part of the proof it is best to regard signs as values in $\mathbb{F}_2$.) Theta oddity means that if $C_1 \cup C_2$ is a theta graph and the third circle in it is $C = C_1 \oplus C_2$, then $\bar\sigma(C) = \bar\sigma(C_1) + \bar\sigma(C_2)$; i.e., theta oddity is literally additivity.

Choose a maximal forest T. We use the fundamental circles to define σ, namely,

$$\sigma(e) := \begin{cases} 0 & \text{if } e \in E(T),\\ \bar\sigma(C_T(e)) & \text{if } e \notin E(T).\end{cases}$$

Thus, for a non-forest edge e, $\sigma(e) = 0$ if $C_T(e) \in \mathcal{B}$ and 1 otherwise. Our task is to prove that $\mathcal{B}(\sigma) = \mathcal{B}$, which means that $\sigma(C) = \bar\sigma(C)$ for every circle. We employ induction on the number of non-forest edges in C.

Case 1: C is a fundamental circle. Since $C = C_T(e)$, by reversing the definition we find that

$$\sigma(C) := \sum_{f \in C} \sigma(f) = \sigma(e) = \bar\sigma(C).$$

Case 2: C is not a fundamental circle. Then $\nu(C) \geq 2$, so by Lemma B.3, $C = C_1 \oplus C_2$, a theta sum in which $\nu(C_1),\ \nu(C_2) < \nu(C)$. By theta additivity and induction on ν,

$$\begin{aligned}\bar\sigma(C) = \bar\sigma(C_1) + \bar\sigma(C_2) &= \sum_{f \in C_1} \sigma(f) + \sum_{f \in C_2} \sigma(f)\\ &= \sum_{f \in C_1 \oplus C_2} \sigma(f) = \sum_{f \in C} \sigma(f) = \sigma(C).\end{aligned}$$

This establishes that $\mathcal{B}(\sigma) = \mathcal{B}$, as we wished. $\square$

In the preceding proof, we obtained an extension of $\bar{\sigma}$ to $\mathbb{F}_2E$; however, this extension (usually) is not unique. Recall that $Z_1 = \langle \delta 1_v : v \in V\rangle^\perp = (B^1)^\perp = \langle E(v, V\setminus v) : v \in V\rangle^\perp$, where δ is the coboundary mapping. (The first equals sign is more an identification than an actual equality.) 2014 Oct 22: Richard Behr

Furthermore, $Z_1 \cap B^1 = \{0\}$. [**WRONG! Counterex.: C_4 where E is both a cycle and a cut.**]

It follows that $\mathbb{F}_2E = Z_1 \oplus B^1$. Thus, the extensions of the function σ' defined on Z_1 are determined by $B^1 = \langle 1_{E(v,V\setminus v)}\rangle$.

We can combinatorialize this algebraic presentation by using switching. Choose $\zeta : V \to \{+,-\}$ and switch to σ^ζ; equivalently, choose $X = \zeta^{-1}(-) \subseteq V$, and switch X; that is, negate the edges of $E(X, X^c)$.[3] This preserves the set of positive circles (Theorem A.6). Observe that switching is essentially the same thing as negating a cut; that is, switching σ by X is the same as taking $\sigma + 1_{E(X,X^c)}$. But every nonzero element of the binary cut space B^1 is a cut (Section **??**), so taking $\sigma + 1_{E(X,X^c)}$ is the same as taking $\sigma + b$ for some $b \in B^1$. Thus, any extension of σ' to $\mathbb{F}_2E$ is a switching of the chosen extension σ, and vice versa. So, if we have σ' defined on Z_1, not only does it extend to σ on $\mathbb{F}_2E$, but also all extensions constitute the coset $\sigma + B^1$ of B^1 in $\mathbb{F}_2E$. [**NEED THE FOLLOWING?**] In other words, if $D = \vec{E}(X, X^c)$, then $1_D = D = 1_{E(X,X^c)}$, and $\sigma + 1_D = \sigma^X$.

Oct 13: Nate Reff

C. Connection

[[LABEL 2.connection]]

C.1. Balanced components. [[LABEL 2.balcomp]]

Suppose we have a signed graph $\Sigma = (V, E, \sigma)$ with some subset $S \subseteq E$. Recall that a path in Σ, $P = e_1e_2\dots e_k$ (not containing any half edges), has a sign $\sigma(P) = \sigma(e_1)\sigma(e_2)\dots\sigma(e_k)$. A circle whose sign is $+$ is said to be *positive* or *balanced*. We say that S is *balanced* if it contains no half edges and every circle is balanced. Recall that we denote by $c(S) = c(V, S) = c(\Sigma\,|\,S)$ the total number of components (that is, node components). We will denote by $b(S) = b(V, S) = b(\Sigma|S)$ the number of balanced components. Recall that

$$\boldsymbol{\pi}(S) = \{\text{vertex sets of components of } S\}.$$

[**THIS IS DUPLICATIVE. Partitions and partial partitions are treated elsewhere.**]

We also write

$$\begin{aligned}\boldsymbol{\pi}_\mathrm{b}(S) &= \{\text{vertex sets of balanced components of } S\}\\ &= \{X \in \boldsymbol{\pi}(S) \mid S{:}X \text{ is balanced}\}.\end{aligned}$$

This may be called the *balanced partial partition of V induced by S*. Then $c(S) = |\boldsymbol{\pi}(S)|$ and $b(S) = |\boldsymbol{\pi}_\mathrm{b}(S)|$.

Let's take a moment to review partitions of a set. A *partition* of V is a class $\{B_1, B_2, \dots, B_k\}$ of disjoint, nonempty sets B_i such that $B_1 \cup B_2 \cup \dots \cup B_k = V$. A *partial partition* of V is a partition of a subset of V; its *support* $\mathrm{supp}(\pi) := \bigcup \pi$ is that subset. (One should not

---

[3] I say the same thing three ways to emphasise the equivalence of the ways. Each has its uses, as we will see.

overlook the unique partition of the empty set: it is the empty partition, $\varnothing$, and it is a partial partition of V.) We denote the class of partitions and partial partitions by Π_V and $\Pi_V^\dagger$, respectively. So as an immediate observation we have $\boldsymbol{\pi}(S) \in \Pi_V$ and $\boldsymbol{\pi}_{\mathrm{b}}(S) \in \Pi_V^\dagger$. Also one should note that $\Pi_n^\dagger \cong \Pi_{n+1}$. This is because a partial partition π is in bijective correspondence with the partition $\pi \cup \{\{0,1,\dots,n\} \setminus \operatorname{supp}(\pi)\}$ of $\{0,1,\dots,n\}$. (The block $\{0,1,\dots,n\} \setminus \operatorname{supp}(\pi)$ is called the "zero block" of π, by those who like to have it. This isomorphism does not give us a new kind of lattice, but instead a new structure to be studied.)

Now we turn our attention to the natural isomorphism $\mathcal{P}(E) \cong \mathbb{F}_2^E$. The latter is a binary vector space (a structure that is equivalent to an abelian group of index 2). We will denote by $\oplus$ the binary vector addition operator. We denote by $\mathcal{C} = \mathcal{C}(\Gamma)$ the class of circles in Γ. Suppose we have three circles $C, C_1, C_2 \in \mathcal{C}$, we say C is the *theta sum* of C_1 and C_2 if $C = C_1 \oplus C_2$ and $C_1 \cup C_2$ is a theta graph.

We know from [**someplace in CHAPTER I**] that, given a maximal forest T of an unsigned ordinary graph, the fundamental system of circles with respect to T, $\{C_T(e) \mid e \notin T\}$, is a basis for the cycle space $Z_1(\Gamma;\mathbb{F}_2)$ and that $C = \bigoplus_{e \in C \setminus E(T)} C_T(e)$. In fact we can rearrange the sum to in terms of sums in theta graphs.

Lemma C.1. *[[LABEL L:1013lemma1]] C can be obtained from $\{C_T(e) \mid e \in E(C) \setminus E(T)\}$ by theta sums.*

Proof. For convenience in the proof we define $Q_C := E(C) \setminus E(T)$. Now we do induction on $|Q_C|$. For the base case, if $|Q_C| = 1$, then $C = C_T(e)$ where $\{e\} = Q_C$. For the induction step, where $|Q_C| > 1$, we give two proofs by two different methods.

First Proof (by a direct argument).

Since $|Q_C| > 1$, $C \setminus Q_C$ is a disconnected graph. This means that T contains a path connecting two vertices in different components of $C \setminus Q_C$. Now suppose P is a minimal such path. Then P is internally disjoint from C, by minimality. Therefore, P is chordal path of C, so $P \cup C$ is a theta graph, and $C = C_1 \oplus C_2$ where C_1 and C_2 are the circles in $P \cup C$ that contain P. Hence, the P we wanted exists.

Second Proof (to illustrate the use of bridges).

We split C into two circles C_1 and C_2 such that $C =$ theta sum of C_1 and C_2 (We will prove that this is possible after the induction is completed.) and $Q_C = Q_{C_1} \mathbin{\dot\cup} Q_{C_2}$. Since $|Q_{C_1}|, |Q_{C_1}| < |Q_C|$, by the induction hypothesis $\{C_T(e) \mid e \in Q_{C_1}\}$ generates C_1 by theta sums, and $\{C_T(e) \mid e \in Q_{C_2}\}$ generates C_2 by theta sums. Therefore the disjoint union $Q_{C_1} \mathbin{\dot\cup} Q_{C_2} = Q_C$ generates the entirety of C by theta sums. This completes the induction argument, so now we turn back to prove the existence of the theta sum.

Suppose that C is drawn as in figure C.1.

We say P is a *chordal path* of C if P is a path which connects two vertices in C but is internally disjoint from C. Equivalently, $C \cup P$ is a theta graph.

In the context of this proof we want to find such a $P \subseteq T$. Notice that all the nontree edges of C_1 are in C_1, all nontree edges of C_2 are in C_2, and all nontree edges of C is the disjoint union on nontree edges in C_1 and nontree edges of C_2, so $Q_C = Q_{C_1} \mathbin{\dot\cup} Q_{C_2}$. In figure C.1 P is a *bridge* of C, as is the red subgraph seen in figure C.1. Every vertex of C is in T, so every edge of C that is not an edge of T is a bridge of T. So $T \setminus E(C)$ splits into bridges of C and isolated vertices that are not bridges. There is at least one bridge that contains

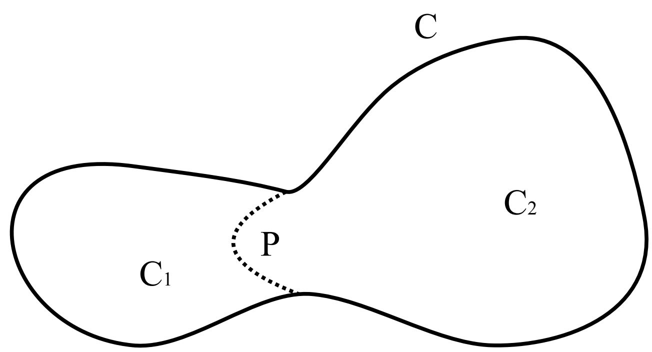

FIGURE C.1. F:1013Figure1

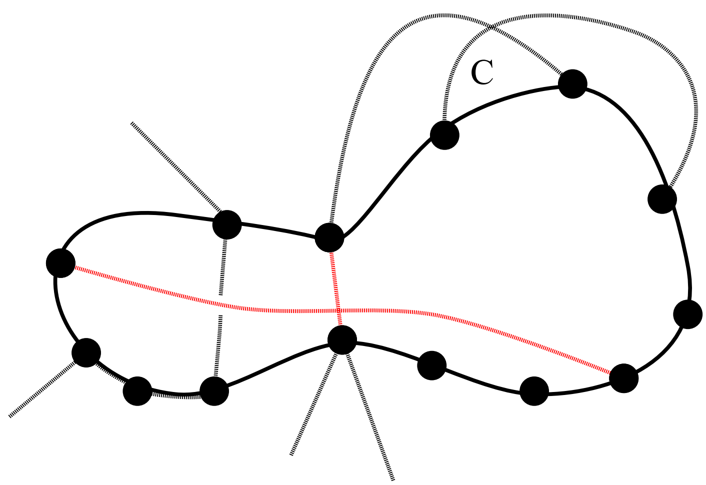

FIGURE C.2. F:1013Figure2

vertices of two components of $T \cap C = C \setminus Q_C$, which is disconnected since $|Q_C| > 1$. This completes the bridge proof. □

Oct 20: Yash Lodha

C.2. **Unbalanced blocks.** [[LABEL 2.blocks]]

C.2.1. *Menger's theorem.*

We take a moment to call to mind Menger's theorem. A *block* of Γ is a maximal inseparable subgraph, which means that every pair of edges in the subgraph is in a common circle of the subgraph. A block graph is a graph that is a block of itself, or in other words, inseparable.

Theorem C.2 (Menger's theorem)**.** *[[LABEL T:1020menger]] In a 2-connected graph Γ, given any two vertex sets X, Y (not necessarily disjoint) such that $|X|, |Y| \geq 2$, there exist two disjoint XY-paths.*

Corollary C.3 (Usual Menger's theorem). *[[LABEL C:1020mengervv]] For a 2-connected graph Γ and any two non-adjacent vertices x and y, there exist two internally disjoint xy-paths.*

Corollary C.4 (Another form of Menger's theorem). *[[LABEL C:1020mengersv]] For a 2-connected graph Γ, any set X of at least two vertices, and any vertex z, there exist two internally disjoint Xy-paths whose endpoints in X are distinct.*

We use Menger's theorem (in whichever form) mainly for $k = 2$. The following Harary-type vertex and edge theorems show the method.

C.2.2. *Vertices and edges in unbalanced blocks.*

In an unbalanced block there are no vertices or edges that don't participate in the imbalance. This is implied by Harary's second theorem and its edge version.

Theorem C.5 (Vertex Theorem (Harary 1955a)). *[[LABEL T:1020unbalblockvertex]] Let Σ be an unbalanced signed block with more than one edge. Then every vertex belongs to a negative circle.*

Proof. Let D be a negative circle. If v is in D we're done. Otherwise, by Menger's theorem there are two paths from v to D, disjoint except that both start at v. Call them $P_1{:}vw_1$ and $P_2{:}vw_2$, and let P be the combined path from w_1 to w_2. Also, let Q and R be the two paths into which w_1 and w_2 divide D. Then $D \cup P$ is a theta graph. As D is negative, one of the two circles $P \cup Q$ and $P \cup R$ must be negative. □

This is not Harary's proof. As is commonly true, the original proof was much longer.

A stronger result is the edge version. I don't know why Harary didn't think of it, but probably because his attention was focussed on the vertices, which represented the persons in a social group to which the theory of signed graphs was intended to apply. The edges themselves were not interesting in that context.

Theorem C.6 (Edge Theorem). *[[LABEL T:1020unbalblockedge]] In an unbalanced block with more than one edge, every edge is in a negative circle.*

There is a short proof of the Edge Theorem, similar to that of the Vertex Theorem but slightly harder due to having two nontrivial cases. The proof is a good homework problem.

Oct 22: Jackie Kaminski

D. IMBALANCE AND ITS MEASUREMENT

[[LABEL 2.imbalance]]

We know what it means for a graph to be balanced, as well as unbalanced, but we could certainly have more information about unbalanced graphs. For example some unbalanced graphs might be 'almost' balanced, with a single unbalanced circle, and others might be 'more unbalanced'.

D.1. **Ways to balance.** [[LABEL 2.balancing]]

We now introduce several definitions that help us us talk about the degree of imbalance of unbalanced graphs.

Definition D.1. [[LABEL D:1022 Bal Vertex]] A *balancing vertex* in a connected signed graph Σ is a vertex such that Σ/v is balanced, but Σ is not.

If Σ has a balancing vertex v, then every negative circle in Σ passes through v. Any graph with a balancing vertex does not contain two vertex-disjoint circles.

Similarly, we define a balancing edge and a balancing edge set. But as these are complicated enough to have two definitions, they get separate lengthy treatment.

D.1.1. *Balancing edges.* [[LABEL 2.baledge]]

A special kind of edge, actually (though not obviously) the nearest signed-graphic analog of an isthmus, is an edge whose deletion increases balance—in a certain precise sense. Dec 5b: Peter Cohen and Thomas Zaslavsky

Definition D.2. [[LABEL D:1205baledge]] A *partial balancing edge* of Σ is an edge e such that $b(\Sigma \setminus e) > b(\Sigma)$.

There are three types of partial balancing edge.

Proposition D.1. *[*[LABEL P:1205baledge]] *An edge e is a partial balancing edge of Σ $\iff$ it is one of the following three types:*

(1) *An isthmus between two components of $\Sigma \setminus e$, of which at least one is balanced. [*[LABEL P:1205baledge isthmus]]
(2) *A negative loop or half edge in an otherwise balanced component of Σ. [*[LABEL P:1205baledge loop]]
(3) *A link $e{:}vw$ added to a component of $\Sigma \setminus e$ that is balanced, e having sign opposite to that of a vw-path in $\Sigma \setminus e$. [*[LABEL P:1205baledge link]]

[This can give an alternative proof of the circuit-closure rank property in Chapter IV, maybe. The idea: It implies that closure preserves rank. I think it's different from my usual proof. Anyway, it might be good in the course notes as a second proof. (If it is a second proof and not merely duplicating what I already have.)]

The proof is a good homework problem. The three cases can be further analyzed. In Case 1, deleting e creates either one balanced component where previously there was an unbalanced component, or two balanced components where there was only one before. In Cases 2 and 3, e is all that prevents the component from being balanced. In Case 3, the vw-path can be any one that does not use e since balance of the component without e implies every such path has the same sign (Theorem **??**). An example of this case is a negative link (not an isthmus) in an otherwise all-positive component of Σ.

For another characterization of a partial balancing edge see Proposition F.7.

Call two edges *series equivalent* if they belong to the same set of circles and they are not isthmi or half or loose edges. An equivalence class is a *series class* of edges.

Proposition D.2. *[*[LABEL P:1205baledgeseries]] *In a connected signed graph Σ, the set of partial balancing edges that are not isthmi is empty or a series class.*

Note that any isthmus might be a partial balancing edge, regardless of what else; it only depends on whether one shore of the isthmus is balanced.

Proof. **[NEEDS PROOF] [Is it possible that no such case arises? If each is a p.b.e., then only one can exist? I feel uneasy about this prop. – TZ]**

□

One could make a different definition. A *total balancing edge* is an edge whose deletion makes an unbalanced signed graph balanced. One might think that, if Σ is connected, a total balancing edge and a partial balancing edge are the same thing; but that isn't so. However, it is true that:

Proposition D.3. *[[LABEL P:1205stronglybaledge]] If Σ is connected, a total balancing edge is an edge of type* (2) *or* (3) *in Proposition D.1.*

This is a two-way result: all total balancing edges are of those types, and any edge of those types is a total balancing edge. The proof is another good exercise for the mental muscles.

D.1.2. *Balancing sets.* [[LABEL 2.balset]]

Oct 22: Jackie Kaminski

With a balancing edge set we find two essentially different concepts.

Definition D.3. [[LABEL D:1022 Bal Set]] An edge set is a *total balancing set* of Σ if its deletion leaves a balanced graph.

An edge set S is a *partial balancing set* of Σ if its deletion increases the number of balanced components; that is, if $b(\Sigma \setminus S) > b(\Sigma)$. A *strict balancing set* is a partial balancing set whose deletion does not increase the number of connected components; that is, it makes one or more existing unbalanced components balanced without breaking any of them apart.

A total balancing set of minimum size has $l(\Sigma)$ edges, by the definition of frustration index.

If Σ is balanced, the empty set is a total balancing set but, obviously, not a partial balancing set. A bond is a minimal partial balancing set but (obviously) not a minimal total balancing set.

A total balancing set makes Σ balanced, while a partial balancing set may not make it balanced but does make it, in a sense, more balanced than it was before. Both kinds of balancing set have to be considered because they serve different purposes. As we shall see, total balancing sets are related to frustration, while partial and strict balancing sets are inovlved with cuts and matroids. We are especially interested in minimal balancing sets, and then there is a simple relationship between the two kinds.

Lemma D.4. *[[LABEL L:1022mintbs]] A total balancing set of Σ consists of a total balancing set of each connected component. An edge set is a minimal total balancing set if and only if it consists of a minimal total balancing set of each unbalanced component.*

Proof. Let $S \subseteq E$ and for each component Σ_i, let $S_i := S \cap E_i$. Then $\Sigma \setminus S$ is balanced if and only if every $\Sigma_i \setminus S_i$ is balanced. That proves the first part and makes the second part obvious. □

Lemma D.5. *[[LABEL L:1022minsbs]] A minimal strict balancing set is a minimal partial balancing set.*

Proof. By Lemma D.4 we may assume Σ is connected. Let B be a minimal strict balancing set. Then $\Sigma \setminus B$ is connected so $b(\Sigma \setminus B) = 1$. By minimality, adding back any edge $e \in B$ makes the graph unbalanced (since it cannot change the number of components), hence $b((\Sigma \setminus B) \cup e) = 0$; in other words, $B \setminus e$ is not a partial balancing set. Thus, B is a minimal partial balancing set. □

The structure of a minimal partial balancing set that is not strict can be rather complicated. It will be developed in our treatment of cuts in Section **??**.

A total, or partial, balancing edge is a total, or partial, balancing set of size 1 (more correctly, the balancing set is $\{e\}$ if e is the balancing edge). A strict balancing edge is also a total balancing set of size 1, provided that Σ is connected (and unbalanced); this is the edge described in Proposition D.1(3). The reader familiar with matroid theory will notice that a partial balancing edge corresponds to a matroid coloop. (See Proposition F.7 for more about this.) **[THAT WILL REQUIRE EXPLANATION ADDED NEAR THE PROP. NAMELY, A BALANCING EDGE OF Σ is a balancing edge of $E \setminus e$.]**

D.2. **A plethora of measures.** [[LABEL 2.plethoraimbalance]]

We now present a list of eight possible measures (generated in class, some by me and some by the students) that one might use to measure the imbalance of a signed graph. This list is in no way meant to be exhaustive. We will follow this this with a discussion of which ones are actually used in certain situations. We would also like to point out that any of the following measurements may be normalized by dividing through by an appropriate quantity.

(1) The minimum number of vertices whose deletion makes Σ balanced. This is the *vertex elimination number* (or "vertex deletion number"), denoted by $l_0(\Sigma)$ [[LABEL R:1022vdeletion]]
(2) The minimum number of edges whose deletion makes Σ balanced. This is the *frustration index*, which we denote by $l(\Sigma)$. (Former or alternative names: line index of balance—whence the letter l; deletion index.) [[LABEL R:1022frustration]]
(3) The minimum number of edges whose negation makes Σ balanced. This is the *negation index*. [[LABEL R:1022negation]]
(4) The maximum number of vertex-disjoint negative circles. [[LABEL R:1022vdnegcircles]]
(5) The maximum number of edge-disjoint negative circles. [[LABEL R:1022ednegcircles]]
(6) The number of negative circles in Σ. Or, the normalized version, which is the proportion of all circles that are negative. [[LABEL R:1022negcirc]]
(7) The minimum number of negative fundamental circles with respect to a maximal forest. (It is not the same for every maximal forest; see below.) [[LABEL R:1022 NFC]]
(8) The minimum number of circles whose successive deletion leaves a balanced graph. [[LABEL R:1022circdeletion]]

The first two have (relatively) standard names. The ones that seem to me to be worth studying are the vertex elimination number (1), the frustration index (2), and the two numbers of disjoint negative circles, (4) and (5).

The frustration index (2) shows up in small-group psychology (usually under Harary's name "line index of balance"), which is where it originated (Abelson and Rosenberg (1958a)) and in physics, especially in spin glass theory (Toulouse (1977a)). Finding the frustration index is NP-hard, because it contains the maximum cut problem, one of the standard NP-complete problems (cf. Akiyama, Avis, Chvátal, and Era)1981a), p. 229); see Corollary D.12. (For a quick explanation of NP, see Section D.3.2.)

The vertex elimination number (1) is NP-hard even when restricted to signed complete graphs—that is, deciding whether it is $\leq k$ is NP-complete (due to Akiyama, Avis, Chvatál, and Era (1981a), p. 232). Evaluating it is also NP-hard, even when restricted to negated line graphs of signed graphs; see Section L. **[Give precise reference to theorem that**

$l_0(-\Lambda(\Sigma) = l(-\Sigma))$ **in line graphs section. Proof: Deleting the edge set S in Σ is the same as deleting the vertex set S in $\Lambda(\Sigma)$. $\Lambda(\Sigma)$ is antibalanced iff Σ is [EXPLAIN: CITE LG section], so $-\Lambda(\Sigma)$ is balanced iff $-\Sigma$ is. Thus, $-\Lambda(\Sigma) \setminus S$ is balanced iff $-\Sigma \setminus S$ is balanced.]**

Although I don't believe (6) actually has a use at present (despite some early consideration in the psychology literature), Tomescu (1976a) and Popescu and Tomescu (1996a et al.) found interesting things to say about it for signed complete graphs.

By the way, the normalized measure in (6) seems to me more meaningful than the unnormalized one because it is comparable between different graphs with different number of circles (or of circles with fixed length). Either measure is equally good if we have a fixed underlying graph—such as K_n.

Example D.1. [[LABEL X:1022negfundcircles]] The value of (7) may in fact differ with the choice of spanning forest T. To see this consider $-K_4$ with T_1 as three edges incident to a single vertex. Then each of the edges in $K_4 \setminus T_1$ has a negative fundamental circle. But if we take T_2 to be a path of length 3, then two edges in $-K_4 \setminus T_1$ have fundamental circles that are triangles and hence negative, but the third edge has a quadrilateral as its fundamental circle, which is positive.

The next lemma tells us that minimal total balancing sets are minimal negative edge sets.

Lemma D.6. *[*[LABEL L:1022minbalset]] *If S is a minimal total balancing set of Σ, then Σ can be switched so that S is its set of negative edges.*

Contrariwise, the negative edge set of a switching Σ^ζ is a total balancing set. It is minimal such if and only if it is minimal among negative edge sets $E^-(\Sigma^\zeta)$ of switchings.

Proof. For the first part, $\Sigma \setminus S$ has the same number of connected components as Σ; otherwise S would not be minimal since one could add to it an edge connecting two of its components that are in the same component of Σ. Take T a maximal forest in $\Sigma \setminus S$; it is also a maximal forest of Σ. By Lemma A.8 we can switch Σ so T is all positive. Then every edge not in S is positive, because its fundamental circle is positive since $\Sigma \setminus S$ is balanced. Every edge in S has to be negative, because if $e \in S$ were positive, $S \setminus e$ would be a smaller total balancing set. Thus, S is the negative edge set of the switched Σ.

I leave the proof of the second part to the reader. ∎

Proposition D.7. *[*[LABEL L:1022 2and4]] *The imbalance measure in* (7) *is not less than the frustration index, and is equal to it for some choice of maximal forest.*

Proof. The number of negative fundamental circles with respect to a maximal forest T equals the number of negative edges when T is switched to be all positive. This number is not less than $l(\Sigma)$.

To prove (7) can be equal to $l(\Sigma)$, take S to be a minimum total balancing set. Then by Lemma D.6 there is a switching in which $E^- = S$. By the proof of that lemma, $\Sigma \setminus S$ contains a maximal forest of Σ, call it T. (7) for this choice of T equals the frustration index. ∎

D.3. **Frustration index.** [[LABEL 2.frustrationindex]]

D.3.1. *Properties.* [[LABEL 2.frustrationindexproperties]]

It seems that frustration index is far the most important measure of imbalance. Here are some of its properties. The first one is an essential property, first stated (in their unique matrix language) by Abelson and Rosenberg (1958a) and then (in more ordinary matrix language) by S. Mitra (1962a). I don't remember who gave the first explicit proof.

Lemma D.8. *[[LABEL L:1022frustrationindex]] There is a switching of Σ in which the number of negative edges equals the frustration index, but no switching has fewer negative edges.*

Proof. This is an immediate consequence of Lemma D.6. The frustration index is, by definition, the size of a minimum total balancing set. Let S be such an edge set and switch Σ so $S = E^-$. Then $|E^-| = l(\Sigma)$.

On the other hand, any set E^- in a switching of Σ is a total balancing set for Σ, so it cannot be smaller than $l(\Sigma)$. □

The first part of the next theorem is due to Harary. The second part is the preceding lemma.

Theorem D.9. *[[LABEL T:1022 Harary]] For a signed graph Σ, the frustration index $l(\Sigma)$ = the negation index of $\Sigma = \min_\zeta |E^-(\Sigma^\zeta)|$, the minimum number of negative edges in any switching.*

Proof. Suppose negating $R \subseteq E$ makes Σ balanced. Then every circle in $\Sigma \setminus R$ is positive, so $\Sigma \setminus R$ is balanced. On the other hand, if S is a minimal total balancing set, switch so it is the negative edge set. Then negating S makes the switched graph balanced. Therefore, negating S makes Σ balanced. That proves the first equation.

For the second, Lemma D.8 states that $l(\Sigma)$ equals the minimum number of negative edges in a switched Σ. □

Lemma D.10. *[[LABEL L:1022frustrateddegree]] If Σ is a signed link graph such that $l(\Sigma) = |E^-|$, then $d^-(v) \leq \frac{1}{2}d(v)$ at every vertex.*

Proof. Suppose $d^-(v) > \frac{1}{2}d(v)$, or equivalently $d^-(v) > d^+(v)$. Then by switching v we reduce the number of negative edges at v while not changing the signs of the other edges. Thus, to minimize $|E^-(\Sigma^\zeta)$ we have at least to switch so every vertex has negative degree no larger than its positive degree. □

The problem of frustration index includes the well known max-cut problem for graphs.

Corollary D.11. *[[LABEL P:1022 negative frustration]] For a graph Γ, the frustration index of $-\Gamma$ is given by*

$$l(-\Gamma) = |E(\Gamma)| - \max_X |E(X, X^c)|,$$

the complement of the maximum cut size in Γ.

Proof. Recall that a cut $E(X, X^c)$ consists of the edges with one endpoint in each set of a bipartition $\{X, X^c\}$ of V. Let $E(X, X^c)$ be a cut of Γ. The cut edges of an all-negative graph will form an all-negative bipartite graph, where all circles are of even length and therefore positive. So the remaining edges (the edges of $E(\Gamma) \setminus$ the cut) are a set whose deletion balances $-\Gamma$.

Certainly, $\max_{X \subseteq V} |E(X, X^c)|$ will minimize the size of $E(\Gamma) \setminus$ a cut.

Lastly, since every balanced subgraph of an all-negative graph must be bipartite, every total balancing set is the complement of a bipartite subgraph. This completes the proof. □

D.3.2. *Computational complexity.* [[LABEL 2.frustrationindexcomplexity]]

2014 Oct 31: T.Z.

Corollary D.11 tells us something important about the computational difficulty of the frustration index. To understand this one has to know the current classification of computational complexity.

A *decision problem* is a question of the following form, where P is a fixed property: "Given an input object I, does it have property P?" For instance, the input might be an ordinary graph together with a number k and the property might be that the maximum cut size is at most k. The answer could be "Yes" or "No". This problem is called the maximum cut problem, "max-cut" for short, and written MAXCUT. (It's conventional to write computational problems in capitals.)

Given a decision problem, we want to know whether a "Yes" answer can be found quickly or not. "Quickly" is somewhat arbitrarily defined as "in time polynomial in the length of the input". For a graph, the length of the input may be a constant times $|V| + |E|$ (depending on the input format). Note that our question, "Is the max cut size in Γ at most k?", is asymmetric; we are only asking for a positive answer. If we wanted a negative answer, we would ask the opposite question, "Is the max cut size in Γ greater than k?" It's possible that one question can be answered quickly and the other cannot.

There are two systems for computing an answer. If straightforward computation, considering every possible input and every possible sequence of trials (in MAXCUT that means we may have to test every set $X \subseteq V$ because it's possible that the only), will always give us the "Yes" answer in polynomial time, the problem belongs to the class P of *Polynomial-time decision problems*. (MAXCUT does not belong to P because we might have to test 2^{n-1} sets X before finding the one whose cut size is $\leq k$.) If straightforward computation, considering every possible input and only the shortest possible sequence of trials (in MAXCUT that means we happen to test first an X whose cut is indeed no bigger than k, so we immediately get the answer "Yes"), gives the answer "Yes", then we say the problem belongs to NP, the class of *Nondeterministically Polynomial-time decision problems*. (The name comes from an equivalent way of modelling the process: we assume we have unlimited parallel processing so we can test every sequence of trials simultaneously.) Every problem in P is also in NP, but not conversely—*maybe!* In fact, probably everyone believes P $\neq$ NP but no one has been able to prove it.

There are decision problems that are not in NP. I will not be concerned with them. There are problems that are in NP and have maximal difficulty, meaning that if we can solve such a problem in a certain time $T(l)$, depending on the input length l, then we can solve every NP problem in time that is a polynomial in $T(l)$. Such problems are *NP-complete*. MAXCUT is one of the best-known NP-complete problems. (If you can prove MAXCUT is, or is not, solvable in polynomial time, you have solved the P vs. NP question and you win a Clay Foundation Millennium Prize.)

One more explanation. A problem that is not a decision problem, such as "What is the maximum cut size in Γ", may be solvable in polynomial time. For instance, the diameter of a graph can be computed in polynomial time. The analog of an NP decision problem for computing a number (instead of making a decision) is the *NP-hard problem*.

Now we can state the complexity consequence of Corollary D.11.

Corollary D.12. *[*[LABEL C:1022 NP]] *The frustration index of signed graphs is an NP-hard problem. The question "Is $l(\Sigma) \leq k$?" is NP-complete.*

Proof. The maximum-cut problem is already NP-hard, and MAXCUT is NP-complete. (See any book on algorithmic complexity. The classic is Garey and Johnson [GJ].) □

In other words, don't expect to find frustration index in polynomial time (but if you do, you have proved P = NP!).

D.3.3. *Computational methods.* [[LABEL 2.frustrationindexcomputations]]

On the other hand, if you are interested in a particular class of signed graphs, it is conceivable that frustration index can be computed "quickly". (In this section we assume the underlying graphs are simple.) Such is the case for signed planar and toroidal graphs (see Katai and Iwai (1978a)) and Barahona (1981a, 1982a)), including the square lattice graphs of interest in physics (see Bieche, Maynard, Rammal, and Uhry (1980a)). However, it is not the case for signed 3-dimensional lattice graphs (according to Barahona (1982a)) or signed bipartite graphs (see ?? (??)); frustration index limited to either class is still NP-complete. I don't know, and I think no one knows, whether it can be computed in general in time much less than $Cn^2 2^{n-1}$, which is how long it takes to test every possible switching for the number of negative edges. (C is a constant. The n^2 is the time it takes to check the sign of every edge of K_n—neglecting the constant of 1/2.) This time is exponential in n and is considered—rightly in this case—slow. 2014 Nov 23: T.Z.

Problem D.1. [[LABEL Pr:1022]] Improve the general upper bound in terms of n on the time required to compute $l(\Sigma)$ for an arbitrary signed simple graph. (To repeat, I do not know the current best bound.)

Here is a brief account of the algorithm of Barahona and Katai–Iwai for finding $l(\Sigma)$ when the underlying graph $|\Sigma|$ is planar. First, embed $|\Sigma|$ in the plane; this is called a *plane graph.* A plane graph has regions and each region has a boundary walk, whose sign product is assigned to the region. That gives the dual graph $|\Sigma|^*$, whose vertices are the regions of $|\Sigma|$ and whose edges join the regions on opposite sides of each edge of $|\Sigma|$, a vertex signature. We now find a matching of negative vertices in the dual graph, where the matching consists of paths between matched vertices, such that the total length of all paths is a minimum. (In that case, no two matching paths can have a common edge.) The edges of Σ that correspond to matching path edges in the dual graph form a minimum negation set; thus, the number of such edges is $l(\Sigma)$. Summarizing:

Theorem D.13 (Katai–Iwai, Barahona)**.** *[[LABEL T:1022planar]] The frustration index of a signed planar graph Σ equals the minimum number of edges in a set of matching paths between negative vertices in the dual graph $|\Sigma|^*$.*

It is interesting that the publication of Katai–Iwai, which preceded Barahona's by three years, was in a psychology journal, where its mathematics was exceptionally high-powered. Barahona published in a physics journal, where the mathematics was not especially high-powered but was of an unusual kind. It may not be a coincidence that Barahona's presentation is much simpler, shorter, and more easily read.

D.4. **Maximum frustration.** [[LABEL 2.maxfrustration]]

Computing the maximum frustration index of any signature of a given graph should be no less difficult than finding the frustration index of a particular signed graph, although I don't know of any proof about this. Nevertheless, there are some theorems. 2014 Oct 31: T.Z.

Definition D.4. [[LABEL D:1022 max frust]]$l_{\max}(\Gamma) := \max_{\sigma:E\to\{+,-\}} l(\Gamma,\sigma)$, the maximum frustration index over all signatures.

This number $l_{\max}$ was introduced by Akiyama, Avis, Chvátal, and Era (1981a). Computing it leads us to an often-rediscovered theorem of Petersdorf.

Theorem D.14 (Petersdorf (1966a)). *[[LABEL T:1022 Petersdorf]] For the complete graph,*

$$l_{\max}(K_n) = l(-K_n) = \left\lfloor \frac{(n-1)^2}{4} \right\rfloor .$$

The signatures whose frustration index achieves the maximum are precisely those in the switching class of $-K_n$.

In other words, antibalance uniquely maximizes the frustration index.

Proof. We have three things to prove: the exact value of $l(-K_n)$, that the maximum frustration index is achieved by $-K_n$, and that no other signature, up to switching, achieves the same frustration index.

Part 1. To see that $l(-K_n) = \lfloor\frac{(n-1)^2}{4}\rfloor$, we observe that by Proposition D.11, $l(-K_n) = |E(K_n)| - |\text{max cut of } K_n|$. An edge cut is just the set of edges with one endpoint in each part of a bipartition of V. In K_n, such a set is a complete bipartite graph $K_{i,n-i}$, which has $i(n-i)$ edges. Therefore, $l(-K_n) = \max_{0\leq i\leq n} i(n-i)$. Since $i(n-i)$ is an increasing function of i for $i < \frac{n}{2}$ and decreasing for $i > \frac{n}{2}$, $\max_{i=0,1,\dots,n} i(n-i) = \lfloor\frac{n}{2}\rfloor(n-\lfloor\frac{n}{2}\rfloor)$. If n is even this is $\frac{n}{2}\cdot\frac{n}{2} = \frac{n^2}{4}$. If n is odd it is $\frac{n-1}{2}\cdot\frac{n+1}{2} = \frac{n^2-1}{4}$. Both cases can be expressed as $\lceil\frac{n^2-1}{4}\rceil$. The frustration index is therefore $\left\lfloor\binom{n}{2} - \frac{n^2-1}{4}\right\rfloor = \lfloor\frac{(n-1)^2}{4}\rfloor$. This gives the value of $l(-K_n)$, which takes care of the first part of the proof.

Part 2. By definition, $l_{\max}(K_n) = \max_{\sigma:E\to\{+,-\}} l(K_n,\sigma)$, which equals the maximum negation index of any (K_n,σ) by Theorem D.9. We assume from now on that (K_n,σ) is already switched so the number of negative edges equals its frustration index. By Lemma D.10 every vertex has negative degree $\leq \lfloor(n-1)/2\rfloor$. Thus, the number of negative edges is at most $\frac{1}{2}n\lfloor(n-1)/2\rfloor$.

If n is even this is $\frac{1}{4}n(n-2) = \lfloor\frac{(n-1)^2}{4}\rfloor$, so $-K_n$ does have maximum frustration.

If n is odd, it is $\lfloor\frac{1}{4}n(n-1)\rfloor$, which is larger than $l(-K_n)$. We must look deeper. Suppose there are two positively adjacent vertices, v and w, both with negative degree $\frac{1}{2}(n-1)$. The total number of negative edges from $\{v,w\}$ to $V\setminus\{v,w\}$ is $n-1$. The total number of edges between $\{v,w\}$ and $V\setminus\{v,w\}$ is $2(n-2)$. Therefore, by switching $\{v,w\}$ we reduce the number of negative edges. **[PICTURE HERE.]** That contradicts the hypothesis that $|E^-|$ equals the frustration index; we conclude that no two vertices with negative degree $\frac{1}{2}(n-1)$ can be positively adjacent. This implies that, if $d^-(v) = \frac{1}{2}(n-1)$ for some vertex v, then all other vertices with the same degree are neighbors of v. Thus, there cannot be more than $\frac{1}{2}(n+1)$ vertices with degree $\frac{1}{2}(n-1)$. The remaining $\frac{1}{2}(n-1)$ vertices have degree at most $\frac{1}{2}(n-3)$. Adding up these degrees, there are no more than $\frac{n+1}{2}\frac{n-1}{2} + \frac{n-1}{2}\frac{n-3}{2} = \frac{(n-1)^2}{4}$ negative edges, the exact value of $l(-K_n)$. Consequently, $-K_n$ has maximum frustration in the odd case.

Part 3. We ask whether there is more than one switching class that has maximum frustration.

In the odd case we get the largest frustration when (K_n,σ) has $\frac{1}{2}(n+1)$ vertices with $d^-(v) = \frac{1}{2}(n-1)$. None of these vertices can be positively adjacent; thus, they form a clique of order $\frac{1}{2}(n+1)$ in the negative subgraph. Each vertex has $\frac{1}{2}(n-1)$ neighbors in the clique, so it cannot be negatively adjacent to any other vertex. Thus, the most negative edges arise when the remaining $\frac{1}{2}(n-1)$ vertices also form a negative clique. This is precisely $-K_n$ with a maximum cut switched to positive. Thus, the only signature on K_n that has maximum frustration is the all-negative one.

In the even case the negative subgraph Σ^- must be $\frac{n}{2}$-regular for maximum frustration. The solution is similar to that for odd n but slightly more complicated. Instead of showing that two vertices of maximum negative degree must be negative neighbors, we prove that no three vertices can be positively adjacent and deduce that no two positively adjacent vertices can have a common negative neighbor.

Suppose first that u, v, w are positively adjacent. Then all their $3(\frac{n}{2}-1)$ negative neighbors are in $V \setminus \{u, v, w\}$. That leaves $3(\frac{n}{2}-2)$ positive edges between $\{u, v, w\}$ and $V \setminus \{u, v, w\}$, so switching $\{u, v, w\}$ reduces the number of negative edges, contradicting the hypothesis on σ. Therefore, no three vertices can be positively adjacent.

Now suppose v, w are positively adjacent. Their negative neighborhoods combined, $N^-(v) \cup N^-(w)$, constitute at most $2(\frac{n}{2}-1) = |V \setminus \{v, w\}|$ vertices. By the preceding paragraph there cannot be a vertex that is positively adjacent to both v and w. Consequently, $N^-(v) \cup N^-(w) = V \setminus \{v, w\}$, from which we deduce that $N^-(v) \cap N^-(w) = \varnothing$. We have proved that, if two vertices are negative non-neighbors, their neighborhoods are disjoint. Restating that, if two vertices have a common negative neighbor, they must be negatively adjacent. Hence, Σ^- is a union of disjoint cliques, each of degree $\frac{1}{2}n-1$, thus of order $\frac{1}{2}n$. So $\Sigma^- = K_{n/2} \mathbin{\dot\cup} K_{n/2}$ and (K_n, σ) is a switching of $-K_n$. That concludes the last part of the proof. □

For completeness' sake we mention that $l(\text{any signed forest}) = 0$, since it has no circles, and therefore $l_{\max}(\text{any forest}) = 0$. We just did $l_{\max}(K_n)$ above.

A next logical graph to consider is $K_{r,s}$; but this is considerably more of a problem than K_n. With K_n, the 'obvious' signing $-K_n$ yields the maximum frustration index. However, with $K_{r,s}$ there is no 'obvious' signature to yield a high frustration index, since the all-negative signature has frustration index 0 and there is no clear substitute. In view of the relatively obscurity of signed graphs within graph theory, it may be surprising that the value of $l_{\max}(K_{r,s})$ has been the subject of several papers. The reason is that it is the 'rectangular' generalization of the Gale–Berlekamp switching game, which has been a challenging problem for the last oh-so-many years (see, i.a., Brown and Spencer (1971a) and Solé and Zaslavsky (1994a)).

The Gale–Berlekamp switching game is played on $K_{r,r}$, or rather, on an $r \times r$ board with a light bulb in each square and a switch for each row and column. Initially, some of the lights are on and some are off. A switch will reverse all the bulbs in its row or column. The goal is to keep switching so as to minimize the number of lit bulbs. The problem is to find the exact upper bound on that number. Transforming the board into a signed $K_{r,r}$ by making edge $v_i w_j$ negative when the bulb in row i and column j is lit, we have the problem of evaluating $l_{\max}(K_{r,r})$.

It follows from coding theory that frustration index of a randomly signed $K_{r,s}$ (for variable r and s) is NP-hard; this leads one to expect that $l_{\max}(K_{r,s})$ is also NP-hard—although I don't know of a proof. Nevertheless, we do know how to solve one general case, that in which $s = k2^{r-1}$, from Garry Bowlin's recent doctoral thesis (2009a). In a signed graph, let $v(N)$ denote a vertex whose negative neighborhood is N.

Theorem D.15. [[LABEL T:1022 bowlin]] *For the complete bipartite graph $K_{r,k2^{r-1}}$ with left set $[r]$, where $r, k > 0$, the signature with largest frustration index is the one that has k right vertices $v(N)$ for each $N \subseteq [r]$ such that $|N| < r/2$ and also (if r is even) for each $N \subseteq [r]$ such that $|N| = r/2$ and $1 \in N$.*

Bowlin (2009a) also shows that there are a systematic construction of reasonably tight bounds for all s, given a fixed value of r; however, he does not prove the exact maximum frustration.

Despite its ineffectiveness on bipartite graphs, the all-negative signature (that is, antibalance) is tempting. I propose the following problem, about whose solution I have no clue:

Problem D.2. [[LABEL Pr:1022 maxfr]] Find necessary, sufficient, or necessary and sufficient conditions on a graph Γ for $l(-\Gamma)$ to equal the maximum frustration $l_{\max}(\Gamma)$.

I close this topic with the statement of Akiyama et al.'s general theorem about maximum frustration, which I will not prove. (The first inequality is the nontrivial part.)

Theorem D.16 (Akiyama, Avis, Chvátal, and Era (1981a))**.** [[LABEL T:1022maxfrust]] *For an ordinary graph Γ with $m := |E|$, $m/2 - \sqrt{mn} \leq l_{\max}(\Gamma) \leq m/2$.*

And to conclude, here is a suggestion for an example that may provide a new exact value. (I don't know of any published value.)

Conjecture D.3. [[LABEL Cj:1022 Krst...]] The maximum frustration of the complete multipartite graph $K_{n_1,n_2,\ldots,n_k}$ where $k \geq 3$ is $l(-K_{n_1,n_2,\ldots,n_k})$.

The conjecture is not complete, however.

Problem D.4. [[LABEL Pr:1022 Krst...]] Evaluate $l(-K_{n_1,n_2,\ldots,n_k})$, assuming $k \geq 3$.

D.5. **Disjoint negative circles.** [[LABEL 2.disnegcircles]]

We now turn our attention to imbalance measure (4) and consider when the maximum number of vertex-disjoint negative circles is 1. The reader familiar with matroid theory will be interested to know that for a 2-connected signed graph, having no two vertex-disjoint negative circles is equivalent to having a binary frame matroid. I state a theorem, first proposed by Lovasz with an incomplete proof, that was finally established by Slilaty.

Theorem D.17 (Slilaty (2007a))**.** [[LABEL T:1022 Slilaty]] *Σ has no vertex-disjoint negative circles if and only if one or more of the following is true:*

(1) Σ is balanced,
(2) Σ has a balancing vertex,
(3) Σ embeds in the projective plane, or
(4) Σ is one of a few exceptional cases.

The proof of this remarkable theorem, as well as a formal definition of a signed graph embedding (technically, "orientation embedding"—see especially Zaslavsky (1992a)), are beyond the scope of this course. But I note that the backward direction of the proof is easier than the forward direction, and that in a signed graph embedding, a circle is negative if and only if it is orientation reversing in the embedding.

---

24 Oct 2008: Simon Joyce

E. Minors of Signed Graphs

[[LABEL 2.minors]]

For a signed graph as for a graph, a *minor* is any result of a sequence of contractions and deletions of edge sets and deletions of vertices. So before we can discuss minors we must define contraction.

But first let us dispose of deletion. We mentioned in Section A.1 that any subgraph of Σ is naturally signed by inheritance from Σ. That takes care of deletion, since deleting edges or vertices from Σ amounts to taking a subgraph.

E.1. **Contraction.** [[LABEL 2.contraction]]

Contraction of edges in a signed graph is substantially more complex than in ordinary graphs. Thus, we develop the notion of contraction in two stages: first we contract a single edge, then an arbitrary set of edges.

E.1.1. *Contracting a single edge.*

If e is a positive link we delete e and identify its endpoints, which is how we normally contract a link in an unsigned graph. If e is a negative link we take a switching ζ of Σ such that e is a positive link in Σ^ζ. Now we contract e in the usual way. We must check that this operation is in some sense well defined.

Lemma E.1. *[[LABEL L:1024 link contraction equivalence]] In a signed graph Σ any two contractions of a link e are switching equivalent. The contraction of a link in a switching class is a well defined switching class.*

Proof. If e is a positive link the result is immediate so let's assume e is a negative link. Let ζ_1 and ζ_2 be any two switching functions of Σ such that e is a positive link in both Σ^{ζ_1} and Σ^{ζ_2}. We want to show Σ^{ζ_1}/e and Σ^{ζ_2}/e are switching equivalent. Since $|\Sigma^{\zeta_1}/e| = |\Sigma^{\zeta_2}/e|$ by theorem A.6 it will suffice to show $\mathcal{B}(\Sigma^{\zeta_1}/e) = \mathcal{B}(\Sigma^{\zeta_2}/e)$.

Let C be a circle in Σ. Since switching does not change the sign of the circle, C has the same sign in both Σ^{ζ_1} and Σ^{ζ_2}. If e is not an edge of C, then contracting e won't affect the sign of C in Σ^{ζ_1}/e or Σ^{ζ_2}/e. If e is an edge of C, since the sign of e is positive in Σ^{ζ_1} and Σ^{ζ_2} contracting it won't affect the sign of C in Σ^{ζ_1}/e or Σ^{ζ_2}/e either. It follows that $\mathcal{B}(\Sigma^{\zeta_1}/e) = \mathcal{B}(\Sigma^{\zeta_1}/e)$. □

When we contract a positive loop or a loose edge e we just delete e.

If e is a negative loop or half edge and v is the vertex of e, we cut out v (as if with scissors) and delete e. This operation may produce several half and loose edges, as can be seen in Figure E.1. Since we are deleting e and v, $V(\Sigma/e) = V(\Sigma) \setminus \{v\}$, and $E(\Sigma/e) = E(\Sigma) \setminus \{e\}$. Also for any edge $f \neq e$ we have $V_{\Sigma/e}(f) = V_\Sigma(f) \setminus \{v\}$. So if f is a link with endpoints v and w it becomes a half edge at w in the contraction. If f is a loop at v, f becomes a loose edge in the contraction. (This is one of two reasons why we have half and loose edges.)

E.1.2. *Contracting an edge set S.*

Contraction of an arbitrary edge set S of a signed graph Σ will also be more complicated than contraction for ordinary, unsigned graphs. The process differs for the balanced and unbalanced components of S. The edge set and vertex set of Σ/S will be as follows:

$$E(\Sigma/S) := E(\Sigma \setminus S),$$
$$V(\Sigma/S) := \{\text{vertex sets of balanced components of } (V, S) = \Sigma|S\}$$

$$= \boldsymbol{\pi}_{\mathrm{b}}(S).$$

To contract we first apply a switching function ζ so the balanced components of S are all positive. Lemma A.8 guarantees we can do this. Once we have switched, we contract each balanced component of S in the usual way.

To contract the unbalanced components of S we cut them out and delete all the edges and vertices of each unbalanced component in a similar process to how we contracted a negative loop or half edge. This may create some half edges or loose edges. If an edge $e \notin S$ is a link and has a single endpoint in an unbalanced component of S, then it becomes a half edge in the contraction. If both endpoints of e are in unbalanced components of S, or if e is a half edge with it's endpoint in an unbalanced component of S, then e becomes a loose edge in the contraction.

The signature $\sigma_{\Sigma/S}$ is the sign function induced by Σ^{ζ}. Then any edge that is a link or loop in Σ/S keeps its (switched) sign. Any half or loose edges have no sign.

To summarise, once we have found such a switching ζ, we have,

$$\Sigma/S = \big(V(\Sigma/S), E(\Sigma/S), \sigma_{\Sigma/S}\big).$$

Notice that contraction of an unsigned graph Γ behaves exactly like contraction of $+\Gamma$ as we have defined it here.

An example of how the process of contraction works is presented in Figure E.2. Here $\Sigma = \pm K_5$ and S is the set of red edges. Since $|E(S)| = 4$ we have $|E(\pm K_5/S)| = 16$. Notice that $\boldsymbol{\pi}_{\mathrm{b}}(S) = \{v_3, v_4\}$ and this will correspond to the only vertex in the contraction. To contract we switch the vertex v_3 and then contract the edge $-e_{34}$. $+e_{34}$ is now a negative loop. Now we cut out the unbalanced component, i.e., we reduce the vertices v_1, v_2, and v_5 and delete the remaining edges of S. All the edges with an endpoint at the new contracted vertex become half edges and all the edges with all endpoints in the unbalanced component become loose edges.

As with contraction of a single edge, we must show this process is in some sense well defined. That is the content of the next result.

Lemma E.2. *[*[LABEL L:1024 contraction equivalence]]

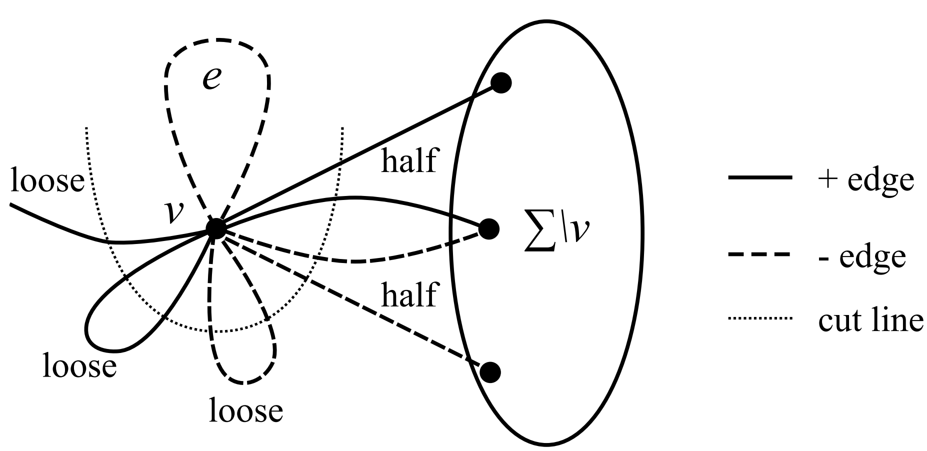

FIGURE E.1. Cutting out v leaves half and loose edges.
[[LABEL F:1024Figure1]]

(a) *Given Σ a signed graph and $S \subseteq E(\Sigma)$, all contractions Σ/S (by different choices of switching Σ) are switching equivalent. Any switching of one contraction Σ/S is another contraction and any contraction Σ^ζ/S of a switching of Σ is a contraction of Σ.*

(b) *If $|\Sigma_1| = |\Sigma_2|$, $S \subseteq E$ is balanced in both Σ_1 and Σ_2, and Σ_1/S and Σ_2/S are switching equivalent, then Σ_1 and Σ_2 are switching equivalent.*

Part (a) tells us that a contraction of a signed graph is not really a signed graph; it is a switching class. It also tells us that contraction really applies to switching classes more than to individual signed graphs. We can summarize this in a formula:

$$[\Sigma]/S = [\Sigma/S].$$

Proof. By theorem A.6, since $|\Sigma^\zeta/S|$ is the same for any switching function, if we can show $\mathcal{B}(\Sigma^\zeta/S)$ does not depend on the switching function ζ, our result will follow. When we contract by S we contract each component of S separately so it will suffice to show the result holds when we contract a single balanced component or unbalanced component.

First assume S is composed of a single balanced component. To contract S we must apply a switching function so that all the edges of S are positive. Again, such switching fuctions exist by Proposition A.9. Let ζ_1 and ζ_2 be two such switching functions. Let x be the vertex corresponding to S in Σ^{ζ_1}/S and let $C \in \mathcal{B}(\Sigma^{\zeta_1}/S)$.

If $x \notin V(C)$, then $C \in \mathcal{B}(\Sigma)$. Since switching does not change the sign of circles it follows that $C \in \mathcal{B}(\Sigma^{\zeta_2}/S)$.

Now suppose $x \in V(C)$. Consider the path $P \in \Sigma^{\zeta_1}$ induced by the edges of C. P is positive since C is balanced. If P is closed, then $C \in \mathcal{B}(\Sigma)$ and so $C \in \mathcal{B}(\Sigma^{\zeta_2}/S)$. Otherwise P has distinct endpoints $v, w \in V(S)$ and $E(P) \cap S = \varnothing$. Since all the edges of S in Σ^{ζ_1} and Σ^{ζ_2} are positive, there is a positive path Q in S with endpoints v and w in both Σ^{ζ_1} and Σ^{ζ_2}. Therefore the circle $P \cup Q \in \mathcal{B}(\Sigma^{\zeta_1})$. It follows that $P \cup Q \in \mathcal{B}(\Sigma^{\zeta_2})$ and since the edges in Q are all positive we get that $C \in \mathcal{B}(\Sigma^{\zeta_2}/S)$.

A similar argument shows that if $C \in \mathcal{B}(\Sigma^{\zeta_2}/S)$, then $C \in \mathcal{B}(\Sigma^{\zeta_1}/S)$, so $\mathcal{B}(\Sigma^{\zeta_2}/S) = \mathcal{B}(\Sigma^{\zeta_2}/S)$.

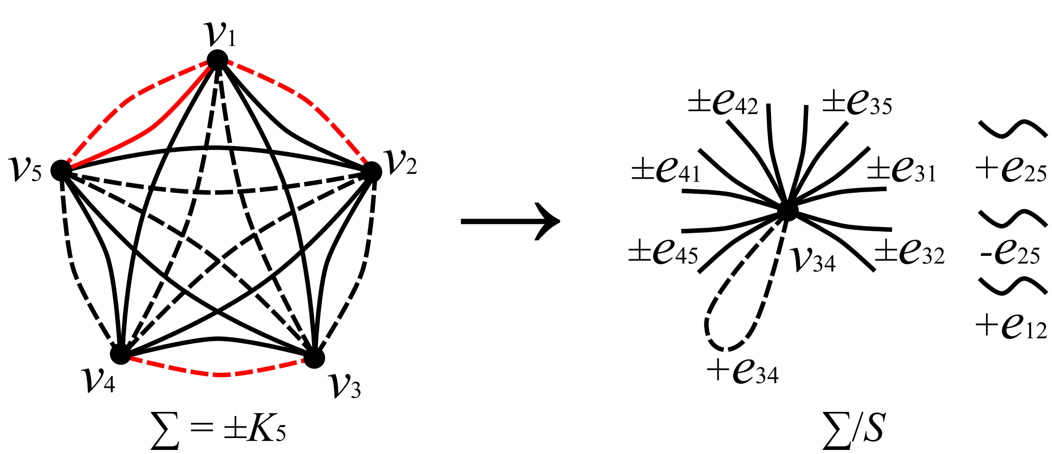

FIGURE E.2. A contraction of $\pm K_5$.
[[LABEL F:1024Figure2]]

Now assume S is composed of a single unbalanced component. Let $C \in \mathcal{B}(\Sigma^{\zeta_1}/S)$. Since no vertex of C can be in S we have that $C \in \mathcal{B}(\Sigma)$, and therefore $C \in \mathcal{B}(\Sigma^{\zeta_2}/S)$. It follows that $\mathcal{B}(\Sigma^{\zeta_1}/S) = \mathcal{B}(\Sigma^{\zeta_2}/S)$.

If ζ is a switching function of Σ/S, then we can define a pullback $\hat{\zeta}$ that is a switching function of Σ by $\hat{\zeta}(v) := \zeta(V_i)$ for $v \in V_i \in \boldsymbol{\pi}_{\mathrm{b}}(S)$ and letting $\hat{\zeta}|_{V_0(S)}$ be arbitrary. Then $\Sigma^{\hat{\zeta}}/S = (\Sigma/S)^{\zeta}$, i.e., $(\Sigma/S)^{\zeta}$ is another contraction of Σ. That Σ^{ζ}/S is a contraction of Σ where ζ is a switching function is immediate.

For part (b) of the Theorem, since Σ_1/S and Σ_2/S are switching equivalent, Σ_2/S is a contraction of Σ_1 by part (a). So there is a switching function ζ_1 such that $\Sigma_1^{\zeta_1}/C = \Sigma_2/C$. Note all the edges of S are positive in $\Sigma_1^{\zeta_1}$ so that we can contract S. Now Σ_2/C is obtained from Σ_2 by applying a switching function ζ_2 to Σ_2 that made all the edges in S positive and then contracting S. This means that the edge signs of $\Sigma_2^{\zeta_2}$ are the same as the edge signs of $\Sigma_1^{\zeta_1}$ and therefore Σ_1 and Σ_2 are switching equivalent. □

[**Proof needs to be checked**]

Part (b) of lemma E.2 fails if S is unbalanced. An example of this is shown in Figure E.3.

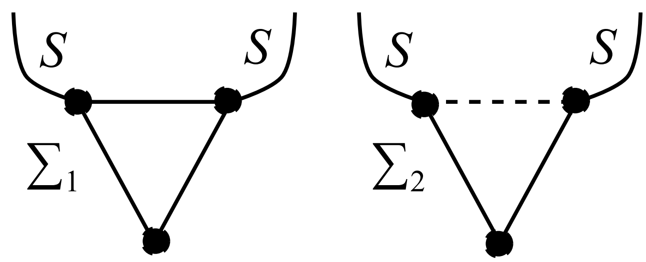

FIGURE E.3. Part (b) of Lemma E.2 fails for unbalanced S.
[[LABEL F:1024Figure3]]

Exercise E.1. [[LABEL Ex:1024conn-contraction]] Suppose Σ is connected. Show that Σ/S is connected if S is balanced, but not necessarily if S is unbalanced.

Compare with Exercise I.**??**.

Exercise E.2. [[LABEL Ex:1024unbal-contraction]] Suppose Σ is unbalanced and connected. Prove that Σ/S is unbalanced if and only if it has at least one vertex.

E.2. **Minors.** [[LABEL 2.minors.minors]]

A minor of a signed graph can, by definition, be constructed—amongst other ways—as follows. First, delete all edges that are supposed to be deleted. Now all vertices to be deleted become isolated; delete those vertices. Finally, contract all edges that are supposed to be contracted. In short, a minor of a signed graph is a contraction of a subgraph. This is only one of the many choices of the order in which to carry out the deletions and contractions that result in a particular minor. We must prove that all choices give the same result.

Theorem E.3. *[*[LABEL T:1024 minors are minors]]

(a) *Every minor of a signed graph is obtainable as a contraction of a subgraph, and also as a subgraph of a contraction.*

(b) *All minors of* Σ *that result from deleting an edge set* S *and a vertex set* X *and contracting an edge set* T *disjoint from* S *are switching equivalent to* $(\Sigma \setminus S \setminus X)/T$ *and to* $(\Sigma/T) \setminus S \setminus \bar{X}$*, where* $\bar{X}$ *is the image of* X *in the contraction.*

Proof of (b). Combine (a) with Lemma E.2. □

Oct 27: Nate Reff

Proof of (a). We have to prove three things: that a subgraph of a subgraph is a subgraph, which is obvious; that a subgraph of a contraction is a contraction of a subgraph; and that (E.1)
$(\Sigma/S)/T = \Sigma/(S \cup T)$ when $S, T \subseteq E$ with $S \cap T = \varnothing$.[[LABEL E:sg contraction contraction]]

(Equality here, as in any repeated contraction of a graph, means identity of edge sets along with a bijection of the vertex sets such that the same edge is incident to corresponding vertices in the two graphs.)

To prove Equation (E.1), remember that $V(\Sigma/S) = \boldsymbol{\pi}_{\mathrm{b}}(V, S)$ and, similarly, that

$$V(\Sigma/(S \cup T)) = \boldsymbol{\pi}_{\mathrm{b}}(V, S \cup T) \in \Pi_V^{\bullet},$$
$$V((\Sigma/S)/T)) = \boldsymbol{\pi}_{\mathrm{b}}(V/S, T) \in \Pi_{V/S}^{\bullet}.$$

Therefore $V(\Sigma/(S \cup T))$ and $V((\Sigma/S)/T))$ cannot really be equal. If we want equality we have to allow vertex bijection but the identity correspondence for the edge sets.

We write $V_0(\Sigma) = \{$vertices of unbalanced components$\}$, and $V_{\mathrm{b}}(\Sigma) = \{$vertices of balanced components$\}$ So we have $\bigcup \boldsymbol{\pi}_{\mathrm{b}}(S) = V_{\mathrm{b}}(S)$.

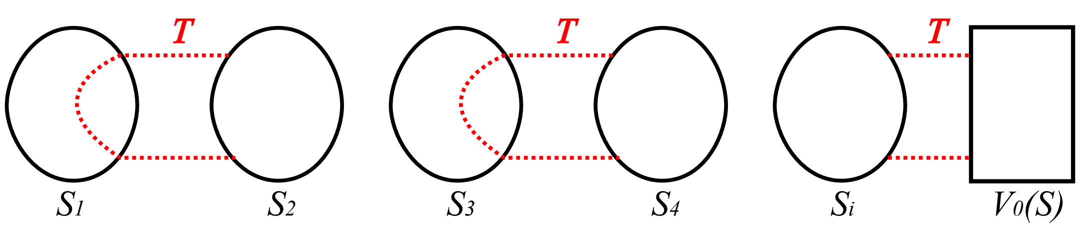

FIGURE E.4. If we throw in T what happens to the unbalanced and balanced components? [[LABEL F:1027Figure1]]

Let $\boldsymbol{\pi}_{\mathrm{b}}(S) = \{B_1, B_2, \ldots, B_k\}$, where $B_j = V(S_j)$ as seen in Figure E.4. Suppose that every balanced component $S{:}B_i$ of S is positive. Looking at the components C_i of $S \cup T$ with $T_i := E(T \cap C_i)$ (the edge set of $T \cap C_i$). Any of these C_i that contains an unbalanced component of S is unbalanced. In Σ/S, C_i becomes loose edges and at least one half edge $\iff$ T_i had an edge with an endpoint outside $V_0(S)$ $\iff$ $N(C_i) \subseteq V_0(S)$.

Table E.1 shows how T affects the components of $\Sigma \setminus T$ and $(\Sigma \setminus T)/S$. There are four cases to examine. Notice that there is a natural bijection between C and C' in Case III. □

If we zoom in our attention to the specific situations we can discuss them a little more clearly with visual aid.

In Figure E.5 we can see that anything connected to an unbalanced component will make an unbalanced component of T trivially.

If we have the situation in Figure E.6, a negative T edge in a balanced component makes the set unbalanced. This is because in the contraction of S this negative edge makes a

The effect of T		
	on $B_i \subseteq V(\Sigma)$	**on** $B_i \in V(\Sigma/S)$
Case I	Connects B_i to $V_0(S)$ so $B_i \subseteq V_0(S \cup T)$.	T makes a half edge at B_i, so $B_i \in V_0(\Sigma/S;T)$.
Case II	T edges are within $V_0(S)$.	T is a loose edge.
Case III	T edges connect up one or more B_i into an unbalanced component of $S \cup T$. Also $B_i \subseteq V_0(S \cup T)$.	T forms an unbalanced component of T in Σ/S. Also these $B_i \in V_0(\Sigma/S;T)$.
Case III	T connects one or more B_i into a balanced component C of $S \cup T$, making C a vertex of $\Sigma/(S \cup T)$. Then $C \subseteq V$, with $C = \bigcup B_i \in \boldsymbol{\pi}_{\mathrm{b}}(S)$, so $C \in V(\Sigma/(S \cup T))$	T connects one or more vertices of Σ/S into a balanced component of C' of T in Σ/S. Then C' is a vertex of $(\Sigma/S)/T$. Then $C' \subseteq V/S$, with $C' = \{B_i \mid B_i \in \boldsymbol{\pi}_{\mathrm{b}}(S)$ and $B_i \subseteq C\}$, so $C' \in V((\Sigma/S)/T))$, where $C' = \{B_i \in \boldsymbol{\pi}_{\mathrm{b}}(S) \mid B_i \subseteq C$ in $\Sigma\}$

TABLE E.1. The effect of T on balanced components in Σ and Σ/S.

[[LABEL Tb:1027T]]

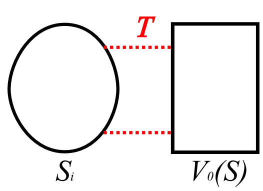

FIGURE E.5.

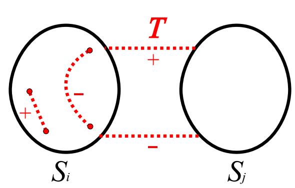

FIGURE E.6.

negative circle in Σ/S with S, and hence unbalanced in $S \cup T$. A positive T edge preserves that S_j is balanced.

Lemma E.4. *[[LABEL L:1027Lemma2]] Let S be balenced in Σ and $T \subseteq E \backslash S$. Then $S \cup T$ is balanced in $\Sigma \iff T$ is balanced in Σ/S.*

Suppose we have the situation in Figure E.7. If the loop (T circle) is negative then it gives an unbalanced component in the contraction. One should note that switching does not change the sign of a circle. Also, contracting a proper subset of circle edges does not change the sign of the circle.

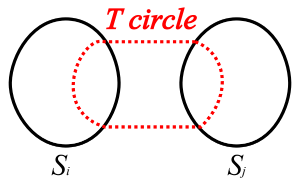

FIGURE E.7.

F. CLOSURE AND CLOSED SETS

[[LABEL 2.closure]]

Closure in a signed graph, while fundamentally similar to that in a graph (very similar, according to Proposition F.3), is certainly more complicated. Oct 29: Peter Cohen and Thomas Zaslavsky

F.1. **Closure operator.** [[LABEL 2.closure.operator]]

The best way to define the closure of an edge set in Σ is in two steps. First we define an operator on balanced sets, then we use it to define the closure of any edge set. Notice that our definition of closure in a signed graph generalizes the characterization of graph closure in Theorem **??** rather than the definition of graph closure. There is a generalization of the latter definition (see Theorem F.6), and it is important, but it is not as simple.

Definition F.1. [[LABEL D:1029closures]] The *balance-closure* of $T \subseteq E$ is

$$\mathrm{bcl}(T) := T \cup \{e \in T^c : \exists \text{ a positive circle } C \subseteq T \cup e \text{ such that } e \in C\} \cup E_0(\Sigma),$$

where $E_0(\Sigma)$ is the set of loose edges in Σ. (The name is not "balanced closure"; $\mathrm{bcl}(T)$ need not be balanced—but see Lemma F.2.)

The *closure* of an edge set $S \subseteq E$ is

$$\mathrm{clos}(S) := (E{:}V_0(S)) \cup \bigcup_{i=1}^{k} \mathrm{bcl}(S_i) \cup E_0(\Sigma),$$

where $S_1, \ldots, S_k$ are the balanced components of S and $V_0(S)$ is the vertex set of the union of all unbalanced components of S, that is, $V_0(S) = V \setminus (B_1 \cup \cdots \cup B_k)$. We can restate this directly in terms of $\boldsymbol{\pi}_{\mathrm{b}}(S)$ (since $S_i = S{:}B_i$ for $B_i \in \boldsymbol{\pi}_{\mathrm{b}}(S)$) as

$$\mathrm{clos}(S) := (E{:}V_0(S)) \cup \bigcup_{B \in \boldsymbol{\pi}_{\mathrm{b}}(S)} \mathrm{bcl}(S{:}B_i) \cup E_0(\Sigma),$$

which has the advantage of not implying that k is finite. In the definitions of the closure, the union with $\cup E_0(\Sigma)$ is only necessary in case $k = 0$, i.e., $\boldsymbol{\pi}_{\mathrm{b}}(S) = \varnothing$.

Lemma F.1. *[*[LABEL L:1029bclpositive]] *If* $T \subseteq E^+(\Sigma)$, *then* $\mathrm{bcl}(T) = \mathrm{clos}_{\Sigma^+}(T)$, *the graph closure of* T *in the positive subgraph of* Σ.

Proof. First suppose $\Sigma = +\Gamma$, all positive. Then, comparing the definition of bcl in Σ with the second definition of clos_Γ in Definition **??**, we see they are the same.

A positive circle contained in $T \cup e$ has sign $\sigma(e)$; thus only a positive edge can be in $\mathrm{bcl}\, T$. That means $\mathrm{bcl}_\Sigma T = \mathrm{bcl}_{+\Sigma^+} T = \mathrm{clos}_{\Sigma^+} T$. □

Lemma F.2. *[*[LABEL L:1029bclbalance]] *If* T *is balanced, then* $\mathrm{bcl}(T)$ *is also balanced, and furthermore* $\mathrm{bcl}(\mathrm{bcl}\, T) = \mathrm{bcl}(T) = \mathrm{clos}(T)$.

Proof. The main step is to assume by switching Σ that T is all positive. Then we apply Lemma F.1. Since $\operatorname{bcl} T$ is again all positive, it is balanced, and that means it was balanced before switching. Furthermore, as $\operatorname{bcl} T$ is all positive, $\operatorname{bcl}(\operatorname{bcl} T) = \operatorname{clos}_{\Sigma^+}(\operatorname{clos}_{\Sigma^+} T) = \operatorname{clos}_{\Sigma^+} T = \operatorname{bcl} T$ by idempotency of graph closure.

The equation of $\operatorname{bcl} T$ and $\operatorname{clos} T$ is obvious from the definition of closure. □

Note that we have not said balance-closure is an abstract closure operator. In fact, it is not. It is increasing and isotonic but it is not idempotent. (*Exercise*: Find a counterexample. It must be unbalanced, of course.)

It's easy to see that balance-closure is a direct generalization of graph closure, as we state formally in the next result (an obvious corollary of Lemma F.1).

Proposition F.3. *[[LABEL P:1029ordinaryclosure]] If Γ is an ordinary graph, then* $\operatorname{clos}_{+\Gamma}(S) = \operatorname{bcl}_{+\Gamma}(S) = \operatorname{clos}_{\Gamma}(S)$.

An interesting observation is that the union of the balance-closures of subsets with no common vertices is the same as the balance-closure of the union of the subsets. That is,

$$\bigcup_{i=1}^{k} \operatorname{bcl}(S_i) = \operatorname{bcl}\big(\bigcup_{i=1}^{k} S_i\big)$$

if the vertex sets $V(S_i)$ are pairwise disjoint. The sets S_i themselves need not be balanced. The reason for this is that balance-closure acts within the components of an edge set. We can formalize this as the first statement in the next lemma.

Lemma F.4. *[[LABEL L:1029balptn]] For an edge set S, whether balanced or not,* $\boldsymbol{\pi}(\operatorname{bcl} S) = \boldsymbol{\pi}(S)$ *and* $\boldsymbol{\pi}_{\mathrm{b}}(\operatorname{clos} S) = \boldsymbol{\pi}_{\mathrm{b}}(\operatorname{bcl} S) = \boldsymbol{\pi}_{\mathrm{b}}(S)$.

Proof. Set $\boldsymbol{\pi}(S) = \{B_1, \ldots, B_k, C_1, \ldots, C_l\}$, where $S{:}B_i$ is balanced while $S{:}C_j$ is unbalanced.

All the sets $\operatorname{bcl}(S{:}B_i)$ in the definition of $\operatorname{bcl} S$ are balanced (by Lemma F.2) and connected; each set $\operatorname{bcl}(S{:}C_j)$ is connected and unbalanced (because it contains the unbalanced component $S{:}C_j$ of S); and these are the components of $\operatorname{bcl} S$. Thus, the partition due to $\operatorname{bcl} S$ is the same as that due to S, and the same is true for the balanced partial partition.

Each $E{:}C_j$ is unbalanced, because it contains $S{:}C_j$. Thus, every component of $E{:}V_0(S)$ is unbalanced, so the balanced components of $\operatorname{clos}(S)$ are the $\operatorname{bcl}(S{:}B_i)$. Therefore, $\boldsymbol{\pi}_{\mathrm{b}}(\operatorname{clos} S) = \boldsymbol{\pi}_{\mathrm{b}}(S)$. □

Proposition F.5. *[[LABEL P:1029closureclosure]] The operator* clos *on subsets of $E(\Sigma)$ is an abstract closure operator.*

Proof. The definition makes clear that $S \subseteq \operatorname{clos} S$ and that $\operatorname{clos} S \subseteq \operatorname{clos} T$ when $S \subseteq T$. What remains to be proved is that $\operatorname{clos}(\operatorname{clos}(S)) = \operatorname{clos}(S)$.

As before, let $\boldsymbol{\pi}_{\mathrm{b}}(S) = \{B_1, \ldots, B_k\}$, so $S{:}B_i$ is balanced. Then $\boldsymbol{\pi}_{\mathrm{b}}(\operatorname{clos} S) = \boldsymbol{\pi}_{\mathrm{b}}(S)$ so also $V_0(\operatorname{clos} S) = V_0(S)$. Thus,

$$\begin{aligned}\operatorname{clos}(\operatorname{clos} S) &= (E{:}V_0(\operatorname{clos} S)) \cup \bigcup_{i=1}^{k} \operatorname{bcl}((\operatorname{clos} S){:}B_i)\\ &= (E{:}V_0(S)) \cup \bigcup_{i=1}^{k} \operatorname{bcl}((\operatorname{bcl} S){:}B_i)\end{aligned}$$

$$= (E{:}V_0(S)) \cup \bigcup_{i=1}^{k} \operatorname{bcl}(S{:}B_i)$$
$$= \operatorname{clos} S. \qquad \square$$

Oct 29: T. Zaslavsky

Closure via frame circuits.

We have defined closure in terms of induced edge sets and balanced circles (through the balance-closure); but we also want a definition in terms of circuits, analogous to that of closure in an ordinary graph.

Theorem F.6. *[*[LABEL T:1029cctclosure]*]For $S \subseteq E$ and $e \notin S$, $e \in \operatorname{clos} S$ iff there is a frame circuit C such that $e \in C \subseteq S \cup e$.*

Proof. We treat a half edge as if it were a negative loop, since they are equivalent in what concerns either closure or circuits.

[**The proof needs figures for the cases.**]

Necessity. We want to prove that if $e \in \operatorname{clos} S$, then a circuit C exists. There are three cases depending on where the endpoints of e are located.

Case 0. A trivial case is where e is a loose edge. Then $e \in \operatorname{clos} S$ and $C = \{e\}$.

Case 1. Suppose e has its endpoints within one component, S'. Then there is a circle C' in $S' \cup e$ that contains e. If S' is balanced, then $e \in \operatorname{bcl} S'$ so there exists a positive C', which is the circuit C. In general, if C' is positive it is our circuit C. (This includes the case of a positive loop e, where $C = \{e\}$.)

Let us assume, therefore, that S' is unbalanced and C' is negative. In S' there is a negative circle C_1. If e is an unbalanced edge at v, there is a path P in S' from v to C_1; then $C = C_1 \cup P \cup e$ is the circuit we want. If e is a balanced edge, it is a link $e{:}vw$ contained in the negative circle C'. There are three subcases, depending on how many points of intersection C' has with C_1. If there are no such points, take a minimal path P connecting C' to C_1 and let $C = C_1 \cup P \cup C'$. If there is just one such point, $C = C_1 \cup C'$. If there are two or more such points, take P to be a maximal path in C' that contains e and is internally disjoint from C_1. Then $P \cup C_1$ is a theta graph in which C_1 is negative; hence one of the two circles containing P is positive, and this is the circuit C.

Case 2. Suppose e has endpoints in two different components, S' and S''. For e to be in the closure, it must be in $E{:}V_0$. Hence, S' and S'' are unbalanced. Each of them contains a negative circle, C' and C'' respectively, and there is a connecting path P in $S \cup e$ which contains e. Then $C' \cup P \cup C''$ is the desired circuit.

Sufficiency. Assuming a circuit C exists, we want to prove that $e \in \operatorname{clos} S$. Again there are three cases, this time depending on C and its relationship with e.

Case 0. C is balanced. Then $e \in \operatorname{bcl} S \subseteq \operatorname{clos} S$.

Case 1. C is unbalanced and e is not in the connecting path. Let C_1, C_2 be the two negative circles and P the connecting path of C, and assume $e \in C_1$. Since $C \setminus e$ is connected, it lies in one component S' of S. Thus, $C_2 \subseteq S'$, whence S' is unbalanced. It follows that $e \in E{:}V_0 \subseteq \operatorname{clos} S$.

Case 2. C is unbalanced and e is in the connecting path. With notation as in Case 1, now $C \setminus e$ has two components, one containing C_1 and the other containing C_2. The components of S that contain C_1 and C_2 are unbalanced. (There may be one such component or two, depending on whether C_1 and C_2 are connected by a path in S.) Therefore, e has both endpoints in V_0, so again, $e \in E{:}V_0 \subseteq \operatorname{clos} S$. □

Dec 8: Yash Lodha

With closure in terms of frame circuits we can offer another characterization of a partial balancing edge.

Proposition F.7 (Balancing Edge Properties)**.** *[[LABEL P:1208BE]] In a signed graph Σ let S be an edge set and e an edge not in S. The following relationships between e and S are equivalent:*

(i) $e \in \operatorname{clos}(S)$.
(ii) There is a frame circuit C such that $e \in C \subseteq S \cup e$.
(iii) $b(S \cup e) = b(S)$.
(iv) e is not a partial balancing edge of $S \cup e$.

Proof. Parts (iii) and (iv) are equivalent by the definition of a partial balancing edge. The equivalence of (i) and (ii) is Theorem F.6. What we need to prove is the equivalence of (ii) and (iii). We treat a half edge as a negative loop.

Case 1. $V(e) \subseteq V(S_1)$ where S_1 is a component of S. If S_1 is unbalanced, it has a negative circle C_i and there is a path in S_1 joining the endpoints of e.

[figures go here]

Therefore if S_1 is unbalanced, C exists as in (ii).(This will be clear from the images) Also $b(S \cup e) = b(S)$. If S_1 is balanced, then either every circle $e \in C \subseteq S_1 \cup e$ is negative or every such circle is positive. This is because we can switch so that all edges in S_1 are positive and so the resulting sign of e is the sign of all the circles $e \in C \subseteq S_1 \cup e$. Therefore $b(S \cup e) = b(S)$ $\iff$ e is positive after switching $\iff$ there exists a frame circuit $e \in C \subseteq S_1 \cup e$ which will be a positive circle.

Case 2. e is an isthmus of $S \cup e$, joining components S_1 and S_2.

[diagram]

If S_1 and S_2 are unbalanced, then e is in a circuit handcuff of $S_1 \cup S_2 \cup e$, and also $b(S \cup e) = b(S)$ because $S_1 \cup S_2 \cup e$ is unbalanced.

[diagram]

Suppose S_2 is unbalanced.

[diagram]

Then e is not in a frame circuit **[I have to check the cases]**, and $b(S \cup e) = b(S) - 1$ (since S_1 is unbalanced implies that one balanced and one unbalanced component, S_2 and S_1 become one unbalanced component $S_1 \cup S_2 \cup e$.)

Case 3. e is a half edge. Treat this as a negative loop, which is Case 1.

Case 4. e is a loose edge. Then $b(S) = b(S \cup e)$ abd $e \in \{e\}$ which is a circuit.

Hence the proposition is proved. □

Oct 29: Peter Cohen and Thomas Zaslavsky

F.2. **Closed sets.** [[LABEL 2.closure.closed]]

Now we look at the closed sets themselves. The first fact is that they form a lattice.

Definition F.2. [[LABEL D:1029lattices]] The lattice of closed sets of Σ is

$$\operatorname{Lat} \Sigma := \{S \subseteq E \mid S \text{ is closed}\}.$$

The semilattice of closed, balanced sets is

$$\mathrm{Lat}^{\mathrm{b}}\,\Sigma := \{S \subseteq E \mid S \text{ is closed and balanced}\}.$$

(Be careful! By closed, balanced edge sets, we mean edge sets that are both closed and balanced. This is completely different from sets that are balance-closed, which need not even be balanced.)

We haven't yet proved that $\mathrm{Lat}\,\Sigma$ is a lattice.

Proposition F.8. *[[LABEL P:1029lattices]]* $\mathrm{Lat}\,\Sigma$ *is a lattice with* $S \wedge T = S \cap T$*, and* $S \vee T = \mathrm{clos}(S \cup T)$.

$\mathrm{Lat}^{\mathrm{b}}\,\Sigma$ *is a meet semilattice with* $S \wedge T = S \cap T$*. It is an order ideal in* $\mathrm{Lat}\,\Sigma$ *(that is, every subflat of a flat in* $\mathrm{Lat}^{\mathrm{b}}\,\Sigma$ *is also in* $\mathrm{Lat}^{\mathrm{b}}\,\Sigma$*).*

$\mathrm{Lat}\,\Sigma$ *is ranked by the rank function* $\mathrm{rk}(S) = n - b(S)$.

Proof. □

In $\mathrm{Lat}\,\Sigma$ there is one maximal closed set: E. Its rank is $n - b(\Sigma)$. All maximal closed, balanced sets have rank $n - c(\Sigma)$. These facts are proved in Section **??**; they are true because, in the matrix, each vertex allows one potential dimension, while each balanced component will have a row dependence relation, reducing the rank by 1. **[This should be proved somewhere and cross-referenced to where it's proved. – TZ]**

Oct 31: Yash Lodha

F.3. **Signed partial partitions.** [[LABEL 2.sppartitions]]

Now we come to a new way of looking at the closed sets of a signed graph: it's the signed-graph version of partitions of the vertex set. Two refinements are required: we need partial partitions, and we need signed blocks.

F.3.1. *Partial partitions.*

A partial partition is a partition of any subset of X. Partitions are found all over combinatorics and other mathematics but partial partitions are unjustly rare. We shall have much to say about them.

Definition F.3. [[LABEL Df:1031ppartition]] A *partial partition* of a set V is defined as $\pi = \{B_1, B_2, \ldots, B_k\}$ where $B_i \subseteq V$, each B_i and B_j are pairwise disjoint, and each $B_i \neq \varnothing$. The B_i's are called the *blocks* or sometimes *parts* of π. The *support* $\mathrm{supp}\,\pi$ is the union of the blocks. The set of all partial partitions of V is written $\Pi^{\dagger}_V$. Note that $\varnothing$ is a partial partition of V—the unique one with no blocks.

A partition of V is therefore a partial partition with the additional condition that $[n] = \bigcup_{i=1}^{k} B_i$. The refinement ordering of the set Π_V of partitions (see Section **??**) clearly agrees with the refinement ordering of partial partitions.

The set of partial partitions of $[n]$ is denoted by $\Pi^{\dagger}_n$. It is partially ordered in the following way: For two partial partitions π and τ, we define $\pi \leq \tau$ if each block of τ is a union of blocks of π. We say π *refines* τ—though the support of π need not be contained in that of τ. The refinement ordering makes $\Pi^{\dagger}_n$ a poset. This poset is a geometric lattice. In fact:

Proposition F.9. *[[LABEL P:1031ppartition]]* $\Pi^{\dagger}_n \cong \Pi_{n+1}$.

Proof. A partial partition $\pi = \{B_1, B_2, \ldots, B_k\}$ naturally corresponds to

$$\pi' := \{B_1, B_2, \ldots, B_k, B_0\}, \text{ where } B_0 := [n+1] \setminus \bigcup_{i=1}^{k} B_i .$$

It is easy to see that this correspondence is order preserving and bijective, hence a poset isomorphism. □

The minimum element of $\Pi_n^\dagger$ is $\hat{0}_n := \{\{i\} : i \in [n]\}$. Its maximum element is the empty partial partition $\varnothing$.

F.3.2. *Signed partial partitions.*

Suppose we have a set B and a sign function $\tau : B \to \{+,-\}$. The pair (B,τ) is a *signed set*. Two signed sets (B,τ_1) and (B,τ_2) are *equivalent* if there is a sign $\varepsilon \in \{+,-\}$ such that $\varepsilon\tau_1 = \tau_2$. We write the equivalence class of (B,τ) with square brackets: $[B,\tau]$. (In a way, an equivalence class is a kind of switching class but defined on vertices rather than edges.)

Definition F.4. [[LABEL Df:1031sppartition]] A *signed partial partition* of V is a set $\theta = \{[B_i,\tau_i]\}_{i=1}^k$, where $\pi(\theta) := \{B_i\}_{i=1}^k$ is a partial partition of V, called the underlying partial partition, and τ_i is a function $B_i \to \{+,-\}$.

The *support* of θ is $\operatorname{supp}(\theta) := \operatorname{supp}(\pi(\theta)) = \bigcup_{i=1}^k B_i$. The set of all signed partial partitions of V is denoted by $\Pi_V^\dagger(\{+,-\})$, or for short, $\Pi_V^\dagger(\pm)$.

Signed partial partitions are partially ordered in the following way: $\theta \leq \theta'$ if $\pi(\theta) \leq \pi(\theta')$ and, whenever $B_i \subseteq B'_j$, we have $\tau_i = \varepsilon\tau_j|_{B_i}$ for some sign ε.

The poset of signed partial partitions of V is denoted by $\Pi_V^\dagger(\{+,-\})$, or for short, $\Pi_V^\dagger(\pm)$. In particular, the set of signed partial partitions of $[n]$ is written $\Pi_n^\dagger(\pm)$. It is a poset, in fact a geometric lattice (as we shall see later); it is the Dowling lattice of the sign group as originally defined by Dowling (1973b).

(Some people think of a signed partial partition as a sort of partially signed partition $\{[B_1,\tau_1],\ldots,[B_k,\tau_k],B_0\}$, where $\{B_1,\ldots,B_k,B_0\}$ partitions $[n]\cup\{0\}$, having a special "zero block" $B_0 \ni 0$ that is not signed. I find this artificial, since the "zero block" is completely different from all other blocks. However, it may have its uses.)

We define a function $\Theta_{\mathrm{b}} : \operatorname{Lat}(\Sigma) \to \Pi_V^\dagger(\pm)$, which will be an order preserving injection.

A *potential function* for $T \subseteq E(\Sigma)$ is a function $\rho : V \to \{+,-\}$ such that

$$\sigma(e_{vw}) = \rho(v)^{-1}\rho(w) \text{ for every edge in } T.$$

(One can equivalently define ρ as a switching function that makes T all positive; but that is not a definition which generalizes to gain graphs; see Chapter III.) If Σ is connected ρ is unique up to negation. If $B_i \in \boldsymbol{\pi}_{\mathrm{b}}(S)$ is the vertex set of a balanced component of (V,S) then ρ is what we want for τ_i. So $\Theta_{\mathrm{b}}(S)$ sends S to $\{B_i,\tau_i\}$ where B_i are the balanced components of $\Sigma|S$ and $\tau_i = \rho(S{:}B_i)$.

Note that we can actually define $\Theta_{\mathrm{b}} : \mathcal{P}(E) \to \Pi_v^\dagger(\pm)$, but it will not be an injection.

Now define

$$\Pi^\dagger(\Sigma) := \{\Theta_{\mathrm{b}}(S) \mid S \subseteq E\},$$

which is a subposet of $\Pi_V^\dagger(\pm)$.

Theorem F.10. *[[LABEL T:1031lattices]]* $\Theta_{\mathrm{b}} : \operatorname{Lat}(\Sigma) \to \Pi^\dagger(\Sigma)$ *is a poset isomorphism.*

Lemma F.11. *[[LABEL L:1031ppartition]]* $\Theta_{\mathrm{b}}(S) = \Theta_{\mathrm{b}}(\operatorname{clos}(S))$.

Proof. The partition $\pi(\Theta_{\rm b}(S))$ is unchanged by taking the closure: $\pi(\Theta_{\rm b}(S)) = \pi(\Theta_{\rm b}(\operatorname{clos}(S)))$ since $\boldsymbol{\pi}_{\rm b}(S) = \boldsymbol{\pi}_{\rm b}(\operatorname{clos}(S))$ by a previous lemma. [**(will put in the name)**] The potential function depends on a spanning tree of $S{:}B_i$ which is still a spanning tree in the closure. So it is clear that $\Theta_{\rm b}(S) = \Theta_{\rm b}(\operatorname{clos}(S))$. Hence the lemma is proved. □

Proof of Theorem F.10. The theorem follows easily from the lemma. □

Example F.1. [[LABEL X:1031dowling]] $\Pi^\dagger(\pm K_n^\circ) \cong \Pi_v^\dagger(\pm)$.

Proof. We proved that $\Pi^\dagger(\pm K_n^\circ) \cong \operatorname{Lat}(\pm K_n^\circ)$. So it will suffice to prove that $\operatorname{Lat}(\pm K_n^\circ) \cong \Pi_v^\dagger(\pm)$.

To prove this we need to look at the flats of $\operatorname{Lat}(\pm K_n^\circ)$. These flats look like $(E{:}X) \cup A$ where $X \subseteq V(\pm K_n)$ and A is a balanced, balance-closed set of $E{:}X^c$. Clearly, the components of the balanced closed set give us a partial partition of the vertex set $V(\Sigma)$ and the signs of each block of this partial partition are exactly the signs that make the balanced, balanced closed set positive. This construction/map gives us an element of the signed partial partition lattice of the vertex set of Σ. This mapping is precisely the function $\Theta_{\rm b}$ defined above.

We will first show that it is order preserving. Let A, B be two flats of $\pm K_n^\circ$ such that $A \leq B$. Let P_1, P_2 be the elements of $\Pi_V^\dagger(\pm)$ be the image of A, B respectively in our map defined above. That $\pi(P_1) \leq \pi(P_2)$ is obvious from the fact that $A \leq B$ because $\pi(A), \pi(B)$ are the underlying partial partitions of the vertex set of Σ with blocks as the vertex sets of the balanced components of $A{:}V$ and $B{:}V$. Given block C_i of $\pi(P_1)$ which is contained in a block D_j of $\pi(P_2)$, it is clear that the edge set $E(B{:}\operatorname{supp}(D_j))$ contains $E(A{:}\operatorname{supp}(C_i))$ so the signs associated with the vertices $\operatorname{supp}(D_j)$ must be switching equivalent to the signs associated with $\operatorname{supp}(D_j)$ restricted to $\operatorname{supp}(C_i)$. Therefore our map is order preserving.

We now show that our map is an injection. For any two different flats A, B we first present the case where the components of $A{:}V$ and $B{:}V$ are different in which case it is obvious that the partial partitions associated with these flats will have different supports. In case of these support being the same we observe that the edge sets of a balanced component of $A{:}V$ and one of $B{:}V$ having the same vertex set have different edge sets, giving us different switching sets for the same vertex sets because had these switching sets been the same, because of balanced closure these flats would be the same. [**THAT SENTENCE NEEDS REWRITING. IT'S IMPENETRABLE.**] So we get different signed partial partitions in the image.

Our map is surjective because the method used to obtain a signed partial partition is reversible. We show an example of such a reverse map. Given a signed partial partition $[A_i, \tau_i]$, for the vertices $a, b \in A_i$ if the signs of a, b are the same, we connect them with a positive edge, and if the signs are opposite them we connect them with a negative edge. And if the sign on a is positive, we add the positive loop at a, and a negative loop if the sign is negative. This way we can obtain the closed set associated with our signed partial partition.

Hence the bijection is established. □

F.3.3. *Signed partitions.* 2014 Nov 16: T.Z.

Some partial partitions of V are partitions; they form the sublattice $\Pi_V \subseteq \Pi_V^\dagger$. Some signed partial partitions are likewise partitions; they form a meet subsemilattice,

$$\Pi_V(\{\pm\}) := \{\theta \in \Pi_V^\dagger(\{\pm\}) : \operatorname{supp}\theta = V\}.$$

There is a corresponding subsemilattice of $\Pi^\dagger(\Sigma)$; it is

$$\Pi(\Sigma) := \{\theta \in \Pi^\dagger(\Sigma) : \operatorname{supp}\theta = V(\Sigma)\}.$$

Recall from [**WHERE? WRITE THAT BIT.**] that $\mathrm{Lat}^{\mathrm{b}}\,\Sigma$ is the set of balanced closed sets.

Theorem F.12. [[LABEL T:1103sgd ptns]] *The natural bijection* $\mathrm{Lat}\,\Sigma \leftrightarrow \Pi^\dagger(\Sigma)$ *restricts to a bijection* $\mathrm{Lat}^{\mathrm{b}}\,\Sigma \leftrightarrow \Pi(\Sigma)$.

Proof. [**NEEDS PROOF**] □

Therefore we may refer to $\Pi(\Sigma)$ as the *lattice of signed partitions of* Σ.

[**ADD the fact (if true) that it's a lattice iff Σ is balanced or exceptions—as an exercise.**]

Nov 3: Jackie Kaminski

F.4. **Examples of signed graphs and closed sets.** [[LABEL 2.closure.examples]]

We have now built up definitions and some machinery about closed sets, balanced edge sets, and closed, balanced edge sets. It will be good to know what these sets are for certain graphs and types of graphs. This information is presented as both a reference and a tool to help the reader build up his or her intuition.

Throughout, $\Gamma = (V, E)$ is an ordinary graph without loops. We recall that $\Gamma^\circ = (V, E^\circ)$ is the unsigned graph with a loop at every vertex, in contrast to $+\Gamma^\circ$ which is a signed graph with a negative loop at each vertex. For $B \subseteq V$, by K_B we mean the complete graph on the vertex set B.

Remember that an edge set is balanced if it has no negative circles or half edges (Definition **??**), that the balance-closure of S is

$$\mathrm{bcl}(S) := S \cup \{e \notin S : \exists\, C \in \mathcal{B}(\Sigma) \text{ with } e \in C \subseteq S \cup e\} \cup \{\text{all loose edges of } \Sigma\},$$

and the closure of S is

$$\mathrm{clos}(S) := (E{:}V_0(S)) \cup \bigcup_{i=1}^{k} \mathrm{bcl}(S_i),$$

where V_0 is the vertex set of the union of the unbalanced components of S and $S_1, \ldots, S_k$ are the balanced components of S. (See Section F.1).

Example F.2. [[LABEL X:clos-pmKnfull]] $\pm K_n^\circ$ (the complete signed graph [not to be confused with a signed complete graph]).

- *Balanced edge sets*: Any switching of a positive edge set of K_n. We note that this is a little imprecise; what we mean is to take any switching of any edge set in $+K_n$. Then for an edge e in this switching if e is positive, take the edge $+e \in \pm K_n^\circ$, otherwise take $-e \in \pm K_n^\circ$.
- *Closed, balanced sets*: Take $\pi \in \Pi_n$, take $E(\pi) := \bigcup_{B\in\pi} E(K_B)$, and assign signs in a balanced way (as above). Notice that E_π, as the union of pairwise disjoint complete graphs, is a closed set in K_n.
- *Closed sets*: To create a closed set S, take any $W \subseteq V$ and a partition π of $V \setminus W$ and let $S := E(\pm K_W^\circ) \cup \bigcup_{B\in\pi}(K_B, \sigma_B)$, where (K_B, σ_B) denotes the complete graph on vertex set B with a balanced signature σ_B.

Example F.3. [[LABEL X:clos-pmGfull]] $\pm\Gamma^\circ$ (the full signed expansion of a graph).

- *Balanced edge sets*: Any switching of an edge set in $+\Gamma$, with the same technical clarification as in the $\pm K_n^\circ$ case.
- *Closed, balanced sets:* A closed edge set in Γ, signed in a balanced way (i.e., take a closed edge set $S \subseteq E$, and take any switching of $+S$).
- *Closed sets*: To create a closed set S, take $W \subseteq V$, and take S^* to be any closed set in $\Gamma \setminus W$. Sign S^* in a balanced way. Then $S := E(\pm[\Gamma{:}W]^\circ) \cup S^*$ is a closed set.

Example F.4. [[LABEL X:clos-full]] Σ° (the filled version of a signed graph Σ).

This generalizes the previous examples.

- *Balanced edge sets*: The balanced edge sets of Σ° are precisely the balanced sets in Σ.
- *Closed, balanced sets*: The closed, balanced edge sets of Σ° are precisely the closed, balanced sets in Σ.
- *Closed sets*: For any $W \subseteq V$, take $E(\Sigma^\circ{:}W) \cup$ a balanced closed set in $\Sigma \setminus W$. (This construction is obvious from the definition of closed sets. A closed set has two parts, an unbalanced part which is the subgraph induced by some vertex set, and a balanced part, in the complementary vertex set. Neither of these parts needs to be connected; also, either one may be void.)

Example F.5. [[LABEL X:clos-pmkn]] $\pm K_n$ (the complete signed link graph).

This is just slightly more complicated than $\pm K_n^\circ$.

- *Balanced edge sets*: The same as in $\pm K_n^\circ$. (Any switching of a positive edge set of K_n.)
- *Closed, balanced sets*: The same as in $\pm K_n^\circ$. (Take $\pi \in \Pi_n$, take $E(\pi)$, and assign signs in a balanced way. In other words, it's the union of pairwise disjoint, balanced complete graphs on subsets of V.)
- *Closed sets*: Similar to $\pm K_n^\circ$. To create a closed set S, take any $W \subseteq V$ and take a partition π of $V \setminus W$ and let $S := E(\pm K_W) \cup \bigcup_{B\in\pi}(K_B, \sigma_B)$, where $|W| \neq 1$ in order to avoid duplication in the construction. (When W is a singleton we get a closed set but it is the same as that obtained through replacing W by $\varnothing$ and adding the singleton set W to π.)

Example F.6. [[LABEL X:clos-pmG]] $\pm\Gamma$ (the signed expansion of a graph).

This is similar to $\pm\Gamma^\circ$, but again, a bit more complicated because there are no loops to identify vertices.

- *Balanced edge sets*: Any switching of an edge set in $+\Gamma$, with the standard technical clarification.
- *Closed, balanced sets*: Take a closed edge set in Γ and sign it in a balanced way (i.e., take a closed edge set $S \subseteq \Gamma$, and choose any switching of $+S$).
- *Closed sets*: To create a closed set S, take W to be any subset of V such that W is not stable (that is, $E{:}W \neq \varnothing$). Take S^* to be a subset of $E(\Gamma \setminus W)$ and sign $S*$ in a balanced way. Then $S = E(\pm\Gamma{:}W) \cup S^*$ is a closed set.

Example F.7. [[LABEL X:clos-all+]] $+\Gamma$ (an all-positive graph).

- *Balanced edge sets*: Any edge set of Γ.
- *Closed, balanced sets*: Any closed edge set of the unsigned graph Γ.
- *Closed sets*: The same as the closed, balanced sets.

Example F.8. [[LABEL X:clos-all+full]]

$+\Gamma^\circ$ (a full all-positive graph).

- *Balanced edge sets*: Any edge set in Γ.
- *Closed, balanced sets*: Any closed set in Γ.
- *Closed sets*: This is similar to Σ°. For any $W \subseteq V$, take $(E^\circ{:}W) \cup$ a closed set of $\Gamma \setminus W$. The set W is identifiable as the set of vertices at which there are unbalanced edges, so any different choice of W results in a different closed set.

  There is another technique that will work here. We could consider the unsigned graph join, $\Gamma \vee K_1$ (Γ plus one new vertex adjacent to every vertex of Γ), then look at the various sets in $\Gamma \vee K_1$ (keeping in mind that being closed has a different definition for $\Gamma \vee K_1$), and then pull back the results to $+\Gamma^\circ$.

Example F.9. [[LABEL X:clos-all-]] $-\Gamma$ (an all-negative graph).

- *Balanced edge sets*: The bipartite edge sets, which are exactly the edge sets where every circle has even length.
- *Closed, balanced sets*: Take a connected partition $\pi \in \Pi(\Gamma)$, and in each block $B \in \pi$, take a maximal cut. Taking any cutset in B will still produce a closed, balanced set; however, taking only maximal cuts has the nice property that for $\pi \in \Pi(\Gamma)$, and S a set consisting of a maximal cut in each block of π, then $\pi(S) = \pi$.
- *Closed sets*: Each closed set has the form $S = E(-\Gamma{:}W) \cup$ a closed, balanced set in $-(\Gamma \setminus W)$, where $W \subseteq V$ is such that $\Gamma{:}W$ has no bipartite components. Notice that if we took a vertex subset W such that $\Gamma{:}W$ had a bipartite component, we would still get a closed set but in more than one way, since the same set is generated by a smaller vertex subset, namely, the one obtained by removing from W the vertices of bipartite components of $\Gamma{:}W$.

Example F.10. [[LABEL X:all-Kn]] $-K_n$ (the all-negative, or antibalanced, complete graph).

This is simpler than $-\Gamma$, because any cut in K_n is a complete bipartite graph.

- *Balanced edge sets*: The bipartite edge sets.
- *Closed, balanced sets*: The union of pairwise-disjoint complete bipartite subgraphs in V.
- *Closed sets*: Take an induced edge set $E{:}W$ together with disjoint complete bipartite graphs in $V \setminus W$. We should require $|W| \neq 1, 2$ in order that each closed set arise uniquely, for if $|W| = 2$ then $E{:}W$ is a complete bipartite subgraph, and if $|W| = 1$ then $E{:}W = \varnothing$; in either case we can restructure S to have empty W.

Example F.11. [[LABEL X:clos-all-full]] $-\Gamma^\circ$ (a full all-negative graph).

- *Balanced edge sets*: The bipartite edge sets (same as $-\Gamma$).
- *Closed, balanced sets*: As in $-\Gamma$, take a connected partition $\pi \in \Pi(\Gamma)$, and in each block of $B \in \pi$, take a maximal cut.
- *Closed sets*: Somewhat as for $-\Gamma$, take any $W \subseteq V$, then let $S = (E^\circ{:}W) \cup$ a closed, balanced set in $-(\Gamma \setminus W)$. We need not restrict W; each different choice of W gives a different closed set since W is identifiable as the set of vertices that support unbalanced edges of S.

Example F.12. [[LABEL X:clos-all-Kn]] $-K_n^\circ$ (the full all-negative complete graph).

This is even simpler than $-\Gamma^\circ$.

- *Balanced edge sets*: The bipartite edge sets.
- *Closed, balanced sets*: Take a partition $\pi \in \Pi_V$, and in each block $B \in \pi$, take a maximal cut, i.e., the edges of a spanning complete bipartite graph.
- *Closed sets*: As with $-\Gamma^\circ$, take any $W \subseteq V$; then each closed set has the form $S = (E^\circ{:}W) \cup$ a disjoint union of complete bipartite graphs on subsets of $V \setminus W$.

There, wasn't that fun!

Nov 3: Jackie Kaminski

G. Incidence and Adjacency Matrices

[[LABEL 2.matrices]]

A signed graph, like a graph, has incidence and adjacency matrices that describe the graph.

G.1. **Incidence matrix.** [[LABEL 2.incidencematrix]]

We now introduce the incidence matrix of a signed graph. Unlike with an unsigned graph, there is only one kind of incidence matrix, the oriented one. As with an unsigned graph, the incidence matrix comes in a family, differing in arbitrary sign choices for the columns.

Definition G.1. [[LABEL D:1103 Incidence Matrix]] An *incidence matrix* of a signed graph Σ is a $V \times E$ matrix $\mathrm{H}(\Sigma) = (\eta_{ve})_{v,e}$ (read 'Eta') whose column indexed by e is shown in Figure G.1, with a zero column for a loose edge. Thus a link has two nonzero elements in

$$\begin{array}{c} \\ \\ \\ v \\ \\ \\ \\ w \\ \\ \\ \\ \end{array}\begin{pmatrix}0\\ \vdots\\ 0\\ \pm1\\ 0\\ \vdots\\ 0\\ \mp\sigma(e)\\ 0\\ \vdots\\ 0\end{pmatrix} \qquad \begin{array}{c} \\ \\ \\ v \\ \\ \\ \\ \\ \end{array}\begin{pmatrix}0\\ \vdots\\ 0\\ \pm1\mp\sigma(e)\\ 0\\ \vdots\\ \vdots\\ 0\end{pmatrix} \qquad \begin{array}{c} \\ \\ \\ v \\ \\ \\ \\ \\ \end{array}\begin{pmatrix}0\\ \vdots\\ 0\\ \pm1\\ 0\\ \vdots\\ \vdots\\ 0\end{pmatrix}$$

a link $e{:}vw$ — a loop $e{:}vv$ — a half edge $e{:}v$

Figure G.1. The columns of the incidence matrix that correspond to each kind of edge.

[[LABEL F:1103column]]

its column, each of which is ±1 and which are the same for a negative link and the same for a positive link (we can state this as the requirement that $\sigma(e)\eta_{ve} + \eta_{we} = 0$); positive loops and loose edges have columns of all zeros; the column of a half edge at v is zero except for ±1 in the row of v; and for a negative loop, the column is all zero except for ±2 in the v row.

Although we say "the" incidence matrix, it is not unique due to the free choice of one sign in each non-zero column.

The incidence matrix is a good descriptor of a graph, but not perfect because it cannot distinguish between positive loops and loose edges, and it doesn't say where loops are located on the graph.

Signed-graphic incidence matrices let us explain the existence of the two kinds of incidence matrix, oriented and unoriented, of a graph. The oriented incidence matrix $\mathrm{H}(\Gamma)$ is just $\mathrm{H}(+\Gamma)$. The unoriented incidence matrix $B(\Gamma)$ is the incidence matrix $\mathrm{H}(-\Gamma)$ with non-negative entries.

Nov 5: Nate Reff

Another way to define an incidence matrix $\mathrm{H}(\Sigma) = (\eta_{ve})_{V\times E}$ is by giving a formula for the (v, e) entry, as follows:

$$\eta_{ve} = \begin{cases} 0 & \text{if } v \text{ and } e \text{ are not incident,} \\ \pm 1 & \text{if } v \text{ and } e \text{ are incident once, so that if } e{:}vw \text{ is a link then } \eta_{ve}\eta_{we} = -\sigma(e), \\ 0 & \text{if } e \text{ is a positive loop at } v, \\ \pm 2 & \text{if } e \text{ is a negative loop at } v. \end{cases}$$

The columns are still defined only up to negation. The reason for that will be explained when we come to orientation, and specifically to incidence matrices of bidirected graphs (Section H.2).

G.2. **Incidence matrix and frame circuits.** [[LABEL 2.incidencecoldep]]

The relation between the incidence matrix and the closure operation is through one of the fundamental structures in a signed graph, the frame circuit.

Definition G.2. [[LABEL Df:1105framecircuit]] A *frame circuit* of Σ is a positive circle, a loose edge, or a pair of negative circles C_1 and C_2 which meet in at most one vertex (and no edges) together with a minimal connecting path P if C_1 and C_2 are vertex disjoint. (When there is a common vertex, we consider it to be a minimal connecting path of length 0.)

The characteristic of a field $\mathbf{F}$ is denoted by $\operatorname{char}\mathbf{F}$. We write $\mathbf{x}_e :=$ the column of e in $\mathrm{H}(\Sigma)$ and $\mathbf{b}_i$ for the ith unit coordinate vector of $\mathbf{F}^n$. When $S \subseteq E$, we denote by $\mathrm{H}(\Sigma)|S$ the matrix consisting of the columns of S from $\mathrm{H}(\Sigma)$ and by $\mathbf{x}_S$ the set of those columns considered as vectors in $\mathbf{F}^n$.

Theorem G.1. *[[LABEL T:1105Theorem1]] Let S be an edge set in Σ and consider the corresponding columns in $\mathrm{H}(\Sigma)$ over a field $\mathbf{F}$.*

(1) *When $\operatorname{char}\mathbf{F} \neq 2$, the columns corresponding to S are linearly dependent $\iff$ S contains a frame circuit.*
(2) *When $\operatorname{char}\mathbf{F} = 2$, the columns corresponding to S are linearly dependent $\iff$ S contains a circle, a loose edge, or a path joining two half edges.*

From a matroid perspective, this means the frame circuits are the circuits of a matroid on a ground set E, and the incidence matrix represents the matroid. This is the *frame matroid*[4] of Σ, which we will study in Chapter IV.

---

[4]Originally and sometimes still called the *signed-graphic matroid* [SG].

Proof of Sufficiency. For (1) it suffices to prove that a frame circuit is dependent. For (2) it suffices to prove that a circle or loose edge, or the path together with the two half edges, is dependent.

Case I: The frame circuit is a loose edge e. Then $\mathbf{x}_e = \mathbf{0}$, which is linearly dependent.

Case II: The frame circuit is a positive loop e incident with v_i. Then $\mathbf{x}_e = \mathbf{0}$, which is dependent.

If e is negative, $\mathbf{x}_e = \pm 2\mathbf{b}_i$, which is independent if $\operatorname{char}\mathbf{F} \neq 2$. If e is a half edge, $\mathbf{x}_e = \pm\mathbf{b}_i$, which is always independent.

Case III: The frame circuit is a circle $C = v_0e_1v_1e_2v_2 \dots e_lv_l$, where $v_0 = v_l$, with $l \geq 2$. Switch so C is all positive with the possible exception of e_1, whose sign is $\sigma(C)$. The incidence submatrix corresponding to C is

$$\begin{array}{cc}
 & \begin{array}{ccccccc} e_1 & e_2 & e_3 & e_4 & \cdots & e_{l-1} & e_l \end{array} \\
\begin{array}{c} v_1 \\ v_2 \\ v_3 \\ \vdots \\ v_{l-1} \\ v_l \\ 0 \\ \vdots \\ 0 \end{array} &
\begin{pmatrix}
-1 & +1 & 0 & 0 & \cdots & 0 & 0 \\
0 & -1 & +1 & 0 & \cdots & 0 & 0 \\
0 & 0 & -1 & +1 & \cdots & 0 & 0 \\
\vdots & \vdots & \vdots & \vdots & \ddots & \vdots & \vdots \\
0 & 0 & 0 & 0 & \cdots & -1 & +1 \\
\sigma(C) & 0 & 0 & 0 & \cdots & 0 & -1 \\
0 & 0 & 0 & 0 & \cdots & 0 & 0 \\
\vdots & \vdots & \vdots & \vdots & \ddots & \vdots & \vdots \\
0 & 0 & 0 & 0 & \cdots & 0 & 0
\end{pmatrix}
\end{array}$$

(where we interpret the sign $\sigma(C)$ as a number, ± 1). The sum of all columns is $(\sigma(C)-1)\mathbf{b}_l$. Hence the vectors are linearly dependent if C is positive. They generate $2\mathbf{b}_l$ if C is negative, whence they are linearly independent iff $\operatorname{char}\mathbf{F} \neq 2$. (For independence one proves that the vectors generate all unit basis vectors $\mathbf{b}_1, \dots, \mathbf{b}_l$ if they generate any one; that follows from connectedness of S_j.)

Case IV: The frame circuit is a handcuff. If there are l vertices, there are $l+1$ edges in dimension l.

In characteristic 2, this covers the case of two half edges with a connecting path of length $l-1$. □

Proof of Necessity. It suffices to prove, for (1), that an edge set that does not contain a frame circuit is independent, and for (2), that an edge set that does not contain a circle is independent. We may consider each component separately since the different components act within disjoint sets of coordinates.

Assume that S is connected and not empty and contains no frame circuit. Then it is a tree (of order at least 2) or a 1-tree. In either case it has a univalent vertex or it is a negative circle. A negative circle is independent by the argument at Sufficiency, Case III, so we may assume S has a vertex v of degree 1. In the incidence matrix of S the row of v has a single 1, that in $\mathbf{x}_e$ where e is the edge at v. The column of e is consequently linearly independent of all other columns of $\mathrm{H}(\Sigma|S)$. We may strip e out of S, leaving a smaller example of the same kind (a tree or 1-tree) whose state of independence is the same as that of S. Continuing in the same manner we arrive at either a negative circle, which is independent in characteristic other than 2, or a half edge or a link, also independent. It follows that S itself is independent.

The proof for characteristic 2 is similar but we have only a tree or a 1-tree that contains a half edge. Alternatively, the incidence matrix is identical to the binary incidence matrix of an unsigned graph so we can appeal to **[WHAT IS THE EXACT RESULT?]** □

Nov 5: Nate Reff

Extending the conclusion of Sufficiency Case III, if we have a closed walk $W = e_1e_2 \dots e_l$ from v_k to v_k, then a suitable linear combination of vectors $\mathbf{x}_{e_1}, \mathbf{x}_{e_2}, \dots, \mathbf{x}_{e_l}$ equals $(\sigma(W) - 1)\mathbf{b}_k$. The precise formula is that

$$\sum_{i=0}^{l-1} \sigma(e_l e_{l-1} \dots e_{l-i}) \mathbf{x}_{e_{l-1}} = (\sigma(W) - 1)\mathbf{b}_k = \begin{cases} \mathbf{0} & \text{if } \sigma(W) = +, \\ -2\mathbf{b}_k & \text{if } \sigma(W) = -. \end{cases}$$

Corollary G.2. *[*[LABEL C:1105Corollary2]] *Assume W is a closed walk which has an edge that appears just once.*

(1) *The vectors of W are linearly dependent if $\sigma(W) = +$ or* $\operatorname{char} \mathbf{F} = 2$.
(2) *The vectors of W generate $2\mathbf{b}_k$ if $\sigma(W) = -1$.*

[This needs explanation/proof!]

Nov 7 (draft): Simon Joyce & T.Z.

Lemma G.3. *[*[LABEL L:1107 negcircle]] *For a negative circle C, $\mathbf{x}_C$ is independent in characteristics other than 2, dependent in characteristic 2.*

Proof. We proved this incidentally when we demonstrated in the proof of Theorem G.1(1) that a positive circle is dependent. □

Recall that we treat half edges as negative loops in this discussion. For instance, "two negative circles" means one or both circles may be half edges instead.

Lemma G.4. *[*[LABEL L:1107 Frame Circuits]] *S contains a frame circuit if and only if it contains a balanced circle or it has two negative circles in the same component.*

Proof. If S contains a frame circuit, that circuit is a balanced circle or contains two connected negative circles.

On the other hand, suppose S contains no positive circle. Then it is a contrabalanced cactus forest (by Exercise **??**), which contains an unbalanced frame circuit if and only if it has a component cactus with at least two negative circles. □

A 1-*tree* is a tree with one extra edge on the same vertices; the extra edge is a half edge or forms a circle.
[MOVE this defn. TO BASICS in Ch. 1. Bal. or unbal. 1-tree in Ch. 2 basics? or maybe better here.]

Theorem G.5. *[*[LABEL T:1107 dep rk]]*Given a signed graph Σ and $S \subseteq E(\Sigma)$. Over a field of characteristic other than 2:*

(1) $\mathbf{x}_S$ *is linearly dependent if and only if S contains a frame circuit.*
(2) $\mathbf{x}_S$ *is linearly independent if and only if each component of S is a tree or a contrabalanced 1-tree.*
(3) *The rank of $\mathrm{H}(\Sigma|S)$ is $n - b(S)$. In particular, $\mathrm{H}(\Sigma)$ has rank $n - b(\Sigma)$ and nullity $|E| - n + b(\Sigma)$. The nullity of its transpose $\mathrm{H}(\Sigma)^T$ is $b(\Sigma)$. [*[LABEL C:1107 matrix rank]]

Proof. Part (1) is Theorem G.1(1).

Part (2) follows from part (1) and Lemma G.4.

As for part (3), the largest possible independent subset of $\mathbf{x}_S$, by part (2), consists of a spanning tree in each balanced component and a contrabalanced 1-tree in each unbalanced component. The size of such a set is $n - b(S)$. It follows that $\dim\langle \mathbf{x}_S \rangle = n - b(S)$. This gives the dimension of the column space of $\mathrm{H}(\Sigma|S)$ and consequently the rank of that matrix. □

We infer the ranks of the oriented and unoriented incidence matrices of an ordinary graph Γ, which are, respectively, $\operatorname{rk}\mathrm{H}(+\Gamma) = n - c(\Gamma)$ and $\operatorname{rk}\mathrm{H}(-\Gamma) = n - c_{\mathrm{bip}}(\Gamma)$, $c_{\mathrm{bip}}(\Gamma)$ denoting the number of bipartite components. Thus we have established the truth of Theorem I.**??**.

Theorem G.6. *[*[LABEL T:1107 span clos]] *Given a signed graph Σ and $S \subseteq E(\Sigma)$, over a field $\mathbf{F}$ of characteristic other than 2 the following properties hold:*

(1) $\operatorname{clos}(S) = \{e \in E : \mathbf{x}_e \in \langle \mathbf{x}_S \rangle\}$*; that is,* $\mathbf{x}_{\operatorname{clos} S} = \mathbf{x}_E \cap \langle \mathbf{x}_S \rangle$.

(2) *S is a closed edge set if and only if $\mathbf{x}_S$ is the intersection of $\mathbf{x}_E$ with a flat of $\mathbf{F}^n$.*

Proof. I will give a graph-theoretic proof based on the definition of signed-graph closure, without using Theorem G.1. That theorem is essentially a linear-algebra and matroid property and not intrinsically graphical.

Let $\boldsymbol{\pi}_{\mathrm{b}}(S) = \{V_1, V_2, \ldots, V_k\}$ and $V_0 := V_0(S)$, and switch so the balanced components $S_i := S{:}V_i$ of S are all positive.

Proof of Part (1). We begin with a lemma. Define $Z := \{\mathbf{x} \in \mathbb{R}^V : \sum_1^n x_i = 0\}$ and $\hat{V}_i := \{\mathbf{x} \in \mathbb{R}^V : x_w = 0 \text{ if } w \notin V_i\}$. The subspaces $\langle \mathbf{x}_{S_i} \rangle$ are independent subspaces, in that $\langle \mathbf{x}_{S_i} \rangle \subseteq \hat{V}_i$. Now we prove that

(G.1) $\mathbf{x}_E \cap \langle \mathbf{x}_S \rangle = (\mathbf{x}_E \cap \langle \mathbf{x}_{S_0} \rangle) \cup (\mathbf{x}_E \cap \langle \mathbf{x}_{S_1} \rangle) \cup \cdots \cup (\mathbf{x}_E \cap \langle \mathbf{x}_{S_k} \rangle)$.[[LABEL E:1107 span clos]]

Notice that $\langle \mathbf{x}_{S_i} \rangle \subseteq Z$.

Suppose f has endpoints $u \in V_i$ and $w \in V_j$ with $i < j$ (so, S_j is all positive). Thus, $\langle \mathbf{x}_{S_j} \rangle \subseteq Z$. Now, the support of $\mathbf{x}_f$ is $\{u, w\}$ with one vertex in V_j; the restriction of $\mathbf{x}_f$ to $\hat{V}_j$ is the signed unit basis vector $\pm\mathbf{b}_w$. The only way S can generate the vector $\mathbf{x}_f$ whose support meets V_j is for S_j to span its restriction to V_j, $\pm\mathbf{b}_w$. However, $\pm\mathbf{b}_w \notin Z$; therefore, $\pm\mathbf{b}_w \notin \langle \mathbf{x}_{S_j} \rangle$ and $\mathbf{x}_f$ cannot be spanned by S.

Similarly, if f is a negative or half edge in $E{:}V_j$, then $\mathbf{x}_f \notin Z$ so $\mathbf{x}_{S_j}$ cannot span $\mathbf{x}_f$. But the only way to generate $\mathbf{x}_f$ from $\mathbf{x}_S$ is for it to be spanned by $\mathbf{x}_{S_j} \subseteq \hat{V}_j$.

If $f{:}uw$ is a positive edge in $E{:}V_j$, there is a uw-path in S_j, say $u = u_0, e_1, u_1, \ldots, e_l, u_l = w$. Taking $\mathbf{x}_{e_i} = \mathbf{b}_{u_i} - \mathbf{b}_{u_{i-1}}$, the sum of these vectors is $\mathbf{b}_w - \mathbf{b}_u$, which is $\mathbf{x}_f$. Thus, $\mathbf{x}_f \in \langle \mathbf{x}_{S_j} \rangle$.

The subspace $\langle \mathbf{x}_{S_0} \rangle$ equals all of $\hat{V}_0$. To see that, note that because each component of S_0 is unbalanced it contains a spanning unbalanced 1-tree. The 1-tree is independent (Theorem G.5(2)) and has as many edges as vertices. The union of all the 1-trees of components of S_0 consequently has cardinality equal to that of V_0; it follows that the vectors of the union of 1-trees form a basis for $\hat{V}_0$. Therefore, $\langle \mathbf{x}_{S_0} \rangle$ contains $\mathbf{x}_f$ for every edge f whose endpoints are in V_0.

In he course of this proof we showed that the edges for which $\mathbf{x}_f \in \langle \mathbf{x}_S \rangle$ are precisely those in the closure of f. That establishes Part (1).

Proof of Part (2). Here is a chain of equivalences:

$$S \text{ is closed} \iff S = \operatorname{clos} S \iff (\text{by part (1)})\ \mathbf{x}_S = \mathbf{x}_E \cap \langle \mathbf{x}_S \rangle.$$

The last property is equivalent to saying that $\mathbf{x}_S = \mathbf{x}_E \cap A$ for some flat A. To prove this, note that $\langle \mathbf{x}_S \rangle$ is one possible flat A if any such flat exists, since $\mathbf{x}_S \subseteq \mathbf{x}_E \cap \langle \mathbf{x}_S \rangle \subseteq \mathbf{x}_E \cap A = \mathbf{x}_S$ by the assumption that an A exists. Thus,

$$S = E{:}V_0 \cup E^+{:}V_1 \cup \cdots \cup E^+{:}V_k \iff \mathbf{x}_S = \mathbf{x}_E \cap \langle \mathbf{x}_S \rangle$$

and the proof is complete. □

Exercise G.1. [[LABEL Ex:1107 span clos]] Prove Theorem G.6 using the characterization of closure in Theorem F.6. (That proof would be a matroid-style proof. The one I presented is essentially graphical.)

There are two other important corollaries, which a reader who is not involved with matroids may ignore. Let us define $\operatorname{Lat} M$, for an $n \times m$ matrix M, to be the family of subspaces of $\mathbb{R}^n$ that are generated by columns of M; for instance, the smallest such space is the zero space, generated by the empty set of columns, and the column space $\operatorname{Col}(M)$is the largest such space. It's well known that $\operatorname{Lat}(M)$ is a geometric lattice (in fact, that's where the name comes from).

Theorem G.7. *[[LABEL C:1107 matroid rank]]In a signed graph Σ, the closure operator is a matroid closure,* rk *is a matroid rank function, and* $\operatorname{Lat}\Sigma$ *is a geometric lattice with rank function $n - b(S)$, isomorphic to* $\operatorname{Lat} \mathrm{H}(\Sigma)$*. Furthermore, $\Pi^\dagger(\Sigma)$ is a geometric lattice with rank function $\operatorname{rk}(\theta) = n - |\pi(\theta)|$.*

Proof. The key is to prove that $\operatorname{Lat}\Sigma$ and $\operatorname{Lat} \mathrm{H}(\Sigma)$ are isomorphic. The specific isomorphism is that $S \in \operatorname{Lat}\Sigma \mapsto \langle \mathbf{x}_S \rangle = \langle \mathbf{x}_e : e \in S \rangle \in \operatorname{Lat} \mathrm{H}(\Sigma)$.

[NEEDS MORE PROOF.]

□

Corollary G.8. *[[LABEL C:1107 cor of cor]] The set $\Pi_n^\dagger(\pm)$ of signed partial partions of $[n]$ is a geometric lattice.*

Proof. We know from the definitions that $\Pi_n^\dagger(\pm) = \Pi^\dagger(\pm K_n^\circ)$, so we apply Theorem G.7. □

2014 Nov 17: T.Z.

G.3. **Adjacency matrix.** [[LABEL 2.adjacencymatrix]]

The *adjacency matrix* of a signed graph Σ tells which vertices are adjacent (including to themselves). It is the matrix $A(\Sigma) = (a_{ij})_{n\times n}$ in which

$$a_{ij} = (\text{number of positive } v_iv_j \text{ edges}) - (\text{number of negative } v_iv_j \text{ edges})$$

when $i \neq j$ and

$$a_{ii} = 2(\text{number of positive loops}) - 2(\text{number of negative loops}) \text{ at } v_i.$$

In terms of the positive and negative subgraphs there is the simple expression

$$A(\Sigma) = A(\Sigma^+) - A(\Sigma^-).$$

Cancellation makes it hard to recover a signed graph from its adjacency matrix. I'll call Σ *reduced* if it has no parallel edges of opposite sign—that includes no parallel loops of opposite sign. If it does have any such parallel pairs, *reducing* Σ is the operation of deleting them; that is, we delete negative digons until none remain, and we delete pairs of a positive and negative loop at the same vertex until none of those remain. The resulting signed graph is obviously reduced. Clearly, all we can recover from $A(\Sigma)$ is the reduced form of Σ. [**MOVE DEFN TO EXAMPLES OF S.G., earlier.**]

Powers of A count walks of given length, as with unsigned graphs, but with cancellation of oppositely signed walks exactly as oppositely signed edges cancel in A.

Theorem G.9. *[*[LABEL T:Apowers]] *In Σ define $w_l^\varepsilon(v,w)$ to be the number of walks of length l with endpoints v and w whose sign is ε. For each $l \geq 0$, the (i,j) element of $A(\Sigma)^l$ equals $w_l^+(v_i,v_j) - w_l^-(v_i,v_j)$.*

The proof is an easy but pleasant exercise in induction that I leave to the reader.

Recall that the degree matrix of the underlying graph, $D(|\Sigma|)$, is the diagonal matrix with $d_{ii} = d_{|\Sigma|}(v_i)$, where a loop counts twice in the degree. (Explanations: A loop $e{:}v_iv_i$ makes v_i adjacent to itself twice, once from each end of e. A half edge does not create an adjacency.)

Nov 10: T.Z. & Peter Cohen

Theorem G.10. *[*[LABEL T:1110amatrix]] *The adjacency matrix of a signed graph satisfies $A(\Sigma) = D(|\Sigma|) - \mathrm{H}(\Sigma)\mathrm{H}(\Sigma)^T$.*

Proof. Since HH^T is the matrix of dot products of the rows of H, we compute those products.

For rows i and j where $i \neq j$, the product is the sum of $+1$ for each negative edge e_{ij} and -1 for each positive edge e_{ij}. That gives the value of a_{ij} for $i \neq j$.

In the product of row i with itself we get 4 for each negative loop, 1 for each half edge, 1 for each link, and 0 for each positive loop incident with v_i. Subtracting from the degree matrix leaves -2 for each negative edge, $+2$ for each positive edge, and 0 for a half edge. □

Corollary G.11. *[*[LABEL C:1110aregular]] *If $|\Sigma|$ is k-regular then all eigenvalues of $A(\Sigma)$ are $\leq k$. The multiplicity of k as an eigenvalue is $b(\Sigma)$.*

Proof. First, some matrix theory. A *Gram matrix* G is the matrix of inner products of a set of vectors. Rephrasing the definition in matrix terms, $G = M^\mathrm{T}M$ for some matrix M; that is, G is the matrix of inner products of the columns of M. If M is real, the Gram matrix G is real and symmetric, so it has only real eigenvalues, and it has n such eigenvalues (with multiplicity). Furthermore, G is positive semidefinite so it has no negative eigenvalues. The rank of $G =$ the rank of M by matrix theory, so the nullity of G, which is the multiplicity of 0 as an eigenvalue of G, equals the nullity of M^T.

Now, $D - A = \mathrm{HH}^\mathrm{T}$ is a Gram matrix (with $M^\mathrm{T} = \mathrm{H}$). By its positive semidefiniteness, all eigenvalues of $D - A$ are non-negative. The multiplicity of 0 as an eigenvalue of $D - A$ is $\operatorname{nul}\mathrm{H}^\mathrm{T}$. By Theorem G.5(3), this is $b(\Sigma)$.

We check what that means for A, remembering that $D = kI$. If λ is an eigenvalue of A with eigenvector $\mathbf{x}$, then $A\mathbf{x} = \lambda\mathbf{x}$, so $(D - A)x = (kI - A)x = (k - \lambda)x$. By the positive semidefiniteness of $D - A = \mathrm{HH}^\mathrm{T}$, $k - \lambda \geq 0$ for every eigenvalue λ and the eigenvalue $\lambda = k$, corresponding to the eigenvalue 0 of $D - A$, has multiplicity $\operatorname{nul}\mathrm{H}^\mathrm{T} = b(\Sigma)$. □

Nov 10 (draft): Peter Cohen

H. Orientation

[[LABEL 2.orientation]]

An oriented signed graph is a bidirected graph; thus, we begin by explaining bidirection.

H.1. **Bidirected graphs.** [[LABEL 2.bidirected]]

Bidirected graphs were introduced by Jack Edmonds (1965a??) to treat matching theory. Our use for them is entirely different.

Definition H.1. [[LABEL D:1110bidirected]] An edge with an independent direction at each end is called a *bidirected edge*. A *bidirected graph* is a graph with an independent direction on each of the ends of each edge; that is, where every edge is bidirected.

Loose edges are bidirected by having no directions, as they have no ends. Half edges are bidirected by having one direction, as they have only one end. A loop has two ends that have the same endpoint, so a loop, like a link, is bidirected by getting two directions.

We may think of the directions pictorially as arrows or algebraically as signs. To denote the signs, we will introduce new notation, $\tau(v,e)$, which is the sign of the end of edge e_k at vertex v_i. The definition of τ in terms of directions is:

$$\tau(v,e) = \begin{cases} +, & \text{if } e \text{ enters } v, \\ -, & \text{if } e \text{ leaves } v, \\ 0, & \text{if } e \text{ and } v \text{ are not incident.} \end{cases}$$

(We often write τ_{ve} for $\tau(v,e)$; and when the edges and vertices are numbered, τ_{ik} for $\tau_{v_i e_k}$.) A direction into a vertex is positive, while a direction out of a vertex is negative. The edge itself is positive if it is balanced; both directions are the same (going from one vertex to the other, so one is positive and the other is negative). A negative edge has either two negative or two positive ends. **[THIS NEEDS TO BE COORDINATED WITH THE NEXT DAY'S EXPLANATION OF** H(B)**.]**

Example H.1. [[LABEL X:1110small]] **[Is this really an example? What is the example? What is it for?]** For our oriented graph, the matrix $\mathrm{H}(\Gamma)$ will have, as an example, for the column of edge e_k:

$$\begin{pmatrix} 0 \\ -1 \\ 0 \\ +1 \\ 0 \end{pmatrix}.$$

where the -1 indicates the edge leaves that vertex, and a $+1$ indicates that the edge enters that vertex. In this example, the edge is a positive edge (in a 5-vertex graph). A positive loop will have a column like

$$\begin{pmatrix} 0 \\ 0 \\ 0 \\ 0 \\ 0 \end{pmatrix}.$$

So the matrix $\mathrm{H}(\Gamma)$ cannot distinguish between a positive loop and a loose edge. A negative loop will have ± 2 in one entry of its column while the other entries are zero.

Definition H.2. [[LABEL Df:1110sgbidir]] The *signed graph associated* with a bidirected graph B is $\Sigma(\mathrm{B}) := (|\mathrm{B}|, \sigma_{\mathrm{B}})$ where $|\mathrm{B}|$ is the underlying graph of B and $\sigma_{\mathrm{B}}(e) := -\tau_{ve}\tau_{we}$ for a link or loop $e{:}vw$. If Σ is the signed graph associated with B, we say that B is an *orientation* of Σ.

Nov 12: Yash Lodha

H.2. **Incidence matrix of a bidirected graph.** [[LABEL 2.orientation.incid]]

We now define the incidence matrix $\mathrm{H(B)} = (\eta_{ik})$ of the bidirected graph $\mathrm{B} = (\Gamma, \tau)$, where Γ is the underlying graph.

For an edge e_k incident to the vertex v_i, $\tau_{ik} = +$ if the direction/orientation at the vertex v_i end is directed towards the vertex v_i and $\tau_{ik} = -$ if the direction/orientation is directed away from the vertex. The column of a link $e = v_iv_j$ has i-th entry τ_{ik} and j-th entry τ_{jk}, and the remaining entries are zero. For a loop v_iv_i the i-th entry equals $\tau_{ik} + \tau'_{ik}$ where each τ'_{ik} is the same as τ_{ik} except for the other end of the loop. For a loose edge all the entries in the corresponding column are zero.

More formally, $\eta_{ik} = \sum_\varepsilon \tau_\varepsilon$, summed over all edge ends ε of e_k incident with v_i.

Notice that an incidence matrix of a bidirected graph B is an incidence matrix of its signed graph $\Sigma(\mathrm{B})$. Conversely, an incidence matrix of Σ is the incidence matrix of an orientation of Σ.

A *source* is a vertex where every edge end departs, ie every $\eta_{ik} \leq 0$ for all edges e_k. A *sink* is a vertex where every edge end enters, i.e. every $\eta_{ik} \geq 0$ for all edges e_k.

A *cycle* in a bidirected graph is an oriented frame circuit with no source or sink. This means that every vertex of degree two in the circuit must have consistent orientation, i.e. the direction/orientation of both edge ends incident to the vertex agree. So a positive circuit has exactly two orientations with no source or sink, and they are opposite. A negative circle must have an orientation with sources or sinks.

Definition H.3. [[LABEL D:1112cyclicacyclic]] We say an oriented signed graph $\vec{\Sigma}$ is *acyclic* if it has no cycles, *cyclic* if it has a cycle, and *totally cyclic* if each edge is in a cycle.

Recall that $\Sigma(\mathrm{B})$ has edge signs $\sigma(e{:}vw) = -\tau_{ve}\tau_{we}$.

Walks and coherence.

In a walk $W = v_0e_1v_1e_2 \cdots v_{l-1}e_lv_l$, the two edge ends (v_i, e_i) and (v_i, e_{i+1}) incident to vertex v_i (when $0 < i < l$) may have either of two interrelations: they may be *coherent* or *consistent* (both terms are used), which means that one of their arrows points into the common vertex and the other points out (in terms of the bidirection function, $\tau(v_i, e_i)\tau(v_i, e_{i+1}) = -$), or they may be *incoherent* or *inconsistent*, which means both arrows point into the vertex or both point out (that is, $\tau(e_{i-1})\tau(e_i) = +$).

Lemma H.1. *[[LABEL L:1112coherentwalk]] Let $W = v_0e_1v_1 \cdots e_lv_l$ be a walk in which each vertex v_i for $0 < i < l$ is consistently oriented in W. Then $(-1)^l\tau_{0l}\tau_{ll} = \sigma(W)$.*

If W is a closed walk, so $v_0 = v_l$, then it is positive if it is consistent at $v_l = v_0$, and negative if it is inconsistent.

Proof. Take the product of the signs of all oriented edge ends ε in W and compute it in two ways.

$$\prod_\varepsilon \tau_\varepsilon = \prod_{i=1}^{l}(\tau_{i-1,i}, \tau_{ii}) = \prod_{i=1}^{l} -\sigma(e_i) = (-1)^l\sigma(W).$$

Also,

$$\prod_\varepsilon \tau_\varepsilon = \tau_{01}(\tau_{11}\tau_{12})(\tau_{22}\tau_{23})...(\tau_{(l-1),(l-1)}\tau_{(l-1),l})\tau_{ll}.$$

Therefore, $(-1)^l\sigma(W) = (-1)^{l-1}\tau_{01}\tau_{ll} \implies \sigma(W) = -\tau_{01}\tau_{ll}$. □

Corollary H.2. *[[LABEL C:1112coherentclosedwalk]] A closed walk W, in which (as above) each vertex v_i for $0 < i < l$ is consistently oriented, is consistent at v_l if and only if $\sigma(W) = +$.*

An application of the corollary is that a positive circle can be oriented consistently and a negative circle can be oriented consistently except for one inconsistent vertex, which is a source or a sink.

[WE NEED DIAGRAMS for all these explanations.]

In a frame circuit with no source or sink, every divalent vertex must be coherent. Therefore we can orient a positive circle cyclically (i.e., to have no source or sink) in only two ways; once we have oriented one edge, every other edge orientation is determined by coherence. Corollary H.2 ensures that it is possible to make every vertex coherent.

A contrabalanced handcuff C likewise has only two cyclic orientations. Each negative circle, $C_i = C_1$ or C_2 must be coherent except at the vertex v_i that lies on the connecting path P. (If the two negative circles share a vertex, we consider that vertex to be the connecting path.) Since v_1 is incoherent, hence a source or sink, in C_1, the orientation of the end (v_1, e_{1P}) of the connecting-path edge e_{1P} is determined by the requirement that v_1 not be a source or sink in the handcuff. (An edge $e_1 \in C_1$ at v_1 is thus coherent with e_{1P}.) The orientations of all edges of P are then determined by coherence in P, and the orientation of (v_2, e_{2P}) determines that of each edge $e_2 \in C_2$ at v_2 and hence everywhere. (If P has length 0 so $v_2 = v_1$, the orientations of the ends (v_1, e_{1P}) determine those of the ends (v_1, e_{2P}).) Summarizing this discussion, we have a proposition:

Proposition H.3. *[[LABEL P:1112cycliccircuit]] A frame circuit has exactly two cyclic orientations, which are negatives of each other.*

Nov 12: Yash Lodha

I. Equations from Edges, and Signed Graphic Hyperplane Arrangements

[[LABEL 2.equations]]

I.1. Equations from edges. [[LABEL 2.edgeequations]]

An equation from an edge is dual to its column vector $\mathbf{x}_e$ from $\mathrm{H}(\Sigma)$. Let $\mathbf{x} = (x_1, ..., x_n)$. So the equation from an edge e will be $\mathbf{x}_e \cdot \mathbf{x} = 0$.

For a positive edge $\mathbf{x}_e \cdot \mathbf{x} = x_i - x_j$ or $x_j - x_i$. So we get $x_i = x_j$.

For a positive loop this says $x_i = x_i$, which gives us the "degenerate hyperplane", $\mathbb{R}^n$.

For a signed edge $x_i = \sigma(e)x_j$ because from $\mathbf{x} \cdot \mathbf{x}_e = 0$ we get $\pm(\mathbf{b}_i - \sigma(e)\mathbf{b}_j) = 0$. So for a negative edge we get $x_i = -x_j$.

For a half edge we get $x_i = 0$.

For a loose edge we get $0 = 0$, which gives us the degenerate hyperplane.

So each edge $e = e_{ij}$ gives us a hyperplane $h_e = h_{ij}^{\sigma(e)}$ where $h_{ij}^{\varepsilon} = \{\mathbf{x} \mid x_i = \varepsilon x_j\}$. For a half edge e_i,, $h_{e_i} = \{\mathbf{x} \mid x_i = 0\}$, which is a coordinate hyperplane, and for a loose edge $e_\varnothing$, $h_{e_\varnothing} = \mathbb{R}^n$, the degenerate hyperplane.

So we get a signed graphic hyperplane arrangement $\mathcal{H}[\Sigma]$, and the intersection lattice of this arrangement, ordered by reverse inclusion, is the poset obtained from the set of flats.

Formally:

$$\mathcal{L}(\mathcal{H}[\Sigma]) = \{A \subseteq \mathbb{R}^n \mid A = \bigcap \mathcal{S} \text{ for } \mathcal{S} \subseteq \mathcal{H}[\Sigma]\} = \{\bigcap_{e \in S} h_e \mid S \subseteq E\}.$$

Theorem I.1. [[LABEL T:1112latticeisom]] $\mathcal{L}(\mathcal{H}[\Sigma]) \cong \mathrm{Lat}(\Sigma)$ *by the correspondence* $A \mapsto \{e \mid h_e \supseteq A\}$.

Proof. By vector-space duality,

$$\mathcal{L}(\mathcal{H}[\Sigma]) \cong \{\text{flats in } \mathbb{R}^n \text{ generated by subsets of the columns of } \mathrm{H}(\Sigma)\},$$

which is isomorphic to $\mathrm{Lat}(\Sigma)$ by Theorem G.7. The exact formula is a matter of tracing the correspondences. □

[MISSING NOTES] 11/14: Kaminski

Nov 17a (draft): Simon Joyce

I.2. **Additive representations: binary and affine.** [[LABEL 2.additiverep]]

For this section we will be working over $\mathbb{F}_2$, whose additive group is isomorphic to the sign group $\{+, -\}$. For a signed graph Σ we can describe the graph in matrix from using $\mathrm{H}(|\Sigma|)$, the unoriented incidence matrix and an extra row containing the signs of the edges. A 0 in this extra row denotes a positive edge and a 1 denotes a negative edge. The is how the augmented incidence matrix is defined.

Definition I.1. [[LABEL D:1117 AGIM]] The *augmented graphic incidence matrix* of a signed graph Σ is $M(\Sigma)$, defined as follows:

$$M(\Sigma) = \begin{array}{c} \\ \\ v_1 \\ \vdots \\ v_n \end{array} \begin{pmatrix} e_1 & e_2 & \cdots & e_m \\ s_1 & s_2 & \cdots & s_m \\ & & & \\ & \mathrm{H}(|\Sigma|) & & \\ & & & \end{pmatrix} \begin{array}{c} \\ x_0 \\ x_1 \\ \vdots \\ x_n \end{array}$$

[This still doesn't look right but it's better.]
where the s_i in row 0 belong to $\{0, 1\} = \mathbb{F}_2$. This matrix is an $\mathbb{F}_2$-matrix (so $1 = -1$ in $\mathrm{H}(|\Sigma|)$).

A loose edge or a positive loop is represented by an all zero column. A negative loop or a half edge is represented by a column with a one in the x_0 row and zeros in all other rows. We denote a column vector of this form by $\mathbf{b}_0$.

We have seen that the minimally dependent edge sets in the signed incidence matrix are the frame circuits. For the augmented incidence matrix we call the minimally dependent edge sets lift circuits. The posible lift circuits are shown in Figure I.1. We treat half edges like negative circles so that a pair of half edges would also be a lift circuit for example.

Definition I.2. [[LABEL D:1117 lift circuits]] In a signed graph Σ, a lift circuit is a positive circle, a contrabalanced tight handcuff, or contrabalanced loose bracelets (i.e., two vertex-disjoint negative circles).

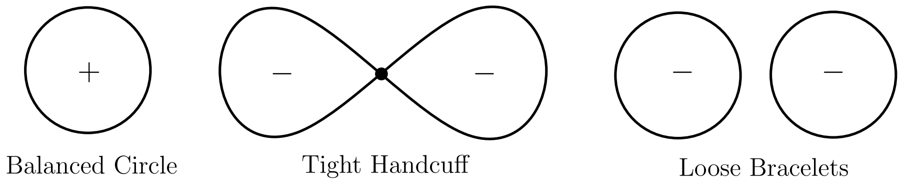

FIGURE I.1. The different kinds of lift circuits.
[[LABEL F:1117liftcircuits]]

Theorem I.2. *[[LABEL T:1117 lift circuits]] A set of columns of $M(\Sigma)$ is linearly dependent if and only if the corresponding edge set contains a lift circuit.*

Proof. To prove this we'll need a few lemmas.

Lemma I.3. *[[LABEL L:1117 lemma1]] The columns of $M(\Sigma)$ corresponding to $S \subseteq E$ generate $\mathbf{b}_0$ if and only if S is unbalanced.*

Proof. If S is balanced, switch so that all the edges of S are positive. Then for all edges $e \in S$, their corresponding columns in M all have 0 in the x_0 row. Therefore $\mathbf{b}_0$ cannot be generated.

Now assume S contains an unbalanced circle C. Let $\mathbf{z}_e$ be the column vector in M corresponding to the edge e. Denote the sign of e in $\mathbb{F}_2$ by z_{e0} and let $\mathbf{z}'_e$ be the column vector corresponding to e in $\mathrm{H}(|\Sigma|)$, so that $\mathbf{z}_e = \begin{pmatrix} z_{e0} \\ \mathbf{z}'_e \end{pmatrix}$. In $\mathrm{H}(|\Sigma|)$ over $\mathbb{F}_2$ we have $\sum_{e \in C} \mathbf{z}'_e = \mathbf{0}$. [**This should have been proved in chapter one but I can't find a reference**] Since C is unbalanced, in M we get

$$\sum_{e \in C} \mathbf{z}_e = \begin{pmatrix} \sum z_{e0} \\ \sum \mathbf{z}'_e \end{pmatrix} = \begin{pmatrix} \sigma(C) \\ \mathbf{0} \end{pmatrix} = \mathbf{b}_0.$$

A half edge is unbalanced set on its own and corresponds to a $\mathbf{b}_0$ column in M so the Lemma is also true in the case where S contains a half edge. ☐

[**Is the next lemma actually needed for the proof?**]

Lemma I.4. *[[LABEL L:1117 lemma2]] Row 0 of M is a linear combination of the other rows if and only if Σ is balanced.*

Proof. We want to prove that

$$\mathrm{rk}[M] = \mathrm{rk}[\mathrm{H}(|\Sigma|)] \iff \Sigma \text{ is balanced} \iff \mathbf{b}_0 \notin \mathrm{Col}(M) \iff \mathrm{rk}[M|\mathbf{b}_0] = \mathrm{rk}[M] + 1.$$

Equivalently, we can show that

$$\mathrm{rk}[M] = \mathrm{rk}[\mathrm{H}(\Sigma)] + 1 \iff \mathrm{rk}[M|\mathbf{b}_0] = \mathrm{rk}[M].$$

By adding $\mathbf{b}_0$ to all the negative edges in M we get that $\mathrm{rk}[M|bfb_0] = \mathrm{rk}[\mathrm{H}(\Sigma)] + 1$ and by Lemma I.3 $\mathrm{rk}[M|\mathbf{b}_0] = \mathrm{rk}[M]$ if and only if Σ is unbalanced. The result follows. ☐

Now it's clear that if an edge set S contains a lift circuit, then the corresponding columns in M are linearly dependent. Conversely, if S is linearly dependent, let T be a minimally dependent subset of S. If no subset of T sums to $\mathbf{b}_0$, then T forms a positive circle since the edges must form a circle to be minimally dependent in $\mathrm{H}(|\Sigma|)$ [**reference?**] and there must be an even number of 1's in the x_0 row. Otherwise there must be a bipartition T_1 and T_2 of

T such that $\sum_{e\in T_1} = \mathbf{b}_0$ and $\sum_{e\in T_2} = \mathbf{b}_0$ since T is minimally dependent. By Lemma I.3 T_1 and T_2 are the edge sets of unbalanced circles. If $V(T_1)\cap V(T_2) \geq 2$, then T will contain a theta graph. So T will contain a positive circle by Theorem B.2, but this contradicts the minimal dependence of T. Therefore T_1 and T_2 form a tight handcuff or loose bracelets. $\square$

[the next part needs fleshing out. We should discuss the content.]

For a signed graph Σ, denote by $\mathcal{A}[\Sigma]$ the affinographic hyperplane arrangement of Σ over $\mathbb{F}_2$. This is the arrangement whose hyperplanes are given by the correspondence

$$e : vw \;\longleftrightarrow\; \begin{bmatrix}\sigma(e)\\ 0\\ 1\\ 0\\ 1\\ 0\end{bmatrix} \;\begin{aligned}&\longleftrightarrow \text{equation } x_j - x_i = \sigma(e)x_0 \text{ in } \mathbb{F}_2^{n+1}\\ &\longleftrightarrow \text{linear hyperplane } \bar{h}_{ij}^{\sigma(e)} \text{ in } \mathbb{F}_2^{n+1}\\ &\longleftrightarrow \text{affine hyperplane } h_{ij}^{\sigma(e)} \text{ in } \mathbb{A}^n(\mathbb{F}_2).\end{aligned}$$

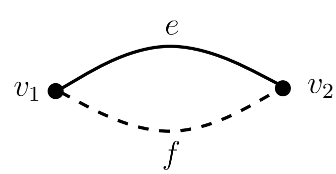

FIGURE I.2.

[[LABEL F:1117hyparr]] For example, in $\mathbb{F}_2^n$ the edge e in Figure I.2 corresponds to the hyperplane where $x_1 - x_2 = 0$ and f corresponds to the hyperplane where $x_1 - x_2 = 1$.

Definition I.3. [[LABEL D:1117 intersection sublattice]] For a signed graph Σ, define the *intersection sublattice*, $\mathcal{L}(\mathcal{A}[\Sigma])$, to be

$$\mathcal{L}(\mathcal{A}[\Sigma]) = \{\bigcap \mathcal{S} \mid \mathcal{S} \subseteq \mathcal{A}[\Sigma],\ \bigcap \mathcal{S} \neq \varnothing\}.$$

Theorem I.5. *[[LABEL T:1117 to prove later]]*

$$\mathcal{L}(\mathcal{A}[\Sigma]) \cong \operatorname{Lat}^{\mathrm{b}} \Sigma.$$

This will be proved in Chapter III because the proof is the same in the greater generality of gain graphs.

Nov 17b (draft): Simon Joyce

J. CHROMATIC FUNCTIONS

[[LABEL 2.chromatic]]
[THIS IS DUPLICATED IN THE NEXT DAY'S NOTES.]

Definition J.1. [[LABEL D:coloration]] Given a signed graph Σ with vertex set V, a k-*coloration* is a mapping $\gamma : V \to \Lambda_k^*$ or $\gamma : V \to \Lambda_k$. Here the *color sets* are $\Lambda_k^* = \{\pm1, \pm2, \ldots, \pm k\}$ and $\Lambda_k = \Lambda_k^* \cup \{0\}$. Colorations of the former type are called *zero free*.

We call an edge $e : vw$ in Σ *proper* (with respect to γ) if $\gamma(v) \neq 0$ and $\gamma(w) \neq \sigma(e)\gamma(v)$. It is *improper* if $\gamma(w) = \sigma(e)\gamma(v)$. We also consider loose edges to be improper. A half edge $e{:}v$ is improper if and only if $\gamma(v) = 0$.

The definition implies that a negative loop $e{:}vv$ is proper when $\gamma(v) \neq -\gamma(v)$, i.e., $2\gamma(v) \neq 0$. That is equivalent to having color 0 at v when working over characteristic other than 2; in characteristic 2 a negative loop, like a positive loop, can never be proper, which is correct because in characteristic 2 $+$ and $-$ are the same.

Nov 19: Nate Reff As with unsigned graphs, I call any function that depends on coloring or that satisfies the algebraic laws of the chromatic polynomial (or the dichromatic polynomial) a *chromatic* (or *dichromatic*) *function*.

J.1. **Coloring a signed graph.** [[LABEL 2.coloring]]
Suppose that $\Sigma = (V, E, \sigma)$ is a signed graph.

Definition J.2. [[LABEL D:1119coloration]] A *k-coloration* is a a mapping $\gamma : V \to \Lambda_k$, where the *color set* is

$$\Lambda_k := \{\pm 1, \pm 2, \ldots, \pm k\} \cup \{0\}.$$

A coloration is *zero free* if it does not use the color 0 (that is, $0 \notin Im(\gamma)$); the *zero-free color set* is

$$\Lambda_k^* := \Lambda_k \setminus \{0\} = \{\pm 1, \pm 2, \ldots, \pm k\}.$$

[DUPLICATION.]
Just as in ordinary unsigned graph coloring, with respect to a particular coloration there are two kinds of edges, proper and improper. An edge $e{:}vw$ is *proper* if $\gamma(w) \neq \sigma(e)\gamma(v)$, or *improper* if $\gamma(w) = \sigma(e)\gamma(v)$. A half edge $e{:}v$ is proper if $\gamma(v) \neq 0$. A loose edge is always improper. A *proper coloration* is a coloration with no improper edges. The *chromatic number* of Σ is

$$\chi(\Sigma) := \min\{k : \exists \text{ a proper } k\text{-coloration}\},$$

and the *zero-free chromatic number* is

$$\chi^*(\Sigma) := \min\{k : \exists \text{ a zero-free proper } k\text{-coloration}\}.$$

If there does not exist a proper coloration (or, equivalently, zero-free coloration) at all, then $\chi(\Sigma) = \infty$ (or, $\chi^*(\Sigma) = \infty$).

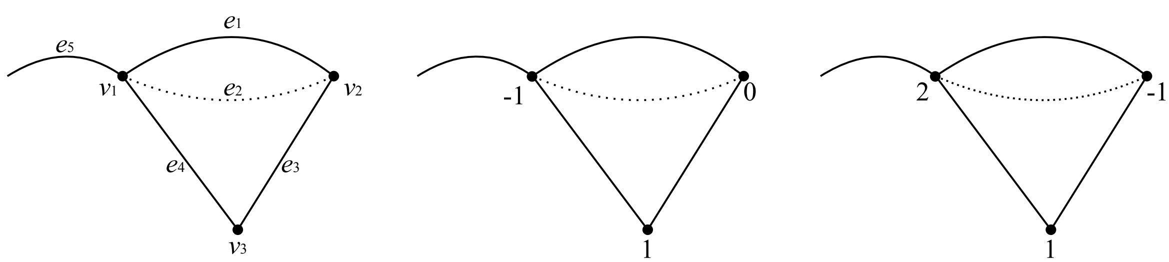

FIGURE J.1. Signed graph Σ, a proper 1-coloration of Σ, a proper zero-free 2-coloration of Σ.
[[LABEL 1119image1]]

Consider the example of a signed graph Σ in Figure J.1. There clearly does not exist a proper 0-coloration. There is, however, a proper 1-coloration as seen in Figure J.1, and so $\chi(\Sigma) = 1$. If we try to find zero-free colorations, it is easy to see that there is no proper zero-free 1-coloration due to the $+K_3$ subgraph present, but there is a proper zero-free 2-coloration as seen also in Figure J.1. Therefore $\chi^*(\Sigma) = 2$.

J.2. **Chromatic numbers.** [[LABEL 2.chromaticnumber]]

Recall that we write $\Sigma^\bullet$ for the signed graph obtained from Σ by adding a negative loop or half edge at every vertex. One can see that, under our definition of proper coloration, $\chi(\Sigma^\bullet) = \chi^*(\Sigma)$.

Let's make a few observations. First,

$$\text{[[LABEL E:1119chromaticnumberineq]]}\chi(\Sigma) \le \chi^*(\Sigma) \le \chi(\Sigma) + 1. \tag{J.1}$$

Furthermore, the lower value obtains if and only if Σ is full, since only then is the color 0 ruled out.

Next, take a look at an all-positive graph:

$$\chi(+\Gamma) = \left\lceil \frac{\chi(\Gamma) - 1}{2} \right\rceil \text{ and } \chi^*(+\Gamma) = \left\lceil \frac{\chi(\Gamma)}{2} \right\rceil.$$

Looking at these two equations we can see that if $\chi(\Gamma)$ is even, then $\chi(+\Gamma) = \chi^*(+\Gamma)$. It is possible that $\chi^*(+\Gamma) > \chi(+\Gamma)$, but Equation (J.1) leaves little room for difference between the two chromatic numbers.

Coloring the complete signed graph $\pm K_n^\bullet$, we can only have zero-free proper colorations due to the negative loop or half edge at each vertex. To ensure a coloration is proper, each vertex must get a different absolute value of color. Thus see that

$$\chi^*(\pm K_n^\bullet) = \chi(\pm K_n^\bullet) = \chi(K_n^\bullet) = \chi(K_n) = n,$$
$$\chi(\pm K_n) = \chi(K_n) - 1 = n - 1, \text{ and } \chi(\pm\Gamma) = \chi(\Gamma) - 1.$$

A general rule is that, if you switch Σ by ζ, you also switch colorations: γ switches to γ^ζ defined by

$$\gamma^\zeta(v) := \zeta(v)\gamma(v).$$

Lemma J.1. *[[LABEL L:1119Lemma1]] e is proper in Σ colored by $\gamma \iff$ it is proper in Σ^ζ colored by γ^ζ.*

Proof. First suppose $e{:}vw$ is a link. Then e is proper in $\Sigma \iff \gamma(w) \ne \sigma(e)\gamma(v) \iff \zeta(v)\zeta(v)\zeta(w)\gamma(w) \ne \zeta(v)\zeta(v)\zeta(w)\sigma(e)\gamma(v) \iff \zeta(w)\gamma(w) \ne \zeta(v)\sigma(e)\zeta(w)\zeta(v)\gamma(v) \iff \gamma^\zeta(w) \ne \sigma^\zeta(e)\gamma^\zeta(v) \iff e$ is proper in Σ^ζ with γ^ζ.

Now suppose $e{:}v$ is a half edge, or $e{:}vv$ is a negative loop. Then e is proper in $\Sigma \iff \gamma(v) \ne 0 \iff \zeta(v)\gamma(v) \ne 0 \iff \gamma^\zeta(v) \ne 0 \iff e$ is proper in Σ^ζ with γ^ζ. □

Proposition J.2. *[[LABEL P:1119Prop1]] Switching does not change chromatic numbers. That is, $\chi(\Sigma) = \chi(\Sigma^\zeta)$ and $\chi^*(\Sigma) = \chi^*(\Sigma^\zeta)$ for all switching functions ζ.*

Proof. Use switching of colors and Lemma J.1. □

J.3. **Chromatic polynomials.** [[LABEL 2.chromaticpoly]]

The archetypical chromatic functions of signed graphs are the counting functions for the two types of proper coloration.

Definition J.3. [[LABEL D:1119chromaticpolys]] Let k be any non-negative integer. We define

$\chi_\Sigma(2k+1) :=$ the number of proper k-colorations and
$\chi^*_\Sigma(2k) :=$ the number of zero-free proper k-colorations.

Obviously, the two functions of k are non-decreasing. Evidently, $\chi(\Sigma)$ is the smallest non-negative integer k for which $\chi_\Sigma(2k+1)$ is not zero, and $\chi^*(\Sigma)$ is the smallest non-negative integer k for which $\chi^*_\Sigma(2k)$ is non-zero.

Notice that $\chi^*_\Sigma(2k) = \chi_{\Sigma^\bullet}(2k+1)$, which reduces χ^*_Σ to $\chi_{\Sigma^\bullet}$. The functions χ_Σ and χ^*_Σ will turn out to be polynomials, but just as with ordinary graph coloring, this is not a trivial fact.

Theorem J.3. *[*[LABEL T:1119Theorem1]] *The chromatic functions $\chi_\Sigma(2k+1)$ and $\chi^*_\Sigma(2k)$ have the following properties:*

Unitarity:

$$\chi_\varnothing(2k+1) = 1 = \chi^*_\varnothing(2k) \textit{ for all } k \geq 0.$$

Multiplicativity:

$$\chi_{\Sigma_1 \,\dot\cup\, \Sigma_2}(2k+1) = \chi_{\Sigma_1}(2k+1)\chi_{\Sigma_2}(2k+1)$$

and

$$\chi^*_{\Sigma_1 \,\dot\cup\, \Sigma_2}(2k) = \chi^*_{\Sigma_1}(2k)\chi^*_{\Sigma_2}(2k).$$

Invariance: *Suppose $\Sigma_1 \cong \Sigma_2$; then*

$$\chi_{\Sigma_1}(2k+1) = \chi_{\Sigma_2}(2k+1) \textit{ and } \chi^*_{\Sigma_1}(2k) = \chi^*_{\Sigma_2}(2k).$$

Switching Invariance: *For every switching function ζ,*

$$\chi_\Sigma(2k+1) = \chi_{\Sigma^\zeta}(2k+1) \textit{ and } \chi^*_\Sigma(2k) = \chi^*_{\Sigma^\zeta}(2k).$$

Deletion-Contraction:

$$\chi_\Sigma(2k+1) = \chi_{\Sigma\setminus e}(2k+1) - \chi_{\Sigma/e}(2k+1)$$

if e is not a half-edge or a negative loop, and

$$\chi^*_\Sigma(2k) = \chi^*_{\Sigma\setminus e}(2k) - \chi^*_{\Sigma/e}(2k).$$

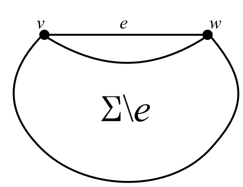

FIGURE J.2.
[[LABEL 1119image2]]

Proof. Unitarity holds true by general agreement about functions with domain $\varnothing$ (the empty function). Nullity and invariance are obvious. To prove switching invariance we use Lemma J.1.

The hard part is to prove the deletion-contraction property.

To prove $\chi_\Sigma(2k+1) = \chi_{\Sigma\setminus e}(2k+1) - \chi_{\Sigma/e}(2k+1)$, we start by coloring $\Sigma \setminus e$ properly in k colors. If $\gamma(v) \neq \gamma(w)$, then Σ is properly colored (and otherwise if $\gamma(v) = \gamma(w)$ Σ is not colored properly). We can contract γ to $\gamma/e : V(\Sigma/e) \to \Lambda_k$ such that $\gamma/e(v_e) = \gamma(v)\gamma(w)$. To prove that γ/e is a proper coloration of Γ/e. An improper edge in Σ/e must be incident with v_e. If it is a link $v_e u$, then it was a link vu or wu, therefore it is proper. If v_e is a loop $v_e v_e$, then it was a loop vv or ww or link vw, therefore it is proper since the endpoint colors are the same in Σ and Σ/e. If v_e is a half edge $f : v_e$, then it was $f : v$ or $f : w$ in Σ, therefore it it proper since the endpoint colors are the same in Σ and Σ/e. If v_e is a loose edge, then it was a loose edge in Σ. Conversely, every proper coloration of Σ/e pulls back to a proper coloration of $\Sigma \setminus e$ where $\gamma(v) = \gamma(w)$. So the number of proper colorations of $\Sigma \setminus e$ equals the sum of the number of proper colorations of Σ and Σ/e. Therefore our formula follows.

[MISSING: proof for χ^*.] □

In signed graph coloring the color 0 is special because it has no sign. Some colorations will include the 0 color, while other colorations, which we will call "zero-free", do not use the color 0. The number of zero-free proper k-colorations of Σ is denoted by $\chi^*_\Sigma(2k)$ or $\chi^{\mathrm{b}}_\Sigma(2k)$; the function $\chi^*_\Sigma(\lambda)$ is called the *zero-free chromatic polynomial* or *balanced chromatic polynomial* of Σ. That it is a polynomial follows from (The reason for the second name will appear in Theorem **??** with amplified explanation in Chapter III [**gains chapter**].) Nov 21 (draft): Peter Cohen and Thomas Zaslavsky

Definition J.4. [[LABEL T:1121full]] A graph is *full* if every vertex supports at least one unbalanced edge. The notation $\Sigma^\bullet$ denotes a signed graph Σ with an unbalanced edge adjoined to every vertex that did not already support one.
[THIS APPEARS EARLIER.]

Theorem J.4. *[[LABEL T:1121dczero-free]] The number of zero-free proper colorations of a signed graph satisfied the deletion-contraction law*

$$\chi^*_\Sigma(\lambda) = \chi^*_{\Sigma\setminus e}(\lambda) - \chi^*_{\Sigma/e}(\lambda).$$

[DOES THIS DUPLICATE A PREVIOUS THEOREM?]

[picture with v, e, sigma e, etc]

Proof. In the case of $\Sigma \setminus e$, vertex v has color $\neq 0 \iff$ it is a proper coloration of Σ. Vertex v has color $= 0 \iff$ it has a proper coloring of $\Sigma \neq \Sigma \setminus e$.
[FIX THIS PROOF.] □

[THIS IS PART OF THE PROOF?]
The zero-free polynomial of Σ is close to the ordinary chromatic polynomial of the full form of Σ. The relationship is that

$$\chi^*_{\Sigma^\bullet}(2k) = \chi_{\Sigma^\bullet}(2k+1) = \chi_{\Sigma^\bullet\setminus e}(2k+1) - \chi_{\Sigma^\bullet/e}(2k+1).$$

Lemma J.5. *[[LABEL T:1121fullcontract]] Any contraction $\Sigma^\bullet/e$ is full.*

Observe that, if Σ itself contains no unbalanced edges, then $\Sigma^\bullet \setminus e$ is full $\iff$ e is a balanced edge. Therefore, if e is a balanced edge,

$$\begin{aligned}\chi^*_\Sigma(2k) &= \chi_{(\Sigma\setminus e)^\bullet}(2k+1) - \chi_{(\Sigma/e)^\bullet}(2k+1)\\ &= \chi^*_{\Sigma\setminus e}(2k) - \chi^*_{\Sigma/e}(2k)\end{aligned}$$

Theorem J.6 (Polynomiality)**.** *[[LABEL T:1121chromatic polyonmialit]] The chromatic and zero-free chromatic functions $\chi^{[*]}_\Sigma(\lambda)$ are polynomial functions of $\lambda = 2k+1$ (if general) or $2k$ (if zero-free), monic, of degree n, of the form $\chi_\Sigma(\lambda) = \lambda^n - a, \lambda^{n-1} + \cdots + (-1)^{n-b(\Sigma)} a_{b(\Sigma)} \lambda^{b(\Sigma)}$ or $\chi^*_\Sigma(\lambda) = \lambda^n - a^*, \lambda^{n-1} + \cdots + (-1)^{n-b(\Sigma)} a^*_{b(\Sigma)} \lambda^{b(\Sigma)}$ where all a_i or $a^*_i > 0$.*

It should be noted that a_1 is the number of edges in Σ, and a_1^* is the number of links in Σ, if Σ is simply signed in the sense that there do not exist any parallel links with the same sign and no vertex has two (or more) unbalanced edges.

Proposition J.7 (Subset Expansion). *[[LABEL T:1121chromaticsubset]] The chromatic polynomials have the subset expansions*

$$\chi_\Sigma(\lambda) = \sum_{S\subseteq E} (-1)^{|S|}\lambda^{b(S)} \quad \textit{and} \quad \chi^*_\Sigma(\lambda) = \sum_{\substack{S\subseteq E\\ \text{balanced}}} (-1)^{|S|}\lambda^{b(S)}.$$

Definition J.5. [[LABEL T:1121unbalcount]] The number of unbalanced components of a graph (or subgraph) is $u(\Sigma)$. This equals the number of components minus the number of balanced components; $u(\Sigma) = c(\Sigma) - b(\Sigma)$.
[**MOVE TO EARLIER; mention here.**]

We can write the ordinary and zero-free chromatic polynomials of a signed graph in terms of $u(S)$ as

$$\chi_\Sigma(\lambda) = \sum_{S\subseteq E} (-1)^{|S|}\lambda^{b(S)}1^{u(S)} \quad \text{and} \quad \chi^*_\Sigma(\lambda) = \sum_{S\subseteq E} (-1)^{|S|}\lambda^{b(S)}0^{u(S)}.$$

This formulation shows there is a comprehensive polynomial to incorporate both. Define the *total chromatic polynomial* as

$$\chi_\Sigma(\lambda, z) := \sum_{S\subseteq E} (-1)^{|S|}\lambda^{b(S)}z^{u(S)}.$$

When $z = 1$ we have the chromatic polynomial and when $z = 0$ we have the zero-free chromatic polynomial.

Proof. [**from del/con, let e be any balanced edge**]

$$\sum_{S\subseteq E} (-1)^{|S|}\lambda^{b(S)} = \sum_{S\subseteq E\setminus e} (-1)^{|S|}\lambda^{b(S)} + \sum_{S\subseteq E/e} (-1)^{|S|}\lambda^{b(S)}.$$

In the zero-free case, where e is not an unbalanced edge,

$$\sum_{S\subseteq E} (-1)^{|S|}\lambda^{b(S)} = \chi^*_{(\Sigma\setminus e)}\lambda + \sum_{T\subseteq E\setminus e} (-1)^{|S|+1}\lambda^{b(T\cup e)}.$$

[**we had a previous lemma that said S is balanced in sigma, iff S-R is balanced in Sigma/R, need to cite it**]

Therefore,

$$\sum_{T\subseteq E\setminus e} (-1)^{|T|}\lambda^{b_{\Sigma/e}(T)} = \sum_{T\subseteq E\setminus e} (-1)^{|T|+1}\lambda^{b_{\Sigma/e}(T)} = \sum_{T\subseteq E\setminus e} (-1)^{|T|+1}\lambda^{b_\Sigma(T\cup e)}$$

as we needed.

The components of $T\cup e$ don't become disconnected when we contract a balanced edge, therefore the number of balanced components of T is the same as the number of balanced components of $T\cup e$; that is $b_{\Sigma/e}(T) = b_\Sigma(T\cup e)$.

Suppose that Σ has only unbalanced edges. Then Σ only contains half edges and negative loops so it has one component per vertex. In other words, $c(\Sigma) = |V|$. All vertices in Σ

are either k_1 or (k_1+e), an unbalanced edge. The (k_1+e) edges are full, by definition. So therefore, the coloration is the sum of the coloration of the k_1's and the (k_1+e)'s;

$$\chi_{\Sigma}^{[*]}(\lambda) = \chi_{k_1}^{[*]}(\lambda)^{n-i} + \chi_{k_1^\bullet}^{[*]}(\lambda)^{i} = \lambda^{n-i} + (\lambda^i \text{ if 0-free}, (\lambda-1)^i \text{ if all colorations}).$$

Being disconnected, the sum of the coloration is the same as the product of the sums;

$$\sum_{S\subseteq E}(-1)^{|S|}\lambda^{b(V,S)} = \prod_{i=1}^{n}\sum_{S_i\subseteq E_i}(-1)^{|S_i|}\lambda^{b(V_i,S_i)}. \qquad \square$$

Nov 24: Jackie Kaminski

J.4. **Counting acyclic orientations.** [[LABEL 2.acycliccount]]

We now take up the generalization to signed graphs of Stanley's theorem, Theorem **??** interpreting the chromatic polynomial at negative arguments.

The sesquijection of acyclic orientations.

The key to everything is the generalization of the sesquijection, or 1:1/2:2 correspondence, of acyclic orientations of a graph (Lemmas **??** and **??**) to a sesquijection between acyclic orientations of Σ and those of $\Sigma \setminus e$ and Σ/e.

Definition J.6. Two walks,

$$W = v_0, e_1, v_1 \dots v_{l-1}e_l v_l \text{ and } W' = v_0', e_1', v_1' \dots v_{l'-1}' e_{l'}' v_{l'}',$$

are *internally disjoint* if each internal vertex of one walk, W or W', is not in the other, respectively W' or W. That is, no $v_j =$ any v_j' except that v_0, v_l may be $v_0', v_{l'}'$.

Recall that $\mathrm{AO}(\Sigma)$ is the set of all acyclic orientations of Σ.

Lemma J.8. *[*[LABEL L:1124aonumber]] $|\mathrm{AO}(\Sigma)| = |\mathrm{AO}(\Sigma\setminus e)| + |\mathrm{AO}(\Sigma/e)|$ *for e not a positive loop or loose edge. [**corrected**].*

Proof. [**This is as much of the proof as we did Monday**]

Let α be an acyclic orientation of $\Sigma\setminus e$ with e not a positive loop or loose edge. This means e is a link or half edge or negative circle. If e is a link, we assume we have used switching so e is positive. We would like to show that there is a 1:1/2:2 correspondence (a *sesquijection*) between $\mathrm{AO}(\Sigma)$ and $\mathrm{AO}(\Sigma\setminus e)\cup \mathrm{AO}(\Sigma/e)$. We will show that the 0, 1, or 2 acyclic extensions of α to Σ are in sesquijective correspondence to α as an element of $\mathrm{AO}(\Sigma\setminus e)$ and possibly $\mathrm{AO}(\Sigma/e)$.

As we consider adding e back to $\Sigma\setminus e$, there are two possible orientations for it, $e{:}v\vec{w}$ and $e{:}w\vec{v}$, and each of these orientations may or may not contain a cycle. This gives us four types of situation, which really reduce to three:

- Type II: both orientations of e produce acyclic orientations of Σ,
- Type I: adding $e{:}v\vec{w}$ produces an acyclic orientation of Σ, but adding $e{:}w\vec{v}$ produces a cyclic orientation of Σ,
- Also Type I: adding $e{:}v\vec{w}$ produces a cyclic orientation of Σ, but adding $e{:}w\vec{v}$ produces an acyclic orientation of Σ,
- Type O: both orientations of e produce cyclic orientations of Σ,

where the middle two cases can be treated identically.

Since α (and α extended to include e in Σ and α "restricted" to Σ/e) are the only orientations in question, we will drop the cumbersome arrows in the notations $\vec{\Sigma}$, $\overrightarrow{\Sigma/e}$, etc.

Type II: Both orientations of e produce acyclic orientations of Σ.

In other words α extends to two acyclic orientation of Σ. Since α is an acyclic orientation of $\Sigma \setminus e$, we simply want to show that α applied to Σ/e is also acyclic. Then we will have a 2:2 correspondence between the two acyclic orientations extending α in $\mathrm{AO}(\Sigma)$ and the two acyclic orientations of $\Sigma \setminus e$ and Σ/e implied by α. We will look at two subcases: when e is a link (which we assume is positive by switching), and when e is a negative loop or half edge.

Subcase A: e is a positive link.

First we note that since e is positive link any consistently oriented walk W containing e will still be consistently oriented in Σ/e. Now, for a proof by contradiction, suppose that Σ/e contains an oriented cycle. Since $\Sigma \setminus e$ is acyclic, this cycle must contain the vertex v_e, let $W = v_e e_1 v_2 \cdots v_{k-1} e_k v_e$ be a closed walk around the oriented cycle in Σ/e beginning at v_e.[5] Now consider the closed walks in Σ. Notice that if e_1 and e_k are both incident to v or both incident to w in Σ, then the closed oriented walk W is an oriented circle in $\Sigma \setminus e$, which contradicts our assumption that $\alpha \in \mathrm{AO}(\Sigma \setminus e)$. So one of e_1 and e_k is incident to v and the other to w, by choice of notation, we choose e_1 incident to v and e_k incident to w.

Now we consider two coherent closed walks in Σ that contain e in opposite orientations, namely,

$$W_1 = w, e{:}v\vec{w}, v, e_1, v_2, \ldots, v_{k-1}, e_k, v$$

and

$$W_2 = w, e{:}w\vec{v}, v, e_1, v_2, \ldots, v_{k-1}, e_k, w.$$

Since W was a walk around a consistently oriented circle **[MORE?]**

If W in Σ/e was oriented so e_1 left v_e and e_{k-1} entered v_e then W_2 is consistently oriented in Σ. Furthermore, since $\sigma(e) = +$ (by assumption), the circle(s) (and paths) of W_2 are still circles(s) (and paths) in $W_2 \cup e{:}w\vec{v}$ with the same sign(s). Therefore $W_2 \cup e{:}w\vec{v}$ is a cycle in Σ, and since it was oriented we have an oriented cycle in Σ, which is contrary to the assumptions of Subcase A. Furthermore, if we don't have $W \in \Sigma/e$ oriented so e_1 left v_e and e_{k-1} entered v_e then $W \in \Sigma/e$ was oriented so e_1 enters v_e and e_{k-1} leaves v_e (since W is consistently oriented in Σ/e these are the only two options). In this case we have an identical argument with $W_1 \cup e{:}v\vec{w}$, and we reach the same contradiction.

Therefore Σ/e does not contain an oriented cycle, and in particular α "restricted" to Σ/e is acyclic. Therefore we have the two acyclic extensions of α to Σ in 2:2 correspondence with the two acyclic orientations α of $\Sigma \setminus e$ and the "restricted" α on Σ/e.

Subcase 2: e is a negative loop or half edge. To simplify this proof we will assume that e is actually a half edge with vertex v.

This subcase is similar to the first. For proof by contradiction we assume that $\alpha \in \mathrm{AO}(\Sigma \setminus e)$ extends to two acyclic orientations of Σ, namely $\alpha \cup e{:}v\vec{w}$ and $\alpha \cup e{:}w\vec{v}$, but that α "extended" to Σ/e contains an oriented cycle.

[5]If the circuit is a handcuff circuit, then this walk will simply repeat the circuit path.

We note that this cycle must use a half edge f created by contracting e, in other words, f was a v, v_1 link in Σ.[6] If this isn't the case it is immediate that we have an oriented cycle in Σ. Now we note that f is itself an unbalanced circle, so the oriented cycle containing f in Σ/e must be of the negative handcuff type. So there exists a circuit path P beginning at v leading to another unbalanced circle C where $P \cup C \cup$ the half edge f is consistently oriented in Σ/e. Now we notice that since f was a v, v_1 link $C \cup P \cup f_\Sigma \cup e$ is a cycle in Σ. Furthermore, this cycle is consistently oriented for $C \cup P \cup$ the half of f at v_1, than no matter which way f is oriented at v, one of the orientations of e is consistent with f at v, meaning we have an oriented cycle in Σ, which contradicts our assumption.

Therefore Σ/e does not contain an oriented cycle, and in particular α "restricted" to Σ/e is acyclic. Therefore we have the two acyclic extensions of α to Σ in 2:2 correspondence with the two acyclic orientations α of $\Sigma \setminus e$ and of α on Σ/e.

Thus we have proved our 2:2 correspondence for Type II.

[**THE PROOF IS IN CASES.**
WHERE TO FIND THE CORRECT PROOFS OF THE CASES (guide for who writes what):
Case I. e is a positive loop or loose edge. (Trivial; see 11/24 or 11/26.) Case II. e is a link (+ by switching). Case III. e is a half edge or negative loop.
Case II has 3 types. We have an acyclic orientation α of $\Sigma \setminus e$. Type Two. $\alpha \cup e$ is acyclic in both orientations of e. Type One. Only in one orientation. Type Zero. Not in any orientation.
I think Types Two, One were dealt with on 11/26 with some supplementation on 12/1.
Type Zero was treated on 11/26 and 12/1. It has three cases. Case 1. P is a path. (Done 11/26.) Case 2. P is a handcuff with e in the connecting path. Case 3. Same with e in one of the circles. These were treated on 12/1.]

[The following should all be redone by 11/26 and 12/1 people, and is provided here just in case it helps:]

Type I: Adding $e{:}\vec{vw}$ produces an acyclic orientation of Σ, but adding $e{:}\vec{wv}$ produces a cyclic orientation of Σ.

Let P be a closed walk in $\Sigma \setminus e$ s.t. $P \cup e{:}\vec{wv}$ is a an oriented circuit in Σ. This implies that P is oriented consistently (within P) from v to w. We would like to show that α is acyclic when extended to Σ, but not acyclic on Σ/e (for either orientation of e), thus giving a 1:1 correspondence.

We now look at 3 subcases,

- Subcase A: $P \cup e$ is a positive circle
- Subcase B: $P \cup e$ is a negative handcuff, with e in the circuit path of the circuit
- Subcase C: $P \cup e$ is a negative handcuff, with e in a negative circle of the circuit

[6]Note that f could not have been a negative loop or half edge at v. If it were a half edge or negative loop, then f together with one of the orientations of e would yield an oriented cycle in Σ. And if f was a positive loop then f (with any orientation) is an oriented cycle in Σ.

Subcase A: We consider α extended to σ/e (with either orientation of e). Note that since e is a positive link by assumption, this contraction makes sense. Furthermore, the contraction doesn't alter the sign of the circle $P \cup e$, by Lemma E.4 (For S balanced in Σ and $T \subseteq E \setminus S$. Then $S \cup T$ is balanced in $\Sigma \iff T$ is balanced in Σ/S.) P is balanced (positive) in Σ/e. Furthermore, since the contraction didn't affect the internal vertices of P, the edges of the circle P is oriented coherently at all vertices except e_v. And since the path P in $\Sigma \setminus e$ was oriented from v to w, the circle P is oriented coherently in Σ/e. Therefore Σ/e is cyclic.

Since α extended to Σ is acyclic for exactly one orientation of e by assumption, we have a one to one correspondence between $\mathrm{AO}(\Sigma \setminus e)$ and $\mathrm{AO}(\Sigma)$, which is in fact one to one between $\mathrm{AO}(\Sigma \setminus e)$ and $\mathrm{AO}(\Sigma) \cup \mathrm{AO}(\Sigma/e)$ for Case 2 C.

For the other subcases, we need a sublemma.

Lemma J.9. *[*[LABEL L:1124 SubLemma]] *For e a positive link, and $P \cup e$ a coherently oriented walk in $\vec{\Sigma}$, then P is a coherently oriented walk in $\overrightarrow{\Sigma/e}$.*

proof of sublemma. On Wed? □

Subcases B & C: (If the SubLemma is true we can treat A, B, C together, otherwise we need to do work here on wednesday)

This concludes Type I.

Type O: The acyclic orientation of $\Sigma \setminus e$ extends only to cyclic orientations of Σ.

We wish to show that this is impossible, that there are no acyclic orientations of $\Sigma \setminus e$ with $\Sigma \cup e{:}\vec{vw}$ and $\Sigma \cup e{:}\vec{wv}$ cyclic orientations of Σ. We will do so by contradiction.

[THIS ENTIRE PROOF (OF CASE 3) IS SUPERSEDED.]

Let $P{:}\vec{wv}$ and $Q{:}\vec{vw}$ be oriented walks in $\Sigma \setminus e$ (oriented by α of course) s.t. $P \cup e{:}\vec{vw}$ is a coherently oriented cycle, and similarly $Q \cup e{:}\vec{wv}$ is a coherently oriented cycle. Furthermore, the concatenation PQ is a coherently oriented closed walk. Now we wish to show that there is subwalk of PQ that is a coherently oriented cycle. To this end, we look 2 cases,

- Subcase A: P, Q are internally disjoint
- Subcase B: P, Q are not internally disjoint

Subcase A: Then we have several subsubcases. (Note that we have omitted the cases where the rolls of P and Q are simply reversed.) In each of these cases we will find a circuit in $\Sigma \setminus e$, giving us a contradiction.

- $P \cup e$ is a positive circle $Q \cup e$ is a positive circle
- $P \cup e$ is a handcuff with e in one of the negative circles and $Q \cup e$ is a positive circle
- $P \cup e$ is a handcuff with e in one of the negative circles and $Q \cup e$ is a handcuff with e in one of the negative circles
- $P \cup e$ is a handcuff with e in one of the negative circles and $Q \cup e$ is a handcuff with e in the circuit path
- $P \cup e$ is a handcuff with e in the circuit path $Q \cup e$ is any circuit

[This all needs pictures and a few words about why the orientations are still coherent.]

Subcase B: By choice of notation, Q meets P internally at some vertex u. So u is a vertex in Q and an internal vertex in P.

Note that this includes the possibility that $u = v$ or $u = w$, since v (or w) could be internal to P, and v (and w) are vertices in Q.

[This is where things got really messy in class. I haven't straightened them out yet. [They probably can't be salvaged without re-doing. – TZ]]

□

Nov 26 (draft): Simon Joyce

Lemma J.10. [[LABEL L:1126 Lemma1]] *Given a link or unbalanced edge e and an orientation β of a signed graph Σ such that $\Sigma \backslash e$ is acyclic. Then β is acyclic if and only if β/e is acyclic.*

Proof. Necessity is the subject of J.11. Sufficiency is the subject of J.12. We prove the contrapositive in both cases. **[Double check that I have used the correct logical statements]**

Lemma J.11. [[LABEL L:1126 Lemma1A]] *Given a link or an unbalanced edge e contained in a cycle C in Σ, then C/e is a cycle in Σ/e.*

Proof. We may assume $\Sigma = C$ since this is the only part of the graph we care about.

The cases where e is a link, which we may assume is positive by switching are shown in Figure J.4. It is clear in all cases that C/e is a cycle.

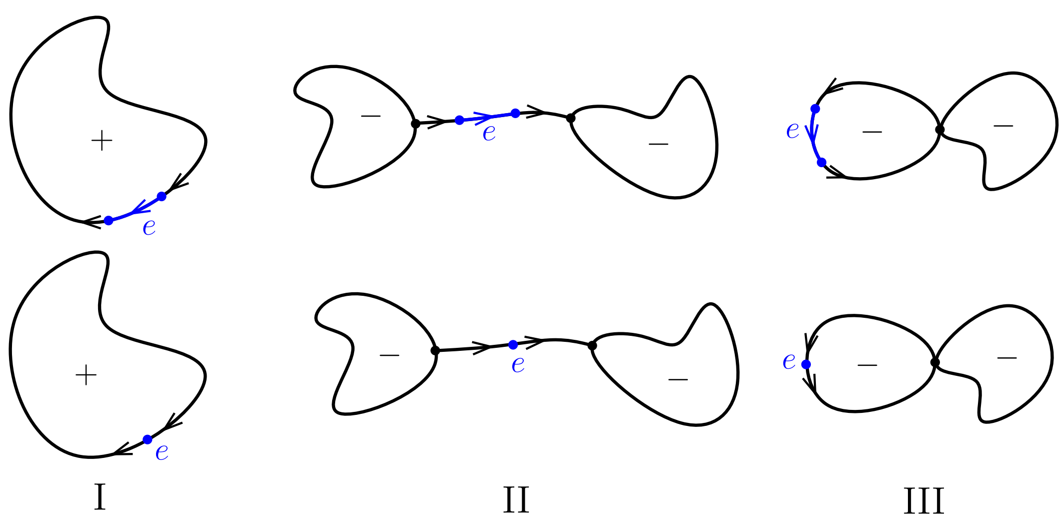

FIGURE J.3.

[[LABEL F:1126link]]

The cases where e is a half edge are shown in Figure J.4. The case where e is a negative loop has exactly the same result. Since e is an unbalanced edge it must be part of a handcuff. In case I there is a non-empty path leading from e to the other negative circle or unbalanced edge in the hand cuff and clearly C/e is also a cycle. In case II e is incident to a vertex in a negative circle. Here again we see C/e results in a cycle. In case III C is made up of two unbalanced edges and the contraction gives a loose edge which is a cycle.

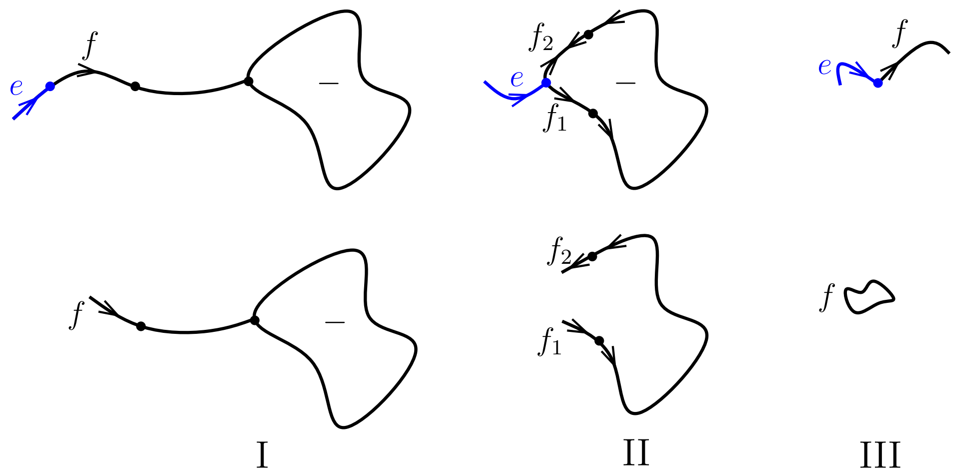

FIGURE J.4.

[[LABEL F:1126ube]]

□

Lemma J.12. *[*[LABEL L:1126 Lemma1B]] *For e a link or unbalanced edge of Σ, if β is acyclic in $\Sigma\backslash e$ and cyclic in Σ/e, then β is cyclic in Σ, or β with the orientation of e reversed is cyclic in Σ (or both).*

Proof. Since $\Sigma\backslash e$ is acyclic, if Σ contains a cycle it must be of the form $e \cup E(C)$ where C is a cycle in Σ/e.

Suppose e is a link which we assume is positive by switching. $C \cup e$ forms a a circle or a handcuff such that all edges in C are oriented consistently in β. The orientation of C determines a unique orientation of e that is consistent C such that $C \cup e$ forms a cycle in Σ.

Now suppose e is an unbalanced edge. Now $C \cup e$ must form a handcuff with e as one of it's unbalanced ends. Suppose $C \cup e$ forms a loose handcuff. Assume the edge adjacent to e is positive by switching. Then the orientation of C determines a unique orientation of e such that $C \cup e$ forms a cycle in Σ.

Now suppose $C \cup e$ forms a tight handcuff and C is not an unbalanced edge in Σ. If we assume all the edges of C that are not adjacent to e in Σ are positive by switching, then one of the edges adjacent to e is positive and one is negative. Since C is consistently oriented this will uniquely determine the orientation of e such that $C \cup e$ forms a cycle in Σ.

If C is an unbalanced edge in Σ, then C is a loose edge in Σ/e, so it has no orientation. Then either orientation of $C \cup e$ forms a cycle as long as C is oriented consistently with e.

□

This completes the proof of J.10.

□

Now suppose both orientations of e give a cyclic orientation of Σ. J.12 shows that Σ/e is also cyclic. To get our 2:2 correspondence in the is case we must show that $\Sigma\backslash e$ must also be cyclic in this case.

Lemma J.13. *[*[LABEL L:1126 Lemma2A]] *If α is an orientation of $\Sigma\backslash e$ such that $\alpha\cup \overrightarrow{e}$ and $\alpha\cup \overleftarrow{e}$ are cyclic on Σ, then α is cyclic on $\Sigma\backslash e$*

Proof. Let v and w be the endpoints of e. Let P and Q be the edge sets in $\Sigma \backslash e$ such that $P \cup e{:}vw$ and $Q \cup e{:}wv$ form cycles in Σ. First we assume $P \cup e$ is a positive circle [**This is case 1 in the next days notes**]. The means we can consider P as a consistently oriented path from w to v since $P \cup e{:}vw$ is a cycle.

First suppose the vertices of P and Q don't intersect except at v and w, then $Q \cup P$ forms a cycle since P is consistently oriented from w to v. The three possible cases are shown in Figure J.4. In case I, $Q \cup e{:}wv$ forms is a consistently oriented positive circle. In cases II and III, $Q \cup e{:}wv$ forms a consistently oriented handcuff where e is in the connecting path in case II and in one of the negative circles in case III.

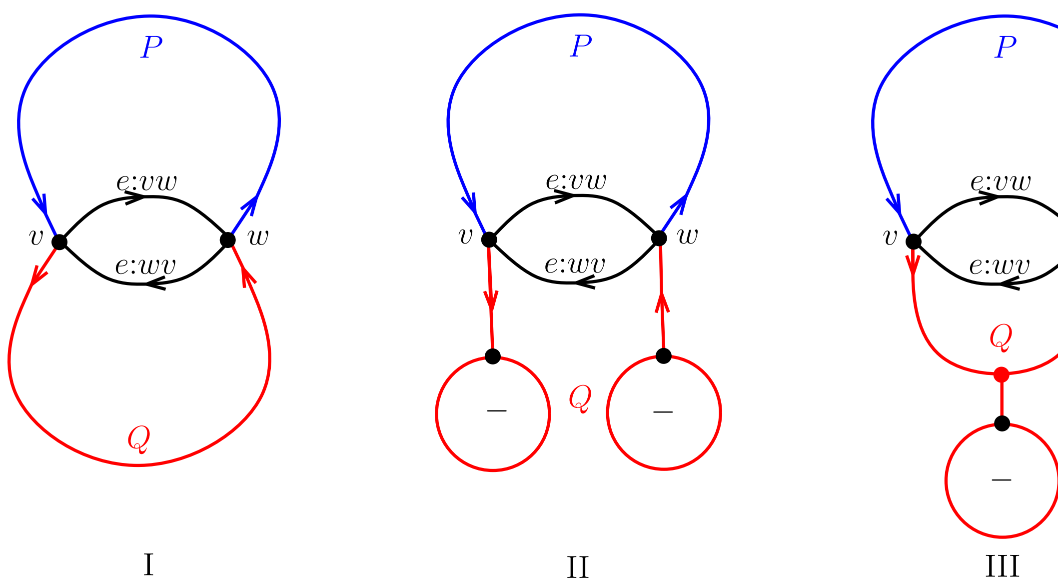

FIGURE J.5.

[[LABEL F:1126Fig1]]

Now suppose the vertices of P and Q do intersect at a vertex other than v and w. Notice the first edge of Q incident with v (or w) cannot be in P since it's orientation is opposite to the edge in P incident with v (or w). Starting at v, follow the edges of Q so that the orientation of the edges is always pointing away from v. Stop at the first vertex x where Q intersects P or when you return to a vertex of Q. If Q intersects with P at x then these edges form a consistently oriented path from v to x. But there will also be a subpath of P that is consistently oriented from x to v. Therefore α contains a consistently oriented positive circle. If you hit a vertex of Q before you intersect with P then go to w and follow the edges of Q so that the orientation of the edges is always toward w. This time you must hit a vertex x of P before a vertex of Q since otherwise P and Q would not intersect and we'd be back in the previous case. These edges of Q form a path consistently oriented from x to w. But there is a subpath of P consistently oriented from w to x. Therefore α again contains a consistently oriented positive circle.

This proof is continued the next day.

□

A *coherent balloon* consists of a negative circle (or half edge) C and a path P of any length (possibly zero) that is disjoint from C except at one end v, oriented so that $C \cup P$ has no Dec 1: Nate Reff source or sink except one. This vertex is called the *tip* of the balloon. It is easy to see that an oriented negative circle (or half edge) cannot be coherent at every vertex; thus, if it is coherent at the largest number of vertices, there is a unique incoherent vertex. This must be the vertex common to P and C; we call it the *jointure* of the coherent balloon. (When P has length 0, we define it to be $P = v$ and the tip is v.) Since C is negative and is coherent everywhere but at v, it cannot be coherent at the v; thus, the orientation of C determines that of P and the tip is the only source or sink. The oriented signed graph seen in Figure J.6 is an example.

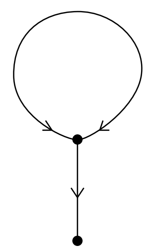

FIGURE J.6.
[[LABEL F:1201image1]]

Lemma J.14. *[*[LABEL L:1201Lemma1ImprovedLemma4]] *Suppose we have a coherently oriented balloon with tip w. Extend coherently from w until you meet the (extended) balloon. Then the extended balloon will contain a cycle.*

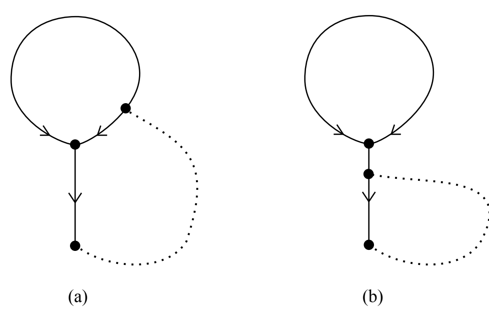

FIGURE J.7.
[[LABEL F:1201image2]]

Proof. Let v be the jointure of the balloon. When the extension of P hits the extended balloon, the configuration looks like one of the two types seen in Figure J.7. (If the hit point is v, we are in the second type.) Each type contains a cycle, as we explain next. The arguments are based on the description in Lemma **??** of a closed walk that is coherent at every internal vertex.

In diagram (a), let x be the point at which the extended path meets the balloon. Follow the circle from x in the direction that makes v coherent; then when we arrive back at v we have a coherent, hence positive, circle, which makes a balanced cycle.

In diagram (b), when we hit the extended path, say at y, we either form a positive circle, which is coherent because it is coherent by definition at every vertex other than y, or a negative circle, which means the entire figure is a handcuff and an unbalanced cycle. □

Proof of Sesquijection Lemma, continued. Suppose that α is an orientation of $\Sigma \setminus e$. We are trying to prove that if $\alpha \cup \vec{e}$ and $\alpha \cup \overleftarrow{e}$ are both cyclic, then α is cyclic (ie "type zero"REF??? does not exist from our previous discussion). We are basically assuming that we have a link $e : vw$. More specifically we assume $e : vw$ is a positive link, $P \cup (e : v\vec{w})$ is a cycle, and $Q \cup (e : \overleftarrow{wv})$ is a cycle.

Case 1: P is a path. (We already did this case.)

Case 2: P is a handcuff and e is in its connecting path. We may also assume that Q is not a path, since that was taken care of in Case 1. Look at Figure J.8. By Q_w we mean the part

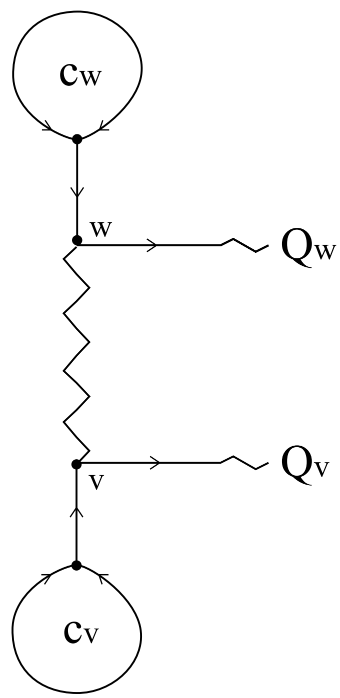

FIGURE J.8.
[[LABEL F:1201image3]]

of Q we can get to by backtracking coherently along Q from w. If Q_w hits P_w, then lemma J.14 tells us that it is a cycle in $\Sigma \setminus e$ (since we are extending coherently and we must hit somewhere). By Q_v we mean to forward track coherently along Q from v. Then the result is similar.

This means that Q must be one of the shapes seen in Figure J.9.

If Q_w does not hit P_w, then $P_w \cup Q_w$ is a cycle.

Case 3: P is a handcuff and e is in one of its circles. We may also assume Q has the same type, or we would be in Case 1 or 2. Looking at Figure J.10, we call the red path P_w, and the blue path Q_w. If we backtrack along Q_w and do not hit P then we will close up and will be back in Lemma J.14. So this is similar to Case 2.

The essential part of Cases 2 and 3 is that P_w and Q_w are two balloons that meet coherently at their tips.

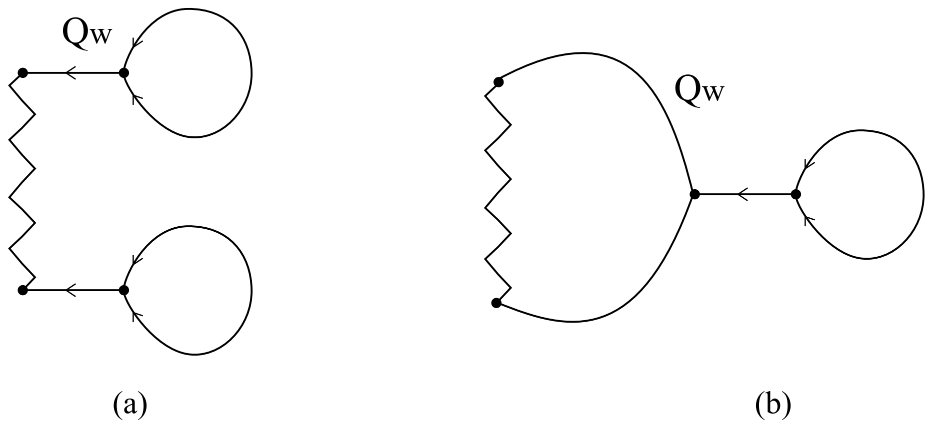

FIGURE J.9.
[[LABEL F:1201image4]]

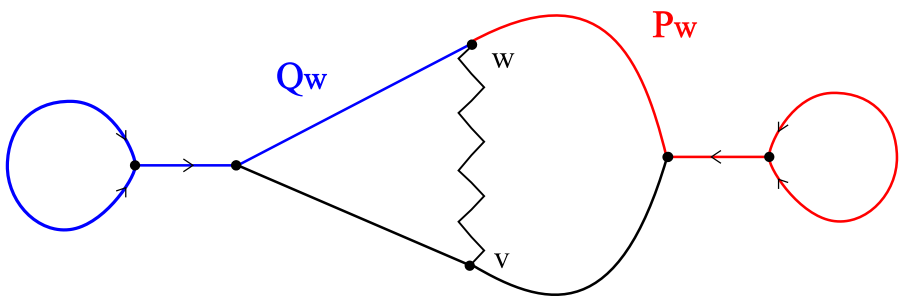

FIGURE J.10.
[[LABEL F:1201image5]]

Corollary J.15. *[*[LABEL C:1201CorollaryToCaseII]] *The union of two coherent balloons, not necessarily internally disjoint, that are joined coherently at their tips, contains a cycle.*

Case III: Suppose e is an unbalanced edge at v, there exists a cycle $P \cup \vec{e}$ in $\alpha \cup \vec{e}$, and there exists a cycle $Q \cup \overleftarrow{e}$ in $\alpha \cup \overleftarrow{e}$ (See Figure J.11).

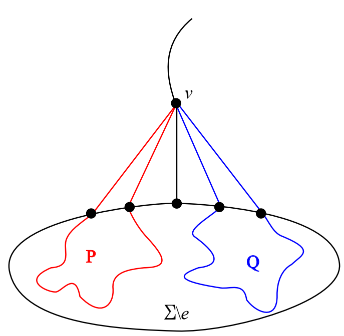

FIGURE J.11.
[[LABEL F:1201image6]]

Possible pairs (γ, α):		
	Proper	Compatible
(a) Extroverted	$\gamma(w) + \gamma(v) > 0$	$\gamma(w) + \gamma(v) \geq 0$
(b) Positive	$\gamma(v) < \gamma(w)$ (This includes a positive loop $v = w$)	$\gamma(v) \leq \gamma(w)$
(c) Introverted	$\gamma(v) + \gamma(w) < 0$ (For a negative loop $2\gamma(v) < 0$	$\gamma(v) \leq \gamma(w)$
(d) Introverted half edge	$\gamma(v) < 0$	$\gamma(v) \leq 0$
(e) Extroverted half edge	$\gamma(v) > 0$	$\gamma(v) \geq 0$

[[LABEL Tb:1201pair types]]

Therefore, P and Q are balloons and meet coherently. Apply Corollary J.15. Thus $\Sigma \setminus e$ has a cycle at v.

This concludes the proof of the Sesquijection Lemma. □

Proper and compatible pairs.

Recall that a coloration of Σ is a mapping $\gamma : V \to \Lambda_k$ where $\Lambda_k := \{0, \pm 1, \pm 2, \ldots, \pm k\}$. A zero-free coloration of Σ is a mapping $\gamma^* : V \to \Lambda_k^*$ where $\Lambda_k^* := \{\pm 1, \pm 2, \ldots, \pm k\}$.

A pair (γ, α) where γ is a coloration and α is an acyclic orientation, can be of any of the types in Figure J.12 defined in the following table.

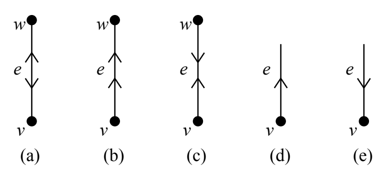

FIGURE J.12. An edge e which is (a) extraverted, (b) positive, (c) introverted, (d) an introverted half edge, (e) an extraverted half edge.
[[LABEL F:1201image7]]

Note: If α exists then there are no loose edges or positive loops.

Let $x(e) :=$ column of e in $\mathrm{H}(\Sigma, \alpha)$. We are saying that $x(e) \cdot \gamma > 0$ (proper), and $x(e) \cdot \gamma \geq 0$ (compatible).

The number of acyclic orientations.

Stanley's theorem on the number of acyclic orientations of a graph (Theorem **??**) extends to signed graphs. The original theorem is the special case of an all-positive signature. For a fixed $k \geq 0$, we write $a(\Sigma) :=$ the number of acyclic orientations, $a_2(\Sigma)$ for the number of compatible pairs, and $a_2^*(\Sigma)$ for the number in which the coloration is zero free. Let's make three important observations:

(1) Given an acyclic orientation α and a coloration γ, an edge e is proper if and only if $x(e) \cdot \gamma > 0$ and it is compatible with γ if and only if $x(e) \cdot \gamma \geq 0$. [**DEFINE** $x(e)$ **before this statement. Give a proof?—a lemma?—cite Table J.4?**]
(2) A proper pair (γ, α) is determined by γ. Therefore the number of proper pairs is equal to the number of proper k-colorations (or zero-free k-colorations), which is equal to $\chi_\Sigma(2k+1)$ (or $\chi_\Sigma^*(2k)$).

(3) If we have an improper compatible pair (γ, α), then γ is improper. In other words, if there exists an e such that $x(e) \cdot \gamma = 0$, then γ is improper. If such an e does not exist, then γ is proper.

Theorem J.16. *[[LABEL T:1201Theorem1StanleyType]] Let k be a non-negative integer. In a signed graph Σ, the number of compatible pairs of an acyclic orientation and a k-coloration is $(-1)^n \chi_\Sigma(-(2k+1))$. The number of compatible pairs with a zero-free k-coloration is $(-1)^n \chi^*_\Sigma(-2k)$.*

Proof. We proceed by induction on the number of links. For zero links we have $a_2(\Sigma) = \prod_v a_2(\Sigma{:}v)$ and $\chi_\Sigma(\lambda) = \prod_v \chi_{\Sigma:v}(\lambda)$. If there exists a link $e{:}vw$ then to prove $a_2(\Sigma) = a_2(\Sigma \setminus e) + a_2(\Sigma/e)$ (and for the zero-free case $a^*_2(\Sigma) = a^*_2(\Sigma \setminus e) + a^*_2(\Sigma/e)$), we may assume e is positive by switching and then deletion and contraction of χ_Σ and χ^*_Σ give us the result.

Let (γ, α) be a compatible pair in $\Sigma \setminus e$.
Case 1: e is proper in γ. Then α extends uniquely to Σ and we get a compatible pair (γ, α_Σ). Also, γ does not color Σ/e because $\gamma(v) \neq \gamma(w)$. Therefore we have a bijection of compatible pairs. In other words one pair in Σ corresponds to one pair in each $\Sigma \setminus e$ and Σ/e.
Case 2: e is improper in γ. Therefore $\gamma(v) = \gamma(w)$ and γ colors Σ/e. If we add e to α, we have

$$\mathrm{AO}(\Sigma) \overset{\text{sesquijection}}{\longleftrightarrow} \mathrm{AO}(\Sigma \setminus e) \cup \mathrm{AO}(\Sigma/e),$$

where the left-hand side is thought of as the number of extensions and the right-hand side is thought of as α applied to $\Sigma \setminus e$ and also to Σ/e.

Since we have a sesquijection this gives us the correct numbers for the compatible pairs with k-colorations. Notice that if γ is zero-free then γ/e is zero-free and vice versa, so the proof is the same. Therefore the theorem is proved. □

12/3: Kaminski

[MISSING NOTES]

Dec 5 (draft): Peter Cohen and Thomas Zaslavsky

J.5. **The dichromatic and corank-nullity polynomials.** [[LABEL 2.dichromatic]]

The algebraic form of the chromatic polynomials, i.e., the subset expansion in Theorem J.7, allows us to generalize greatly. The dichromatic polynomials of a signed graph, like that of a graph, are two-variable generalizations of the chromatic polynomials that have combinatorial properties of their own. A modification, the corank-nullity polynomials, have slightly but significantly different properties.

Dichromatic polynomials.

We begin with the algebraic definitions of three dichromatic polynomials.

Definition J.7. [[LABEL D:1205dichromatic]] The *(ordinary) dichromatic polynomial* of a signed graph Σ is

$$Q_\Sigma(u,v) := \sum_{S \subseteq E} u^{n-b(S)} v^{|S|-n+b(S)}.$$

The *balanced dichromatic polynomial* is

$$Q^*_\Sigma(u,v) := \sum_{\substack{S \subseteq E \\ \text{balanced}}} u^{n-b(S)} v^{|S|-n+b(S)}.$$

The *total dichromatic polynomial* is

$$Q_\Sigma(u,v,z) := \sum_{S\subseteq E} u^{n-b(S)} v^{|S|-n+b(S)} z^{c(S)-b(S)}.$$

The definitions are concocted so that $Q_\Sigma(u,v) = Q_\Sigma(u,v,1)$ and $Q^*_\Sigma(u,v) = Q_\Sigma(u,v,0)$. The purpose of the total dichromatic polynomial is to give a common expression to the ordinary and balanced polynomials, but I do not know of any interpretation of it for values of z other than 1 and 0.

The definitions show that, for an ordinary graph, $Q_{+\Gamma}(u,v) = Q^*_{+\Gamma} = Q_\Gamma(u,v)$. That is, we are generalizing the dichromatic polynomial of a graph. That, of course, is the point.

The chromatic polynomials can be expressed as $\chi^{[*]}_\Sigma(\lambda) = (-1)^n Q^{[*]}_\Sigma(-\lambda,-1)$. That follows from the algebraic forms of the chromatic polynomials (Theorem J.7).

Theorem J.17 (Theorem Q)**.** *[[LABEL T:1205Qdc]] Let e be an edge in the signed graph Σ. If e is neither a balanced loop nor a loose edge, then*

$$\begin{aligned} Q_\Sigma(u,v,z) &= Q_{\Sigma\setminus e}(u,v,z) + Q_{\Sigma/e}(u,v,z) \text{ if } e \text{ is a link},\\ Q_\Sigma(u,v) &= Q_{\Sigma\setminus e}(u,v) + Q_{\Sigma/e}(u,v),\\ Q^*_\Sigma(u,v) &= Q^*_{\Sigma\setminus e}(u,v) + Q^*_{\Sigma/e}(u,v) \text{ if } e \text{ is a link}. \end{aligned}$$

If e is a balanced loop or a loose edge, then

$$Q_\Sigma = Q_{\Sigma\setminus e} + vQ_{\Sigma/e}.$$

Proof. Clearly, if the first formula holds for three variables, it will hold for any specialization of those three variables. Setting z to 0, the formula will simplify to

$$Q^*_\Sigma(u,v) = Q_\Sigma(u,v,0) = Q_{\Sigma\setminus e}(u,v,0) + Q_{\Sigma/e}(u,v,0) = Q^*_{\Sigma\setminus e}(u,v) + Q^*_{\Sigma/e}(u,v),$$

with a similar argument for $Q(u,v)$, so the second and third parts of the theorem are valid for any link, contingent of course on the first part.

For the remaining proof, let's write $u_\Sigma(S) := c(S) - b(S)$, for short; this is the number of unbalanced components of $\Sigma|S$. The definition gives

(J.2)
[[LABEL E:1205Qsimp]]
$$Q_\Sigma(u,v,z) = \sum_{S\subseteq E} u^{b_\Sigma(S)} v^{|S|-n+b(S)} z^{u_\Sigma(S)} = v^{-n}\sum_S (uv)^{b_\Sigma(S)} v^{|S|} z^{u_\Sigma(S)}$$

(a very handy simplification in many computations) and, for $\Sigma \setminus e$ with this simplification,

$$Q_{\Sigma\setminus e}(u,v,z) = v^{-n}\sum_{S\subseteq E\setminus e} (uv)^{b_{\Sigma\setminus e}(S)} v^{|S|} z^{u_{\Sigma\setminus e}(S)} = v^{-n}\sum_{S\subseteq E\setminus e} (uv)^{b_\Sigma(S)} v^{|S|} z^{u_\Sigma(S)}.$$

By subtraction,

(J.3) [[LABEL E:1205Qdiff]]
$$Q_\Sigma(u,v,z) - Q_{\Sigma\setminus e}(u,v,z) = v^{-n}\sum_{S\subseteq E: e\in S} (uv)^{b_\Sigma(S)} v^{|S|} z^{u_\Sigma(S)}.$$

This is valid for any edge e.

Now there are three cases. The edge e may be a link, or it may be unbalanced (a negative loop or a half edge), or it may be a positive loop or a loose edge.

The easiest case first. Suppose e is a positive loop or a loose edge. What distinguishes such an edge is that then $\Sigma/e = \Sigma\setminus e$, as we saw in Section E.1. Also, it's easy to see that

$b(T \cup e) = b(T)$, $u(T \cup e) = u(T)$, and $|T \cup e| = |T| + 1$ for any set $T \subseteq E \setminus e$. Applying these facts in Equation (J.3), we have

$$\begin{aligned} Q_\Sigma(u,v,z) - Q_{\Sigma \setminus e}(u,v,z) &= v^{-n} \sum_{S \subseteq E : e \in S} (uv)^{b_\Sigma(S)} v^{|S|} z^{u_\Sigma(S)} \\ &= v^{-n} \sum_{T \subseteq E \setminus e} (uv)^{b_\Sigma(T)} v^{|T|+1} z^{u_\Sigma(T)} \\ &= v \cdot v^{-n} \sum_{T \subseteq E \setminus e} (uv)^{b_{\Sigma \setminus e}(T)} v^{|T|} z^{u_{\Sigma \setminus e}(T)} \\ &= v Q_{\Sigma \setminus e}(u,v,z) = v Q_{\Sigma/e}(u,v,z). \end{aligned}$$

Therefore, $Q_\Sigma = Q_{\Sigma \setminus e} + v Q_{\Sigma/e}$ if e is a balanced loop or a loose edge.

If e is any other kind of edge, then Σ/e has one vertex less than either Σ or $\Sigma \setminus e$. The next step in the proof is to write out the simplified definition (J.2) for Σ/e; it is

$$Q_{\Sigma/e}(u,v,z) = v^{-(n-1)} \sum_{T \subseteq E \setminus e} (uv)^{b_{\Sigma/e}(T)} v^{|T|} z^{u_{\Sigma/e}(T)}.$$

If e is a link, we can rewrite this as

$$Q_{\Sigma/e}(u,v,z) = v^{-n} \sum_{T \subseteq E \setminus e} (uv)^{b_\Sigma(T \cup e)} v^{|T \cup e|} z^{u_\Sigma(T \cup e)}$$

because $b_{\Sigma/e}(T) = b_\Sigma(T \cup e)$ by Lemma **??** [**CAN'T FIND IT ANYWHERE!**] and $c_{\Sigma/e}(T) = c_\Sigma(T \cup e)$ by Lemma I.**??** [**CAN'T FIND IT ANYWHERE!**], and of course $|T \cup e| = |T| + 1$. This is the same as $Q_\Sigma - Q_{\Sigma/e}$, so we have the familiar equation $Q_\Sigma = Q_{\Sigma \setminus e} + Q_{\Sigma/e}$.

But suppose e is a negative loop or a half edge? Then we cannot predict the number of components of the contraction. However, the rest is as before: $b_{\Sigma/e}(T) = b_\Sigma(T \cup e)$ (by Lemma **??** [**FIND IT!**]) and $|T \cup e| = |T| + 1$. Thus, if we set $z = 1$ to eliminate the effect of $c(T)$, we get a valid identity,

$$Q_{\Sigma/e}(u,v) = v^{-n} \sum_{T \subseteq E \setminus e} (uv)^{b_\Sigma(T \cup e)} v^{|T \cup e|}.$$

This is the case $z = 1$ of the expression in (J.3); so we have the desired reduction formula for $Q(u,v)$.

Another way to eliminate the effect of $c(T)$ is to set $z = 0$, which means we are talking about Q^*. Sad to say, this doesn't help. Because Q^* restricts the sum to balanced edge sets, we can no longer compare the sum in $Q^*_{\Sigma/e}$, which is over balanced sets $T \subseteq E(\Sigma/e)$, to the sum in $Q^*_\Sigma - Q^*_{\Sigma \setminus e}$, which is over balanced sets $S \subseteq E(\Sigma)$ that contain e. But no such sets exist! That is why we are satisfied to prove the reduction formula for Q^* only when e is a link. □

Corank-nullity polynomials. [[LABEL 2.crn]]

The corank-nullity polynomial is most easily defined in terms of the dichromatic polynomial, by the following formulas. There are two important corank-nullity polynomials, which can be combined into one by the addition of a third variable—exactly as with the dichromatic polynomials.

Definition J.8. [[LABEL D:1205crn]] The *corank-nullity polynomial* (or *rank generating polynomial*) of a signed graph is

$$R_\Sigma(u,v) := u^{-b(\Sigma)} Q_\Sigma(u,v).$$

The *balanced corank-nullity polynomial* is

$$R^*_\Sigma(u,v,z) := u^{-b(\Sigma)} Q^*_\Sigma(u,v).$$

The *total corank-nullity polynomial* is

$$R_\Sigma(u,v,z) := u^{-b(\Sigma)} Q_\Sigma(u,v,z).$$

Thus, $R_\Sigma(u,v) = R_\Sigma(u,v,1)$ and $R^*_\Sigma(u,v) = R^*_\Sigma(u,v,0)$.

[VERIFY THE STATEMENT of this result:] Dec 8: Yash Lodha

Theorem J.18 (Theorem R). *[[LABEL T:1208R]] The corank-nullity polynomials of a signed graph have the following properties:*

(1) $R_\Sigma(u,v,z) = R_{\Sigma\setminus e}(u,v,z) + R_{\Sigma/e}(u,v,z)$ if e is a link and not a balancing edge of Σ.
(2) $R_\Sigma(u,v) = R_{\Sigma\setminus e}(u,v) + R_{\Sigma/e}(u,v)$ if e is not a balancing edge and not a positive loop or loose edge.
(3) $R^*_\Sigma(u,v) = R_{\Sigma\setminus e}(u,v) + R_{\Sigma/e}(u,v)$ if e is a link but not a balancing edge.

Proof. Use "Theorem Q" (Theorem J.17) and Proposition F.7.
[We need details here! WHAT IS THE PROP?]

□

Theorem J.19 (Theorem QRM). *[[LABEL T:1208QRM]] $Q_\Sigma(u,v,z)$ and $R_\Sigma(u,v,z)$ satisfy the following identities.*

(M) Multiplicativity:
$$Q_{\Sigma_1 \mathbin{\dot\cup} \Sigma_2} = Q_{\Sigma_1} Q_{\Sigma_2},$$
$$R_{\Sigma_1 \mathbin{\dot\cup} \Sigma_2} = R_{\Sigma_1} R_{\Sigma_2},$$
$$Q_{\Sigma_1 \cup_v \Sigma_2} = Q_{\Sigma_1} Q_{\Sigma_2}.$$
(U) Unitarity:
$$Q_{K_1} = u,\ R_{K_1} = 1,\ Q_\varnothing = R_\varnothing = 1,$$
$$Q_{K_1^\circ} = u + z = R_{K_1^\circ}.$$
(I) Invariance:
$$\Sigma_1 \cong \Sigma_2 \implies Q_{\Sigma_1} = Q_{\Sigma_2} \text{ and } R_{\Sigma_2} = R_{\Sigma_2}.$$
(BE) If e is a balancing edge of Σ_1 which is not an isthmus, then
$$Q_\Sigma = (u+1) Q_{\Sigma\setminus e},$$
$$R_\Sigma = (u+1) R_{\Sigma\setminus e}.$$

Proof. The proofs are an exercise. One should consult Section I.**??** for guidance. □

J.6. **Counting colorations.** [[LABEL 2.allcolorations]]

Recall that $I(\gamma) :=$ set of improper edges of γ. Define

$$X_\Sigma(k,w) := \sum_{\gamma : V \to \Lambda_k} w^{|I(\gamma)|},$$

which is the generating function of all k-colorations by the number of improper edges, and

$$X^*_\Sigma(k,w) := \sum_{\gamma : V \to \Lambda^*_k} w^{|I(\gamma)|},$$

which is the generating function of all zero-free k-colorations.

Theorem J.20. *[[LABEL T:1208allcolorations]]*

$$X_\Sigma^{[*]}(k,w) = (-1)^{b(\Sigma)}(w-1)^n Q_\Sigma^{[*]}(\frac{-\lambda}{w-1}, w-1)$$

where $\lambda = 2k+1$ *if all colors are allowed and* $2k$ *if* 0*-free. (*$\lambda =$ *size of the color set ,* Λ_k *or* Λ_k^**.)*

Lemma J.21 (Lemma A)**.** *[[LABEL L:1208A]] For a coloration* γ*,* $I(\gamma)$ *is closed, and it is balanced if* γ *is* 0*-free.*

Proof. Exercise. □

Lemma J.22 (Lemma B)**.** *[[LABEL L:1208B]]* $\gamma|_{V_0(I(\gamma))} \equiv 0$.

Proof. Recall $V_0(S) = \{$ Vertices of unbalanced components $\} = V \setminus \bigcup \pi_b(S)$. Look at an unbalanced component of $I(\gamma)$. It contains a negative circle or a half edge. A negative circle of improper edges [**diagramcomes here**] generates an equation $2\gamma_i = 0$. (From $[1-\sigma(C)]\gamma_i = 0$.)

Therefore $\gamma(v_i) = 0$ if $v_i \in V_0(I(\gamma))$. Hence proved. □

This means that $V_0(I(\gamma))$ together with $\gamma|_{V\setminus V_0(I(\gamma))}$ completely determine γ.

12/10: Joyce

[**MISSING NOTES**]

Dec 12: Nate Reff

K. Signed Complete Graphs

[[LABEL 2.complete]]

Signed complete graphs $\Sigma = (K_n, \sigma)$ have especially nice properties due, in part, to the existence of adjacencies between all vertices, and in further part, to the fact that the adjacency matrix is zero only on its diagonal. We can regard a signed K_n as determined by its negative subgraph Σ^-. From this point of view we like to write it as $\Sigma = K_\Gamma$ where Γ is a simple graph of order n; this signed graph is $-\Gamma \cup +\Gamma^c$; that is, $\Sigma^- = \Gamma$ and $\Sigma^+ = \Gamma^c$, the complementary graph. Then $K_{\Gamma^c} = -K_\Gamma$.

The trivial examples are $+K_n = K_{(V,\varnothing)} = K_{K_n^c}$ and $-K_n = K_{K_n}$. The nontrivial examples are those in which $\varnothing \subset E(\Gamma) \subset E(K_n)$, so they have edges of both signs.

K.1. **Coloring.** [[LABEL 2.complete.coloring]]

How does this relate to signed graph coloring? Let's look at a zero-free coloration γ. What makes it proper? Looking at Figure K.1 we see that $\gamma^{-1}(\pm i)$ must be properly colored for each i. This leads to two observations. The first is that $K_\Gamma{:}\gamma^{-1}(\pm i)$ has to be antibalanced. Here recall Harary's Balance Theorem A.2: Σ is balanced iff the negative edges are a cut. Thus, Σ is antibalanced iff the positive edges are a cut. The second is that there are 2 ways to put vertex signs on $\gamma^{-1}(\pm i)$, because it induces a connected subgraph of K_Γ.

These observations suggest a three-step coloring procedure.

(1) Choose a partition of V into antibalanced sets $B_1, \ldots, B_l$ (in other words, $K_\Gamma{:}B_i$ is antibalanced; equivalently, $\Gamma^c{:}B_i$ is complete bipartite).
(2) Assign $+$ and $-$ to the two halves of each B_i (there are 2^l ways to do this because each B_i induces a connected subgraph).

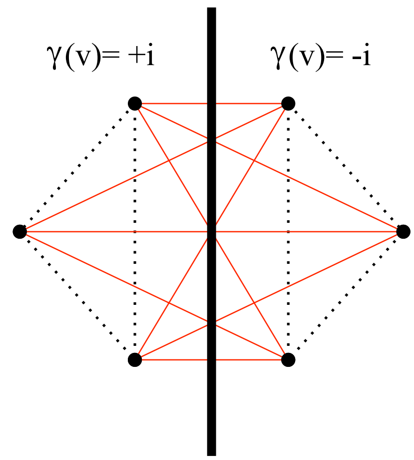

FIGURE K.1. Assigning signs to the vertices of $\gamma^{-1}(\pm i)$ in a signed K_n. The diagram shows the case in which there are 6 vertices colored $\pm i$. The positive edges (in red) are complete bipartite.[**NATE, NOTE: Figure uses too much white space!**]

[[LABEL 1212image1]]

(3) Assign l distinct labels from $[k]$ to the B_i's (there are $(k)_l = l!\binom{k}{l}$ ways to do this).

Suppose we have a definite signed graph Σ. Let's define a partition of V to be *antibalanced* if every part induces an antibalanced signed graph. Our coloring procedure leads to the following description of the chromatic polynomial of a signed K_n, or indeed (by the same proof) of any signed graph that is complete in the sense that each pair of vertices is joined by one or more edges.

Theorem K.1. *[[LABEL T:1212antibalanced chromatic]] If Σ is a signed graph in which all vertices are adjacent, then $\chi^*_\Sigma(\lambda) = \sum_\pi 2^{|\pi|}(k)_{|\pi|}$, where $\lambda = 2k$ and the sum is taken over all antibalanced partitions of V.*

This means the zero-free chromatic polynomial encodes the number of partitions into antibalanced sets.

Corollary K.2. *[[LABEL C:1212minantiptn]] For any signed graph Σ in which all vertices are adjacent, $\chi^*(\Sigma)$ = the minimum size of a partition of V into antibalanced sets.*

A *clique* is a vertex set that induces a complete subgraph. In the next corollary we include $\varnothing$ as a clique, i.e., K_0 as a complete subgraph, since one part of a bipartition may be empty. The corollary gives a structural interpretation, in terms of Γ or its complement, of the zero-free chromatic number of K_Γ.

Corollary K.3. *[[LABEL C:1212mincliquepairptn]] $\chi^*(K_\Gamma)$ = the minimum size of a partition of V into induced complete bipartite subgraphs of Γ^c, which also = the minimum size of a partition into pairs of nonadjacent cliques in Γ.*

We can apply this to get a (less satisfactory) interpretation of the chromatic number.

Corollary K.4. *[[LABEL C:1212Corollary4]] $\chi(K_\Gamma) = \min_{v\in V} \chi^*(K_{\Gamma\setminus v})$.*

Proof. You can use the color 0 only once since all vertices are adjacent. $\square$

Open questions on coloring of signed complete graphs.

(1) What is $\max_\Gamma \chi^*(K_\Gamma)$, over all graphs Γ of order n? Tom thinks $+K_n$ should maximize with $\chi^*(+K_n) = \lceil \frac{n}{2} \rceil$, and $-K_n$ should minimize. Also, $\chi^*(-K_n) = 1$ since they can all be the same color.
(2) Similarly, what is $\max_\Gamma \chi(K_\Gamma)$, over all graphs Γ of order n?
(3) Are the graphs that achieve the maxima unique (up to switching)?

I wrote a short paper, Zaslavsky (1984a), on chromatic number that looked at the very easiest questions of this kind. There is certainly much more to be accomplished by anyone who is interested.

K.2. **Two-graphs.** [[LABEL 2.twographs]]

A *two-graph* is a set of triples chosen from V, in other words $\mathcal{T} \subseteq \mathcal{P}^{(3)}(V)$, such that every quadruple from V contains an even number of triples of $\mathcal{T}$. $\mathcal{T}$ is *regular* if every pair v_iv_j is in the same number of triples of $\mathcal{T}$.

Observe that $\mathcal{T}^c$ is a two-graph if $\mathcal{T}$ is, and moreover that $\mathcal{T}^c$ is regular if $\mathcal{T}$ is.

A signed complete graph K_Γ generates a two-graph $\mathcal{T}(K_\Gamma)$ by the rule:

$$\begin{aligned}\mathcal{T}(K_\Gamma) &:= \{ \text{ vertex sets of negative triangles } \} \\ &= \{ \text{ triples of vertices that support an odd number of edges in } \Gamma \}.\end{aligned}$$

Lemma K.5. *[*[LABEL L:1212swclasstg]] *The class $\mathcal{T}(K_\Gamma)$ is a two-graph, and the whole switching class $[K_\Gamma]$ generates the same two-graph.*

Proof. A nice elementary exercise for the reader. □

Theorem K.6. *[*[LABEL T:1212tggraph]] *Every two-graph is a $\mathcal{T}(K_\Gamma)$ for some graph Γ, which is unique up to switching.*

Proof. We construct Γ from $\mathcal{T}$ as follows: (1) Choose any vertex v. (2) Define all v-edges to be positive. (3) Define the edge uw to be $-$ (negative) if $vuw \in \mathcal{T}$ and $+$ (positive) if not. Then check that this definition is consistent, i.e., that $\mathcal{T} = \mathcal{T}(K_\Gamma)$. [**THIS IS WHAT YOU SHOULD DO in the write-up!**]

To prove uniqueness notice that you can switch any graph so everything agrees on a spanning tree. [**NATE: EXPLAIN HOW THIS PROVES UNIQUENESS.**] □

Graph switching.

Switching originated in the work of J.J. Seidel, who studied equiangular lines, which are sets of lines that all make the same angle with each other. (See van Lint and Seidel (1966a) in [JJS].) We'll see in Chapter III [**GEOMETRY**] that equiangular lines are cryptomorphic [*sic*] to signed complete graphs. Seidel described switching in terms of the graph Γ, not signed graphs; consequently I call switching a graph *Seidel switching*, or simply *graph switching*. This switching means taking Γ and reversing the adjacencies between $X \subseteq V$ and X^c, for some vertex set $X \subseteq V$. From the definitions it is plain to see that switching K_Γ corresponds to (graph-) switching Γ; specifically, that $(K_\Gamma)^X = K_{\Gamma^X}$. (This is how I came to the notion of switching a signed graph.)

The *Seidel adjacency matrix* of Γ is what we are calling $A(K_\Gamma)$. Seidel introduced this matrix early (cf. Seidel (1968a) in [JJS]), strictly in terms of the graph Γ; he called it the $(0, -1, +1)$-adjacency matrix of Γ. It turned out to be a powerful tool because of its eigenvalue theory (cf. Seidel (1976a) in [JJS]). From the perspective of this matrix, switching either Γ or K_Γ corresponds to conjugating $A(K_\Gamma)$ by a diagonal ± 1-matrix.

Lemma K.7. *[[LABEL L:1212Lemma5]] Switching does not change the eigenvalues of $A(K_\Gamma)$.*

Proof. Similar matrices have the same eigenvalues. □

We write $A(\mathcal{T}) :=$ any $A(K_\Gamma)$ such that $K_\Gamma \leftrightarrow \mathcal{T}$. Thus, $A(\mathcal{T})$ is well defined only up to conjugation by a diagonal ± 1-matrix, but that is sufficient to make its spectrum (its eigenvalues and their multiplicities) well defined.

Lemma K.8. *[[LABEL L:1212tga]] Any adjacency matrix A of a two-graph $\mathcal{T}$ satisfies*

(K.1) [[LABEL E:1212tga]] $$A^2 = (n-1)I + (n-2)A - 2(\sigma_{ij}t_{ij})_{ij},$$

where $t_{ij} :=$ the number of triples on v_iv_j.

Notice that this is, properly, a statement about signed complete graphs that is invariant under switching. That is why we can formulate it in terms of a two-graph, which corresponds to a switching class of signed K_n's.

Proof. Note that the incidence numbers t_{ij} satisfy $0 \leq t_{ij} \leq n-2$. We write $\sigma_{ij} := \sigma(v_iv_j)$.

On the diagonal, $(A^2)_{ii} = n-1$, since A has $n-1$ ± 1's in each row and 0's along the diagonal. This accounts for the diagonal elements of all the matrices in Equation (K.1) Thus, we only have to examine an off-diagonal element (i,j) where $i \neq j$.

In A^2, the entry is $(A^2)_{ij} = \sum_{k=1}^n a_{ij}a_{jk} = \sum_{k\neq i,j} \sigma_{ik}\sigma_{jk}$.

Suppose $\sigma_{ij} = +$. Then $v_iv_jv_k$ is a triple in $\mathcal{T} \iff a_{ik}a_{jk} = -1$. So, $t_{ij} =$ the number of triples on v_iv_j that are in $\mathcal{T} =$ the number of negative paths $v_iv_kv_j$. Since $n-2-t_{ij}=$ the number of triples on v_iv_j that are not in $\mathcal{T}$ = the number of positive paths $v_iv_kv_j$, $(A^2)_{ij} = (n-2-t_{ij}) - t_{ij} = n-2-2t_{ij}$.

Suppose on the contrary that $\sigma_{ij} = -$. Then $v_iv_jv_k$ is a triple in $\mathcal{T} \iff a_{ik}a_{jk} = +1 \iff \sigma(v_iv_kv_j) = +$. So $t_{ij} =$ the number of positive paths $v_iv_kv_j$. Meanwhile, $n-2-t_{ij} =$ the number of negative paths $v_iv_kv_j$. Therefore, $(A^2)_{ij} = t_{ij} - (n-2-t_{ij}) = -(n-2-2t_{ij})$.

We conclude that $(A^2)_{ij} = \sigma_{ij}(n-2-2t_{ij})$ off the diagonal. With our calculation of the diagonal, we have proved Equation (K.1). □

Proposition K.9. *[[LABEL P:1212rtga]] Any adjacency matrix A of a regular two-graph with t triples on each pair of vertices satisfies*

(K.2) [[LABEL E:1212rtga]] $$A^2 = (n-1)I + (n-2-2t)A.$$

Conversely, if some adjacency matrix of a two-graph $\mathcal{T}$ satisfies a quadratic equation, then it satisfies (K.2) *and $\mathcal{T}$ is regular with t triples on each vertex pair.*

Proof. The first part is direct from Lemma K.8. The second part follows from comparing the presumed quadratic equation $A^2 = \beta I + \alpha A$ with (K.1). We deduce from the diagonal that $\beta = n-1$ and from the off-diagonal that $\sigma_{ij}(n-2-2t_{ij}) = a_{ij}\alpha$. But we also know that $a_{ij} = \sigma_{ij} \neq 0$, hence every $t_{ij} = \frac{1}{2}(n-2-\alpha)$, a constant. Hence, $\mathcal{T}$ is regular. Comparing with (K.2), this constant is t. □

Theorem K.10. *[[LABEL T:1212Theorem8]] For $n \geq 3$, $\mathcal{T}$ is regular $\iff$ $A(\mathcal{T})$ has at most 2 eigenvalues. Moreover, $A(\mathcal{T})$ cannot have only one eigenvalue.*

Proof. We write $A := A(\mathcal{T})$. Now, $\mathcal{T}$ is regular $\iff$ A satisfies a quadratic equation, specifically Equation (K.2) $\iff$ A has at most two eigenvalues (by matrix theory).

For A to have just one eigenvalue, it must have a linear annihilating polynomial, that is, $A - \alpha I = O$. This is impossible since A is non-zero off the diagonal and $n > 1$. □

The multiplicity trick.

There is a standard but clever and effective trick used in the analysis of integral symmetric matrices, especially the adjacency matrices of graphs, which uses basic facts about the eigenvalue multiplicities. We'll apply this trick to signed complete graphs with two eigenvalues, a.k.a. regular two-graphs. (Again, my account is based on papers by Seidel in [JJS]; see especially Seidel (1976a).) Let the eigenvalues be ρ_1 and ρ_2 with multiplicities μ_1 and μ_2.

By Proposition K.9, $A^2 - (n-2-2t)A - (n-1)I = O$ is an annihilating polynomial of A. It is the minimal polynomial since A cannot have only one distinct eigenvalue. Hence, the eigenvalues are the two zeros of $\rho^2 - (n-2-2t)\rho - (n-1) = 0$. Specifically,

$$\rho_1, \rho_2 = \frac{n-2-2t \pm \sqrt{(n-2-2t)^2 + 4(n-1)}}{2} = \frac{\alpha \pm \sqrt{\Delta}}{2},$$

where for simplicity I write

$$\Delta := (n-2-2t)^2 + 4(n-1) = (n-2t)^2 + 8t$$

for the discriminant and $\alpha := n-2-2t$. Because $(n-2-2t)^2 \geq 0$ and (since $n \geq 3$) $4(n-1) > 0$, the discriminant is positive. Therefore the eigenvalues are real (and distinct, as we knew already).

The multiplicity trick depends on three basic facts:

(1) The multiplicities are whole numbers.
(2) $\mu_1 + \mu_2 = n$.
(3) $\mu_1\rho_1 + \mu_2\rho_2 = \operatorname{tr}(A) = 0$.

In the simplified notation property (3) becomes

$$\mu_1 \frac{\alpha + \sqrt{\Delta}}{2} + \mu_2 \frac{\alpha - \sqrt{\Delta}}{2} = 0.$$

Thus, the multiplicities are

$$\mu_1, \mu_2 = \frac{n}{2}\left(1 \mp \frac{n-2-2t}{\sqrt{(n-2t)^2 + 8t}}\right) = \frac{n}{2\sqrt{\Delta}}(\sqrt{\Delta} \mp \alpha).$$

Case 1: Δ is not a square. Then the eigenvalues are irrational. We can separate their rational and irrational parts to deduce that

$$\mu_1 \frac{\alpha}{2} + \mu_2 \frac{\alpha}{2} = 0$$

and

$$\mu_1 \frac{\sqrt{\Delta}}{2} - \mu_2 \frac{\sqrt{\Delta}}{2} = 0.$$

The first equation tells us that $\alpha = 0$ and the second tells us that $\mu_1 = \mu_2$. Therefore the eigenvalues are $\pm\sqrt{\Delta}/2 = \pm\sqrt{n-1}$, each with multiplicity $n/2$, and $t = \frac{n}{2} - 1$. Evidently, $n-1$ must be odd and not a perfect square.

Case 2: Δ is a square. Then the eigenvalues are rational; by Eisenstein's theorem of number theory, since they are rational zeroes of a monic, integral polynomial, they are integers.

Let $\Delta = q^2$, where $q \in \mathbb{Z}$. Because $q^2 = (n-2t)^2 + 8t$, $q \equiv n \pmod 2$. Write $q = n - 2r$, so $q^2 = (q+2r)^2 - 4t(q+2r-2-t)$. Solve for q:

(K.3) [[LABEL E:1212q]] $q(t-r) = r^2 - 2rt + 2t + t^2 = (t-r)^2 + 2t.$

We conclude that either $t = r$ or

$$q = t - r + \frac{2t}{t-r}.$$

If $t = r$ then (K.3) implies $2t = 0$, so in this case $t = r = 0$. That corresponds to a trivial case: the all-positive complete graph, or Γ with no edges. Let's rule out the trivial cases; we'll look for properties of interesting regular two-graphs with rational eigenvalues. That means $t \neq r$ and $0 < t < n-2$.

If $t \neq r$, let $s = t - r$. Then $q = s + 2 + 2r/s$ and $s|2r$. Directly in terms of r and s,

$$\begin{aligned} t &= r + s, \\ q &= s + 2 + \frac{2r}{s}, \\ n &= s + 2 + 2r + \frac{2r}{s}. \end{aligned}$$

The eigenvalues are

$$\rho_1, \rho_2 = \frac{\alpha \pm q}{2} = \frac{1}{2}\left[\frac{2r - s^2}{s} \pm \left(s + 2 + \frac{2r}{s}\right)\right] = \begin{cases} 1 + \dfrac{2r}{s}, \\ -(s+1) \end{cases}$$

(the upper value is ρ_1, the lower is ρ_2) and the multiplicities are

$$\mu_1, \mu_2 = \frac{n}{2q}(q \pm \alpha) = \frac{n}{2}\,\frac{q \mp (n - 2t - 2)}{q} = \begin{cases} (s+1)\left(1 + \dfrac{2rs}{s^2 + 2(r+s)}\right), \\ 1 + \dfrac{2r}{s} + \dfrac{2r(s+2r)}{s^2 + 2(r+s)}. \end{cases}$$

We can therefore express n and t (the *parameters* of $\mathfrak{T}$) and the eigenvalues and their multiplicities in terms of r and s, and the problem is to find which values of r and s are numerically feasible. After that, the real problem is to find examples of regular two-graphs with feasible parameters, or to show none exist (which is sometimes the case due to more sophisticated reasons). That takes us into group theory and design theory, and I stop here—save for a not-so-short digression on strongly regular graphs.

K.3. **Strongly regular graphs.** [[LABEL 2.twographs.srg]] 2008 Dec 12: Zaslavsky

Let's take a little digression into strongly regular graphs. A simple graph Γ is called *strongly regular* if it is regular—every vertex has degree k—and there are constants λ and μ such that each pair of adjacent vertices has exactly λ common neighbors and each pair of nonadjacent vertices has exactly μ common neighbors. We say Γ is an $\mathrm{SRG}(n, k, \lambda, \mu)$, n denoting the number of vertices; the four numbers are the *parameters.* Strongly regular graphs are used, for instance, to represent finite simple groups, which puts them in combinatorial design theory. Seidel discovered remarkable connections between regular two-graphs and strongly regular graphs through the eigenvalues of the Seidel adjacency matrix of Γ, i.e., $A(K_\Gamma)$. I will give some of the flavor of his ideas here.

First, we'll take the easy way to find a strongly regular graph in a regular two-graph. Then we'll glance at the matrix method. In each case we start with a two-graph $\mathfrak{T}$ of order

n, which has the form $\mathcal{T}(K_\Gamma)$ for various switching-equivalent graphs Γ. Choosing the right Γ is part of the method.

The combinatorics of a detached vertex.

Assume $\mathcal{T} = \mathcal{T}(K_\Gamma)$ is regular with t triples on each pair of vertices. We can pick any vertex u and switch as necessary so it is isolated in Γ^ζ. (This determines ζ uniquely.) Write $\Gamma' := \Gamma^\zeta \setminus u = (V', E')$. Then we can draw a few conclusions, summarized as:

Proposition K.11. *[[LABEL P:1212srg n-1]] If $\mathcal{T}$ is a regular two-graph on V, then $t \geq n/3$, $t \equiv n \pmod 2$, and for each $u \in V$, switching so u is isolated in Γ^ζ, then $\Gamma^\zeta \setminus u$ is a strongly regular graph SRG$(n-1, t, t - \frac{1}{2}(n-t), \frac{1}{2}t)$.*

Conversely, if u is isolated in Γ and $\Gamma \setminus u$ is a strongly regular graph SRG$(n-1, t, t - \frac{1}{2}(n-t), \frac{1}{2}t)$, then $\mathcal{T}(K_\Gamma)$ is a regular two-graph.

Proof. We just count carefully. First, Γ' is a t-regular graph, because each edge $vw \in E'$ makes a triple $uvw \in \mathcal{T}$ while each non-edge vw makes a triple $uvw \notin \mathcal{T}$.

Consider an adjacent pair $vw \in E'$. Let $a_{\alpha\beta}$ be the number of vertices in $\Gamma' \setminus \{v, w\}$ that are adjacent to v iff $\alpha = 1$ and to w iff $\beta = 1$. Thus, $a_{11} + a_{10} + a_{01} + a_{00} = n - 3$. Also, $a_{11} + a_{00} = t - 1$, because the triples xvw that are in $\mathcal{T}$, besides uvw, are those for which $x \in V' \setminus \{v, w\}$ is adjacent to both v and w or to neither. Finally, $a_{11} + a_{10} = d'(v) - 1 = t - 1$ (because one neighbor of v is w) and similarly $a_{11} + a_{01} = t - 1$. These four equations can be solved; one finds that $a_{11} = \frac{1}{2}(3t - n)$. Thus, a_{11} is independent of the particular vw, and we have that part of strong regularity which says λ exists and equals $t - \frac{1}{2}(n - t)$.

Since a_{11} counts something it can't be negative, hence $3t - n \geq 0$. Indeed, if $\mathcal{T}$ is nontrivial, then $3t - n > 0$.

Now consider a nonadjacent pair $vw \in E'$. This time the necessary equations are $a_{10} + a_{01} = t$, because the triples xvw that are in $\mathcal{T}$ are those for which $x \in V' \setminus \{v, w\}$ is adjacent to exactly one of v and w, and $a_{11} + a_{10} = d'(v) = t$ and similarly $a_{11} + a_{01} = t$. The solution is that $a_{11} = t/2$, independently of the pair vw, and we have that part of strong regularity which says μ exists and equals $t/2$. $\square$

Example K.1. [[LABEL X:1212pentagon]] Seidel's favorite example for illustrating the ideas of two-graphs was what he called "the pentagon". It is the two-graph $\mathcal{T}$ obtained from $\Gamma = K_1 \cup C_5$, in other words, the pentagon (naturally) with an extra isolated vertex. It's clear from Proposition K.11 that $\mathcal{T}$ is regular with $n = 6$ and $t = 2$. The adjacency matrix is

$$A = \begin{pmatrix} 0 & 1 & 1 & 1 & 1 & 1 \\ 1 & 0 & -1 & 1 & 1 & -1 \\ 1 & -1 & 0 & -1 & 1 & 1 \\ 1 & 1 & -1 & 0 & -1 & 1 \\ 1 & 1 & 1 & -1 & 0 & -1 \\ 1 & -1 & 1 & 1 & -1 & 0 \end{pmatrix}.$$

The eigenvalues and multiplicities are

$$\rho = \pm\sqrt{5}, \qquad \mu_1 = \mu_2 = 3.$$

Since the eigenvalues are irrational they are negatives of each other and their multiplicities are equal; we're in Case 1 of the multiplicity trick.

The matrix of a detached vertex.

Since u is isolated in Γ, its row and column in $A := A(K_\Gamma)$ are all 1 off the diagonal. Thus, writing $A' := A(K_{\Gamma'})$ and $\mathbf{j}$ for the all-ones vector of order $n-1$,

$$A = \begin{pmatrix} 0 & \mathbf{j}^{\mathrm{T}} \\ \mathbf{j} & A' \end{pmatrix}.$$

Let's put this into the nonzero terms of Equation (K.2):

$$A^2 - (n-2-2t)A - (n-1)I = \begin{pmatrix} 0 & \mathbf{0}^{\mathrm{T}} \\ \\ \mathbf{0} & J + (A')^2 - (n-2-2t)A' - (n-1)I \end{pmatrix},$$

because $\mathbf{j}\mathbf{j}^{\mathrm{T}} = J$, the all-ones square matrix of order $n-1$. The two-graph is regular if and only if this is zero, in other words, if and only if

$$(A')^2 = (n-2-2t)A' + (n-2)I - (J-I).$$

The diagonal part of this equation is satisfied automatically because A' has $n-2$ nonzeros, all ± 1, in each row and column. The interesting part is therefore off the diagonal. One can analyze the off-diagonals to prove Γ' is strongly regular; the best way is to write down the equation satisfied by the Seidel matrix of a strongly regular graph; but I will omit this as we already tested Γ' for strong regularity by combinatorics.

From a strongly regular graph.

The two-graph $\mathcal{T}(K_\Gamma)$ associated with a strongly regular graph may happen to be regular itself.

Proposition K.12. *[*[LABEL P:1212srgtg]] *If Γ is a strongly regular graph with parameters (n, t, λ, μ), then $\mathcal{T}(K_\Gamma)$ is regular if and only if $\lambda + \mu = 2k - \frac{1}{2}n$. Then n is even, $k \geq \frac{1}{4}n$, and $t = 2(k-\mu)$.*

Proof. Like the proof of Proposition K.11, this is simply a matter of counting up edges and triangles. Define $a_{\alpha\beta}$ for Γ just as for Γ' in the proof of Proposition K.11.

Consider first adjacent v, w. The number of common neighbors is $a_{11} = \lambda$. The number of neighbors of v not neighbors of w is $a_{10} = k - 1 - \lambda$ since the total number of neighbors is k and w is one of them. Similarly, $a_{01} = k - 1 - \lambda$. This leaves $a_{00} = (n-2) - \lambda - 2(k-1-\lambda) = n - 2k + \lambda$. The number of triples on $v_i v_j$ is then $t_{ij} = a_{11} + a_{00} = n - 2k + 2\lambda$.

Now suppose v, w are nonadjacent. The number of common neighbors is $a_{11} = \mu$. v has $a_{10} = k - \mu$ neighbors that are not adjacent to w, and of course $a_{01} = k - \mu$ also. Then $t_{ijj} = a_{10} + 1_{01} = 2k - 2\mu$.

For $\mathcal{T}(K_\Gamma)$ to be regular, t_{ij} must be a constant, regardless of whether v and w are adjacent or not. Thus, we have a regular two-graph iff $n - 2k + 2\lambda = 2k - 2\mu$, or $2(\lambda + \mu) = 4k - n$, which is therefore a non-negative integer. ☐

To a strongly regular graph.

The natural next question is the converse: whether, when $\mathcal{T}(K_\Gamma)$ is a regular two-graph, Γ can be switched to become strongly regular. Not always!

Part of the reason comes from applying Proposition K.11 in reverse, which shows that t would have to be even. Another obstacle might be that it's impossible to switch Γ to be regular; an example is the "pentagon" two-graph of Example K.1 (Exercise!).

One can deduce a lot from the eigenvalues and multiplicities. Assume we have a regular two-graph $\mathcal{T}(K_\Gamma)$ where Γ is strongly regular, and let $A := A(\mathcal{T}(K_\Gamma))$. The eigenvalue of A associated with eigenvector $\mathbf{j}$ is $\rho_0 = n-1-2k$, and all other eigenvectors are orthogonal to $\mathbf{j}$ (by matrix theory). The combinatorial definition of strong regularity implies that

$$A(\Gamma)^2 = kI + \lambda A(\Gamma) + \mu A(\Gamma^c),$$

where $A(\Gamma)$ is the standard $(0,1)$-adjacency matrix. As $A(\Gamma^c) = J - I - A(\Gamma)$, we have

$$A(\Gamma)^2 = (\lambda-\mu)A(\Gamma) + (k-\mu)I + \mu J.$$

One can easily calculate that the two-graph's adjacency matrix is $A = J - I - 2A(\Gamma)$. Thus, A satisfies the somewhat quadratic equation

(K.4)
[[LABEL E:1212srgquadratic]] $A^2 - 2[\lambda-\mu+1]A - [2(\lambda+\mu)+1-4k]I = [n-4k+2(\lambda+\mu)]J.$

I say "somewhat" because the J term on the right makes (K.4) not a polynomial in A. We use Equation (K.4) in two ways. Postmultiplying by the eigenvector $\mathbf{j}$ we get a quadratic equation in the eigenvalue ρ_0; since we already know ρ_0, this gives a quadratic equation in n, k, λ, μ which constrains those parameters. Any other eigenvector $\mathbf{x}$, corresponding to an eigenvalue ρ, is orthogonal to $\mathbf{j}$, whence $J\mathbf{x} = \mathbf{0}$. Thus, postmultiplying by $\mathbf{x}$ gives a quadratic equation in ρ,

$$\rho^2 - 2[\lambda-\mu+1]\rho - [2(\lambda+\mu)+1-4k] = 0.$$

The two roots, ρ_1 and ρ_2, and their multiplicities can be treated with the multiplicity trick to extract even more information about the parameters. I will skip further discussion and only mention a conclusion, along with the elementary facts we noticed:

Proposition K.13. *[[LABEL P:1212tgsrg]] Suppose $\mathcal{T}(K_\Gamma)$ is a regular two-graph with eigenvalues ρ_0 (associated with $\mathbf{j}$), ρ_1, and ρ_2, and that Γ is strongly regular with parameters (n,k,λ,μ). Then $\rho_0 = n-1-2k$; t is even; and either $\mu = \lambda+1$, or else ρ_1 and ρ_2 are odd integers.*

All this, once again, is based on Seidel in (1976a) and other papers reprinted in [JJS].

2009 Jan 29: Jackie Kaminski

L. LINE GRAPHS OF SIGNED GRAPHS

[[LABEL 2.lg]]

Now we come to one of the more exciting topics: the line graph of a signed graph, and how it extends the notion of a line graph in ways that are important even beyond signed graphs themselves.

L.1. What are line graphs for? [[LABEL 2.lg.review]]

We begin by reviewing the definition and properties of the line graph of an unsigned graph. For an ordinary link graph Γ, the line graph is $L(\Gamma) = (V(L), E(L))$, where $V(L) = E(\Gamma)$ and $E(L)$ is the set of adjacencies of edges in Γ. Figure L.1 shows a graph Γ and its line graph $L(\Gamma)$.

When Γ is a simple graph, $E(L)$ can be described as $\{ef : e, f \text{ are adjacent in } \Gamma\}$. However the edges of a line graph of a multigraph can't be described any more concisely than as the adjacencies of edges in Γ. We do point the readers attention to Figure L.2 which illustrates that if e, f are two parallel edges in Γ, then they are adjacent twice, which is reflected in

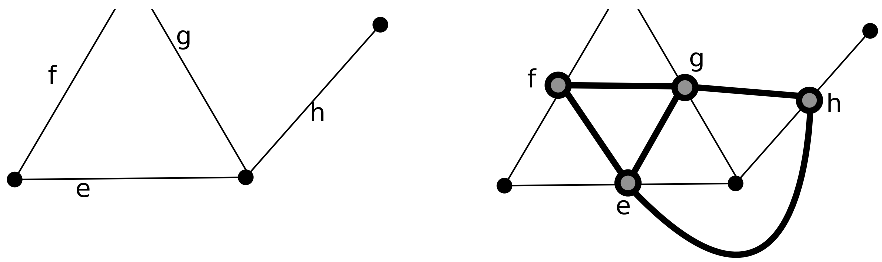

FIGURE L.1. A simple graph Γ, and Γ with its line graph $L(\Gamma)$ superimposed in heavy lines.

[[LABEL F:0129 Line Graph]]

$L(\Gamma)$ as the two edges between vertices e, f. Loops make things very messy, which is why we are restricting our attention to link graphs.

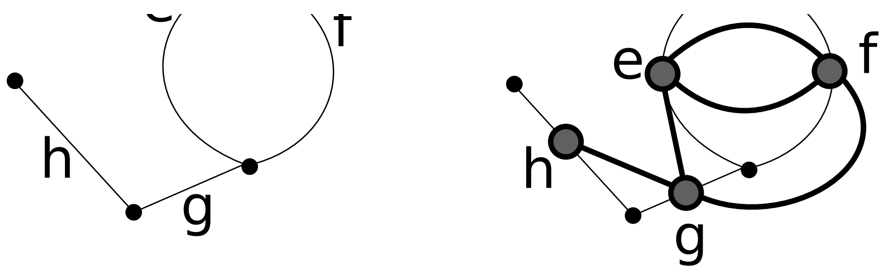

FIGURE L.2. A graph Γ with parallel edges, and Γ with its line graph $L(\Gamma)$ in heavy lines.

[[LABEL F:0129 Line Graph Parallel]]

As further motivation for the line graph of a signed graph, we back up and recall that $B(\Gamma)^{\mathrm{T}}B(\Gamma) = 2I + A(L)$, where $B(\Gamma)$ is the unoriented incidence matrix of Γ and $A(L)$ is the adjacency matrix of the line graph. Furthermore, we recall the corollary, Theorem **??**, that all eigenvalues of a line graph are greater than or equal to -2.

We can't interpret $\mathrm{H}(\Sigma)^{\mathrm{T}}\mathrm{H}(\Sigma)$ (where $\mathrm{H}(\Sigma)$ is the oriented incidence matrix of Σ) for line graphs, we need signed graphs.

L.2. **Ideas for the line graph of a signed graph.** [[LABEL 2.lg.defs]]

It is time to look at possibilities for defining the line graph of a signed graph. Let Σ be a simply signed link graph. Recall that being simply signed means that there are no parallel edges with the same sign. We definitely want our line graph $\Lambda(\Sigma)$ to satisfy $|\Lambda(\Sigma)| = L(|\Sigma|)$, in other words, we want our line graph to have the same underlying graph as the line graph of $|\Sigma|$. Presuming that we want $\Lambda(\Sigma)$ to be a signed graph, we need to decide how to sign the edges of $\Lambda(\Sigma)$. Let's review two ideas that have been tried.

Two previous definitions.

One natural idea would be that for $e' \in E(\Lambda)$, with endpoints $e, f \in V(\Lambda)$, $\sigma_\Lambda(e') = \sigma_\Sigma(e) \cdot \sigma_\Sigma(f)$. However, once we notice that every cycle in Λ is balanced (since every vertex e of the cycle, $e \in V(L)$, e contributes $\sigma_\Sigma(e) \cdot \sigma_\Sigma(e)$ to the cycle sign), we see that this method is trivial: it only gives us line graphs that are balanced, i.e., switching equivalent to $+L(|\Sigma|)$, which means we've lost all the sign information from Σ. We must look for a better idea. (Nevertheless, this line graph has been written about by some people.)

Another signature function for $\Lambda(\Sigma)$ was proposed by Behzad and Chartrand. For an edge ef between e, f, $\sigma_{BC}(ef)$ is $-$ when both $\sigma_\Sigma(e)$ and $\sigma_\Sigma(f)$ are both $-$, and $+$ otherwise. There is literature based on this definition, but as far as I know it has no useful properties. (It doesn't allow us to recover the signs in Σ from the line graph, nor does it preserve the signs of circles, nor does it have eigenvalue properties, etc.)

The definition through bidirection.

The fact is that eigenvalue properties are the main properties that make line graphs interesting (to us, at least, and to many graph theorists). For unsigned graphs we know that $B^{\mathrm{T}}B = 2I + A(L)$, and we know that $\mathrm{H}(\Sigma)\mathrm{H}(\Sigma)^{\mathrm{T}} = \Delta(|\Sigma|) + A(\Sigma)$.

So let's consider $\mathrm{H}(\Sigma)^{\mathrm{T}}\mathrm{H}(\Sigma)$. Recall from Section G.2**??** that the oriented incidence matrix of a signed graph is $\mathrm{H}(\Sigma) = (\eta_{ve})_{V\times E}$, where

$$\eta_{ve} = \begin{cases} 0 & \text{if } v \text{ and } e \text{ are not incident,} \\ \pm 1 & \text{if } v \text{ and } e \text{ are incident once, so that if } e{:}vw \text{ is a link then } \eta_{ve}\eta_{we} = -\sigma(e), \\ 0 & \text{if } e \text{ is a positive loop at } v, \\ \pm 2 & \text{if } e \text{ is a negative loop at } v. \end{cases}$$

So $\mathrm{H}(\Sigma)^{\mathrm{T}}\mathrm{H}(\Sigma)$ is an $E \times E$ matrix, and we notice that row e of $\mathrm{H}(\Sigma)^{\mathrm{T}}$ dot itself is $+2$, since we are only considering link graphs. The dot product will look like $0^2 + \cdots + 0^2 + (\pm 1)^2 + 0^2 + \cdots + 0^2 + (\pm 1)^2 + 0^2 + \cdots + 0^2 = 2$. For the off-diagonal entries of $\mathrm{H}(\Sigma)^{\mathrm{T}}\mathrm{H}(\Sigma)$, row e (of H^{T}) dot column f (of H, which is also row f of H^{T}) gives 0 if e, f are nonadjacent edge (since they will have no vertices in common, there are no positions where both have nonzero entries). If e, f are adjacent, nonparallel links, then the e, j entry of $\mathrm{H}(\Sigma)^{\mathrm{T}}\mathrm{H}(\Sigma)$ is ± 1, depending on how e, f were signed in $\mathrm{H}(\Sigma)$.

To speak more precisely, for this discussion we should be looking at $\vec{\Sigma} = (\Sigma, \tau)$, not just Σ. And we have shown that $H^{\mathrm{T}}(\Sigma, \tau)\mathrm{H}(\Sigma, \tau) = 2I \pm A(\Lambda)$ (for some still unknown convention on signing Λ). And since reversing the orientation of an edge corresponds to switching vertex e in the line graph. So, in some sense we really care about defining $\Lambda(\Sigma, \tau)$ for a switching class of signed graphs, and moreover, since writing the matrix $A(\Lambda)$ necessitates choosing a bidirection for Σ, that's what we should really be looking at. So rather than try to define the line graph of a signed graph, we will define the line graph of a bidirected graph, noting that we can always read signs from a bidirected graph, and if we ever feel compelled to ignore some of the information in our line graph, we have that ability. In summary the basic object on which to take notes is a bidirected graph B[7] (not to be confused with $B(\Gamma)$, the unoriented incidence matrix of Γ). And reorienting B corresponds to switching $\Lambda(B)$.

So now we look at possibilities for how to create the (bidirected) line graph from a bidirected graph $\vec{\Sigma}$. Consider Figure L.3. For a half edge $e{:}v$ in $\vec{\Sigma}$ (where e is the edge and v is the vertex) we have two choices for how to orient the half edge at vertex e in The line graph. Option 1 looks better the way we've drawn it, but we notice that while the half edge $e{:}v$ in $\vec{\Sigma}$ was oriented into the vertex, the corresponding half edge in Λ is oriented out of the vertex. Option 2 is just the opposite. It looks like we're switching the arrows to be backward, however, the half edge that was oriented into the vertex in $\vec{\Sigma}$ is still oriented into the vertex

[7]Note that a bidirection of the unsigned graph Γ does in fact have a sign on each edge, so it is an orientation of a signing of Γ, $\Sigma = (\Gamma, \sigma)$, and when it is convenient we can refer to B as $\vec{\Sigma}$.

in Λ (although the vertex is now e in Λ). Since the matroid theory works out better with Option 2, Option 2 is the right way to create a bidirected line graph from a bidirected graph. Lastly we notice that if we begin with an all negative, all extraverted graph, the line graph (taken with option 1) is all negative, but all introverted, unlike $\vec{\Sigma}$. However, the line graph taken with Option 2 will be an all negative, all extraverted graph, which is the same kind of object as we started with; and this seems preferable.

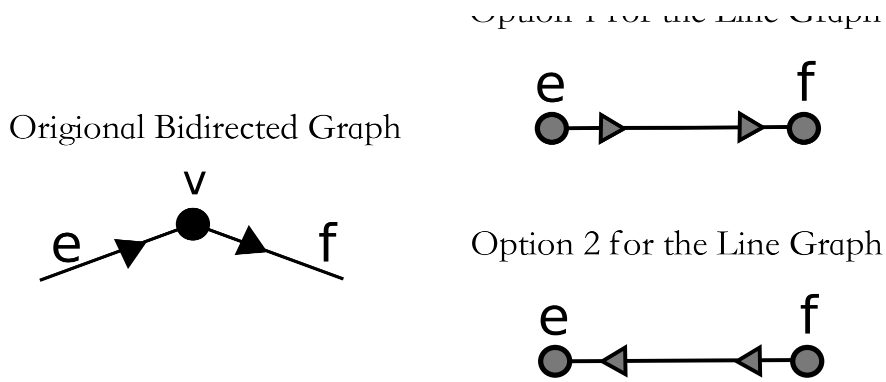

FIGURE L.3. Creating a bidirected line graph from a bidirected graph. [[LABEL F:0129 Line Graph 2]]

Notice that $L(|\vec{\Sigma}|) = |\Lambda(\vec{\Sigma})|$, as desired. So, we know how to create the line graph of a bidirected graph: first we create the line graph of the underlying graph, then we bidirect the edges as above. More formally:

Definition L.1. [[LABEL D:0129 BiDir Line Graph Defn]] The line graph of a bidirected graph $\vec{\Sigma}$ is $\Lambda(\vec{\Sigma})$, whose underlying graph is $|\Lambda| = L(|\Sigma|)$ and whose bidirection is $\tau_\Lambda(e, ef) = \tau_{\vec{\Sigma}}(v, e)$ (where v is the common vertex of e and f).

Notice that we can determine the sign of an edge between vertices e, f of $\Lambda(\vec{\Sigma})$. The formula is

$$\sigma_\Lambda(ef) = -\tau_\Lambda(\varepsilon, e)\tau_\Lambda(\varepsilon', f) = -\tau_\Lambda(\varepsilon)\tau_\Lambda(\varepsilon').$$

where ε' is the end of f at v in Γ (v is between e, f in Γ).

We want to point out also that Option 1 and Option 2 give the same signed graph but the orientations of the edge ends are exactly opposite: $\tau_{\text{Option 1}} = -\tau_{\text{Option 1}}$. In fact, switching $\vec{\Sigma}$ doesn't change the signs of the line graph, i.e., it gives the same the signed line graph $\Sigma(\Lambda(\vec{\Sigma}))$. Therefore, the switching class of the bidirected graph $\vec{\Sigma}$ gives us a signed line graph. On the other hand, the line graph of a signed graph is a switching class of bidirected graphs. Combining these two observations, we can say the line graph of a switching class $[\Sigma]$ of signed graphs is a switching class $[\Lambda(\Sigma)]$ of signed graphs.

[**MISSING NOTES**] 2/3: Lodha

2009 Feb 5: Nate Reff

M. CIRCUITS, COCIRCUITS, AND THEIR SPACES

[[LABEL 2.cyclescircuitsspaces]]

M.1. **Unsigned graphs.** [[LABEL 2.]]

M.1.1. *Cycles, cuts, circuits and bonds.* [[LABEL 2.]]

Suppose we have an ordinary graph $\Gamma = (V, E)$. A *circuit* will be what a circle was in our previous discussions. If $\{X, X^c\}$ is a partition of the vertex set V, then the set of edges, denoted $E(X, X^c)$, which have one end in X and the other in X^c will be called a *cut*. A *bond* is a minimal cut. Recall that every cut is a disjoint union of bonds [[**??**]].

A *binary* (over $\mathbb{F}_2$) *set sum of circuits* is a symmetric difference of circuits. A *binary cycle space* $:= \{$ set sum of circuits $\} = Z_1(\Gamma; \mathbb{F}_2) \subseteq \mathbb{F}_2^E$ or $\mathcal{P}(E)$. For the nonbinary case (over a field $\mathbf{K}$ with $\operatorname{char}\mathbf{K} \neq 2$, or $\mathbf{K} = \mathbb{Z}$) we need to work with indicator vectors, which are defined on orientations (in the binary case this disappears).

M.1.2. *Directed cycles and cuts. Indicator vectors.* [[LABEL 2.]]

Suppose we have a circle $C = e_1 e_2 \cdots e_{l-1} e_l e_1$ in Γ. The *characteristic vector*, or *characteristic function*, is defined as:

$$1_C(e) = \begin{cases} 1 & \text{if } e \in C, \\ 0 & \text{if } e \notin C. \end{cases}$$

Equipped with $1_C + 1_D = 1_{C \oplus D} \mod 2$, where $C \oplus D$ means the set sum of C and D. Note that in characteristic 2 this relation also holds.

Suppose now that we have a fixed orientation of Γ. Let's denote this by $\vec{\Gamma} = (\Gamma, \tau)$ where τ is a bidirection (orient each edge end). With reference to this orientation we define the *indicator vector*, or *indicator function*, of C:

$$I_C(e) = \begin{cases} 1 & \text{if } e \in C \text{ and } \vec{e} \text{ agrees with a chosen direction of } C, \\ -1 & \text{if } e \in C \text{ and } \vec{e} \text{ disagrees with a chosen direction of } C, \\ 0 & \text{if } e \notin C, \end{cases}$$

where $\vec{e}$ means the directed edge e. So I_C and $-I_C$ are the only two indicator vectors of C. We write $\vec{C}$ for a directed C and $I_{\vec{C}}$ for its indicator vector. (We think of a function and a vector as the same thing except for the point of view.)

Observe that C is a cycle (that is, cyclically oriented) if and only if $I_C \geq 0$ or $I_C \leq 0$. This is because the edges have to all agree or all disagree with C.

It is important to notice the circle orientation is independent of edge orientations. Note: we can direct any walk, including a path and a circle. Therefore we can have an indicator vector of a path or a circle or a trail (or a walk, where you add up multiple appearances).

Consider the theta graph in figure M.1. If $\vec{C}_1$ and $\vec{C}_2$ disagree on $\vec{C}_1 \cap \vec{C}_2$ and $\vec{C}_3$ agrees with $\vec{C}_1$, $\vec{C}_2$ on the common path then $I_{\vec{C}_1} + I_{\vec{C}_2} = I_{\vec{C}_3}$.

Proof. If all paths P_{ij} are directed from v_1 toward v_2 then

$$I_{\vec{C}_1} = I_{P_{13}} - I_{P_{12}},$$
$$I_{\vec{C}_2} = I_{P_{12}} - I_{P_{23}},$$
$$I_{\vec{C}_3} = I_{P_{13}} - I_{P_{23}}.$$

This is the proof since we can choose path directions as we like. (We need the minus sign so we can represent signed graphs later on. They cannot be described modulo 2.) $\square$

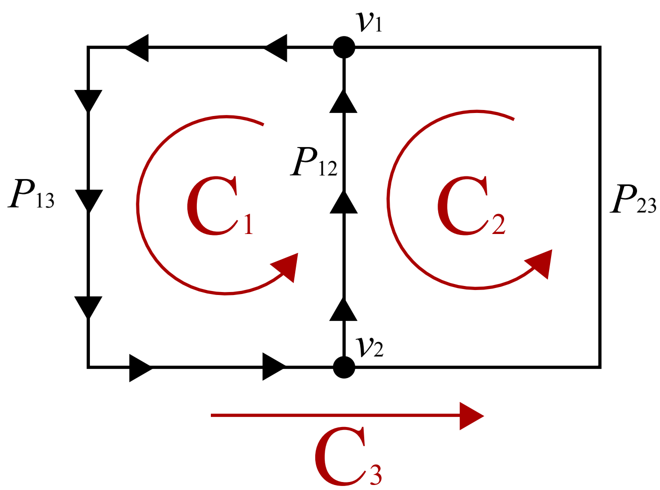

FIGURE M.1.
[[LABEL 0205image1]]

The *cycle space* over $\mathbf{K}$ is the subspace of $\mathbf{K}^E$ generated by all indicator vectors of circuits (circles). We write $Z_1(\Gamma;\mathbf{K})$ for the cycle space over $\mathbf{K}$.

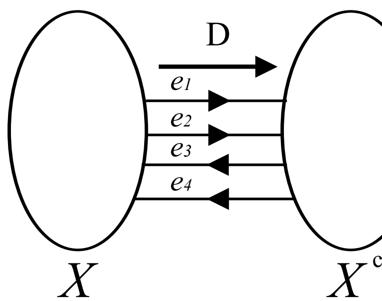

FIGURE M.2.
[[LABEL 0205image2]]

A *directed cut* $\vec{D}$ is the cut $D = E(X, X^c)$ with a direction specified from X to X^c or vice versa. In other words it is directed out of X or into X. See figure M.2. Therefore

$$1_D(e) = \begin{cases} 1 & \text{if } e \in D, \\ 0 & \text{if } e \notin D. \end{cases}$$

and also

$$I_{\vec{D}}(e) = \begin{cases} 1 & \text{if } e \in \vec{D} \text{ and } \vec{e} \text{ agrees with } \vec{D}, \\ -1 & \text{if } e \in \vec{D} \text{ and } \vec{e} \text{ disagrees with } \vec{D}, \\ 0 & \text{if } e \notin \vec{D}. \end{cases}$$

For example, look at figure M.2. Here we have $I_{\vec{D}}(e_1) = 1 = I_{\vec{D}}(e_2)$ and $I_{\vec{D}}(e_3) = -1 = I_{\vec{D}}(e_4)$.

Note that this requires a fixed orientation of Γ. Therefore we have the following relation: $I_{\overrightarrow{D\oplus D'}} = I_{\vec{D}} \pm I_{\vec{D'}}$, where the $\pm$ depends on how $\vec{D}$, $\vec{D'}$ and $\overrightarrow{D \oplus D'}$ are directed. Remember that $\vec{D} \oplus \vec{D}'$ is a cut, otherwise it is $\varnothing$. So this is similar to the theta graph property. The signs present make it possible to work outside of characteristic 2.

The *cut space* over $\mathbf{K}$ is $B^1(\Gamma; \mathbf{K}) := \langle I_D : D \text{ is a cut}\rangle$, the span in $\mathbf{K}^E$ of all indicator vectors of cuts.

M.2. **Signed graphs.** [[LABEL 2.]]

The theory of cycle and cut spaces of signed graphs is largely due to the recent paper by Chen and Wang [CW].

Here the circuits are what are properly called frame circuits. (Lift circuits will be mentioned and later will be suppressed.) The three kinds of circuit look like the following: A

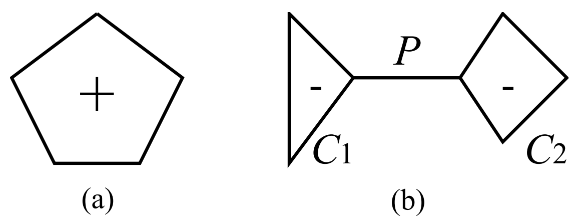

FIGURE M.3. (a) Positive circle, or Type I circuit, (b) Handcuffs, or Types II and III circuits.

[[LABEL 0205image3]]

tight handcuff or Type II circuit is a handcuff where the *circuit path* (*connecting path*) P has length zero. A *loose handcuff* or type III circuit is a handcuff whose circuit path P has length greater than zero.

A *direction* of a circuit is a cyclic orientation (that is, an orientation that has no sources or sinks). This means that we cannot just give a circle an arbitrary orientation as before. Recall that if we do not want sources or sinks then the orientation must be coherent. A divalent vertex is necessarily coherent to avoid being a source or a sink. Therefore orienting one edge forces the rest of the edges present to be oriented in a specific fashion. This means that there exists exactly two different cyclic orientations (directions) of a positive circle. The same will be true for a handcuff, and therefore for all three circuit types.

Indicator vector of C

A *circuit walk* is a minimal closed walk around C.

Given a fixed orientation of $\vec{\Sigma}$, a directed circuit $\vec{C}$ and for each appearance of e in a *circuit walk* around C:

$$I_{\vec{C}}(e) = \begin{cases} 1 & \text{if } e \in \vec{C} \text{ and } \vec{e} \text{ agrees with } \vec{C}, \\ -1 & \text{if } e \in \vec{C} \text{ and } \vec{e} \text{ disagrees with } \vec{C}, \\ 0 & \text{if } e \notin \vec{C}. \end{cases}$$

RESTATE:

$$I_{\vec{C}}(e) = \begin{cases} \pm 1 & \text{if } e \in \vec{C} \text{ and } e \text{ is not in a connecting path } \vec{C}, \\ \pm 2 & \text{if } e \in \vec{C} \text{ and } e \text{ is in a connecting path } \vec{C}, \\ 0 & \text{if } e \notin \vec{C}. \end{cases}$$

Take a walk $W = v_0 e_1 v_1 e_2 \cdots e_l v_l$ in a signed graph Σ. The direction of W gives us an orientation of the edges in W such that each v_i is coherent in W. Call this oriented walk $\vec{W}$. 2009 Feb 10 (draft): Simon Joyce

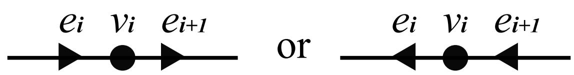

FIGURE M.4. F:0210 The two kinds of coherent edges you could have at $v_i \in \vec{W}$.

If $\vec{\Sigma} = (\Sigma, \tau)$ is a bidirected graph, then each edge $\vec{e}_i \in \vec{W}$ is oriented the same or opposite to the corresponding edge in $\vec{\Sigma}$, so for each e_i we get a + or - depending on whether the orientations of $\vec{\Sigma}$ and $\vec{W}$ agree or not.

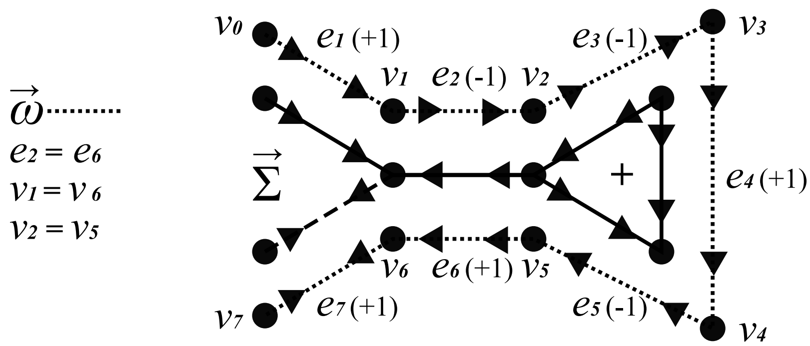

FIGURE M.5. F:0210 Edge signs of $\vec{W}$ for $\vec{\Sigma}$.

$\tau_{\vec{\Sigma}} = \tau$ orients Σ where τ is a bidirection. We can think of τ as a map where τ : $\{\text{edge ends}\} \to \{+, -\}$. $\tau_{\vec{\Sigma}}$ orients edge ends in W, where $\tau_{\vec{\Sigma}}(v_i, e_j)$ depends on i and j where $j = i$ or $j = i + 1$. Note that $\tau_{\vec{W}} = -\tau_{\overleftarrow{W}}$.

M.3. **Indicator vector of a directed walk.**

For a directed frame circuit $\vec{C}$ we define the indicator vector:

$$I_{\vec{C}}(e) = \begin{cases} 0 & \text{if } e \notin C, \\ \pm 1 & \text{for a loose edge or an edge in a circle of } C, \\ \pm 2 & \text{for a half edge or a link in the connecting path of a handcuff.} \end{cases}$$

[**these need to be checked and possibly more added.**]

Definition M.1. [[LABEL D:0210 indicator vector]] Given $\vec{\Sigma}$ a bidirected graph and $\vec{W}$ and a directed walk $\vec{W}$ in $\vec{\Sigma}$, define the *indicator vector*, $I_{\vec{W}}$ to be a map $I_{\vec{W}} : E \to \mathbb{Z}$ such that

$$I_{\vec{W}}(e) = \sum_{e_i = e \in W} \tau_{\vec{W}}(v_i, e_i) \tau_{\vec{\Sigma}}(v_i, e_i).$$

For an abelian group A an *A-flow* is an oriented function $E \to A$ that is conservative at every vertex. [**this sentence needs some attention.**] (We're working over a unital commutative ring $\mathbf{K}$ such that $2 \neq 0$, and possibly we need 2 to be invertible.) [**REMEMBER TO revise this when we figure out what we really need.**]

[**the caption may need attention**]

Ridiculous research questions.

(a) Can there be a matroid on $E(\Sigma)$ whose circuits are the C_3's, the positive circles (including

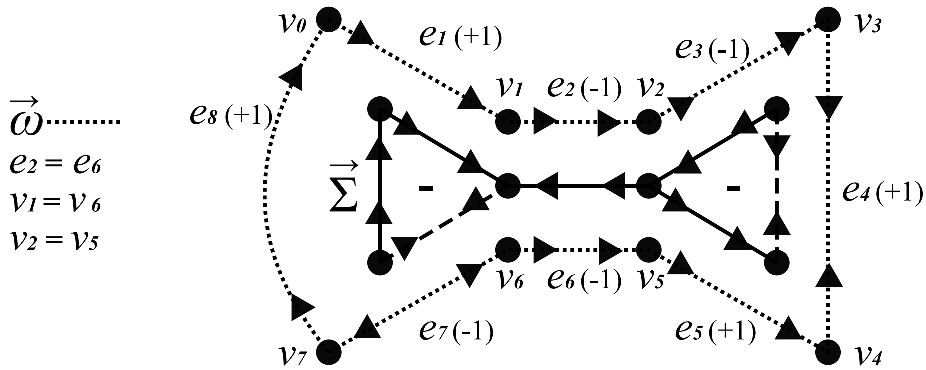

FIGURE M.6. F:0210 Indicator vectors on the graph $\vec{\Sigma}$.

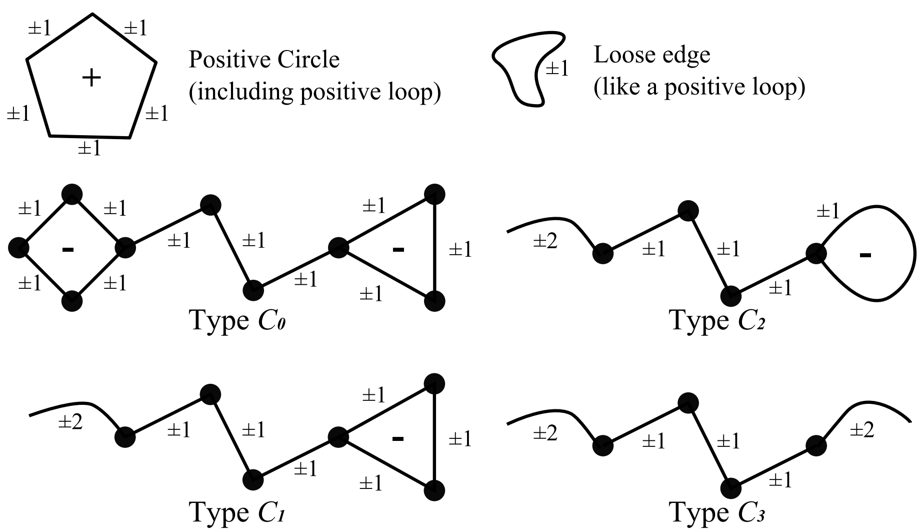

FIGURE M.7. F:0210 Indicator vectors on frame circuits.

loose edges), the ± 1 edges in each C_0, C_1 and C_2 (I don't think so).
(b) Roughly speaking, if not then C_3's should possibly have ± 1's.
(c) Does this help decide between ± 1's on C_3 and ± 2's on C_3.

Hopeful conjecture: We basically get $G(|\Sigma|)$. If ± 1's on C_3 we get $G(|\Sigma| + v_0)$ where v_0 is incident to every half edge, but this might need a half edge to be true.

2009 Feb 12: Jackie Kaminski

M.4. **Flows and Cycles.** [[LABEL 2.cyclespaces]]

M.4.1. *Flows.* [[LABEL 2.flows]]

We begin with the definition of a flow. Throughout this section, we will assume Σ is an oriented signed graph, that is, a bidirected graph. R is a commutative ring.

Definition M.2. [[LABEL D:0212 flow]] An *R-flow* on Σ (also known as a *1-cycle* over R) is a function $f : \vec{E} \to R$ such that at every vertex v,

$$\partial f(v) := \sum_{\varepsilon : v(\varepsilon)=v} f(e(\varepsilon)) \cdot \tau_{\vec{\Sigma}}(\varepsilon) = 0,$$

where the sum is over edge ends ε of $|\Sigma|$, $v(\varepsilon)$ denotes the vertex of the edge end ε, and $e(\varepsilon)$ denotes the edge containing the edge end ε. The *cycle space* or *flow space* of Σ over R is the set of all R-valued flows (or 1-cycles), denoted by $Z_1(\Sigma; R)$.

The condition that $\partial f(v) = 0$ is often stated colloquially as 'the flow is conserved at vertex v', and a flow is called *conservative* if it is conserved at every vertex. The notation Z_1 is chosen to be consistent with that of algebraic topology and homological algebra.

Although we need an orientation on Σ to talk about flows, mostly it's just as a reference point.

Proposition M.1. *[[LABEL P:0212 Z]] $Z_1(\Sigma; R)$ = the null space $\mathrm{Nul}(\mathrm{H}(\Sigma))$ over R.*

Proof. We can think of $f : \vec{E} \to R$ as an $|E| \times 1$ column vector $\vec{f}$ with entries in R. (This is similar to how any function from a finite set of size n can be thought of as an n-tuple.) Now $\vec{f} \in \mathrm{Nul}\,\mathrm{H}(\Sigma)$ if and only if $\mathrm{H}(\Sigma)\vec{f} = \vec{0}$, by definition of the null space. Now $\mathrm{H}(\Sigma)\vec{f} = \vec{0}$ if and only if each row of $\mathrm{H}(\Sigma) \cdot \vec{f}$ is 0, which it is if and only if for each row v of $\mathrm{H}(\Sigma)$,

$$\sum_{e\in E} \eta_{v,e} \cdot f(e) = 0, \iff \sum_{e\in E} \Big(\sum_{\substack{\varepsilon: \\ e(\varepsilon)=e \ \& \ v(\varepsilon)=v}} \tau_{\vec{\Sigma}}(\varepsilon) \Big) \cdot f(e) = 0.$$

Combining into a single summation over all edge ends incident with v, we see that the above is true if and only if

$$\sum_{\varepsilon : v(e)=v} f(e(\varepsilon)) \cdot \tau_{\vec{\Sigma}}(\varepsilon) = 0,$$

which is of course the definition of $\partial f(v) = 0$ for all v.

Therefore $\vec{f} \in \mathrm{Nul}(\mathrm{H}(\Sigma)$ if and only if f is an R-flow. □

Since negating a row doesn't alter the null space of $\mathrm{H}(\Sigma)$, switching a vertex (in both the graph and the flow) doesn't alter a flow. Furthermore, if we negate a column of $\mathrm{H}(\Sigma)$, and then negate the corresponding edge in f, we haven't altered anything about the flow. So in some sense we're considering switching classes yet again. And more importantly we can see that in some ways we really are only using the bidirection in Σ to know whether $f(e)$ is a or $-a$ (for $a \in R$), so it will be nice if we can set things up to have the same orientation on the flow as on Σ.

Further, orthogonality is unaltered by negating the flow value on an edge as well as by negating a column of $\mathrm{H}(\Sigma)$. So if the information we are really interested in is orthogonality, switching doesn't matter at all.

[These two paragraphs should be a general remark about the effect of reorientation and switching on the various spaces.]

Definition M.3. [[LABEL D:0212 circuit space]] The *circuit space* of Σ, $Z(\Sigma; R)$, is the span over R of the indicator vectors of circuits.

It is clear that $Z(\Sigma; R) \subseteq Z_1(\Sigma; R)$, but although there is sometimes equality, they may disagree, for instance when $R = \mathbb{Z}$.

We now return to the argument of what value we want the indicator vectors to have on circuits of the form of two half edges with a connecting path between.

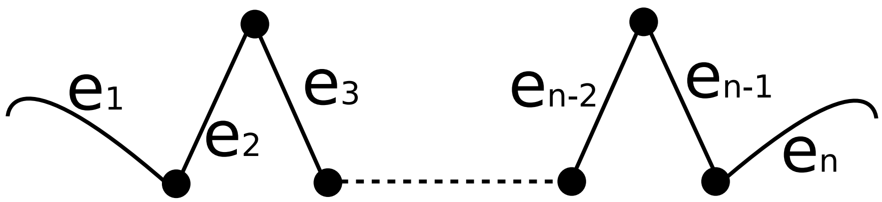

FIGURE M.8. A circuit in $\vec{\Sigma}$, with orientations omitted.
[[LABEL F:0212 Circuit with Half Edges]]

We have our definitions for $I_C(e)$ in circuits as given in **??**, but it's unclear what value we would like the indicator vector to have on the edges of the connecting path of the circuit of this type. Arguments can be made for either ±1's or for ±2's, where the $\pm$ is determined—it simply depends on whether the given orientation of Σ agrees or disagrees with the chosen directed circuit walk.

The arguments in favor of having ±2's is that this is consistent with a circuit path for circuits consisting of two negative circles connected by a circuit path. Additionally, it make it clear that for circuits consisting of one half edge and one negative circle with a circuit path, there is no ambiguity or confusion about what values $I_C(e)$ should have.

As an argument for ±1's, we notice that when we look at the circuit structure of Σ a signed graph (no restrictions, half edges and loose edges allowed), we could get the same information from looking at $\Sigma + v_0$ under the following construction, $V(\Sigma + v_0) = V(\Sigma) + v_0$, and $E(\Sigma + v_0) = \{e | e \text{ is a link or loop in } \Sigma\} \cup \{e^-{:}vv_0 | e \text{ is a half edge in } \Sigma \text{ incident to } v\} \cup \{e^+{:}v_0v_0 | e \text{ is a loose edge in } \Sigma\} \cup \{e^-{:}v_0v_0\}$. Colloquially, keep all links and loops of Σ, then add a new vertex, v_0 with a negative loop. Then replace every half edge (at vertex v), with a negative edge from v to v_0. Finally, Replace every loose edge with a positive loop at v_0.

When $\Sigma = +\Gamma$, readers familiar with matroid theory will notice that the matroid for $\Sigma + v_0$ (as defined above) is isomorphic to the matroid for Σ. Therefore, finally meandering around to our point, we notice that circuits of the form in Figure M.8 turn into positive circles, and the indicator vector of an edge in a positive circle has value ±1.

This leads us to the proposition (the justification of which has already been given).

Proposition M.2. *[*[LABEL P:0212 matroid stuff]] *For Σ a signed graph, with $|\Sigma| = \Gamma$, $\Sigma \cong ((\Gamma^\pm + v_0) \cup e{:}v_0)/\{e{:}v_0\}$.*

In matrix terms, this says $\mathrm{H}(+\Gamma) = \mathrm{H}(\Gamma^\pm + v_0)$, and in matroid terms $G(+\Gamma) = G(\Gamma^\pm + v_0)$.

We recall that $G(+\Gamma)$ means the frame matroid of $+\Gamma$, and we notice that $G((\Gamma^\pm + v_0) \cup e{:}v_0) = G(+\Gamma) \oplus h_0$ coloop. Finally, we close this section with the comment that in graph theory (meaning unsigned graph theory) Z and Z_1 are the same, since $\mathrm{H}(\Gamma)$ is a totally unimodular matrix.

M.5. **Cuts.** [[LABEL 2.cuts]]

Before we even state the definition of a cut in a signed graph, we want to clearly point out that a cut in Σ is *not* always a cut in $|\Sigma|$—and vice versa.

Definition M.4. [[LABEL D:0212 cut]] A *cut* in a signed graph is a nonempty set U of the form $U = E(X, X^c) \cup U_X$ where $X \subseteq V$, and U_X is a minimal total balancing set of $\Sigma{:}X$.

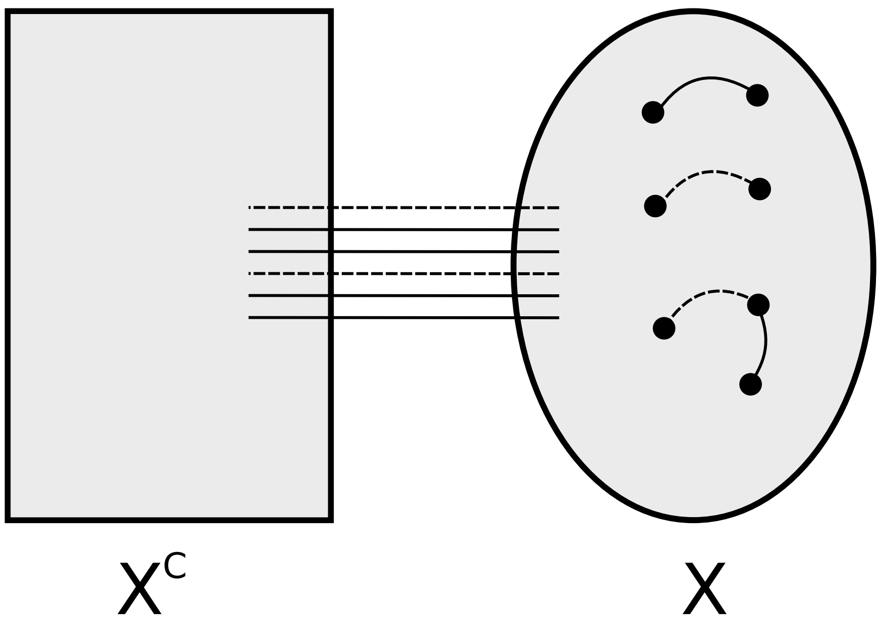

FIGURE M.9. A Cut is Σ
[[LABEL F:0212 cut]]

In Figure M.9, we see the edges of a cut indicated. The rectangle represents $\Sigma{:}X^c$, the oval represents $\Sigma{:}X$. The edges between the two is $E(X, X^c)$ and are part of the cut U. The other edges in $\Sigma{:}X$ represent a minimal balancing set (edges whose removal makes $\Sigma{:}X$ balanced), these are they U_X edges, and they are also part of the cut U.

Although the (unsigned) graph cuts $E(X, X^c)$ and $E(X^c, X)$ are identical (both are the same edge set), in a signed-graph cut, reversing the roles of X and X^c almost always changes the cut, because it changes which set we need to balance, and consequently where the edges in U_X are taken from.

Definition M.5. [[LABEL D:0212 bond]] A *bond* is a minimal cut.

Bonds are, in a vector space sense, dual to circuits, although this relationship is very difficult to express in graph terms. Although for the purpose of justification we point out an example in (unsigned) graph theory. A minimal cut (bond) in a planar graph, is a circuit in the planar dual graph. And, although we are not getting into details here, the subset of the vector space F^E spanned by the circuits of Γ is dual to the vector subspace spanned by the bonds (which is the same subspace spanned by the cuts).

We now define a directed cut in a signed graph. It is an admittedly messy definition.

Definition M.6. [[LABEL D:0212 directed cut]] If $E{:}X \setminus U_X$ is all positive, direct U as follows. Orient each edge of U so that its ends in X satisfy $\tau(\varepsilon) = +1$ (orient the ends into the vertex), or so that for all ends of edges in U that are incident with a vertex in X, $\tau(\varepsilon) = -1$ (orient all edge ends out of the vertices in X). (These two conventions are completely opposite to each other.)

If $E{:}X \setminus U_X$ is not all positive, then switch so that $E{:}X \setminus U_X$ is all positive. This is always possible since every balanced graph is switching equivalent to an all positive graph. Now direct the edges of U as above.

Finally, switch back to the original signature function on Σ, using the same switching as above. Then U is a *directed (signed) cut*.

Notice that, if ζ is a switching function that makes $E{:}X \setminus U_X$ all positive, then $-\zeta$ also does so. Thus, we have a choice of two switching functions, one the negative of the other. If we apply $-\zeta$ with the convention that ends in U are oriented into X, then we get the same directed cut as if we had applied ζ with the opposite convention on orientation. Thus, we only need to define a directed cut with the first convention; the opposite alternative exists of necessity.

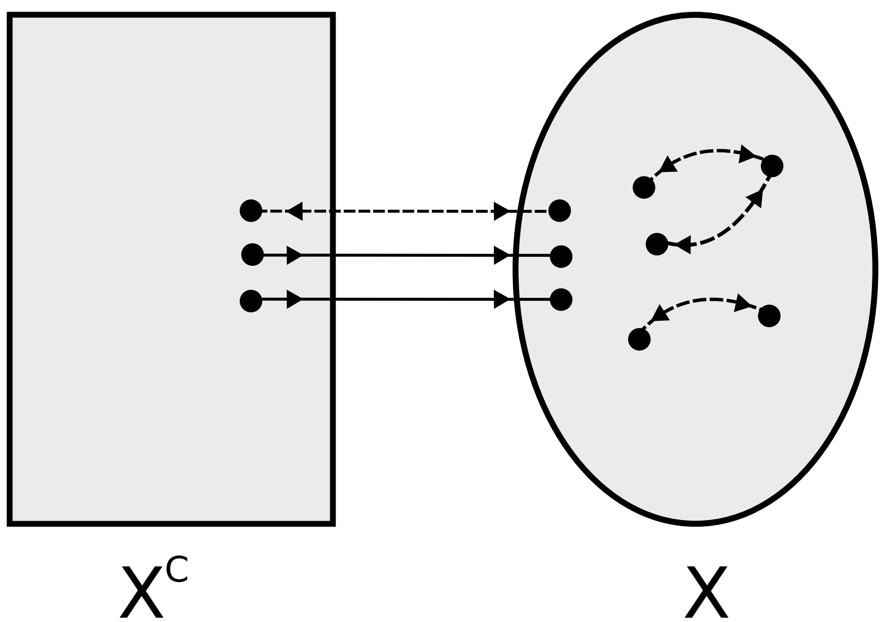

FIGURE M.10. A directed cut in Σ; notice that $(\Sigma{:}X) \setminus U_X$ is balanced [[LABEL F:0212 directed cut]]

Figure M.10 shows a directed cut, where $(\Sigma{:}X) \setminus U_X$ is balanced. Here we have chosen the convention of directing our cut edges into X, but the exact oposite direction is also a directed cut. We notice that since we assume $(\Sigma{:}X) \setminus U_X$ is balanced, and that U_X is a *minimal* balancing set, all U_X edges are negative. Thus the consequence of directing all edges into X is consistent with the edge signs. For the $E(X, X^c)$ edges, regardless of their sign, we direct the ends incident to X into X, then the other end of each edge is directed consistently with its sign.

Finally, we end this section by introducing the indicator vector of a directed cut.

Definition M.7. [[LABEL D:0212 cut indicator]] Let $\vec{U}$ be a directed cut, and $\tau_{\vec{U}}(\varepsilon)$ be the direction of ε in $\vec{U}$ (for ε an edge end in U), finally assume $\vec{U}$ is directed into X. Furthermore, assume that the direction of an edge $e \in \Sigma$ agrees with the direction of e in the cut. Then

$$I_{\vec{U}}(e) = \sum_{\substack{\varepsilon:e(\varepsilon)=e,\\ v(\varepsilon)\in X}} \tau_{\vec{U}}(\varepsilon).$$

Since we have assumed that the directions of the edges in Σ agree with their directions in $\vec{U}$,

$$I_{\vec{U}}(e) = \begin{cases} 0 & \text{if } e \notin U, \\ 1 & \text{if } e \in E(X, X^c), \\ 1 & \text{if } e \in U_X \text{ is a half edge,} \\ 2 & \text{if } e \in U_X \text{ is a link or loop} \end{cases}$$

If there is an edge whose orientation in Σ disagrees with its orientation in $\vec{U}$, we just have a negative value for the indicator vector. On a similar note, if we reverse the orientation of every edge in a cut, we simply negate $I_{\vec{U}}(e)$.

M.6. **The three types of cut.**

2009 Feb 19: Nate Reff

Two kinds of balancing set.

Recall from Definition D.3 that a *partial balancing set* S is a set such that $b(\Sigma \setminus S) > b(\Sigma)$. A *total balancing set* S is a set such that $\Sigma \setminus S$ is balanced.

Notice that if S is a total balancing set then it is not necessary that S be a partial balancing set. Consider the set $S = \varnothing$ where Σ is balanced; then $\Sigma \setminus S$ is balanced but $b(\Sigma \setminus \varnothing) = b(\Sigma)$, hence S is not a partial balancing set. Further, a partial balacing set is not necessarily a total balancing set because you only are increasing the number of balanced components in the deletion and Σ might not be balanced.

Cuts.

There are two kinds of minimal total balancing set S, distinguished by how they change the components of Σ:

$$\text{(i) } c(\Sigma \setminus S) = c(\Sigma),$$

$$\text{(ii) } c(\Sigma \setminus S) > c(\Sigma).$$

Type (i) does not separate components after deletion, but Type (ii) increases the number of components after deletion.

Recall that a *cut* in a signed graph is a nonempty set U of the form $U = E(X, X^c) \cup U_X$ where $X \subseteq V$, and U_X is a minimal total balancing set of $\Sigma{:}X$. Also remember that a *bond* is a minimal cut. See Figure M.11 for an illustration of the general form of a cut.

Here is an easy but important lemma.

Lemma M.3. *[[LABEL L:0219components]] $\pi\big((\Sigma{:}X) \setminus U_X\big) = \pi(\Sigma{:}X)$, or equivalently, $c\big((\Sigma{:}X) \setminus U_X\big) = c(\Sigma{:}X)$.*

Proof. [**I can add a short proof, or you can.**] □

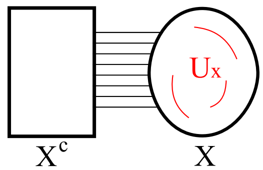

FIGURE M.11. A typical cut in Σ
[[LABEL F:0219image1]]

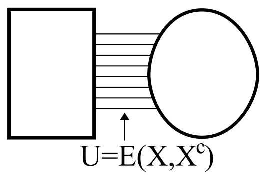

FIGURE M.12. A cut of Type I in Σ.
[[LABEL F:0219image2]]

It's necessary to distinguish three kinds of cut, depending on which of $E(X, X^c)$ or U_X may happen to be empty. Chen and Wang call them "Types I, II, and III".

Type I: A graph cut. In other words, $U_X = \varnothing$. See Figure M.12.

Type II: A cut that is a strict balancing set. In other words, $E(X, X^c) = \varnothing$. This means that $\Sigma{:}X^c$ is a union of components of Σ, and $\Sigma{:}X$ is a union of components of Σ. See Figure M.13.

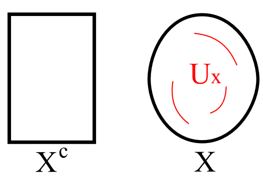

FIGURE M.13. A cut of Type II in Σ.
[[LABEL F:0219image3]]

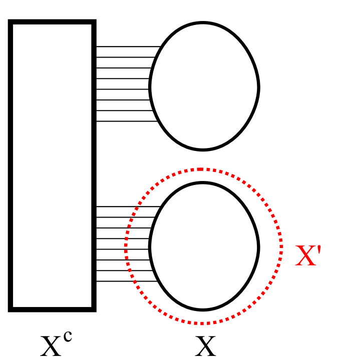

FIGURE M.14. Type I, not a bond.
[[LABEL F:0219image4]]

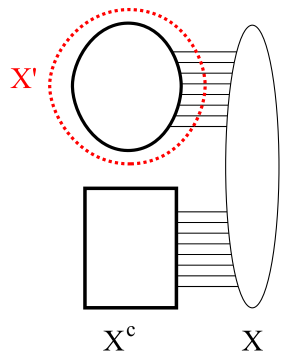

FIGURE M.15. Type I, not a bond.
[[LABEL F:0219image5]]

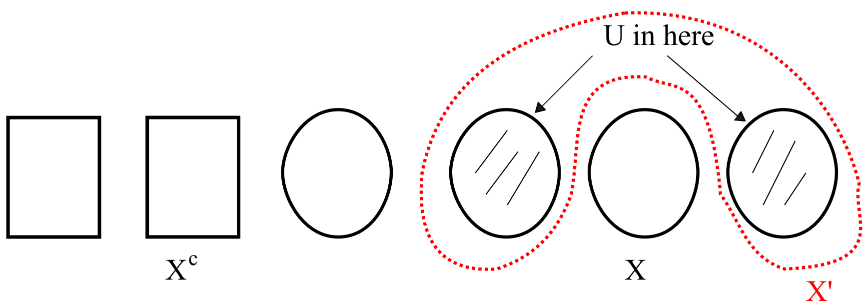

FIGURE M.16. Type II.
[[LABEL F:0219image6]]

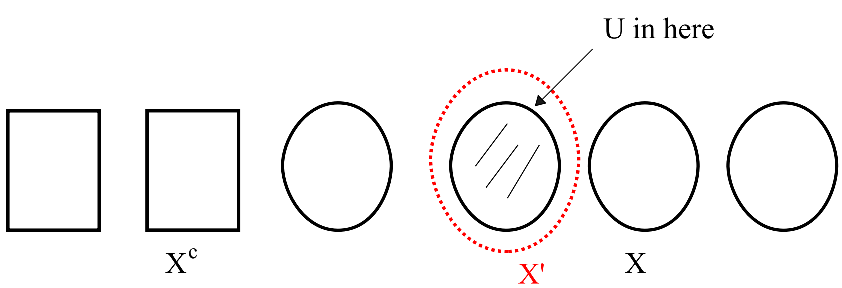

FIGURE M.17. Type II.
[[LABEL F:0219image7]]

In Figures M.14, M.15, M.16 and M.17 we have the two cut types which are not a bond. If we instead choose X' to be our set X then the result would be a bond. [**This needs more explanation. How did we come up with these figures? What are the significant points about the figures? What is the reason for these statements about them?**]

Lemma M.4. *[[LABEL L:0219bondII]] In a type II cut, if U is a bond then we can choose X to be the vertex set of one component of Σ.*

Proof. Choose the vertex set of the union of the vertices of the U_X's in the components of X. □

Sublemma M.5. *[*[LABEL L:0219sublemma1]] *If Σ has a cut U and a component Σ', then $U \cap E'$ is empty or a cut of Σ'.*

Lemma M.6. *[*[LABEL L:0219cutcomponent]] *A bond of Σ is a bond of a component, and a cut is the disjoint union of cuts of one or more components.*

This lemma will allow us to work component by component.

Type III: A *mixed cut*, where $U_X \neq \varnothing$, $E(X, X^c) \neq \varnothing$, and (of necessity) $X, X^c \neq \varnothing$.

Lemma M.7. *[*[LABEL L:0219cut]] *If Σ is balanced then a cut is the same as a graph cut and a bond is the same as a graph bond.*

Proof. The balancing set has to be empty, $U_X = \varnothing$. □

If Σ is unbalanced then we have one of the three types of cuts as described above. What is a bond, then? A bond is either:

(1) A minimal partial balancing set of Σ, which is not a graph cut.
(2) A graph bond of $|\Sigma|$, $E(X, X^c)$ such that $\Sigma{:}X$ is balanced but $\Sigma{:}X^c$ is not.
(3) A graph bond that creates no balanced components, with $E{:}X$ connected, $b(\Sigma{:}X^c) = 0$, together with a minimal total balancing set of $\Sigma{:}X$.

[INSERT PICTURE]

Suppose U is a bond. If one component $\Sigma{:}X_1$ of $\Sigma{:}X$ is balanced then $E(X_1, X_1^c) = U_1 \subseteq U$. Therefore, no component of $\Sigma{:}X$ is balanced. If $\Sigma{:}X$ is not connected then $E(X_1, X_1^c) = E(X_1, X^c)$ because $E(X_1, X_2) = \varnothing$.

Lemma M.8. *[*[LABEL L:0219balset]] ***[DOES the assumption apply to all three parts?]***

(1) If Σ is connected and unbalanced, then a a total partial balancing set is a partial balancing set.
(2) A minimal total balancing set is not a graph cut.
(3) A minimal partial balancing set is either a graph cut or a total balancing set.

2009 Feb 24: Yash Lodha

M.7. **Spaces and orthogonality.**

In the following treatment of edge spaces and subspaces, $\mathbf{K}$ is a field or $\mathbb{Z}$ or an integral domain.

The *edge space* is $\mathbf{K}^E = \{f : E \to \mathbf{K}\}$. The edge space, its members, and its subspaces are always defined with respect to an arbitrary fixed orientation $\vec{\Sigma}$ of Σ. I will omit the orientation from the notation, but don't forget about it!

The *vertex space* is $\mathbf{K}^V$.

Flows and 1*-cycles.*

Definition M.8. [[LABEL Df:0224conserv]] A function in the edge space of Σ is *conservative at $v \in V$* if

$$\sum_{\varepsilon : v(\varepsilon) = v} f(e(\varepsilon)) \tau_\Sigma(\varepsilon) = 0.$$

Here ε denotes an incidence; $v(\varepsilon)$ is its vertex and $e(\varepsilon)$ is its edge. It is *conservative* if it is conservative at every vertex. We call f a *flow*, or a 1*-cycle*, if it is conservative at every vertex.

The 1-*boundary* operator $\partial : \mathbf{K}^E \to \mathbf{K}^V$ is defined by

$$(\partial f)(v) := \sum_{\varepsilon : v(\varepsilon) = v} f(e(\varepsilon))\tau_\Sigma(\varepsilon).$$

Thus, f is a flow iff it lies in the kernel of the boundary operator. (We rarely if ever use other boundary operators, so I will normally omit the "1".)

The *cycle space*, or *flow space*, is the set of all flows:

$$Z_1(\Sigma; \mathbf{K}) := \{f \in \mathbf{K}^E : \partial f = 0\}.$$

Lemma M.9. *[*[LABEL L:0224boundarymap]] *For an edge function f regarded as a column vector, $\partial f = \mathrm{H}(\Sigma) f$.*

That is, H is the matrix of ∂ with respect to the canonical bases of $\mathbf{K}^E$ and $\mathbf{K}^V$.

Proof. ?? □

Proposition M.10. *[*[LABEL P:0224]] $Z_1(\Sigma; \mathbf{K}) = \operatorname{Nul} \mathrm{H}(\Sigma)$.

Proof. By Lemma M.9, an edge function f is conservative iff $f \in \operatorname{Nul} \mathrm{H}(\Sigma)$. □

The *circuit space* $Z(\Sigma; \mathbf{K})$ is the subspace of the edge space $\mathbf{K}^E$ generated by indicator vectors I_C of directed circuits.

Lemma M.11. *[*[LABEL L:0224]] $Z \subseteq Z_1$.

Proof. We defined the indicator vector so it is conservative at every vertex, thus $\partial I_C = 0$. □

That lemma is valid over a commutative, unital ring $\mathbf{K}$, because it only requires that there be a multiplicative identity. The theorem, however, is not as general.

Theorem M.12. *[*[LABEL T:0224zz1]] *Over a field $\mathbf{K}$, $Z = Z_1$.*

Proof. We want to show that the null space, $\operatorname{Nul} \mathrm{H}(\Sigma)$, is generated by circuit indicator vectors.

Recall that the minimal dependent sets of columns are the sets corresponding to frame circuits. (Provided the characteristic of $\mathbf{K}$ is not 2. For characteristic equal to 2 everything is in $|\Sigma|$; the incidence matrix is $\mathrm{H}(|\Sigma|)$, the minimal dependent sets correspond to circles, and so forth. We treated this in Section I.**??**.)

Therefore, if we take a maximal circuit-free set B of columns in H, every other column is generated by those columns via indicator vectors of circuits. To be specific, for each edge $e \notin B$, let $C(e)$ be the unique circuit contained in $B \cup e$. (The existence of this circuit is guaranteed by matroid theory. I will leave that step aside.) The column of e, x_e, is generated by using $I_{C(e)}$ to form a linear combination of the columns from $C(e)$. In the indicator vector, $I_{C(e)}(e) = \alpha_e$, which is ± 1 or ± 2. $I_{C(e)}(f) = \pm 1$ or ± 2 if $f \in B \cup e$, 0 if $f \notin B \cup e$. We use the equation

$$\alpha_e x_e + \sum_{f \in B} I_{C(e)}(f) x_f = \vec{0}.$$

We can solve for x_e by dividing by α_e.

Write $B := \{e_1, e_2, \ldots, e_m$. Let's rearrange the incidence matrix into a convenient form [**(Diagram missing here)**]. In the edge space:

$$I_{C_1} \text{ is such that } I_{C_1}(e_1) \neq 0,\ I_{C_1}(e_2) = 0, \ldots,$$

$$I_{C_2} \text{ is such that } I_{C_2}(e_2) \neq 0,\ I_{C_2}(e_1) = 0, \dots,$$
$$I_{C_3} \text{ is such that } I_{C_3}(e_3) \neq 0,\ I_{C_3}(e_1) = 0, \dots,$$
$$\cdots$$
$$I_{C_m} \text{ is such that } I_{C_m}(e_m) \neq 0,\ I_{C_m}(e_1) = 0, \dots.$$

These vectors are linearly independent and they span $\operatorname{Nul} \mathrm{H}(\Sigma)$. Therefore, $\operatorname{Nul} \mathrm{H}(\Sigma) \supseteq Z(\Sigma; \mathbf{K})$. □

M.7.1. *Cuts.*

Next, let's look at the signed analogs of cuts. We need the dual of the boundary operator. The 0-*coboundary* operator $\delta : \mathbf{K}^V \to \mathbf{K}^E$, which takes a vertex vector $g \in \mathbf{K}^V$ to an edge vector $\delta(g) \in \mathbf{K}^E$, is defined by

$$\delta(g)(e) = \begin{cases} g(w) - g(v) & \text{if } v >> w, \\ g(w) + g(v) & \text{if } v <> w, \\ -g(w) - g(v) & \text{if } v >< w, \\ -g(w) + g(v) & \text{if } v << w. \end{cases}$$

(The <> etc. show the orientation of edge $e{:}vw$ at the endpoints.) **[They are to be replaced by diagrams.]**

Definition M.9. [[LABEL Df:0224B1]] $B^1(\Sigma; \mathbf{K}) := \{\delta g : g \in \mathbf{K}^V\}$. Thus $\delta(g) = \mathrm{H}(\Sigma)^T g$, so B^1 is the row space of $\mathrm{H}(\Sigma)$.

The *cut space* is $B(\Sigma; \mathbf{K}) =$ the span (over $\mathbf{K}$) of indicator vectors of cuts.

Notice that $I_{\{u\}} = \delta(g)$ if we define, for a half edge $e{:}v$,

$$I_{\{u\}}(e) := \begin{cases} \pm 1 \text{ when } u = v, \\ 0 \text{ when } u \neq v, \end{cases}$$

and we treat $I_{\{u\}}$ as the vector (in $\mathbf{K}^E$) of its values on the edges.

Lemma M.13. *[*[LABEL L:0224BinB1]] $B \subseteq B^1$.

Proof. **[PROOF?]** □

Theorem M.14. *[*[LABEL T:0224B1]] $B = B^1$ *over a field* $\mathbf{K}$.

Proof. Exercise. Possibly a dimension argument. **[PROOF?]** □

Theorem M.15. *[*[LABEL T:0224]] Z_1 *and* B^1 *are orthogonal complements in the edge space over a field.*

Proof. The row and null spaces of a matrix are orthogonal complements. $B^1 = \operatorname{Row} \mathrm{H}$, $Z_1 = \operatorname{Nul} \mathrm{H}$. □

Cuts and minimal cuts.

Now here are some contrasting facts.

For graph cuts: The set sum of graph cuts is a graph cut (or ∅). For signed graph cuts: that is false.

For graphs: Every cut is a disjoint union of bonds. For signed graphs: Not even a set sum of bonds.

For signed graphs: A directed bond is a minimal directed cut, but a minimal directed cut need not be a bond.

[**Now comes an example graph with a table of cuts, bonds etc.)**]

Theorem M.16 (Chen and Wang [CW]). *[*[LABEL T:0224dicutunion]] *In a signed graph, every directed cut is a disjoint union of minimal directed cuts.*

I refer to Chen and Wang's important paper for the proof. We just don't have time for it! (Alas.)

Readings and Bibliography

References in the form Name (yearx) will be found in:

[BSG] Thomas Zaslavsky, 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.

The most up-to-date version is on my Web site at

`http://www.math.binghamton.edu/zaslav/Bsg/`

A. Background

[Pearls] Nora Hartsfield and Gerhard Ringel, *Pearls in Graph Theory.* Academic Press, Boston, 1990. MR 91f:05001. Zbl 703.05001. Reprinted by Dover, Mineola, N.Y., 2003.

Ch. 1, §§ 2.1–2, § 3.1, Ch. 8 for elementary background in graph theory. Very readable. Not always precisely correct, so make sure you understand the proofs.

[BM] J.A. Bondy and U.S.R. Murty, *Graph Theory with Applications.* American Elsevier (North-Holland), New York, 1976. MR 54 #117.

An advanced introduction, conceptually oriented. The best graph theory textbook to find background for this course. In connection with our course, see especially Chapters 1–3, 6, 8, 9, 12, and some of 10. The table of contents and links to the entire book are on the Web at

`http://www.ecp6.jussieu.fr/pageperso/bondy/bondy.html`

[Tbook] W.T. Tutte, *Graph Theory.* Encyc. Math. Appl., Vol. 21. Addison–Wesley, Reading, Mass., 1984. MR 87c:05001. Zbl 554.05001. Reprinted by Cambridge Univ. Press, Cambridge, Eng., 2001. MR 2001j:05002. Zbl 964.05001.

A rigorous advanced introduction from the viewpoint of one of the great graph theorists. A large part of Tutte's research, presented in readable form—if you pay close attention. Includes Tutte's definition of connectivity, his decomposition theory for graphs that are not 3-connected, the basics of his treatment of graph invariants, and much more.

[PG] Martin Charles Golumbic, *Algorithmic Graph Theory and Perfect Graphs.* Academic Press, New York, 1980. MR 81e:68081. Zbl 541.05054.

Readable treatment of special topics, e.g., chordal graphs and comparability graphs.

[Cox] H.S.M. Coxeter, *Introduction to Geometry*, second ed. Wiley, New York, 1969. MR 49 #11369. Zbl 181, 481a (e: 181.48101).

Ch. 12–14, §§ 15.1–3. Basic geometrical background for use in the course.

[Oxley] J.G. Oxley, *Matroid Theory.* Oxford Univ. Press, Oxford, 1992. MR 94d:05033. Zbl 784.05002.

The standard place to learn matroid theory, to which I refer occasionally.

[FUTA] Thomas Zaslavsky, *Facing Up to Arrangements: Face-Count Formulas for Partitions of Space by Hyperplanes.* Mem. Amer. Math. Soc., No. 154 (= Vol. 1, Issue 1). American Mathematical Society, Providence, R.I., 1975. MR 50 #9603. Zbl 296.50010.

Real hyperplane arrangements.

[AH] Peter Orlik and Hiroaki Terao, *Arrangements of Hyperplanes*. Grundl. math. Wiss., Vol. 300. Springer-Verlag, Berlin, 1992. MR 94e:52014. Zbl 757.55001.
Complex hyperplane arrangements.

B. Signed Graphs

[SG] T.Z., Signed graphs. *Discrete Appl. Math.* **4** (1982), 47–74. Erratum. *Discrete Appl. Math.* **5** (1983), 248. MR 84e:05095. Zbl 503.05060.
Basic (but not necessarily elementary) concepts and properties, including the frame matroid (that is optional), examples, and incidence and adjacency matrices (Section 8).

[SGC] T.Z., Signed graph coloring. *Discrete Math.* **39** (1982), 215–228. MR 84h:05050a. Zbl 487.05027.
Basic (but not necessarily elementary) concepts and properties of coloring and the chromatic polynomial.

[OSG] T.Z., Orientation of signed graphs. *European J. Combin.* **12** (1991), 361–375. MR 93a:05065. Zbl 761.05095.
Oriented signed graphs (bidirected graphs) and their hyperplane geometry.

[MTS] T.Z., Matrices in the theory of signed simple graphs. *Proc. International Conference on Discrete Mathematics* (Mysore, 2009), to appear.
Introductory survey.

[CW] Beifang Chen and Jue Wang, The flow and tension spaces and lattices of signed graphs. *European J. Combin.* **30** (2009), 263–279.

C. Geometry of Signed Graphs

[GRS] T.Z., The geometry of root systems and signed graphs. *Amer. Math. Monthly* **88** (1981), 88–105. MR 82g:05012. Zbl 466.05058.
A readable introduction to some of the connections between graph theory and geometry.

[CGSS] 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 [JJS, pp. 208–230].
A classic research paper in the background of part of our topic. Not introductory.

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

[Dfs] 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.

[E8] G.R. Vijayakumar, Representation of signed graphs by root system E_8. *Graphs Combin.* 10 (1994), 383–388. MR 96a:05128. Zbl 821.05040.
These are the three main papers on vertex representation of signed graphs.

[JJS] J.J. Seidel, *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.

[GR] Chris Godsil and Gordon Royle, *Algebraic Graph Theory*. Grad. Texts in Math., Vol. 207. Springer, New York, 2001. MR 2002f:05002. Zbl 968.05002.
Textbook. Ch. 10 on strongly regular graphs (§§ 10.1–2), Ch. 11 on two-graphs and equiangular lines, and Ch. 12 (and § 1.7) on line graphs (§§ 12.1–3). Presents much

of [CGSS], with its background, in a somewhat more systematic but more advanced way.

D. Gain Graphs

[BG1] T.Z., Biased graphs. I. Bias, balance, and gains. *J. Combin. Theory Ser. B* **47** (1989), 32–52. MR 90k:05138. Zbl 714.05057.

§ 5: Fundamentals of gain graphs.

[BG3] T.Z., Biased graphs. III. Chromatic and dichromatic invariants. *J. Combin. Theory Ser. B* **64** (1995), 17–88. MR 96g:05139. Zbl 857.05088.

§ 5: General theory of coloring gain graphs.

E. Biased Graphs and Gain Graphs

[BG1] T.Z., Biased graphs. I. Bias, balance, and gains. *J. Combin. Theory Ser. B* **47** (1989), 32–52. MR 90k:05138. Zbl 714.05057.

Fundamentals of gain graphs and biased graphs.

[BG2] T.Z., Biased graphs. II. The three matroids. *J. Combin. Theory Ser. B* **51** (1991), 46–72. MR 91m:05056. Zbl 763.05096.

The closure operations that are basic to the theory, in terms of frame matroids (called bias matroids), in §§ 2, 3. Important open problems.

[BG3] T.Z., Biased graphs. III. Chromatic and dichromatic invariants. *J. Combin. Theory Ser. B* **64** (1995), 17–88. MR 96g:05139. Zbl 857.05088.

Chromatic polynomial et al., with and without colorings.

[BG4] T.Z., Biased graphs IV: Geometrical realizations. *J. Combin. Theory Ser. B* **89** (2003), no. 2, 231–297. MR 2005b:05057. Zbl 1031.05034.

Extensive treatment of the geometry of gain graphs. Fundamentals in §§ 2, 4. More advanced topics throughout, especially in § 7.

F. General References

[GS] Garey and Johnson.

[TLect] W.T. Tutte, Lectures on matroids. *J. Res. Nat. Bur. Standards (U.S.A.) Sect. B* **69B** (1965), 1–47. MR 31 #4023. Zbl 151, 338a (e: 151.33801).