Graphs, Gain Graphs, and Geometry

a.k.a. Signed Graphs and their Friends

Course notes for

Math 581: Topics in Graph Theory: Graphs and Geometry

Binghamton University (SUNY)

Fall, 2008 – Spring-Summer, 2009 – Spring, 2010

Version of May 3, 2010

Lectures by

Thomas Zaslavsky

Notes by

Peter Cohen,
Simon Joyce,
Jackie Kaminski,
Yash Lodha,
Nathan Reff,
and
Thomas Zaslavsky

Redaction by

Thomas Zaslavsky

Contents

Chapter O. Background and Introduction

These are the notes of an extended course on signed graphs and (eventually) their generalizations to gain graphs and biased graphs. Briefly, a signed graph is a graph whose edges are labelled from the sign group, a gain graph has edges labelled (invertibly) from an arbitrary group, and a biased graph is a combinatorial abstraction of the latter that still preserves many of its interesting properties without any algebra. Aug 25: Zaslavsky

This course is not comprehensive. It is a personal selection of the parts of the theories of unsigned, signed, and more general graphs that interest me particularly, seem suitable for an introductory course, and fit into the theme of linear-algebraic structures and geometrical interpretations. Matroids, which are abstractions of both linear algebra and geometry, lie behind many of our ideas and results, but they will not be an explicit part of the course. **[Until later, I hope.]**

The first part of the course, Chapter I, presents graphs from this point of view. The main purpose is to show those parts of graph theory that will be generalized in Chapter II. That chapter is intended to show that and how signed graphs are just like graphs, only more general. (This statement should not be taken too literally.) For instance, a signed graph has incidence and adjacency matrices that directly generalize those of an unsigned graph. Chapter III discusses some of the purely geometrical aspects of signed graphs. Chapter IV concludes the notes with vast generalizations.

A. Day One

This is a fast overview of graphs, signed graphs, and their equations and hyperplanes. Aug 25: Simon Joyce

We begin with a few different definitions of a graph, for discursive purposes. All these definitions are popular, but in decreasing order. (Of course, the one we use is the least popular—and the most complicated. We can't help it.)

Insert picture(s) of graphs here for instructional purposes.

Figure A.1. Pictures of some graphs.
[[LABEL F:0825g]]

Definition A.1. [Simple Graph][[LABEL D:0825simplegraph]] A *graph* is a pair $\Gamma = (V, E)$, where V is a set and E is a subset of $\mathcal{P}_2(V)$, the class of unordered pairs of (distinct) elements of V.

This definition doesn't account for things like loops, whose endpoints coincide, or parallel edges, which are edges with the same endpoints as each other, so we need to extend it for our purposes.

Definition A.2. [Multigraph][[LABEL D:0825multigraph]] A *graph* is a pair $\Gamma = (V, E)$, where V is a set and E is a multisubset of $\mathcal{P}_2(V)$.

However, this definition still doesn't account for loops.

The following definition for a graph is what we will use.

Definition A.3. [Graph][[LABEL D:0825graph]] A *graph* is a triple $\Gamma = (V, E, I)$, where V and E are sets and I is an incidence multirelation between V and E in which each edge has incidences of total multiplicity 2.

Definition A.4. [Signed Graph][[LABEL D:0825signedgraph]] A *signed graph* is a graph whose edges have signs, $+$ or $-$. Formally, $\Sigma = (\Gamma, \sigma) = (V, E, I, \sigma)$, where $\sigma : E \to \{+, -\}$.

Insert picture(s) of signed graphs here.

FIGURE A.2. Pictures of some signed graphs.
[[LABEL F:0825sg]]

Starting at a vertex on a graph we can move along one of its incident edges to another vertex and repeat the process from the new vertex any number of times, to move around the graph in any way we please. To describe different ways of moving around a graph we use the following terms:

- A *path* has no repeated edges or vertices.
- A *trail* has no repeated edges.
- A *walk* may have repeated edges and vertices.
- A *circle* is a closed path, that is, it has no repeated vertices or edges except that the initial and the final vertex are the same.

Each edge of a graph implies an equation. The variables correspond to the vertices and an edge with endpoints v_i, v_j corresponds to the equation $x_i = x_j$ in $\mathbb{R}^n$. The family of all hyperplanes corresponding to all edges, $\mathcal{H}[\Gamma]$, called the hyperplane arrangement generated by Γ, divides up $\mathbb{R}^n$ into regions that have a remarkable combinatorial meaning. For a signed graph, a positive edge $+v_iv_j$ has hyperplane $x_i = x_j$ and a negative edge $-v_iv_j$ has hyperplane $x_i = -x_j$. We'll study the geometry of these arrangements of hyperplanes, both to learn more about the graph and to use the graph in order to understand the hyperplane arrangement.

B. AS THINGS COME UP

All kinds of basic background information will be added during the lectures, both in the beginning of Chapter I and when needed as the lectures progress.

Chapter I. Graphs

In this chapter we meet graphs, to develop the understanding and the technical background for signed graphs. Most of what we say about graphs will generalize later, to the more advanced topics of signed graphs, gain graphs, and even biased graphs. Aug 27a: Nate Reff

A. Basic Definitions

[[LABEL 1.defs]]

Here we meet the basic concepts and vocabulary of our version of graph theory.

A.1. **Definition of a graph.** [[LABEL 1.defsgraph]]

We give a formal definition in terms of incidence between vertices and edges. This is rather heavy on notation, so we'll tend to ignore the technical statement in practice, but it's what we mean even when we don't mention it.

- An *incidence multi-relation* I between sets V and E is a multi-subset of $V \times E$.
- A *graph* $\Gamma = (V, E)$ is an ordered pair consisting of sets V and E with an incidence multi-relation I between them such that every edge is incident to at most 2 vertices (not necessarily distinct).
- The elements of V are called the *vertices* of the graph Γ.
- The elements of E are called the *edges* of the graph Γ.

An example:

In the figure edge q is incident to vertex v_4 twice, so $2 \cdot (v_4, q) \in I$. This is consistent with our definition since we do not need edges to be incident to distinct vertices.

Valuable notation:

- Always, $n := |V|$.
- Sometimes, $m := |E|$.
- $V(e)$ is the multiset of vertices of the edge e.
- Suppose $S \subseteq E$; then $V(S)$ is the set of endpoints of edges in S.

A.2. **Types of edge.** [[LABEL 1.edge]]

In the most general definition, there are four kinds of edge in a graph.

- A *loop* is an edge with two equal endpoints. A notation we often use is $e{:}vv$. Another is e_{vv}.
- A *link* is an edge with two distinct endpoints. A notation is $e{:}vw$. Another is e_{vw}.
- A *half edge* is an edge with one endpoint. A notation is $e{:}v$.
- A *loose edge* (or *free loop*) is an edge with zero endpoints. A notation is $e{:}\varnothing$.
- An *ordinary edge* is a link or a loop.

The set of ordinary edges is $E^* := \{$ links and loops $\}$

A.3. **Types of graph.** [[LABEL 1.graphtypes]]

There are three essential kinds of graph:

- A *simple graph* is a graph in which all edges are links and there are no parallel edges (edges with the same endpoints).
- An *ordinary graph* is a graph with no half edges or loose edges; that is, all edges are ordinary.

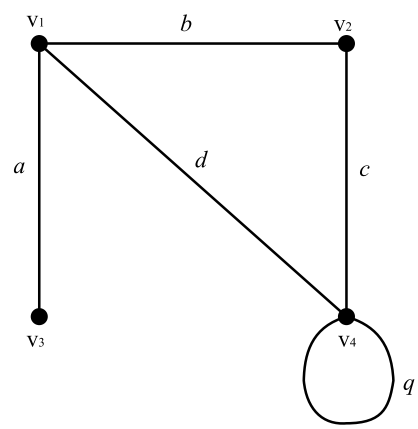

FIGURE A.1. A graph with a loop, but no multiple edges. It is not simple, because of the loop.

[[LABEL F:0827graph]]

- A *link graph* is a graph whose edges are links. A simple graph is a link graph and a link graph is an ordinary graph, but not vice versa, obviously.

Most graph theorists would call these the only kinds of graph, ignoring half and loose edges. We will need those edges later when we generalize to signed graphs and even further; but in this chapter, graphs will be ordinary graphs unless we indicate otherwise.

I should confess that a link graph is usually called a *multigraph* when it isn't simply called a "graph", but both names are also applied to ordinary graphs, where loops are permitted, so I'm adopting the more specific name. I won't confuse the reader by listing any other of the names applied to different kinds of graph.

Special graphs. [[LABEL 1.specialgraphs]]

Aug 29c: Jackie Kaminski

Here are some of the main examples of graphs. All of them are simple.

- A *complete graph*, written K_n, is a simple graph in which every pair of vertices is adjacent. We write K_V when we want a complete graph on a specified vertex set V.
- A *bipartite graph* is a graph whose vertex set has a bipartition $V = V_1 \cup V_2$ such that every edge has one endpoint in V_1 and the other in V_2. It need not be simple.
- A *complete k-partite graph* has vertices partitioned into k (non-empty) parts, and for vertices v, w, if v, w are in the same part, then are no vw-edges. And if v, w are

in different parts, there is a vw edge. A complete k-partite graph with part sizes $n_1, n_2, \ldots, n_k$ is denoted by $K_{n_1,n_2,\ldots,n_k}$.

Figure C shows a complete tripartite graph with tripartition $\{x_1\}, \{v_1, v_2\}, \{w_1, w_2\}$.

Complementation. [[LABEL 1.complements]]

There are three complementations in graph theory: of graphs, of vertex sets, and of edge sets. I will use a superscript c for all of them, as well as for the complement of an arbitrary set within a larger set. Aug 27: Zaslavsky

- The *complement of a simple graph* $\Gamma = (V, E)$ is Γ^c, whose vertex set is V and whose edge set $E(\Gamma^c) := \{vw \mid v, w \in V;\ v \neq w;\ vw \notin E)\}$. That is, $E(\Gamma^c)$ is the set of edges of K_V that are not in Γ.

  Only a simple graph has a complement. There is no absolute notation of complementation for a graph with loops or multiple edges, although one could define the complement of a subgraph within a graph (but we won't).
- The notation X^c, when $X \subseteq V$, denotes the complementary vertex set, $V \setminus X$.
- The notation S^c, when $S \subseteq E$, denotes the complementary edge set, $E \setminus S$.

A.4. **Degree.** [[LABEL 1.degree]]

An edge has a certain number of *ends*: two for a link or loop, one for a half edge, and none for a loose edge. To avoid getting lost in notation, we don't formally define edge ends, but the reader's intuition should make the meaning clear. The important points are that a loop, though it has only one vertex, has two ends, and that the number of ends is the difference between a loop and a half edge. Aug 29c: Jackie Kaminski

Definition A.1. [[LABEL D:0829degree]] The *degree* or *valency* of a vertex, denoted by $d(v)$, is the number of edge ends incident with v.

Hence, a loop adds 2 to the valency (because it has two ends at the same vertex) and a link or half edge adds 1 to the valency of each endpoint.

See Figure A. [**ADD FIGURES**]

Notice that the common definition of the valency of v as the number of neighbors of v is only adequate for simple graphs.

Definition A.2. [[LABEL D:0829isolated]] An *isolated vertex* is a vertex that has no incident edges; i.e., a vertex of degree 0.

Definition A.3. [[LABEL D:0829regular]] A *k-regular* graph is a graph where every vertex has degree k.

A.5. **Types of subgraph.** [[LABEL 1.subgraphtypes]]

There are, of course, subobjects in graph theory; not only subgraphs in general, but also several special kinds. Aug 27a: Nate Reff

- A *subgraph* of Γ is Γ' such that $V' \subseteq V$, $E' \subseteq E$, has the same incidence multi-relation between V and E, every endpoint of every edge in E' is in V', and each edge retains its type.
- A *spanning subgraph* is a subgraph Γ' such that $V' = V$. (Γ' need not have any edges; it just must have all the vertices.)
- $\Gamma \setminus e := (V, E \setminus e)$.

- The *deletion of a vertex set*, denoted by $\Gamma \setminus X$ where $X \subseteq V$, is the subgraph with
$$V(\Gamma \setminus X) := V \setminus X \text{ and } E(\Gamma \setminus X) := \{e \in E \mid V(e) \subseteq V \setminus X\}.$$
The subgraph $\Gamma \setminus X$ includes all the loose edges, if there are any.
- An *induced subgraph* of Γ is a subgraph of the following special form: Let $X \subseteq V$. The subgraph induced by X is
$$\Gamma{:}X := (X, E{:}X), \text{ where } E{:}X := \{e \in E \mid \varnothing \neq V(e) \subseteq X\}.$$
We often write $E{:}X$ as shorthand for $(X, E{:}X)$. In other words, induced subgraphs only contain the inducing vertices, not all the vertices of Γ.
Notice that an induced subgraph has no loose edges; this is the difference between $\Gamma{:}X$ and $\Gamma \setminus X^c$.
- Similarly, $S{:}X$ is the set of edges in S that have all of their endpoints in the vertex set X. We often write $S{:}X$ as shorthand for the subgraph $(S, S{:}X)$.

A.6. **Vertex sets.** [[LABEL 1.specialvsets]]

Aug 29c: Jackie Kaminski

Two kinds of vertex subsets are especially important.

- A *stable* or *independent* set of vertices is a vertex set that induces the empty set of edges; that is, $W \subseteq V$ such that $E{:}W = \varnothing$.
In figure C, $\{x_1\}, \{v_1, v_2\}, \{w_1, w_2\}$ are five stable sets.
- A *clique* is a vertex set whose members are pairwise adjacent.

A.7. **Contraction of an edge.** [[LABEL 1.edgecontraction]]

Aug 27a: Nate Reff

Contraction, intuitively, means shrinking an edge to a point. The two endpoints of a link therefore become one vertex. Oddly, this intuition fails when it comes to contracting half or loose edges—which is why I'll define their contraction here, although it becomes important mainly in connection with signed graphs in Chapter II. The following descriptions cover the basics of contracting an edge. (We'll treat contraction of a set of edges later, in Section C.2.)

The graph Γ with an edge e contracted is denoted by Γ/e.

- *Case 1*: For a link e with vertices v and w, Γ/e has v and w identified to a single vertex and e deleted. Sometimes the identified vertex will be denoted by v_e.
- *Case 2*: For a loop or loose edge, $\Gamma/e = \Gamma \setminus e$.
- *Case 3*: For a half edge e incident to vertex v, to get Γ/e we remove v and e but keep all other edges. A link $f{:}vw$ becomes a half edge $f{:}w$. A loop $f{:}vv$ or a half edge $f{:}v$ becomes a loose edge $f{:}\varnothing$. All other edges remain as they were in Γ.
Intuitively, I think of contracting a half edge $e{:}v$ as like cutting the v end off each edge incident with v, with scissors, and deleting those ends as well as e and v.

Aug 29a: Jackie Kaminski

B. Basic Structures

[[LABEL 1.importantkinds]]

We introduce here some general kinds of graph and some structures within a graph that are essential to graph theory.

We start with some definitions. Recall that $V(e)$ is the set of endpoints of the edge e.

B.1. **Walks, trails, and paths.** [[LABEL 1.walks]]

There are several different ways to get from one place to another in a graph.

- A *walk* is a sequence $v_0 e_1 v_1 \cdots e_l v_l$ where $V(e_i) = \{v_{i-1}, v_i\}$ and $l \geq 0$.
- The *length* of a walk is the number of edges in it, counted as many times as they appear. A walk of length zero is just a vertex.
- A *closed walk* is a walk where $l \geq 1$ and $v_0 = v_l$.
- A *trail* is a walk with no repeated edges.
- A *path* is a trail with no repeated vertex. Sometimes it is called an *open path* to distinguish it from a closed path.
- A *closed path* is a closed trail with no repeated vertex other than that the last vertex is the first one. Despite the name, a closed path is not a path.

B.2. **Connection.** [[LABEL 1.connection]]

Two vertices are said to be *connected* if there exists a path between them. The fundamental property is this:

Theorem B.1. [[LABEL T:0829connequiv]] *The relation of being connected is an equivalence relation on $V(\Gamma)$.*

The proof, which is basic graph theory and is left to the reader, makes use of the next proposition.

Proposition B.2. [[LABEL P:0829walkconn]] *Vertices v, w are connected by a walk $\iff$ they are connected by a path.*

The proof is also basic graph theory and is left to the reader.

Now we explore the implications for graph structure. Here we are assuming, as we normally do throughout this chapter, that the graph is ordinary, that is, without half or loose edges. In this connection (pun intended), half edges are not a problem but loose edges require careful, special treatment.

- A *connected component* (or *vertex component*, or simply *component*) of Γ is the subgraph induced by an equivalence class of the connectedness relation on V.
- We write $c(\Gamma) :=$ the number of components (i.e., vertex components). Often, we write $c(E{:}X)$ as shorthand for $c(X, E{:}X)$, the number of connected components in the subgraph induced by X, and similarly $c(S{:}X) = c(X, S{:}X)$ when $S \subseteq E$.
- We say that Γ is *connected* if the relation of connection on V has exactly one equivalence class.
- The empty graph, $\varnothing := (\varnothing, \varnothing)$ (that is, the graph with no vertices and no edges), is not connected.

  This may seem strange, occasionally even to experienced graph theorists, but it's logically correct: the empty graph does not have exactly one connection equivalence class of vertices.

Aug 27b: Nate Reff

Connection when there are loose and half edges.

The preceding definitions and properties apply to graphs without loose edges. If we want to allow loose edges we need more powerful definitions. Here is one approach.

- A *generalized path* is a sequence $x_0x_1 \dots x_k$ where the x_i's are alternately vertices and edges. If x_i is a link or a loop with endpoints v and w then $\{x_{i-1}, x_{i+1}\} = \{v, w\}$ (note that these are multisets). If x_i is a half edge $e{:}v$, it is x_0 or x_k so that $x_0x_1 = ev$ or $x_{k-1}x_k = ve$. Lastly, if x_i is a loose edge the path is simply x_i.

  Similarly, there are generalized walks and trails.

  Be aware that a generalized path is not necessarily a path, and also that it is unconventional. I introduce it only to explain how elements of a graph that are not necessarily vertices can be considered connected to each other.
- Two elements of Γ, x and y (each of which may be a vertex or edge), are *connected* if there exists a generalized path containing both.

The form Theorem B.1 takes in the more general situation is this:

Theorem B.3. *The relation of being the same or connected is an equivalence relation on $V \cup E$.*

The proof is similar to that of Theorem B.1 so I omit it.

Here are the generalized definitions of connectedness and components:

Definition B.1. [[LABEL Df:0827topcomponent]] A *topological component* of a graph Γ is an equivalence class of $V \cup E$ under the relation of generalized connection.

A *component* (or *vertex component*) of Γ is the subgraph induced by an equivalence class of the connectedness relation on V.

Definition B.2. [[LABEL Df:0827topconn]] We say that Γ is *topologically connected* if the relation of connection on $V \cup E$ has exactly one equivalence class. Equivalently, Γ is topologically connected if it has exactly one (vertex) component and no loose edges, or it is a loose edge.

An alternate definition of a topological component of Γ is as a maximal topologically connected subgraph. Then a component is a topological component that has at least one vertex.

According to our definitions, a loose edge is not a component. This is admittedly strange. Sometimes we might want a loose edge to be a component; we defined topological components to prepare for that possibility, should it ever arise.

Bridges, cutpoints, and blocks.

Bridges are an important concept in connectivity and decomposition of graphs.

Definition B.3. [[LABEL D:1008 bridge]] Let Δ be a subgraph of a graph Γ. A *bridge of* Δ *in* Γ is a maximal subgraph of Γ that is entirely connected without passing through vertices or edges of Δ.

Definition B.4. [[LABEL D:1008 block]] A *cutpoint* of Γ is a vertex with more than one bridge. A *block of* Γ is a maximal subgraph of Γ that has no cutpoints. A *block* (or *block graph*) is a graph that has only one block.

Obviously, each block of a graph is a block graph. Indeed, the blocks of Γ are precisely the maximal block subgraphs.

According to our definition, a vertex is a cutpoint if it supports a loop or half edge and is incident to any other edge. This is not the only existing definition. In fact, the usual one is that a cutpoint is a vertex whose deletion, together with that of every incident edge, increases the number of connected components. That is equivalent to our definition if one first removes all loops and half edges; but our definition is better because it has the important property given in Theorem B.4.

Aug 29b: Jackie Kaminski

B.3. **Circles and pairs of circles.** [[LABEL 1.circles]]

Definition B.5. [[LABEL Df:0829circle]] A *circle* of Γ is a connected 2-regular subgraph of Γ which has at least one vertex, or its edge set. Another definition (equivalent to the first) is that a circle is the graph, or edge set, of a closed path.

For example, any loop is a circle, as is Figure B [**FIGURE NEEDED**]. We require the subgraph to have a vertex in order to exclude loose edges as circles.

Please note that a closed path and the graph of a closed path are not quite the same thing. A closed path has a direction as well as a beginning point. The graph of a closed path has neither. In another direction there is real ambiguity in our use of the term 'circle', as sometimes we mean the edge set, sometimes the graph; but the context should always make the meaning clear.

We denote by $\mathcal{C}(\Gamma)$ the set of circles of a graph Γ.

A main theorem of graph theory concerns the relation of belonging to a common circle.

Theorem B.4. [[LABEL T:1008 blocks and circles]] *Given a graph Γ and $e_1, e_2 \in E(\Gamma)$, e_1 and e_2 are in the same block of Γ if and only if there is a circle in Γ that contains both e_1 and e_2.*

The smallest graphs with two circles are two vertex-disjoint circles, and two circles whose intersection is a single vertex. There is a third kind of graph that, in a sense, has only two independent circles, namely, a *theta graph*, which is the union of three internally disjoint paths between two distinct vertices. This graph has three circles, but any one of them is the set sum of the other two. Theta graphs have an absolutely fundamental role in the entire theory of signed graphs and their graphic generalizations.

FIGURE OF THETA GRAPH

Figure B.1. A theta graph.

B.4. **Trees and their relatives.** [[LABEL 1.trees]]

Graphs without circles, or with a unique circle, will play a large role in our work, the latter especially in the later chapters. Some basic definitions:

- A *tree* is a connected graph which does not contain a circle (as a subgraph).
- A *forest* is a graph which does not contain a circle (as a subgraph).
  Or equivalently, we can define a forest as a graph whose components are all trees. Please refer to Figure C.
- Observations:
  A tree is a connected forest.
  An empty graph (no vertices or edges) is a forest, but is *not* a tree. Recall that a connected graph must have exactly one connected component.
- A *spanning forest* is a spanning subgraph of Γ which is a forest.
  Observe that for any $\Gamma = (V, E)$, the graph $(V, \varnothing)$ is a spanning forest for Γ.

- Similarly, a *spanning tree* is a spanning subgraph of Γ which is a tree.
  Disconnected graphs do not contain any spanning trees.
- A *maximal forest* is a forest which is not properly contained in any other forest.
  Please refer to Figure C.

As an aside, please don't confuse *maximal*, which means not properly contained in any other object (or set) of the same type, with *maximum*, which means having the most elements. For forests in a graph, however, they come to the same thing.

Theorem B.5. [[LABEL T:0829maxforest]] *All maximal forests in Γ have the same number of edges, namely $n - c(\Gamma)$, where $n = |V|$.*

This theorem is elementary, yet not so easy to prove. For a proof see any graph theory textbook. (If you know matroid theory, notice that it is equivalent to the fact that every basis of the graphic matroid has the same size.) Usually, Theorem B.5 is combined with other fundamental properties of maximal forests, as in the following list:

Theorem B.6. [[LABEL T:0829forest]] *For an edge set S in Γ, the following properties are equivalent:*

(i) S is a maximal forest (a maximal edge set that contains no circles).
(ii) S is a minimal edge set that connects everything within each component of Γ.
(iii) S has $n - c(\Gamma)$ edges and connects everything within each component of Γ.
(iv) S has $n - c(\Gamma)$ edges and contains no circles.

Furthermore:

(1) Γ contains a spanning tree $\iff$ it is connected.
(2) A maximal forest consists of a spanning tree of each component of Γ.

The proof is left to the reader—or, see any graph theory textbook.

By definition, the edges not in a maximal forest are the ones that make the circles in Γ. Thus, the number of non-forest edges is, in a sense that can only be made precise through the binary cycle space (Section J.1), the number of independently generated circles of the graph. We call this number the *cyclomatic number* of Γ; that is,

$$\begin{aligned} \xi(\Gamma) &:= |E| - |E(T)| \text{ where } T \text{ is any maximal forest} \\ &= |E| - n + c(\Gamma). \end{aligned}$$

The cyclomatic number of an edge set S is that of the subgraph (V, S), thus $|S| - n + c(S)$.

Tree-like graphs. [[LABEL 1.treelike]]

Other tree-like graphs are:

- A *1-tree* is a tree with one extra edge (not a loose edge). That is, it is a connected graph with cyclomatic number 1 (except that we allow half edges in this definition). See Figure D.
- A *1-forest* is a graph where every component is a 1-tree.
- A *pseudotree* is a graph which is a tree or a 1-tree.
- A *pseudoforest* is a graph in which every component is a pseudotree.

C. Deletion, Contraction, and Minors

A subgraph is "contained" in the graph in the sense of subsets. There are several other ways a graph can "contain" another. The most important is called "containment as a minor". We say Γ_1 contains Γ_2 as a minor if we can get Γ_2 from Γ_1 by any process of repeatedly taking subgraphs and contracting edge sets. Taking a subgraph, which is the same thing as deleting edges and vertices (so it is often called "deletion"), is easy; the complicated part of minors is the operation of contraction.

C.1. **Deletion.** [[LABEL 1.deletionreview]]

We saw several kinds of deletion in Section A.5. The most important for minors is deletion of an edge e or an edge set S, written $\Gamma \setminus e$ or $\Gamma \setminus S$. There is also deletion of an isolated vertex. We can get any subgraph of Γ by first deleting the edges not in the subgraph and then deleting any isolated vertices that are not in the subgraph; every remaining vertex, including all non-isolated vertices, must be in the subgraph.

C.2. **Contraction.** [[LABEL 1.contractionbyset]]

[(Most notes are from an earlier class.)]

We are now restricting ourselves to ordinary graphs again.

- We already defined how to contract a link, loop, half edge, and loose edge.
  Refer to Figure D for a visual representation of contraction by a single edge.
- The contraction of Γ by an edge set $S \subseteq E$ is denoted by $\Gamma/S = (V/S, E \setminus S)$. It is equivalent to a sequence of edge contractions by the edges in S. It can be shown that the resulting graph is the same regardless of the order in which the edges are contracted (provided you aren't too pedantic about the naming of vertices in the resulting graph). Proving this certainly takes some work but is left to the reader.
  See Figure E for an example.
- For a graph Γ, let $\pi(S) :=$ the partition of V such that each block is the vertex set of a (connected) component of (V, S). (Partitions and their blocks are defined in Section D.1.) In other words, $V(\Gamma/S)$ is $\pi(S)$. We will let $[v]$ denote the block of $\pi(S)$ containing the vertex v.
  See Figure F.
- An edge f of the contraction Γ/S is $f \in E \setminus S$, and for $V(f) = \{v, w\}$, f in Γ/S has endpoints $[v], [w]$.

Sept 3a: Yash Lodha

C.3. **Minors.** [[LABEL 1.minors]]

A *minor* of Γ is defined as a contraction of a subgraph of Γ. It turns out that the order of contracting and taking subgraphs makes no difference.

Theorem C.1. *Any graph obtained from a graph Γ by a series of edge contractions and deletions and vertex deletions is a minor of Γ.*

We'll prove more general theorems later, in Chapters II and IV **[GAINS chapter]**, so I omit the proof here.

The following theorem is one of the main ways in which minors are used. It characterizes the graphs that embed in a surface in terms of forbidden minors. Each successive part is much harder to prove. The general name for these results is "Kuratowski-type theorems".

Chap 1. B (Basic Structures)

Aug 29
1 of 2

Fig A

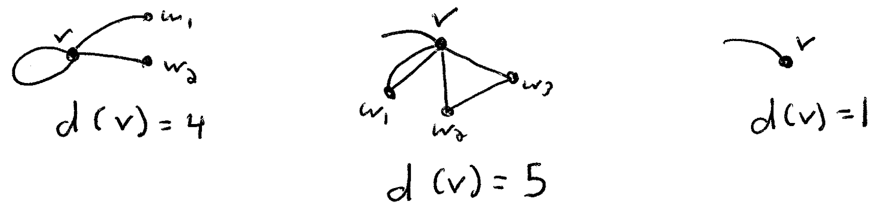

Fig B

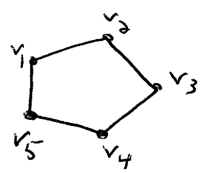

Fig C

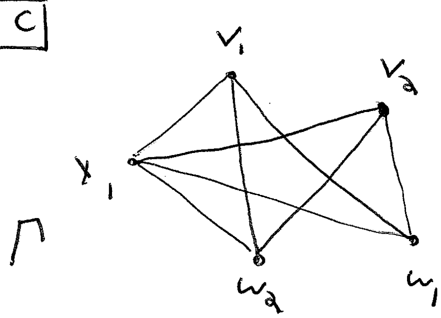

a Forest of Γ

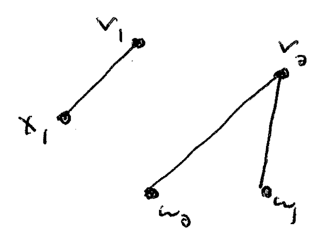

a maximal forest of Γ

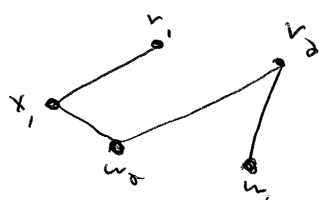

Fig D Contraction

For e a link

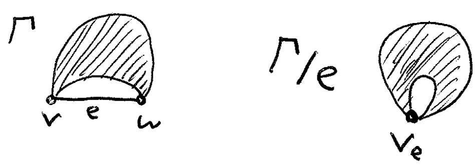

For e a loop

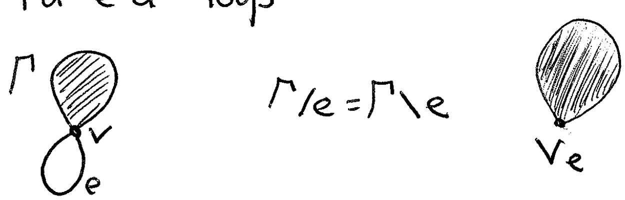

Theorem C.2 (Kuratowski-type theorems)**.** [[LABEL T:0903kuratowski-type]] *Let* Γ *be a graph.*

(1) [Kuratowski (mainly) and Wagner] Γ *is planar iff* Γ *does not contain either* K_5 *or* $K_{3,3}$ *as minors.*
(2) [Archdeacon, Glover, and Huneke] Γ *is projective planar iff* Γ *does not contain as a minor any of a list of* 35 *graphs.*
(3) [Robertson and Seymour] Γ *embeds in a surface* S *iff* Γ *does not contain as a minor any of a finite list of graphs, which depends on* S*.*

D. Closure and Connected Partitions

[[LABEL 1.closure]]

One of the chief ideas in our treatment of graphs is the closure of an edge set, which corresponds to objects in graph invariants and graphical geometry.

D.1. **Partitions.** [[LABEL 1.partitions]]

A *partition* of a set V is a class π of subsets of V, called the *blocks* or (sometimes) *parts* of π, such that

(1) the union $\bigcup_{B\in\pi} B$ equals V,
(2) any two blocks are disjoint, and
(3) each block $B \in \pi$ is nonvoid.

(The last property means that, if we want to allow empty blocks, we do not have a partition.) The *size* of π is the number of blocks, k. The only partition of the empty set, $X = \varnothing$, is $\pi = \{\}$. In all other cases $k \geq 1$. The partition $\hat{1} = \hat{1}_V := \{X\}$ with one block is called the *trivial partition*; the partition in which every block is a singleton, $\hat{0} = \hat{0}_V := \big\{\{x\} : x \in X\big\}$, is the *total partition*. Every set except $\varnothing$ has a trivial partition; every set has a total partition.

It is unfortunate that the term 'block' conflicts with the graph-theoretic usage, but it's too late to change so we'll have to live with it.

We define

$$\Pi_V := \{\text{all partitions of } V\} \quad \text{and} \quad \Pi_n := \Pi_{[n]}.$$

Partitions of V are partially ordered by refinement, namely, $\pi \leq \tau$ (π *refines* τ) if each block of π is contained in a block of τ. In Π_V, the unique minimum element is $\hat{0}$ and the unique maximum element is $\hat{1}$ (if $V \neq \varnothing$).

Now let V be the vertex set of a graph Γ. We say $\pi \in \Pi_V$ is *connected* (in Γ) if each block $B \in \pi$ induces a connected subgraph. Let

$$\Pi(\Gamma) := \text{ the set of connected partitions of } V.$$

We define the *partition* of V *induced by an edge set* S as $\boldsymbol{\pi}(S) := \boldsymbol{\pi}(V, S) :=$ the partition of V into the subsets which are the vertex sets of the connected components of S, that is, of (V, S). That is, the blocks of the partition are the equivalence classes of the connection relation of V in the subgraph (V, S). Closed sets are intimately related to connected partitions.

D.2. **Abstract closure.** [[LABEL 1.abstractclosure]]

We now remind ourselves of the definition of an abstract closure operator on E.

Definition D.1. [[LABEL D:0903closure]] An *abstract closure operator* is a function $\mathcal{P}(E) \to \mathcal{P}(E)$, written $S \mapsto \overline{S}$, such that the following axioms hold for subsets S and T of E:

(1) Increase: $S \subseteq \overline{S}$. [[LABEL R:0903clos1]]

(2) Isotonicity: $S \subseteq T \implies \overline{S} \subseteq \overline{T}$. [[LABEL R:0903clos2]]

(3) Idempotency: $\overline{\overline{S}} = \overline{S}$. [[LABEL R:0903clos3]]

A set $S \subseteq E$ is called *closed* if $S = \overline{S}$.

The closed sets when ordered by inclusion form a partially ordered set (*poset*) which is closed under set intersection and includes the universe E. (Those two properties characterize classes of abstractly closed sets.) This poset is a lattice which is described precisely by a standard result.

Proposition D.1. [[LABEL P:0908meetjoin]] *For any abstract closure operator on E, the class of closed subsets forms a lattice with meet and join defined as follows: For S, T closed subsets of E, $S \wedge T = S \cap T$ and $S \vee T = \operatorname{clos}(S \cup T)$.*

D.3. **Graph closure.** [[LABEL 1.graphclosure]]

There is a natural operation of closure on the edges of a graph.

Definition D.2. [[LABEL D:0903graphclosure]] In an ordinary graph Γ, for $S \subseteq E$, the *closure* of S is

$$\operatorname{clos} S := S \cup \{e : \text{the endpoints of } e \text{ are joined by a path in } S\}.$$

Equivalently, there is a circle $C \subseteq S \cup e$ such that $e \in C$. We say $S \subseteq E$ is *closed* if $\operatorname{clos}(S) = S$.

Notice that it is redundant to list S in the definition of $\operatorname{clos}(S)$, since the endpoints of an edge of S are always connected in S. The restatement in terms of circles, though easy to prove, is more fundamental than might appear at first sight, as we shall see in Chapters II and IV [**GAINS chapter**].

One should keep in mind that Proposition D.1 holds for the closure operator in a graph.

The graph closure operator obeys, besides the abstract closure properties (1–3), a very important fourth property, the *exchange property*:

(4) Let S be a closed subset of E. If $e, f \notin S$ and $e \in \overline{f \cup S}$, then $f \in \overline{e \cup S}$.

(The proof is a nice exercise.) Those familiar with matroids will know that the exchange property is what makes the closure operator on edges a matroid closure.

Recall that $S{:}B$ is the set of edges in S with all of their endpoints in the vertex set B.

Theorem D.2. [[LABEL T:0903indclosure]] *For $S \subseteq E$,* $\operatorname{clos} S = \bigcup_{B \in \boldsymbol{\pi}(S)} E{:}B$.

Proof. An edge e is in $E{:}B$ for some $B \in \boldsymbol{\pi}(S) \iff e$ has both endpoints within one block B of $\boldsymbol{\pi}(S) \iff$ the endpoints of e are connected by $S \iff e \in \operatorname{clos} S$. This establishes the partition formula for closure. □

Theorem D.3. [[LABEL T:0903closedptns]] *The poset of closed sets ordered by inclusion is isomorphic to the poset $\Pi(\Gamma)$ of connected partitions of Γ ordered by refinement.*

Proof. Theorem D.2 presents a bijection between closed edge sets and connected partitions of Γ. It is clear from the definitions of partition ordering and connected partitions that the bijection is order preserving. □

Sept 8a:
Jackie
Kaminski

D.4. **Edge sets induced by partitions.** [[LABEL 1.partitionsedges]]

Recall that $\Pi(\Gamma)$ is the set of all connected partitions of V, i.e., for $\pi \in \Pi(\Gamma)$ and $B \in \pi$, any two vertices in B are connected in $\Gamma{:}B$. We notice immediately that for any $S \subseteq E$, the partition $\boldsymbol{\pi}(S) \in \Pi(\Gamma)$. This observation allows us to define a function $\boldsymbol{\pi} : \mathcal{P}(E) \to \Pi(\Gamma)$ by $S \mapsto \boldsymbol{\pi}(S)$. We now present a definition followed by a lemma about $\boldsymbol{\pi}$.

Definition D.3. [[LABEL D:0908Epi]] For any partition π of V, $E(\pi) := E{:}\pi := \bigcup_{B\in\pi} E{:}B$.

We apply this definition mainly to connected partitions, because when π is not a connected partition some of the terms in the union are not connected and some may be empty. The following lemma is only for connected partitions.

Lemma D.4. [[LABEL L:0908clos]] *For each $\pi \in \Pi(\Gamma)$, $\boldsymbol{\pi}(E{:}\pi) = \pi$. Furthermore, $E{:}\boldsymbol{\pi}(S) = \operatorname{clos}(S)$.*

Thus, from $\boldsymbol{\pi}(S)$ we can't in general recover S, but we can always recover $\operatorname{clos}(S)$.

Corollary D.5. [[LABEL C:0908piofclos]] $\boldsymbol{\pi}(\operatorname{clos}(S)) = \boldsymbol{\pi}(S)$.

Proof. Let $\boldsymbol{\pi}(S) = \{B_1, \ldots, B_k\}$. From Theorem D.2, $\operatorname{clos}(S) = \bigcup_{i=1}^{k} E{:}B_i$. Each part in $\boldsymbol{\pi}(\operatorname{clos}(S))$ will be the vertex set of a maximal connected component of $\bigcup_{i=1}^{k} E{:}B_i$. These are precisely the sets B_i. $\square$

Corollary D.6. [[LABEL C:0908piEpi]] *For any $S \subseteq E$, $\boldsymbol{\pi}(E{:}\boldsymbol{\pi}(S)) = \boldsymbol{\pi}(S)$.*

Proof. By definition $E{:}\boldsymbol{\pi}(S) = \bigcup_{B\in\boldsymbol{\pi}(S)} E{:}B$, and similarly $\boldsymbol{\pi}(E{:}\boldsymbol{\pi}(S)) = \boldsymbol{\pi}(\bigcup_{B\in\boldsymbol{\pi}(S)} E{:}B)$, which is precisely $\boldsymbol{\pi}(S)$ since each $E{:}B$ is connected. $\square$

We supplement Theorem D.2 with two further characterizations of closed sets, which follow immediately from that theorem and Corollary D.5.

Corollary D.7. [[LABEL C:0908indclosure]] *An edge set S is closed $\iff$ it equals $E{:}\pi$ for some $\pi \in \Pi(\Gamma)$ $\iff$ it equals $E{:}\pi$ for some $\pi \in \Pi_V$.* $\square$

D.5. **Lattices.** [[LABEL 1.lattices]]

Whenever $S \subseteq S' \subseteq E$, then $\boldsymbol{\pi}(S)$ is a refinement of $\boldsymbol{\pi}(S')$, that is to say, each of the parts of $\boldsymbol{\pi}(S)$ is contained in a part of $\boldsymbol{\pi}(S')$. Readers familiar with partitions of a set V will think of the last statement as $\boldsymbol{\pi}(S) \le \boldsymbol{\pi}(S')$; this defines a partial ordering of partitions called the *refinement ordering*. It is well known that the set $\Pi(V)$ of all partitions of V with the refinement ordering forms a lattice. It is left to the reader to check that the set of connected partitions also forms a lattice in which the meet operation is the same as in $\Pi(V)$ and the join operation is $\tau \vee \tau' = \bigwedge\{\pi \in \Pi(\Gamma) : \pi \ge \tau, \tau'\}$.

When τ, τ' are two partitions of V such that $\tau \le \tau'$ (ordered by refinement), then $E{:}\tau \subseteq E{:}\tau'$. Here we remind the reader that $\mathcal{P}(E)$, ordered by set inclusion, is also a lattice with the intersection and union operations. These observations and the following definition lead us to our next theorem.

Definition D.4. [[LABEL D:0908lattice]] $\operatorname{Lat}(\Gamma)$ is the class whose members are the closed edge sets of Γ, ordered by containment.

Theorem D.8. [[LABEL T:0908LatIso]] $\Pi(\Gamma) \cong \operatorname{Lat}(\Gamma)$. *Specifically, $\boldsymbol{\pi} : \operatorname{Lat}(\Gamma) \to \Pi(\Gamma)$ is an order isomorphism.*

Proof. We already noted that $S \subseteq S' \subseteq E \implies \boldsymbol{\pi}(S) \leq \boldsymbol{\pi}(S')$ and that for $\tau, \tau' \in \Pi(\Gamma)$, $\tau \leq \tau' \implies E{:}\tau \subseteq E{:}\tau'$. So all that's left to show is that $\boldsymbol{\pi}$ is a bijection between the connected (vertex) partitions of Γ and the closed (edge) subsets of Γ.

To see that $\boldsymbol{\pi}$ is injective, let S, S' be closed subsets of E, and assume $\boldsymbol{\pi}(S) = \boldsymbol{\pi}(S')$. By Theorem D.2, $S = E{:}\boldsymbol{\pi}(S) = E{:}\boldsymbol{\pi}(S') = S'$. To see that $\boldsymbol{\pi}$ is surjective, we notice that for τ a (connected) partition of Γ, $E{:}\tau$ is closed by Theorem D.2. This completes our proof. $\square$

Sept 3b:
Yash Lodha

E. Incidence and Adjacency Matrices

[[LABEL 1.matrices]]

Incidence and adjacency matrices let graph theory benefit from the use of matrix theory.

E.1. **Incidence matrices.** [[LABEL 1.incidmg]]

An incidence matrix describes the incidence relation between vertices and edges. A graph has two kinds of incidence matrix.

Definition E.1. [[LABEL D:0903orincidencematrix]] An *oriented incidence matrix* of a graph is a $V \times E$ matrix $\mathrm{H}(\Gamma)$ (pronounced 'Eta') which has, for each edge e, in the column labelled by e, an entry $\eta_{ij} = +1$ at the row of one endpoint and an entry $\eta_{ij} = -1$ at the other endpoint, with 0s elsewhere. If e is a loop incident with v_i, the entry $\eta_{ij} = 0$ (yes, the whole column is 0).

There are many different oriented incidence matrices of a graph, in fact, $2^{m'}$ where m' is the number of links (and half edges, if allowed).

Definition E.2. [[LABEL D:0903unorincidencematrix]] The *unoriented incidence matrix* $B(\Gamma)$ is a $V \times E$ matrix. The entry $b_{ij} = 0$ if the edge e_j is not incident with the vertex v_i, and $b_{ij} = 1$ if e_j is incident with v_i. If e is a loop incident with v_i, the entry $\eta_{ij} = 2$.

The incidence matrix most commonly seen in graph theory is the unoriented one. However, its proper place is with signed graphs, as we shall see in Section II.G. For our line of interest, the right graphical incidence matrix is the oriented one. Most of the reason is the relationship between linear dependence of columns in the matrix and graph structure to be developed in Lemmas G.3 and G.4. The fact, which the reader will have noticed, that the oriented incidence matrix is uniquely defined only up to negating columns, does not affect column linear dependence.

The oriented incidence matrix has rank $n - c(\Gamma)$, hence nullity $|E| - n + c(\Gamma)$, the cyclomatic number; we shall prove this as a special case of a signed-graph theorem in Section II.G. [**Make reference more specific when the theorem gets into the notes. NB: MISSING S.G. THEOREM!**]

E.2. **Adjacency and valency (degree) matrices.** [[LABEL 1.adjmatrix]] [**FORMERLY 1.adjmg**]

Let $V(\Gamma) = \{v_1, v_2, ..., v_n\}$.

Definition E.3. [[LABEL D:0903adjacencymatrix]] The *adjacency matrix* $A(\Gamma)$ is the $n \times n$ matrix (a_{ij}) defined by the rules:

- For a simple graph, the entry $a_{ij} = 1$ if v_i and v_j are adjacent and 0 if they are not. Thus $a_{ii} = 0$.

- For an arbitrary ordinary graph, a_{ij} is the number of edges that join v_i with v_j, but with a loop counting twice, once for each end.

The *degree matrix* or *valency matrix* $D(\Gamma)$ is a $V \times V$ diagonal matrix where the entry d_{ii} is the degree of the vertex v_i, while the off-diagonal entries are 0. Remember that a loop counts 2 in the degree, while a half edge counts 1. The next theorem is not quite correct if there are half edges.

Theorem E.1. [[LABEL T:0903incidence-adjacency]] *The adjacency, degree, and incidence matrices are related by the formula* $A(\Gamma) = D(\Gamma) - \mathrm{H}(\Gamma)\mathrm{H}(\Gamma)^T = B(\Gamma)B(\Gamma)^T - D(\Gamma)$.

Proof. To prove that $\mathrm{HH}^{\mathrm{T}} = D - A$ we check the cases $i \neq j$ and $i = j$ separately when multipying the ith row of H with the jth column of H^{T}. One should pay special attention to the diagonal when there are loops.

The proof for B is similar. □

We'll have a more detailed proof when we get to the signed-graph generalization, Theorem G.10 in Section G.3.

Sept 5: Peter Cohen and Yash Lodha

F. Orientation

[[LABEL 1.orientation]]

F.1. **Orienting a graph.** [[LABEL 1.orienting]]

If we have a graph Γ, we *orient* it by giving every edge a direction. We write $\vec{\Gamma}$ for an orientation of Γ.

The notation for an oriented edge can be a bit tricky. We could write v_1v_2, or $\overrightarrow{v_1v_2}$,but this presents a problem with parallel edges. We will call an oriented edge $e_i{:}\overrightarrow{v_1v_2}$.

The direction of an edge is drawn as an arrow on the edge. Formal notation is more complicated and is best explained in terms of the incidence matrix; see Section F.2.

I distinguish between an "oriented edge" and a "directed edge", although in many ways they are the same. An orientation is not inherent in the edge but is imposed on it for some purpose. In a directed edge the direction is inherent. Especially, a directed edge can only be traversed in the direction of the edge, but an oriented edge can be traversed in either direction, with or against its orientation. An *oriented graph* is a graph whose edges happen to be oriented in some way that may vary; a *directed graph* or *digraph* is a graph where each edge has a fixed direction.

A key concept is coherence of an oriented walk, especially a circle. A walk in an oriented graph is *coherently oriented* if every two consecutive edges are coherent. Two consecutive edges, incident at a vertex v, are *coherent* or *consistent* if their directions agree, i.e., one of them is directed into v and the other is directed out of v. If the walk is closed, its last and first edges are considered consecutive at the initial vertex; we say it is a *coherently oriented closed walk* if its orientation is consistent at that vertex as well as at all others.

Definition F.1. [[LABEL Df:0905cycle]] A *cycle* in an oriented graph is a circle that is oriented so each vertex is consistent. An orientation of Γ is *acyclic* if it has no cycles, *cyclic* if it has at least one cycle, and *totally cyclic* if every edge belongs to a cycle.

Directing a circle means giving the circle as a whole a direction. This is a completely separate property of the circle from directions on the edges.

Suppose we linearly order the vertex set V, e.g., by numbering the vertices from 1 to n. We get an orientation of Γ by directing each edge $e{:}vw$ from the lower to the higher endpoint. (A loop is oriented either way.) This orientation is acyclic if Γ has no loops. It is obviously uniquely determined by the linear ordering of V; on the other hand, different linear orderings may yield the same acyclic orientation.

Theorem F.1. [[LABEL T:0903acyclic]] *Every acyclic orientation arises from a linear ordering of the vertices.*

Hence there is an equivalence in that statement above.

In an oriented graph there are two special kinds of vertices. A *sink* is a vertex with only entering edges. A *source* is a vertex with only departing edges. The extreme case is an isolated vertex, which is both a source and a sink.

Lemma F.2. [[LABEL L:0903sourcesink]] *Every acyclic orientation has a source and a sink.*

Proof. We start on an edge and walk along a path following edge directions. If we repeat a vertex we form a cycle, which contradicts the assumption that our graph is acyclic. If we never repeat a vertex in our path, then since $|V|$ is finite we must end our path at a vertex that only has entering edges. This proves the existence of a sink.

To prove the existence of a source, reverse the orientations of all edges. A sink in the reversed graph is a source in the original orientation. Alternatively, apply the previous argument in reverse. $\square$

Proof of Theorem F.1. We perform induction on $|V|$. If $\vec{\Gamma}$ is acyclic, then it must have a sink s. Then by our inductive hypothesis $\vec{\Gamma}\backslash s$ is acyclic and has an ordering $v_1 < v_2 < \cdots < v_{n-1}$ The ordering for $\vec{\Gamma}$ is $v_1 < v_2 < \cdots < v_{n-1} < s$. $\square$

A total ordering of V is not necessary for constructing an acyclic orientation. [**I will discuss an example once I learn graphics in Tex.**] In fact a partial ordering can suffice for providing us with a corresponding acyclic orientation of the graph.

Theorem F.3. *For each $\vec{\Gamma}$, there exists a smallest partial ordering of V that gives the orientation $\vec{\Gamma}$. The linear orderings that give $\vec{\Gamma}$ are precisely the linear extensions of that smallest partial ordering.*

An important example is the complete graph.

Example F.1. [[LABEL X:0903kn]] *Acyclic orientations of K_n.* Every partial ordering of V that gives K_n as its comparability graph is a chain (a total ordering). There are $n!$ of these, one for each permutation of V.

Corollary F.4. *The acyclic orientations of K_n correspond bijectively to the permutations of V in a natural way.*

Proof. The correspondence is that a total ordering of V implies an orientation of each edge from lower to higher.

Conversely, suppose K_n is acyclically oriented. Then there is a corresponding partial ordering of V, but it is a total ordering because every pair of vertices is comparable. $\square$

Example F.2. [[LABEL X:0903compar]] A *comparability graph* is the graph of all comparability relations in a poset. This means that the vertex set is the set of elements of the poset, and we connect elements u, v with an edge if u, v are comparable.

There is an extensive literature on comparability graphs. A good, readable source is Golumbic's [PG].

Thus, we can think of an acyclic orientation of a graph as a generalization of a permutation. This point of view gives interesting insights into the regions of the hyperplane arrangement associated with a graph. See Section G.3.

An orientation that is not acyclic is called *cyclic*. But we can also have a *totally cyclic orientation*, where every edge is in a cycle. (Totally cyclic orientations are dual to acyclic orientations; but to explain this properly we want either planar graph duality or the theory of oriented matroids, which are outside our scope.)

Proposition F.5. *Γ has an acyclic orientation iff it has no loops. Γ has a totally cyclic orientation iff it has no isthmi.*

Partial proof. We prove the first part. A loop is necessarily a cycle. Conversely, if there are no loops, we get an acyclic orientation from any linear ordering of V. □

F.2. **Incidence matrix.** [[LABEL 1.omatrix]]

An oriented graph, in contrast to an unoriented graph, has a unique incidence matrix, because the orientation of an edge tells us how to determine the signs in its column of the matrix..

Definition F.2. [[LABEL D:0903orincidence]] An *incidence matrix of an orientation* of a graph has, for each edge e, in the column denoted by e, an entry of $+1$ at the row of its head vertex and an entry of -1 at the tail.

Thus, an incidence matrix $\mathrm{H}(\vec{\Gamma})$ of an orientation of Γ is one of the oriented incidence matrices of Γ, and an oriented incidence matrix of Γ is the incidence matrix of some orientation of Γ.

Sept 8b: Jackie Kaminski

G. Equations and Inequalities from Edges

G.1. **Arrangements of hyperplanes.**

Now we think of the edge set of Γ as $\{v_1, \dots v_n\}$, and we begin by considering only ordinary graphs Γ. We define

$$h_{ij} := \{x \in \mathbb{R}^n \mid x_i = x_i\}.$$

When $i \neq j$, h_{ij} is clearly a hyperplane (a codimension-1 linear subspace) of $\mathbb{R}^n$. We will refer to h_{ii}, which is all of $\mathbb{R}^n$ since it corresponds to the equation $x_i = x_i$, as the "degenerate hyperplane", because it will be convenient later to allow it as one of a family of hyperplanes.

Definition G.1. [[LABEL D0908hyp]] An *arrangement of hyperplanes* is a finite set (or multiset) of hyperplanes in $\mathbb{R}^n$.

Definition G.2. [[LABEL D0908HypGamma]] $\mathcal{H}[\Gamma]$, the hyperplane arrangement induced in $\mathbb{R}^n$ by Γ, is the multiset of hyperplanes $\{h_{ij} \mid e{:}v_iv_j \in E\}$. (Recall that $n = |V|$.)

We notice that each loop in Γ corresponds to the degenerate hyperplane. And furthermore we note the obvious correspondence between the multiset $\mathcal{H}[\Gamma]$ and the edges of Γ. In fact there are many equivalent points of view we can take, as we notice the following (bijective) correspondences, that we describe on elements, but they extend naturally to their respective sets.

- The edge $e{:}v_iv_j \longleftrightarrow$ the equation $x_i = x_j$.
- $x_i = x_j \longleftrightarrow$ the hyperplane h_{ij} in $\mathbb{R}^n$, by geometry.
- $e{:}v_iv_j \longleftrightarrow$ column c_e in $\mathrm{H}(\Gamma)$. (Recall that $\mathrm{H}(\Gamma)$ is the incidence matrix of Γ.) This correspondence is immediate from the definition of $\mathrm{H}(\Gamma)$.
- Column c_e in $\mathrm{H}(\Gamma) \longleftrightarrow$ the equation $x_i = x_j$, by vector space duality.

Before looking into further correspondences, we set up a bit more terminology.

Definition G.3. [Region of $\mathcal{A}$] [[LABEL D:0908region]] For an arrangement $\mathcal{A}$ of hyperplanes in $\mathbb{R}^n$, a *region* of $\mathcal{A}$ is a connected component of $\mathbb{R}^n \setminus \bigcup_{A\in\mathcal{A}} A$. Thus, if there is a degenerate hyperplane in $\mathcal{A}$, then $\mathcal{A}$ has no regions.

Now we define

$$\mathcal{L}(\mathcal{A}) := \{\textstyle\bigcap \mathcal{S} \mid \mathcal{S} \subseteq \mathcal{A}\},$$

which we will later see is a lattice, and we will later have a theorem saying $\mathcal{L}(\mathcal{H}[\Gamma]) \cong \mathrm{Lat}(\Gamma) \cong \Pi(\Gamma)$, where the lattice isomorphisms are all natural. This will eventually allow us to switch between the perspectives of geometry, lattices, and graphs. Furthermore we can think of any of the correspondences above as correspondences between subsets instead of between individual elements.

Finally, we close with two lemmas that we will revisit later.

Lemma G.1. [[LABEL L:0908closspan]] *For $e \in \mathrm{clos}(S)$, $c_e \in \langle c_f : f \in S\rangle$.*

Lemma G.2. [[LABEL L:0908hypintersection]] *For $S \subseteq E$, $\bigcap \mathcal{H}[S] = \bigcap \mathcal{H}[\mathrm{clos}(S)]$.*

This second lemma is the vector dual of the first.

Sept 10:
Nate Reff

G.2. **Graphic hyperplane arrangements and the intersection lattice.** [[LABEL 1.graphichyp]]

Lemma G.3. [[LABEL L:0910lemma1a]] $e \in \mathrm{clos}(S) \implies c_e \in \langle c_f : f \in S\rangle$.

Proof. Let's draw a nice picture to see how things work.

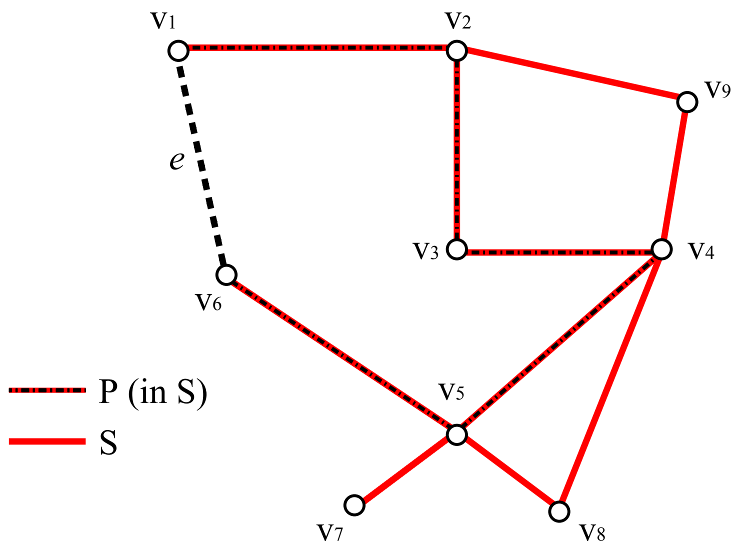

The red lines denote edges of $S \subseteq E$ in a graph $\Gamma = (V, E)$. If $e \in (\text{clos}(S) \backslash S)$ as in the picture, then there exists a path $P \subseteq S$ such that there is a circle. We will show that $\langle c_f : f \in S \rangle$. Because $e \in (\text{clos}(S) \backslash S)$ and thus $e \in \text{clos}(S)$, there exists a path $P = v_1 v_2 \cdots v_l$ connecting the two endpoints of e. Now let's label the vertex set in such a way that we start at v_1, one endpoint of e and traverse P until we reach the other endpoint of e, v_l (in our particular example, v_6). Then arbitrarily assign the remaining vertices. If we do this then the columns of $P \cup e$ are the following:

$$\begin{bmatrix} 1 & 0 & 0 & \dots & 0 & 1 \\ -1 & 1 & 0 & \dots & 0 & 0 \\ 0 & 1 & 1 & \dots & \vdots & \vdots \\ \vdots & \vdots & \ddots & \ddots & \vdots & \vdots \\ 0 & 0 & 0 & \dots & -1 & -1 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots & \vdots \end{bmatrix},$$

where the columns of the matrix correspond to $\{e_1, e_2, \dots, e_l, e\}$ and the rows correspond to $v_1, v_2, \dots, v_l, v_{l+1}, \dots$.

Then $c_e = c_{e_2} + c_{e_3} + \cdots + c_{e_l}$, so c_e is spanned by the column vectors of edges in S. $\square$

Lemma G.4. [[LABEL L:0910lemma1b]] $c_e \in \langle c_f : f \in S \rangle \implies e \in \text{clos}(S)$.

Proof. Suppose $e \notin \text{clos}(S)$. Then the endpoints of e belong to different components of (V, S), simply because there is no path in S connecting the endpoints.

Now, for a working example, let's consider the following graph Γ:

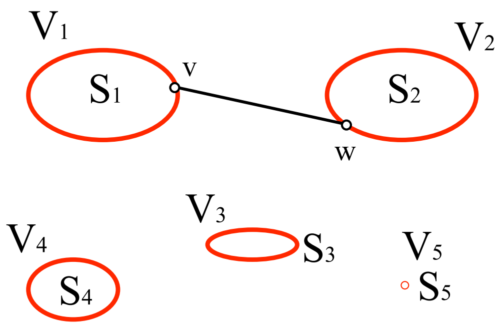

The incidence matrix $\mathrm{H}(\Gamma)$ looks like this, where O is a zero matrix, and $\mathbf{0}$ is a column vector of zeros:

$$\begin{array}{cc}
 & \begin{array}{ccccccc} (S_1) & (S_2) & (S_3) & (S_4) & (S_5) & (e) & (S^c \setminus e) \end{array} \\
\begin{array}{c} (V_1) \\ \\ (V_2) \\ \\ (V_3) \\ (V_4) \\ (V_5) \end{array} &
\left[\begin{array}{ccccccc}
\mathrm{H}(S_1{:}V_1) & O & O & O & O & \begin{pmatrix} 0 \\ \vdots \\ 1 \\ \vdots \\ 0 \end{pmatrix} & * \\
O & \mathrm{H}(S_2{:}V_2) & O & O & O & \begin{pmatrix} 0 \\ \vdots \\ -1 \\ \vdots \\ 0 \end{pmatrix} & * \\
O & O & \mathrm{H}(S_3{:}V_3) & O & O & \mathbf{0} & * \\
O & O & O & \mathrm{H}(S_4{:}V_4) & O & \mathbf{0} & * \\
O & O & O & O & \mathrm{H}(S_5{:}V_5) & \mathbf{0} & *
\end{array}\right]
\end{array},$$

where the columns of the matrix are indexed by the edges of S_1, S_2, S_3, S_4, S_5, e, and $S^c \setminus e$; the rows of the matrix are indexed by the sets V_1, V_2, V_3, V_4, V_5; and the column of $*$'s stands for $\mathrm{H}(S^c \setminus e)$. The nonzero entries in column c_e are, in the rows of V_1, in row v, and in the rows of V_2, in row w.

Now we return to the general proof. Suppose $e{:}vw$ has $v \in V_1$ and $w \in V_2$, and that there is a sum $\sum_{e_i \in S} \alpha_i c_{e_i} = c_e$. The edges in a component S_j of S which doesn't contain an endpoint of e have to add up to zero in the sum, so they can be ignored. Thus, looking only at the rows of V_1,

$$\sum_{e_i \in S_1} \alpha_i c_i + \sum_{e_i \in S_2} \alpha_i c_i = c_e,$$

where for brevity we write c_i for the column of e_i.

Looking only at the rows of V_1, we note two facts. First, let c_i' and c_e' denote just the V_1 rows of c_i and c_e. Then

$$\text{[[LABEL E : 0910S1]]} \sum_{e_i \in S_1} \alpha_i c_i' = \begin{pmatrix} 0 \\ \vdots \\ 1 \\ \vdots \\ 0 \end{pmatrix}. \tag{G.1}$$

Second, all columns in S_1, restricted to the rows of V_1, have entries that sum to zero, so if we add up all the rows in Equation (G.1), the left-hand side of the equation sums to 0 and the right-hand side sums to 1. This is a contradiction! Hence there does not exist a linear combination which is equal to e. Therefore we can say that $e \in \mathrm{clos}(S)$. ☐

Lemma G.5. [[LABEL L:0910lemmaSubLemma]] *For a hyperplane $H_e \in \mathcal{H}[\Gamma]$, $\bigcap \mathcal{H}[S] \subseteq H_e \iff e \in \mathrm{clos}(S)$.*

Lemma G.6. [[LABEL L:0910lemma2]] $\bigcap \mathcal{H}[S] = \bigcap \mathcal{H}[\mathrm{clos}(S)]$.

Proof. Use Lemma G.5, and dualize Lemmas G.3 and G.4. ☐

We define a subset $S \subseteq E$ to be *dependent* if there exists an $e \in S$ such that $e \in \mathrm{clos}(S \setminus e)$.

Proposition G.7. *S is independent $\iff$ S is a forest.*

Proof. This is immediate from the definition of closure. ☐

Theorem G.8. [[LABEL T:0910thm1]] *Let $S \subseteq E$. S is independent in Γ (so S is a forest) $\iff$ the columns of S in $\mathrm{H}(\Gamma)$ are linearly independent.*

Proof. Immediate corollary of Lemmas G.3 and G.4. ☐

We define a *linearly closed set of columns* to be the intersection of $\{c_e : e \in E\}$ with a subspace of F^n.

Corollary G.9. [[LABEL C:0910cor1]] *The closed edge sets $\longleftrightarrow$ the linearly closed sets of columns of $\mathrm{H}(\Gamma)$.*

Theorem G.10. [[LABEL T:0910thm2]] *There are natural isomorphisms $\Pi(\Gamma) \cong \mathrm{Lat}(\Gamma) \cong$ {linearly closed sets of columns} $\cong \mathcal{L}(\mathcal{H}[\Gamma])$.*

Proof. This follows from the relationships we've already seen among the various lattices and closures. ☐

Corollary G.11. [[LABEL C:0910geomlattice]] *$\mathrm{Lat}\,\Gamma$ and $\Pi(\Gamma)$ are geometric lattices.*

Proof. The intersection lattice is dual to the lattice of vector spaces spanned by columns of the incidence matrix, which is known to be a geometric lattice. (See [Oxley], for instance.) ☐

G.3. **Regions and Orientations.** [[LABEL 1.regions]]

An orientation of Γ defines a positive side of each hyperplane $h_{ij} \in \mathcal{H}[\Gamma]$, called the *positive half-space* of the hyperplane. If we orient $e{:}v_iv_j$ from v_i to v_j, the positive half-space is the set $\{x \in \mathbb{R}^n : x_i < x_j\}$. For each orientation, therefore, there is a family of positive half-spaces.

Lemma G.12. [[LABEL L:0910lemma3]] *A cyclic orientation of* Γ *gives an empty intersection of positive half-spaces.*

Proof. Suppose that a graph Γ has a cycle on edges $e_1{:}v_1v_2,\ e_2{:}v_2v_3,\ \ldots,\ e_l{:}v_lv_{l+1}$, where $v_{l+1} = v_1$. We may assume e_j is oriented from v_j to v_{j+1}. Then the corresponding positive half-space for each e_j is the set $\{x \in \mathbb{R}^n : x_j < x_{j+1}\}$. Therefore the intersection of all the positive half spaces is $\{x \in \mathbb{R}^n : x_1 < x_2 < \ldots < x_l < x_1\} = \varnothing$. $\square$

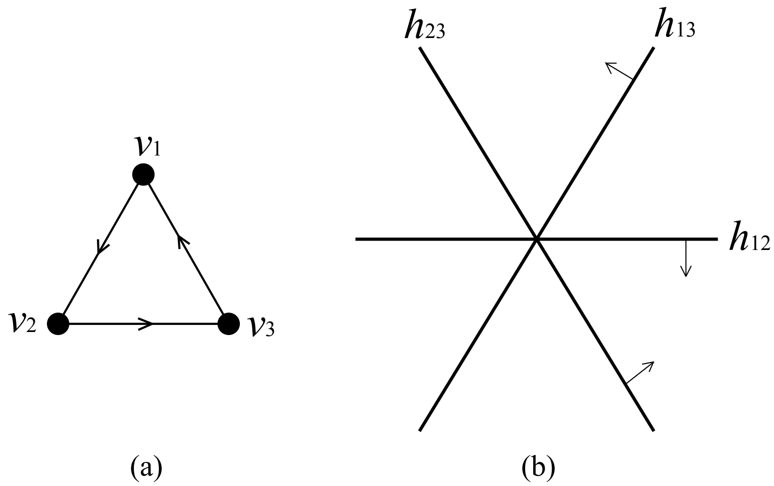

FIGURE G.1

An example illustrates the proof. Suppose that the graph $\Gamma = K_3$ is oriented cyclically, as in (a) of Figure G.1. The corresponding orientation on each hyperplane is shown in (b). By the definition of the positive half space, the corresponding intersection of all the positive half spaces is $\{x \in \mathbb{R}^3 : x_3 > x_2 > x_1 > x_3\} = \varnothing$.

Thus, any region is the intersection of positive half-spaces in a unique orientation of Γ, which is necessarily acyclic.

Theorem G.13. [[LABEL T:0910thm3]] *The intersection of positive half-spaces of an orientation of* Γ *is empty if the orientation is cyclic, but it is a region of* $\mathcal{H}[\Gamma]$ *if the orientation is acyclic.*

Proof. In the cyclic case we just use Lemma G.12. In the acyclic case the orientation corresponds to a linear ordering of vertices, say $v_1 < v_2 < \ldots < v_n$. Then $(1, 2, \ldots, n)$ will be in the intersection of positive half-spaces. Therefore the intersection is nonempty, and in fact a region. $\square$

Sept 12a:
Simon Joyce

H. CHROMATIC FUNCTIONS

[[LABEL 1.chromatic]]

Coloring a graph has inspired all kinds of remarkable developments. We'll concentrate on counting colorations and how it leads to algebraic properties that apply more widely than to coloring, but first we have to know what it means to color a graph.

To explain the section title: I call a *chromatic function* (or a *dichromatic function*, a term that will be explained later in this section) any function that depends on coloring or that satisfies the main algebraic laws that apply to the chromatic polynomial (another term that will be explained in this section).

H.1. **Coloring.** [[LABEL 1.coloring]]

Given a graph Γ, a *coloration* (or *coloring*) of Γ in k colors is a function $\gamma : V \to \Lambda$, a set of k colors. It doesn't matter for the definition exactly which k-element set Λ is, but often enough it is best to choose it to be the set $[k] := \{1, 2, \ldots, k\}$ of the first few positive integers.

An edge $e{:}vw$ is *proper* if $\gamma(v) \neq \gamma(w)$ and a coloration is *proper* if every edge is proper. For example, a graph with a loop can't ever be properly colored. Any coloration γ of a graph Γ has a set of proper edges and a set of improper edges. We will call the set of improper edges $I(\gamma)$.

H.2. **Chromatic number.** [[LABEL 1.chromaticnumber]]

We say a graph is *k-colorable* if there exists a proper coloration in k colors.

Definition H.1. [[LABEL D:0912 chrom num]] For a graph Γ we define its *chromatic number* to be

$$\chi(\Gamma) = \min\{k : \Gamma \text{ is } k\text{-colorable}\}.$$

For instance, $\chi(K_n) = n$ and $\chi(\bar{K}_n) = 1$ for $n \geq 1$. For a forest F with at least one edge, $\chi(F) = 2$. In fact, for any bipartite graph that has at least one edge, $\chi(\Gamma) = 2$. At the opposite extreme, $\chi(\Gamma) = \infty$ if, and only if, Γ has a loop.

H.3. **The chromatic polynomial.** [[LABEL 1.chromaticpoly]]

We now turn our attention to counting functions related to coloring and to the structural properties of those functions.

First is the number of proper colorations of a graph Γ in λ colors. We define the quantity

$$\chi_\Gamma(\lambda) := \text{the number of proper colorations of } \Gamma \text{ in } \lambda \text{ colors},$$

where λ is a positive integer. Obviously, the first non-negative integer for which $\chi_\Gamma(\lambda)$ is not zero is the chromatic number. (I refrain from writing this fact in an inscrutable formula.)

In order to prove results about $\chi_\Gamma(\lambda)$ let's define the set $P_\Gamma = \{\text{proper colorations of } \Gamma\}$. The first property is the famous (believe me!) deletion-contraction identity.

Lemma H.1. [[LABEL L:0912 chrom dc]] *For any edge e in Γ we have*

$$\chi_\Gamma(\lambda) = \chi_{\Gamma\setminus e}(\lambda) - \chi_{\Gamma/e}(\lambda),$$

where $\lambda \in \mathbb{Z}_{>0}$.

Proof. If e is a loop the result is clear because the left-hand side equals 0 and on the right-hand side $\Gamma \setminus e = \Gamma/e$. If e is a link, first observe that $P_\Gamma \subseteq P_{\Gamma\setminus e}$. Consider the set $P_{\Gamma\setminus e} \setminus P_\Gamma$:

$$\begin{aligned} P_{\Gamma\setminus e} \setminus P_\Gamma &= \{\text{proper colorations of } \Gamma \setminus e \text{ which are improper for } \Gamma\} \\ &= \{\text{proper colorations of } \Gamma \setminus e \text{ in which} \\ &\qquad \text{the endpoints of } e \text{ have the same color}\}. \end{aligned}$$

So there is a natural bijection from the set $P_{\Gamma\setminus e} \setminus P_\Gamma$ to the set $P_{\Gamma/e}$, under which $v_e \in \Gamma/e$ gets the same color as that of both endpoints of $e \in \Gamma \setminus e$. We conclude that $|P_{\Gamma\setminus e}| = |P_\Gamma| + |P_{\Gamma/e}|$ and the result follows. ∎

Lemma H.2. [[LABEL L:0912 chrom mult]] *For any positive integer* λ*,*

$$\chi_{\Gamma_1 \,\dot\cup\, \Gamma_2}(\lambda) = \chi_{\Gamma_1}(\lambda)\chi_{\Gamma_2}(\lambda).$$

Proof. By definition, λ is a positive integer k. There is an obvious one-to-one correspondence between colorations $\gamma : V \to [k]$ and coloration pairs (γ_1, γ_2) where $\gamma_i : V_i \to [k]$ (where $i = 1, 2$) are colorations of Γ_1 and Γ_2. Furthermore, because every edge of Γ is contained within V_1 or V_2, γ is proper if and only if γ_1 and γ_2 are both proper. The lemma follows by the multiplication principle. □

Theorem H.3. [[LABEL T:0912chromaticpoly]] *Given a graph* Γ *with no loops, then* $\chi_\Gamma(\lambda)$ *is a polynomial of degree* n *of the form,*

$$\chi_\Gamma(\lambda) = \lambda^n - a_1\lambda^{n-1} + a_2\lambda^{n-2} - \ldots \pm a_{c(\Gamma)}\lambda^{c(\Gamma)}$$

where $a_i > 0$*. If* Γ *is simple,* $a_1 = |E|$*. If* Γ *has a loop, then* $\chi_\Gamma(\lambda) \equiv 0$*.*

Proof. This is easy to prove inductively on the number of edges by means of Lemmas H.1 and H.2. We assume λ is a positive integer k. We need two obvious examples: $\chi_{K_1}(k) = k$ and $\chi_\varnothing(k) = 1$. It is also obvious that a loop makes the number of proper colorations 0, no matter the number of colors.

If Γ has no edges, then $\chi_\Gamma(k) = \chi_{K_1}(k)^n = k^n$, a monic polynomial satisfying the description in the theorem.

Suppose Γ has an edge e but no loops. By induction, $\chi_{\Gamma\setminus e}(k)$ is a monic polynomial of degree n and $a_1' = |E'| = |E| - 1$, where for convenience we write $\Gamma' := \Gamma \setminus e$. If Γ is simple, then Γ/e has no loops, so $\chi_{\Gamma/e}(k)$ is a monic polynomial of degree $n-1$; consequently, $a_1 = a_1' + 1$ by deletion-contraction. In any case, either $\chi_{\Gamma/e}(k)$ is a polynomial of degree $n-1$ or is identically zero; in each case χ_Γ is a polynomial, monic because its leading term is the same as that of $\chi_{\Gamma\setminus e}$. □

Because of Theorem H.3 we are entitled to call $\chi_\Gamma(\lambda)$ by its right name.

Definition H.2. [[LABEL Df:0912chromaticpoly]] The *chromatic polynomial* of Γ is $\chi_\Gamma(\lambda)$.

We may substitute any number for λ; indeed, we should regard λ as an indeterminate that may take on any value in any commutative ring. The key identities of Lemmas H.1 and H.2 are now polynomial identities, because they are valid for infinitely many values $\lambda \in \mathbb{R}$ and χ_Γ has integral coefficients.

Proposition H.4. [[LABEL P:0912 gen chrom poly]] *The chromatic polynomial has the subset expansion*

$$\chi_\Gamma(\lambda) = \sum_{S\subseteq E} (-1)^{|S|}\lambda^{c(S)}.$$

Proof. This, like many other results, follows from Lemma H.1 by induction on the number of edges that are not loops.

For no edges, $\chi_\Gamma(\lambda) = \lambda^n$. Since $S = \varnothing$ only, the proposition is correct.

For a graph with a loop e, the chromatic polynomial equals 0, and the sum equals

$$\sum_{S\subseteq E\setminus e} \left[(-1)^{|S|}\lambda^{c(S)} + (-1)^{|S\cup e|}\lambda^{c(S\cup e)}\right] = \sum_{S\subseteq E\setminus e} \left[(-1)^{|S|}\lambda^{c(S)} + (-1)^{|S|+1}\lambda^{c(S)}\right] = \sum_{S\subseteq E\setminus e} [0],$$

which is the correct value.

For a graph with no loops and at least one link, say e is one of the links. By Lemma H.1 and induction on the number of edges,

$$\begin{aligned}\chi_\Gamma(\lambda) &= \chi_{\Gamma\setminus e}(\lambda) - \chi_{\Gamma/e}(\lambda)\\ &= \sum_{S\subseteq E\setminus e} (-1)^{|S|}\lambda^{c_{\Gamma\setminus e}(S)} - \sum_{S\subseteq E\setminus e} (-1)^{|S|}\lambda^{c_{\Gamma/e}(S)}\\ &= \sum_{S\subseteq E\setminus e} (-1)^{|S|}\lambda^{c_\Gamma(S)} + \sum_{S\subseteq E\setminus e} (-1)^{|S\cup e|}\lambda^{c_\Gamma(S\cup e)}\\ &= \sum_{S\subseteq E}(-1)^{|S|}\lambda^{c_\Gamma(S)},\end{aligned}$$

which is the proposition. □

Sept 12b: Simon Joyce

H.4. **Spanning trees.** [[LABEL 1.treecount]]

The chromatic polynomial is not the only graph function with algebraic properties like those stated in Lemmas H.1 and H.2. Define

$$t(\Gamma) := \text{ the number of spanning trees in a graph } \Gamma.$$

Lemma H.5. [[LABEL L:0912 tree dc]] *The number of spanning trees in a graph has the deletion-contraction property*

$$t(\Gamma) = t(\Gamma \setminus e) + t(\Gamma/e)$$

for any edge e that is not a loop, and

$$t(\Gamma_1 \cup \Gamma_2) = t(\Gamma_1)t(\Gamma_2).$$

[WE NEED A PROOF. Could one be based on the proof for forests given next time? Is there a simple direct proof?]

The spanning tree number is really different from the chromatic polynomial, because no evaluation of the latter can give the former. We show that, with two small examples.

Example H.1. [[LABEL X:0912treechrom]] Consider K_1 versus $K_1^\bullet$, a single vertex with a loop, and K_2 versus $2K_2$, a pair of parallel links.

In the smallest possible example, $t(K_1^\bullet) = 1$ but $\chi_{K_1}(\lambda) = 0$, so evaluating the chromatic polynomial cannot give the spanning tree number. But perhaps this example, whose distinguishing characteristic is that it has a loop, is too trivial.

For a counterexample without loops, consider the fact that K_2 and $2K_2$ have the same chromatic polynomials (from the definition), but $t(K_2) = 1$ while $t(2K_2) = 2$.

But perhaps the reader wants only simple graphs? I'm sure there are known simple graphs with the same chromatic polynomial but different numbers of spanning trees, but I can't giv an example.

Sept 12c: Simon Joyce

H.5. **The dichromatic and corank-nullity polynomials.** [[LABEL 1.dichromatic]]

There is a function that encompasses both the chromatic polynomial and the spanning-tree number, and has the algebraic properties of both of them. That is the dichromatic polynomial.

The dichromatic polynomial.

The dichromatic polynomial generalizes the chromatic polynomial to two variables.

Definition H.3. [[LABEL D:0912 dichrom poly]] The *dichromatic polynomial* of a graph is

$$Q_\Gamma(u,v) = \sum_{S \subseteq E} u^{c(S)} v^{|S|-n+c(S)}.$$

The subset expansion in Proposition H.4 is what tells us the dichromatic polynomial does specialize to the chromatic polynomial; specifically, $\chi_\Gamma(\lambda) = (-1)^n Q_\Gamma(-\lambda, -1)$. Many of the algebraic properties of the chromatic polynomial also generalize; to begin with, the fundamental deletion-contraction identity.

Proposition H.6. [[LABEL P:0912 dichrom dc]] *The dichromatic polynomial of a graph satisfies*

$$Q_\Gamma(u,v) = Q_{\Gamma \setminus e}(u,v) + Q_{\Gamma/e}(u,v)$$

for any edge e that is not a loop.

Proof. There is a standard way to prove this sort of identity. We divide the defining sum of Q_Γ into a part without e and a part with e. The former part is obviously $Q_{\Gamma \setminus e}$ and the latter part is $Q_{\Gamma/e}$, but that is not as obvious.

Here is the calculation:

$$\begin{aligned} Q_\Gamma(u,v) &= \sum_{S \subseteq E} u^{c(S)} v^{|S|-n+c(S)} \\ &= \sum_{S \subseteq E \setminus e} u^{c(S)} v^{|S|-n+c(S)} + \sum_{T \cup e \subseteq E} u^{c(T \cup e)} v^{|T \cup e|-n+c(T \cup e)} \end{aligned}$$

where T is assumed to be $\subseteq E \setminus e$,

$$= Q_{\Gamma \setminus e}(u,v) + \sum_{T \subseteq E(\Gamma/e)} u^{c_{\Gamma/e}(T)} v^{|T|-|V(\Gamma/e)|+c_{\Gamma/e}(T)}$$

because $|V(\Gamma/e)| = n-1$ and contracting an edge does not change the number of components,

$$= Q_{\Gamma \setminus e}(u,v) + Q_{\Gamma/e}(u,v).$$ □

Sept 12: Zaslavsky

Example H.2. [[LABEL X:0912smallQ]] Let's use the definition to do the smallest examples. The empty graph $K_0 = \varnothing$ gives

$$Q_\varnothing = 1$$

since there is only one edge set, $S = \varnothing$. For the same reason,

$$Q_{K_n} = u^n.$$

For a single edge, K_2, we have

$$Q_{K_2} = u^2 + u$$

from the sets $S = \varnothing$ and E.

For a loop, that is, a circle C_1 of length 1, we have

$$Q_{C_1} = u + uv.$$

For a digon C_2, apply the deletion-contraction law, Proposition H.6. For any edge e in C_2, $C_2 \setminus e = K_2$ and $C_2/e = C_1$.

$$Q_{C_1} = Q_{K_2} + Q_{C_1} = (u^2 + u) + (u + uv) = u^2 + 2u + uv.$$

Next, we calculate two larger examples by means of, respectively, the definition and Proposition H.6.

Example H.3. [[LABEL X:0912tree]] A forest F_{nm} of order $n \geq 1$ with m edges has

$$Q_{F_{nm}} = u^{n-m}(u+1)^m.$$

In particular, for a tree T_n of order n,

$$Q_{T_n} = u(u+1)^{n-1}.$$

We prove the forest formula by observing that a subset of E gives a forest with the same order and fewer edges. There are $\binom{m}{k}$ sets $S \subseteq E$ of k edges, each of which has $n-k$ connected components. From the definition, therefore,

$$Q_{F_{nm}} = \sum_{k=0}^{m} \binom{m}{k} u^{n-k} v^0 = u^{n-m} \sum_{k=0}^{m} \binom{m}{k} u^{m-k} = u^{n-m}(u+1)^m.$$

Example H.4. [[LABEL X:0912circle]] For $n \geq 1$,

$$Q_{C_n} = (u+1)^n - 1 + uv.$$

To prove this we may use induction on n, with a single edge contraction that reduces C_n to C_{n-1} and a deletion that reduces it to a tree, actually a path, T_n. The initial case $n = 1$ is in Example H.2. For higher n,

$$Q_{C_n} = Q_{F_{n,n-1}} + Q_{C_{n-1}} = u(u+1)^{n-1} + (u+1)^{n-1} - 1 + uv = (u+1)^n - 1 + uv.$$

Example H.5. [[LABEL X:0912multiedge]] At this point, interested readers may compute Q_{mK_2} for themselves, where mK_2 consists of m parallel edges joining two vertices, and compare it to Q_{C_n}. The comparison is interesting.

Sept 15a: Yash Lodha

The corank-nullity polynomial.

In the definition of the dichromatic polynomial,

$$Q_\Gamma(u, v) = \sum_{S \subseteq E} u^{c(S)} v^{|S|-n+c(S)},$$

there are two quantities which are significant for the graph. The *corank* of $S \subseteq E$ is defined as $c(S) - c(\Gamma)$ and its *nullity* is defined as $|S| - n + c(S)$ (the names come from matrix theory; see Section E.1). Since $c(S)$ is obviously at least as large as $c(\Gamma)$, and $n - c(S) = |T| \leq |S|$ for a maximal forest $T \subseteq S$, both the corank and nullity are nonnegative. The definitions motivate the name of the following polynomial, called the *rank generating polynomial* or *corank-nullity polynomial*, which is

$$R_\Gamma(u, v) := \sum_{S \subseteq E} u^{c(S)-c(\Gamma)} v^{|S|-n+c(S)} = u^{-c(\Gamma)} Q_\Gamma(u, v).$$

The nullity of Γ equals the cyclomatic number of Γ (Section B.4).

Proposition H.7. [[LABEL P:0915BR]] *The corank-nullity polynomial satisfies the relation* $R_\Gamma = R_{\Gamma\setminus e} + R_{\Gamma/e}$ *for an edge* e *that is not a loop or an isthmus.*

Proof. We use the standard method, splitting the defining sum into two parts according to whether e is or is not in S. Thus,

$$\begin{aligned} R_\Gamma &= \sum_{S\subseteq E} u^{c(S)-c(\Gamma)} v^{|S|-n+c(S)} \\ &= \sum_{S\subseteq E\setminus e} u^{c(S)-c(\Gamma)} v^{|S|-n+c(S)} + \sum_{e\in S\subseteq E} u^{c(S)-c(\Gamma)} v^{|S|-n+c(S)} \\ &= \sum_{S\subseteq E\setminus e} u^{c(S)-c(\Gamma\setminus e)} v^{|S|-n+c(S)} + \sum_{T\subseteq E\setminus e} u^{c(S)-c(\Gamma)} v^{|T\cup e|-n+c(T\cup e)} \end{aligned}$$

because e is not an isthmus so $c(\Gamma\setminus e) = c(\Gamma)$,

$$= R_{\Gamma\setminus e} + \sum_{T\subseteq E\setminus e} u^{c(T\cup e)-c(\Gamma)} v^{|T\cup e|-n+c(T\cup e)}$$

through replacing $S \ni e$ by $T\cup e$ where $T \subseteq E\setminus e$. The task now is to express the remaining summation in terms of Γ/e. To do this we make a structural comparison between $T\cup e$ in Γ and T in Γ/e. The essential facts are that $c(T\cup e$, the component count in Γ, equals $c_{\Gamma\setminus e}(T)$, and that $|V(\Gamma\setminus e)| = n-1$ since e is not a loop. Now we continue the previous calculation:

$$\begin{aligned} R_\Gamma - R_{\Gamma\setminus e} &= \sum_{T\subseteq E\setminus e} u^{c_{\Gamma\setminus e}(T)-c(\Gamma/e)} v^{|T|+1-n+c_{\Gamma\setminus e}(T)} \\ &= \sum_{T\subseteq E\setminus e} u^{c_{\Gamma\setminus e}(T)-c(\Gamma/e)} v^{|T|-|V(\Gamma/e)|+c_{\Gamma\setminus e}(T)} \\ &= R_{\Gamma/e}. \end{aligned}$$

by the definition of the corank-nullity polynomial. ☐

(Sometimes I like to write S/e, instead of T, to mean the set $S\setminus e$ in the graph Γ/e. That makes the proof look cuter but it can be confusing.)

Sept 17a (draft): Peter Cohen et al.

More properties of the dichromatic and corank-nullity polynomials.

[NOTE: Needs graphs. The source code has notes where the graphs and diagrams will go in.]

Proposition H.8. [[LABEL P:0917C]] *The dichromatic polynomial has the multiplicative property* $Q_{\Gamma_1 \,\dot\cup\, \Gamma_2} = Q_{\Gamma_1} Q_{\Gamma_2}$.

Proof. The proof is similar to that of the next proposition. ☐

The *vertex amalgamation* of two graphs is defined to be

$$\Gamma_1 \cup_v \Gamma_2 := \Gamma_1 \cup \Gamma_2,$$

where Γ_1 and Γ_2 share a vertex v and have no other vertex or edge in common. This frequently occurs, e.g. when isthmi are contracted.

[add in graph later]

Proposition H.9. [[LABEL P:0917CR]] *The corank-nullity polynomial has the multiplicative properties* $R_{\Gamma_1 \mathbin{\dot\cup} \Gamma_2} = R_{\Gamma_1 \cup_v \Gamma_2} = R_{\Gamma_1} R_{\Gamma_2}$.

Proof. Consider the case of a vertex amalgation, $\Gamma = \Gamma_1 \cup_v \Gamma_2$. Then, first of all, $n = n_1 + n_2 - 1$; secondly, $c(\Gamma) = c(\Gamma_1) + c(\Gamma_2) - 1$, because one component of Γ_1 merges with one component of Γ_2 in the amalgamation; and thirdly, the same relationship holds for any spanning subgraph (V, S) if $S_1 = S \cap E_1$ and $S_2 = S \cap E_2$. So,

$$\begin{aligned}
R_\Gamma &= \sum_{S \subseteq E_1 \cup E_2} u^{c(S)-c(\Gamma)} v^{|S|-n+c(S)} \\
&= \sum_{S_1 \subseteq E_1} \sum_{S_2 \subseteq E_2} u^{c(S_1 \cup S_2)-c(\Gamma_1 \cup \Gamma_2)} v^{|S_1 \cup S_2|-n+c(\Gamma_1 \cup \Gamma_2)} \\
&= \sum_{S_1 \subseteq E_1} \sum_{S_2 \subseteq E_2} u^{[c(S_1)+c(S_2)-1]-[c(\Gamma_1)+c(\Gamma_2)-1]} v^{[|S_1|+|S_2|]-[n_1+n_2-1]+[c(\Gamma_1)+c(\Gamma_2)-1]}
\end{aligned}$$

by the preceding remarks, and then by simplifying and rearranging the exponents and separating the two summations,

$$\begin{aligned}
&= \sum_{S_1 \subseteq E_1} u^{c(S_1)-c(\Gamma_1)} v^{|S_1|-n_1+c(\Gamma_1)} \sum_{S_2 \subseteq E_2} u^{c(S_2)-c(\Gamma_2)} v^{|S_2|-n_2+c(\Gamma_2)} \\
&= R_{\Gamma_1} R_{\Gamma_2}.
\end{aligned}$$

The proof for disjoint unions is similar, but simpler since $n = n_1 + n_2$, $c(\Gamma) = c(\Gamma_1) + c(\Gamma_2)$, and $c(S) = c(S_1) + c(S_2)$. □

Sept 15b: Yash Lodha

H.6. **Counting maximal forests and spanning trees.** [[LABEL 1.maximalforests]]

Let $f(\Gamma)$ be the number of maximal forests of Γ and $t(\Gamma)$ the number of spanning trees. To understand the polynomials discussed above, we calculate them for the graphs $\varnothing, K_1, K_2, \bar{K}_2$ and compare them with the values of the functions f and t for these graphs.

Γ	$Q_\Gamma(u,v)$	$R_\Gamma(u,v)$	$t(\Gamma)$	$f(\Gamma)$
$\varnothing$	1	1	0	1
K_1	u	1	1	1
K_2	u^2+u	$u+1$	1	1
$\bar{K}_2$	$Q_{K_1}(u,v)^2 = u^2$	1	0	1

Lemma H.10. [[LABEL L:0915F]] *The following equations hold for the number of maximal forests of a graph:*

$$f(\Gamma_1 \mathbin{\dot\cup} \Gamma_2) = f(\Gamma_1) f(\Gamma_2),$$
$$f(\Gamma) = f(\Gamma \setminus e) + f(\Gamma / e)$$

if e is not a loop or an isthmus, and

$$f(\varnothing) = 1.$$

Proof. The first equation follows simply from the fact that the maximal forests of $\Gamma_1 \mathbin{\dot\cup} \Gamma_2$ are in bijective correspondence with pairs of maximal forests in Γ_1 and Γ_2. So the result follows from the multiplication principle.

For the second result, we assume that e is a link. There are two kinds of maximal forest of Γ, the ones that contain e and the ones that do not contain e. The ones that contain e are in bijection with the maximal forests of Γ/e and the ones that do not contain e are in bijection with the maximal forests of $\Gamma \setminus e$. This proves the second equation.

For the third equation, we need only keep in mind the empty forest! $\square$

Theorem H.11. [[LABEL T:0915F]] *The number of maximal forests in Γ is $f(\Gamma) = R_\Gamma(0,0)$.*

Proof. Initially, we assume that Γ is connected. We proceed by induction on $|E|$. There are three cases—not mutually exclusive.

Case I: Γ has a loop e. Then $\Gamma = (\Gamma \setminus e) \cup_v K_1^\circ$. By Proposition H.9, $R_\Gamma = (1+v)R_{\Gamma\setminus e}$. So,

$$R_\Gamma(0,0) = 1 \cdot R_{\Gamma\setminus e}(0,0) = 1 \cdot f(\Gamma \setminus e) = f(\Gamma)$$

since e is a loop.

Case II: Γ has no loop and every edge is an isthmus. Then Γ is a tree. By inspection we can see that $f(\Gamma) = 1 = R_\Gamma(0,0)$.

Case III: Γ has a circle C of length greater than one. Let $e \in C$. Then e is not a loop or isthmus, so by Proposition H.7 and Lemma H.10,

$$R_\Gamma(0,0) = R_{\Gamma\setminus e}(0,0) + R_{\Gamma/e}(0,0) = f(\Gamma \setminus e) + f(\Gamma/e) = f(\Gamma).$$

This proves the theorem when Γ is connected.

If Γ has more than one component, we proceed by induction on the number of components of Γ. Let $\Gamma = \Gamma_1 \mathbin{\dot\cup} \Gamma_2$, where our theorem holds for Γ_1 and Γ_2 is a connected graph. Then by Proposition H.9 and Lemma H.10 we get our inductive step:

$$R_\Gamma(0,0) = R_{\Gamma_1}(0,0)R_{\Gamma_2}(0,0) = f(\Gamma_1)f(\Gamma_2) = f(\Gamma). \quad \square$$

Here are a few examples that illustrate the theorem.

$$\begin{aligned} R_\varnothing(0,0) &= 1 = f(\varnothing), \\ R_{K_1}(0,0) &= 1 = f(K_1), \\ R_{K_1^\circ}(0,0) &= 1 = f(K_1^\circ), \\ R_{K_2}(0,0) &= 1 = f(K_2). \end{aligned}$$

Sept 22a: Jackie Kaminski

Theorem H.12. [[LABEL T:0922 tree forest polys]] *The number of spanning trees of a graph $\Gamma \neq \varnothing$ is*

$$t(\Gamma) = u^{-1}Q_\Gamma(u,v)|_{(0,0)} = \frac{\partial}{\partial u}Q_\Gamma(0,0).$$

We begin with a lemma before we prove this theorem.

Lemma H.13. [[LABEL L:0922 rank poly I/L]] *For Γ a graph containing only isthmi and loops, $R_\Gamma(u,v) = (u+1)^{\#\text{ of isthmi}}(v+1)^{\#\text{ of loops}}$.*

Proof. First notice that a graph with only isthmi and loops is a forest with loops. We first introduce notation. For a given edge set S, let S_0 be the set of all loops in S, and S_1 be the set of all isthmi in S. Since we are restricted to the case where our graphs contain only loops and isthmi, $S_0 \uplus S_1 = S$. Now we recall that by definition,

$$R_\Gamma = \sum_{S \subseteq E} u^{c(S)-c(\Gamma)} \cdot v^{|S|-n+c(S)}$$

which, since loops don't affect $c(S)$ but adding an isthmus to a graph decreases the number of connected components by exactly 1, is

$$\begin{aligned}
&= \sum_{S \subseteq E} u^{(n-|S_1|)-c(\Gamma)} \cdot v^{|S|-n+(n-|S_1|)} \\
&= \sum_{S \subseteq E} u^{n-|S_1|-c(\Gamma)} \cdot v^{|S|-|S_1|} \\
&= \sum_{S \subseteq E} u^{n-|S_1|-c(\Gamma)} \cdot v^{|S_0|}
\end{aligned}$$

in which, letting E_0 be the set of loops of Γ and E_1 the set of isthmi of Γ, we can reindex:

$$\begin{aligned}
&= \sum_{S_0 \subseteq E_0} \sum_{S_1 \subseteq E_1} u^{n-|S_1|-c(\Gamma)} \cdot v^{|S_0|} \\
&= u^{n-c(\Gamma)} \Bigg(\sum_{S_0 \subseteq E_0} v^{|S_0|} \Bigg) \Bigg(\sum_{S_1 \subseteq E_1} \left(\frac{1}{u}\right)^{|S_1|} \Bigg) \\
&= u^{n-c(\Gamma|E_1)} \Bigg(\sum_{S_0 \subseteq E_0} v^{|S_0|} \Bigg) \Bigg(\sum_{S_1 \subseteq E_1} \left(\frac{1}{u}\right)^{|S_1|} \Bigg) \\
&= u^{|E_1|} (v+1)^{|E_0|} \left(\frac{1}{u} + 1 \right)^{|E_1|} \\
&= (v+1)^{|E_0|} (1+u)^{|E_1|}. \qquad \square
\end{aligned}$$

Proof of Theorem H.12. In preparation, recall from Lemma H.5 that $t(\Gamma) = t(\Gamma \setminus e) + t(\Gamma/e)$ for all links e, and from Lemma H.10 that $f(\Gamma) = f(\Gamma \setminus e) + f(\Gamma/e)$ for all edges that are not links or loops. When at least one of Γ_1 and Γ_2 is not $\varnothing$, then $t(\Gamma_1 \uplus \Gamma_2) = 0$. This is immediate since spanning trees must be connected. Lastly, by inspection we see that $t(\varnothing) = 0$.

It is easy to see that $u^{-1} Q_\Gamma(u,v)|_{(0,0)} = \frac{\partial}{\partial u} Q_\Gamma(0,0)$. From the fact that Q_Γ is a polynomial with no constant term, it follows immediately that $u^{-1} Q_\Gamma(u,v)|_{(0,0)} = \frac{\partial}{\partial u} Q_\Gamma(0,0)$. Alternatively, an inductive proof of the result that $t(\Gamma) = \frac{\partial}{\partial u} Q_\Gamma(0,0)$ is almost identical to the proof that $t(\Gamma) = u^{-1} Q_\Gamma(u,v)|_{(0,0)} = \frac{\partial}{\partial u} Q_\Gamma(0,0)$, the difference being that the last step is due to linearity of the derivative instead of Proposition H.6.

The proof of the theorem is by induction on $|E|$. First we look at the base cases, where $|E| \leq 1$; then Γ is a single link with 2 vertices ($\Gamma = K_2$), a single loop with one vertex ($\Gamma = K_1^\circ$), or a single vertex ($\Gamma = K_1$).

Notice that $t(K_1) = t(K_1^\circ) = t(K_2) = 1$ since K_1 or K_2 is itself the unique spanning tree. From the defining formula of the dichromatic polynomial it follows that $u^{-1} Q_{K_2}(0,0) =$

$0+1=1$ and $u^{-1}Q_{K_1}(0,0) = u^{-1}Q_{K_1^\circ}(0,0) = 0+1 = 1$. So $t(\Gamma) = u^{-1}Q_\Gamma(u,v)|_{(0,0)}$ holds for the base cases.

Now let Γ be a graph with at least two edges, and assume the theorem holds for all graphs on fewer edges. We will handle all graphs with a non-isthmus link by showing that both sides of the respective formulas satisfy the same deletion-contraction recursion, with the same initial conditions. We will then look at the remaining graphs, which contain only isthmi and loops.

Case 1: Γ contains an edge e which is a link but not an isthmus.

Then $t(\Gamma) = t(\Gamma \setminus e) + t(\Gamma/e)$ by Lemma H.5, and since both $\Gamma \setminus e$, Γ/e have fewer edges than Γ, by the inductive hypothesis we may conclude that

$$\begin{aligned} t(\Gamma \setminus e) + t(\Gamma/e) &= u^{-1}Q_{\Gamma\setminus e}(u,v)|_{(0,0)} + u^{-1}Q_{\Gamma/e}(u,v)|_{(0,0)} \\ &= u^{-1}Q_\Gamma(u,v)|_{(0,0)} \end{aligned}$$

by Proposition H.6, since e was not a loop. Therefore $t(\Gamma) = u^{-1}Q_\Gamma(u,v)|_{(0,0)}$.

Therefore we have proven the theorem for all graphs that have a non-isthmus link, assuming that it holds for graphs with only loops and isthmi.

Case 2: Γ contains only loops and isthmi.

Here, since Γ is a forest with loops, $t(\Gamma) = 1$ if Γ is connected, and otherwise $t(\Gamma) = 0$. Furthermore, evaluating $u^{-1}Q(u,v)$ at $(0,0)$, the only possibility for a non-zero term is when there is a subset S such that $c(S) = 1$ and $|S| + c(S) = n$, i.e., there is a spanning, connected subgraph (V,S) with $n-1$ edges. That is the case where Γ is connected. As Γ contains only isthmi and loops, there is at most one non-zero term in the sum; thus $t(\Gamma) = u^{-1}Q_\Gamma(u,v)|_{(0,0)}$.

We have proven Theorem H.12. □

Sept 22b: Jackie Kaminski

H.7. **The number of improper colorations.** [[LABEL 1.improper]]

Just as the chromatic polynomial gives the number of proper colorations, the dichromatic polynomial counts all colorations, grouped by the number of improper edges. In technical language, the dichromatic polynomial, with a change of variables and normalization, is the generating function of all colorations by the size of the improper edge set.

Definition H.4. [[LABEL D:0922 improper]] For a k-coloration $\gamma : V \to [k]$, we say $e{:}vw$ is *improper* if $\gamma(v) = \gamma(w)$. We let $I(\gamma)$ denote the set of improper edges of γ.

Definition H.5. For $k \in \mathbb{Z}_{>0}$ we define $X_\Gamma(k,\cdot)$ to be the *generating function of k-colorations by the number of improper edges*; that is,

$$X_\Gamma(k,w) = \sum_\gamma w^{|I(\gamma)|} = \sum_{i=0}^{|E|} m_i w^i,$$

where m_i is the number of k-colorations with exactly i improper edges.

By definition X_Γ is a polynomial in w. We would also like to show that it is a polynomial in k. But first notice that $X_\Gamma(k,0) = \chi_\Gamma(k)$, since χ_Γ counts the number of k-colorations that are proper. We now prove a theorem of Tutte's which implies that X is also a polynomial in k.

Theorem H.14 (Tutte)**.** [[LABEL T0922 Tutte]] *For a graph* Γ*, the generating function of colorations by the number of improper edges satisfies*

$$X_\Gamma(k,w) = (w-1)^n Q_\Gamma\big(\frac{k}{w-1}, w-1\big).$$

Proof. First we reformulate the dichromatic polynomial:

$$\text{[[LABEL E : Qreduced]]}Q_\Gamma(u,v) = \sum_{S\subseteq E} u^{c(S)} v^{|S|-n+c(S)} = v^{-n}\sum_{S\subseteq E}(uv)^{c(S)} v^{|S|}. \tag{H.1}$$

Now we look at X_Γ as a sum over all colorations γ:

$$\begin{aligned} X_\Gamma(k,w) &= \sum_{\gamma:V\to[k]} w^{|I(\gamma)|} \\ &= \sum_{S\subseteq E} w^{|S|}\cdot \#(\gamma \text{ such that } I(\gamma)=S). \end{aligned}$$

Here we apply the fact that the number of k-colorations γ whose set of improper edges is precisely S equals the number of proper k-colorations of Γ/S. This follows from the fact that the color on a component of S is constant, because the component is connected through links whose two endpoints have the same color. On the other hand, any edge not in S must have different colors at each end. This is the definition of a proper coloration of Γ/S. (Note that if S is not closed, there are no colorations with it as improper edge set; while, most conveniently, $\chi\Gamma/S(k)=0$.) It follows that

$$X_\Gamma(k,w) = \sum_{S\subseteq E} w^{|S|}\chi_{\Gamma/S}(k),$$

which by Proposition H.4

$$\begin{aligned} &= \sum_{S\subseteq E} w^{|S|} \sum_{T\subseteq E\setminus S}(-1)^{|T|}k^{c_{\Gamma/S}(T)} \\ &= \sum_{S\subseteq E} w^{|S|} \sum_{T\subseteq E\setminus S}(-1)^{|T|}k^{c(S\cup T)} \end{aligned}$$

since $(V, S\cup T)$ has the same number of components as $(V(\Gamma/S), T)$. Now we reindex with $S\cup T = R$ so $T = R\setminus S$:

$$\begin{aligned} &= \sum_{R\subseteq E}\sum_{S\subseteq R} w^{|S|}(-1)^{|R\setminus S|}k^{c(R)} \\ &= \sum_{R\subseteq E}\Big(\sum_{S\subseteq R} w^{|S|}(-1)^{|R\setminus S|}\Big)k^{c(R)} \end{aligned}$$

and by the binomial formula this simplifies:

$$= \sum_{R\subseteq E}(w-1)^{|R|}\ k^{c(R)}.$$

Now multiplying by $\frac{w-1}{w-1}$ in some clever places we get

$$X_\Gamma(k,w) = \sum_{R\subseteq E} \Big(\frac{k}{w-1}(w-1)\Big)^{c(R)}(w-1)^{|R|}$$

and by Equation (H.1) this is

$$= (w-1)^n Q_\Gamma(\frac{k}{w-1}, w-1). \qquad \square$$

We end with an example.

Example H.6. [[LABEL P:0922 Cycle]] For $n \geq 1$,

$$\begin{aligned} X_{C_n}(k,w) &= (w+k-1)^n + (k-1)(w-1)^n \\ &= \sum_{i=0}^{n} w^i \binom{n}{i} \Big[(k-1)^{n-i} + (-1)^{n-i}(k-1)\Big]. \end{aligned}$$

This follows from Example H.4 and Theorem H.14.

The coefficient of w^0 is the chromatic polynomial, so $\chi_{C_n}(k) = (k-1)^n + (-1)^n(k-1)$.

The coefficient of w is the number of k-colorations with exactly one improper edge; that is $n\big[(k-1)^{n-1} + (-1)^{n-1}(k-1)\big]$. Think of this number combinatorially; the one improper edge implies there is one edge whose endpoints have the same color, and contracting that edge gives a proper coloration. There are n choices for the edge and $\chi_{C_{n-1}}(k)$ choices for the proper coloration; thus, the coefficient of w should be $n\chi_{C_{n-1}}(k)$, which is precisely what we found.

Sept 17b (draft): Peter Cohen et al.

H.8. **Acyclic orientations and proper and compatible pairs.** [[LABEL 1.acyclicpairs]]

We defined the chromatic polynomial by its values at positive integers, and extended it to all real (or complex) numbers by proving it is a polynomial. Now we use the deletion-contraction property to establish a combinatorial meaning for the values of that polynomial at *negative* integral arguments.

Define $\mathrm{AO}(\Gamma)$ to be the set of acyclic orientations of Γ. In an oriented graph, the notation $\vec{P}{:}\overrightarrow{vw}$ means a directed path from v to w.

Acyclic vs. cyclic orientations.

[next part needs the graph]

Lemma H.15. [[LABEL L:0917ao-dc]] *Consider an orientation $\vec{\Gamma}$ of Γ and an edge $e{:}\overrightarrow{v_1v_2}$ in $\vec{\Gamma}$.* ***[This is not necessarily the right lemma. We need to check the notes.]***

(1) If there exists $\vec{P}{:}\overrightarrow{v_1v_2}$ in $\vec{\Gamma}\setminus e$ then the orientation is cyclic.
(2) If there does not exist $\vec{P}{:}\overrightarrow{v_1v_2}$ in $\vec{\Gamma}\setminus e$ for any e, then the orientation is acyclic.

[again, graph]
[small simple circle graph]

Proof. [**(needs simple graphs to show, will add soon)**]

[**ARE THESE PROOF CASES RIGHT?** $\vec{P}$ **can't be a cycle because** P **is not a circle. Or, do you mean it's a cycle in the contraction?**]

Case 1: $\vec{\Gamma}/e$ *is oriented as in* $\vec{\Gamma} \setminus e$. In this case, $\vec{P}$ is not a cycle. [**PROOF NEEDED.**]

Case 2: $\vec{P}$ *doesn't exist, so* $\vec{\Gamma}/e$ *is acyclic.* [**PROOF NEEDED.**]

[**IS THIS enough cases?**]

□

Let $a(\Gamma)$ denote the number of acyclic orientations of Γ, i.e., $a(\Gamma) := |\operatorname{AO}(\Gamma)|$. There is a deletion-contraction formula for this number.

Lemma H.16. *For any edge* e, $a(\Gamma) = a(\Gamma \setminus e) - a(\Gamma/e)$.

[**MUST DISCUSS. CAN YOU GET ADDITIONAL NOTES FROM JACKIE OR NATE OR SIMON?**]

[**STANLEY'S THEOREM** $a(\Gamma) = (-1)^n \chi_\Gamma(-1)$.]

(Notes end, need clarification)

We will denote a color set, (α, γ), by $[K]$.

For two vertices, v_1 and v_2, $\gamma(v_1) \leq \gamma(v_2)$

[**needs a lot of work**]

Sept 19: Simon Joyce

Proper and compatible pairs.

A pair (α, γ) consisting of an acyclic orientation and a coloration of Γ with color set $[k]$ is called *proper* if, for each edge $e{:}vw$ such that α orients e from v to w, then $\gamma(v) < \gamma(w)$. The pair is *compatible* if under those conditions $\gamma(v) \leq \gamma(w)$.

Lemma H.17. [[LABEL L:0919 AO]] *Given a graph* Γ *and a link* $e \in E(\Gamma)$, *there is a bijection*

$$\operatorname{AO}(\Gamma) \cup \operatorname{AO}(\Gamma/e) \leftrightarrow \operatorname{AO}(\Gamma \setminus e).$$

Proof. If $\alpha_0 \in \operatorname{AO}(\Gamma \setminus e)$ is such that α_0 is also an acyclic orientation of Γ/e, then α_0 can be extended to an acyclic orientation of Γ by adding $e{:}\overrightarrow{vw}$ or $e{:}\overrightarrow{wv}$. If α_0 is not an acyclic orientation of Γ/e then only one of these is a valid extension to Γ. □

Given a graph Γ and a number of colors k, let

$$p_\Gamma(k) := \#\text{ of compatible pairs.}$$

Lemma H.18. [[LABEL L:0919 c pairs DC]] *Given a graph* Γ *and a link* $e \in E(\Gamma)$, *then*

$$p_{\Gamma\setminus e}(k) = p_\Gamma(k) + p_{\Gamma/e}(k).$$

Proof. For a fixed $k \geq 0$, we will prove there exists a 1:1/2:2 correspondence, or *sesquijection*, between $\operatorname{CP}(\Gamma) \cup \operatorname{CP}(\Gamma/e)$ and $\operatorname{CP}(\Gamma \setminus e)$. Fix $\alpha_0 \in \operatorname{AO}(\Gamma \setminus e)$.

First we assume α_0 orients both $\Gamma \setminus e$ and Γ/e acyclically. The latter means there is no directed path from v to w where v and w are the endpoints of e. Consider γ, a k-coloration of $\Gamma \setminus e$ that is compatible with α_0, so $(\alpha_0, \gamma) \in \operatorname{CP}(\Gamma \setminus e)$.

Either $\gamma(v) = \gamma(w)$ or not. In the former case γ properly colors Γ/e but not Γ, and γ is compatible with both $e{:}\overrightarrow{vw}$ and $e{:}\overrightarrow{wv}$. In the latter case γ doesn't properly color Γ/e but it does properly color Γ, and γ is compatible with exactly one extension of α_0 since $\gamma(v) < \gamma(w)$ or vice versa.

If there exists an oriented path from v to w, we may assume $\gamma(v) < \gamma(w)$. Then α_0 extends by $e{:}\vec{vw}$ and since $\gamma(v) < \gamma(w)$ this extension is unique. Calling this extension α we have $(\alpha_0, \gamma) \leftrightarrow (\alpha, \gamma)$. □

Stanley's famous theorem.

We are now ready to prove our main result.

Theorem H.19 (Stanley)**.** [[LABEL T:0919 Stanley's]] *Given a graph* Γ*,* $\alpha \in \mathrm{AO}(\Gamma)$ *and* γ *a* k*-coloration of* Γ *then,*

$$(-1)^n \chi_\Gamma(-k) = p_\Gamma(k).$$

Proof. If Γ has no links then,

$$(-1)^n \chi_\Gamma(-k) = \begin{cases} 0 & \text{if } \Gamma \text{ contains a loop,} \\ (-1)^n(-k)^n & \text{otherwise.} \end{cases}$$

Also

$$p_\Gamma(k) = \begin{cases} 0 & \text{if } \Gamma \text{ contains a loop,} \\ k^n & \text{otherwise.} \end{cases}$$

So in this case we have equality.

If Γ contains a link then we use lemma H.18, deletion-contraction of χ_Γ and induction. □

The geometry of proper and compatible pairs.

Going back to the idea of coloring, if we take γ to be a k-coloration of a graph Γ we have $\gamma : V \to [k] \subseteq \mathbb{R}$, so we can think of $\gamma \in [k]^n \subseteq \mathbb{R}^n$. So if γ_i and γ_j are the ith and jth coordinates of γ then $\gamma_i \neq \gamma_j$ if $\exists\ e_{ij} \in E(\Gamma)$, i.e., $\gamma \notin h_{ij} = \{x : x_i = x_j\}$ for every $e_{ij} \in E(\Gamma)$, i.e., $\gamma \notin \bigcup \mathcal{H}[\Gamma]$. So we can redefine a proper k-coloration as

$$\gamma \in \mathbb{Z}^n \setminus \bigcup \mathcal{H}[\Gamma] \text{ such that } \gamma \in (0, k+1)^n.$$

This can be restated as

$$\frac{\gamma}{k+1} \in \frac{1}{k+1}\mathbb{Z}^n \text{ and } \frac{\gamma}{k+1} \in (0,1)^n \setminus \bigcup \mathcal{H}[\Gamma].$$

The number of these points is given by a polynomial function of k, $E^\circ(k+1)$, known as the *open Ehrhart polynomial* of $([0,1]^n, \mathcal{H}[\Gamma])$. (See [IOP, Section 5].)

Suppose we have $\gamma : V \to \{0, 1, \ldots, k-1\}$ and $\gamma \in \mathbb{Z}^n \cap [0, k-1]^n$. Each $\alpha \in \mathrm{AO}(\Gamma)$ corresponds to a region $R(\alpha)$ of $\mathcal{H}[\Gamma]$, defined by $x_i < x_j$ when $\exists\ \vec{v_i v_j}$ in α. Also we have $\gamma_i \leq \gamma_j$ when $\exists\ \vec{v_i v_j}$ in α. This defines the *closure* $\overline{R}(\alpha)$ of $R(\alpha)$, called the *closed region of* α. So given γ the number of compatible pairs $(\alpha, \gamma) =$ the number of closed regions of $\mathcal{H}[\Gamma]$ that contain γ.

Given a graph Γ and $x \in \mathbb{R}^n$ we define

$$m(x) := \text{ number of closed regions of } \mathcal{H}[\Gamma] \text{ that contain } x.$$

Now we define the *closed Ehrhart polynomial* to be

$$E(k-1) = \sum_{\gamma \in \mathbb{Z}^n \cap [0,k-1]^n} m(\gamma).$$

Then (by some calculation, which I omit, based on Stanley's theorem), $E^0(t) = (-1)E(-t)$ for $[0,1]^n$ and $\mathcal{H}[\Gamma]$. That is, we really have only one geometric polynomial that counts both proper and compatible pairs. This is no surprise if we already know Stanley's theorem, but it hints at vast generalizations (which are outside the scope of this course, but which apply to signed graphs and will be mentioned when we get to the signed-graphic Stanley-type theorem in Section K.4).

Sept 24: Nate Reff

H.9. **The Tutte polynomial.** [[LABEL 1.tuttepoly]]

The Tutte polynomial is a universal function that satisfies the relations we've been discovering for the corank-nullity polynomial and other polynomials. Let's review these relations.

Tutte–Grothendieck invariants.

We found:

- Deletion-Contraction Property:
  $Q_\Gamma = Q_{\Gamma\backslash e} + Q_{\Gamma/e}$ if e is not a loop.
  $R_\Gamma = R_{\Gamma\backslash e} + R_{\Gamma/e}$ if e is not a loop or isthmus.
- Disjoint Graph Multiplicativity:
  $Q_{\Gamma_1 \mathbin{\dot\cup} \Gamma_2} = Q_{\Gamma_1} Q_{\Gamma_2}$ and $R_{\Gamma_1 \mathbin{\dot\cup} \Gamma_2} = R_{\Gamma_1} R_{\Gamma_2}$.
- Multiplicativity:
  Disjoint Graph Multiplicativity, and $R_{\Gamma_1 \cup_v \Gamma_2} = R_{\Gamma_1} R_{\Gamma_2}$.
- Empty-Graph Unitarity:
  $Q_\varnothing = 1 = R_\varnothing$.
- Unitarity:
  Empty-Graph Unitarity, and $R_{K_1} = 1$.
- Invariance:
  $\Gamma_1 \cong \Gamma_2 \implies Q_{\Gamma_1} = Q_{\Gamma_2}$ and $R_{\Gamma_1} = R_{\Gamma_2}$.

We call a *Tutte–Grothendieck invariant of graphs* any function F on graphs that satisfies all these properties. Let's restate them precisely, in the generality of an arbitrary function F defined on graphs:

(DC) **Deletion-Contraction Identity:**
$F(\Gamma) = F(\Gamma\backslash e) + F(\Gamma/e)$ if e is not a loop or isthmus.

(M) **Multiplicativity:**
$F(\Gamma_1 \mathbin{\dot\cup} \Gamma_2) = F(\Gamma_1 \cup_v \Gamma_2) = F(\Gamma_1)F(\Gamma_2)$.

(U) **Unitarity:**
$F(\varnothing) = F(K_1) = 1$.

(I) **Invariance:**
$\Gamma_1 \cong \Gamma_2 \implies F(\Gamma_1) = F(\Gamma_2)$.

Now let's look at what it means for a function to satisfy these properties, and head toward answering Tutte's question of what are all such functions. First of all, in order for all the properties to make sense, F has to have values in a commutative ring with unity. Next, because of the mulitiplicativity property (M), $F(\Gamma)$ = the product of F(blocks). Due to the property of invariance (I), F(loop) = a value y that is the same for all loops, and also F(isthmus) = a value x that is the same for all isthmi. Lastly, there is a simple form for a basic special case.

Lemma H.20. [[LABEL L:0924lemma1 Loop Isthmus Lemma]] *Suppose Γ has l loops and i isthmi and no other edges. Then $F(\Gamma) = x^i y^l$.*

As another side comment we note that, if the codomain of F is an integral domain, then (U) is almost superfluous; that is, it can be deduced from the other properties, except for a small number of functions F. (This is left as a homework exercise. Hint: Derive (U) from (M) and (I); find the exceptional cases.)

Returning to H.20, let's look as some simple examples. Suppose we define a graph G as in figure H.1. Let C_2 denote the digon graph, which is a circle of length 2; it consists of two vertices and two parallel edges between those vertices. Now for the calculation of F using the deletion-contraction method (as seen in figure H.1) we get the following:

$$\begin{aligned} F(G) &= F(G\backslash e) + F(G/e) \\ &= F((G\backslash e)\backslash a) + F((G\backslash e)/a) + F((G/e)\backslash a) + F((G/e)/a) \\ &= (x^3) + F(K_3) + xF(C_2) + yF(C_2) \\ &= (x^3) + (x^2 + x + y) + x(x + y) + y(x + y) \\ &= x^3 + 2x^2 + x + 2xy + y + y^2. \end{aligned}$$

Theorem H.21. [[LABEL T:0924Theorem1 Main Theorem]] *Suppose F is a Tutte–Grothendieck invariant of graphs. Let $x = F(isthmus)$, and let $y = F(loop)$. Then*

(1) $F(\Gamma) = R_\Gamma(x-1, y-1)$, a polynomial function of x and y,
(2) the polynomial has nonnegative integral coefficients, and
(3) any evaluation of $R_\Gamma(x, y)$ gives a Tutte-Grothendieck invariant of graphs.

Proof. One proves the first two statements by induction on $|E|$, using (DC) and (M). The third statement follows from the fact that R_Γ itself is a Tutte-Grothendieck invariant. □

Corollary H.22. [[LABEL T:0924Corollary1 Main Corollary]] *A Tutte–Grothendieck invariant F is well defined given any choices of $x = F(isthmus)$ and $y = F(loop)$ and is uniquely determined by those choices.*

Proof. This is an immediate corollary of Theorem H.21. □

The Tutte polynomial.

We define the *Tutte polynomial* of Γ as the polynomial obtained by reducing a general $F(\Gamma)$ to x's and y's using the properties defining a Tutte-Grothendieck invariant of graphs. We denote the Tutte polynomial by $T_\Gamma(x, y)$. Our main theorem, Theorem H.21, tells us that $T_\Gamma(x, y) = R_\Gamma(x-1, y-1)$. Using previous results we can now also write $T_\Gamma(1, 1) = R_\Gamma(0, 0) = f(\Gamma)$, the number of maximal forests, and $T_\Gamma(1-\lambda, 0) = R_\Gamma(-\lambda, -1) = (-1)^n \chi_\Gamma(\lambda)$ as well as many other such forms.

Theorem H.23. [[LABEL T:0924Theorem2]] *$T_\Gamma(x, y)$ is a polynomial, with no constant term if $|E| > 0$. The degree of x equals $\mathrm{rk}(\Gamma) = n - c(\Gamma)$, and the degree of y equals the nullity of Γ, that is, $|E| - n + c(\Gamma)$.*

Proof. This is an immediate corollary of Theorem H.21. **[NOT SO IMMEDIATE. Needs some indication of proof.]** □

Let's take another look at the subset expansion of the corank-nullity polynomial:

$$R_\Gamma(u, v) = \sum_S u^{c(S)-c(F)} v^{|S|-n+c(F)} = \sum_{k,l} a_{kl} u^k v^l, \text{[[LABEL E : 0924Tutte1]]} \tag{H.2}$$

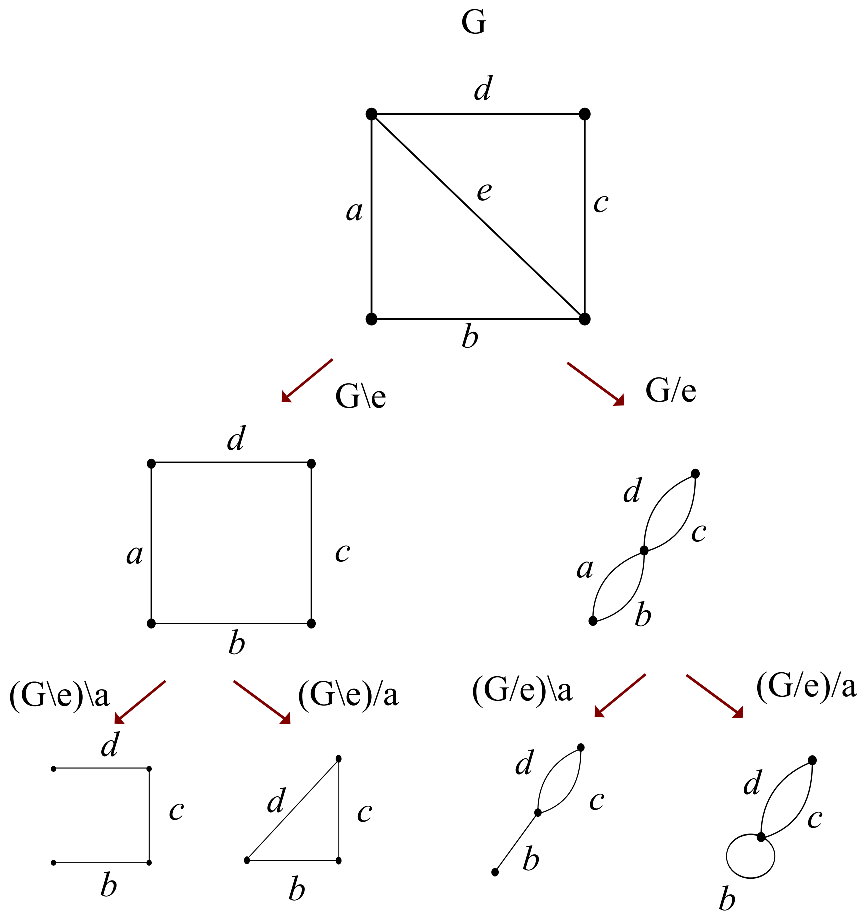

FIGURE H.1. Calculation of F by the method of deletion and contraction. [[LABEL F:0924Figure1]]

where a_{kl} is the coefficient of $u^k v^l$, that is, the number of subsets $S \subseteq E$ that have rank $k = c(S) - c(\Gamma)$ and nullity $l = |S| - n + c(S)$. Write

$$T_\Gamma(x, y) = \sum_{i,j \geq 0} b_{ij} x^i y^j.$$

The significant fact here is:

Proposition H.24. [[LABEL P:0924tuttecoefficients]] *All* $b_{ij} \geq 0$.

Proof. This can be proved by induction. **[NATE: Can you write a proof? OPTIONAL. If not, I'll do it.]** □

We deduce from the correspondence between the Tutte polynomial and the corank-nullity polynomial that

$$R_\Gamma(u, v) = T_\Gamma(u + 1, v + 1)$$

$$
\begin{aligned}
&= \sum_{i,j\geq 0} b_{ij}(u+1)^i(v+1)^j \\
&= \sum_{i,j\geq 0} b_{ij} \sum_k \binom{i}{k} u^k \sum_l \binom{j}{l} v^l \\
&= \sum_{i,l\geq 0} u^k v^k \underbrace{\sum_{i,j\geq 0} b_{ij}\binom{i}{k}\binom{j}{l}}_{\text{coefficients of } u^k v^k \text{ in (H.2)}} \\
&= \sum_{k,l} a_{kl} u^k v^l.
\end{aligned}
$$

This string of equalities shows that:

Proposition H.25. [[LABEL P:0924chromatic-cnp]] $a_{kl} = \sum_{i,j\geq 0} b_{ij}\binom{i}{k}\binom{j}{l}$. □

The proposition allows us to get good lower bounds for certain graph quantities by looking at the coefficients of the Tutte polynomial. In particular, we infer not only that $a_{kl} \geq 0$, but stronger positivity due to the fact that a_{kl} is a positive combination of nonnegative integers b_{ij}. [**TZ will add something here: For instance, ...**]

Properties of the Tutte polynomial.

Here are some significant properties of the Tutte polynomial, that we will not prove. A graph is said to be *separable* if it is not 2-connected or it has a loop. A *series-parallel graph* is a graph such that each block is derived from a single edge by repeatedly subdividing edges and adding parallel edges. Assuming $|E(\Gamma)| \geq 2$, we can say that:

- $b_{01} = b_{10}$.
- $b_{01} = 0 \iff \Gamma$ is separable.
- $b_{01} = 1 \iff \Gamma$ is a series-parallel graph.

Properties of the chromatic polynomial.

Now let's take a look at the chromatic polynomial. We define w_i, called the *Whitney numbers of the first kind* of Γ, to be the coefficients of powers of λ in the chromatic polynomial: $\chi_\Gamma(\lambda) = \sum_{i=0}^n w_i \lambda^{n-i}$. Now we can say that:

$$
\begin{aligned}
\sum_{k=0}^{n} (-1)^{n-k} w_{n-k}\lambda^k &= (-1)^n \chi_\Gamma(-\lambda) \\
&= T_\Gamma(1+\lambda, 0) \\
&= Q_\Gamma(\lambda, -1) \\
&= \sum_{i,j\geq 0} (1+\lambda)^i 0^j b_{ij} \\
&= \sum_i (1+\lambda)^i b_{i0} \\
&= \sum_k \lambda^k \sum_i \binom{i}{k} b_{i0}.
\end{aligned}
$$

Therefore, $w_{n-k} = (-1)^{n-k} \sum_i \binom{i}{k} b_{i0}$. The sum is nonnegative; thus we have the following theorem.

Theorem H.26. [[LABEL T:0924Theorem3 Alternating Sign Theorem]] *The Whitney numbers w_i alternate in sign, with $w_0 = 1$ and $(-1)^i w_i \geq 0$.*

This tells us that the coefficients of the chromatic polynomial alternate in sign. More can be said about the Whitney numbers with further study involving the Tutte polynomial, but we stop here.

Sept 26: Yash Lodha

I. Line Graphs

[[LABEL 1.lg]]

The *line graph* of Γ, denoted by $L(\Gamma)$, is defined as follows:

$$V(L(\Gamma)) = E(\Gamma),$$
$$E(L(\Gamma)) = \{ef \mid e, f \text{ are adjacent in } \Gamma\}.$$

(Recall that edges are *adjacent* when they have a common vertex.) This is the simple definition, valid for simple graphs Γ.

The definition of line graphs raises a few important questions regarding them. First of all, which graphs are line graphs? Secondly, are there graphs that are isomorphic to their line graphs? Thirdly, how many non isomorphic graphs can produce the same line graph? We now provide a few examples:

(1) $L(K_3) \cong K_3$.
(2) $L(K_{1,3}) \cong K_3$.

According to a theorem of Whitney's, these are the only two connected (simple) graphs that have the same line graph.

[I then go on to describe graphically what happens with double edges and loops with graphics.] **[THIS IS NEEDED!]**

Let Γ be a simple graph. Let B be the unoriented incidence matrix of Γ (defined in Section E), and let H be the oriented matrix of Γ. Then the entry $x_{i,j}, i \neq j$, of BB^{T} is the number of ij edges for vertices $i, j \in V(\Gamma)$ and entry $x_{i,i}$ of BB^{T} is the degree valency of the vertex i. It is clear from this that $BB^{\mathrm{T}} = D + A$ where D is the degree matrix or the diagonal $V \times V$ matrix with degrees of vertices in its diagonal entries and A is the adjacency matrix. The entry $x_{i,j}, i \neq j$ of HH^{T} is negative of the number of vw edges, and the entry $x_{i,i}$ of HH^{T} is the degree of the vertex i.

Theorem I.1. [[LABEL T:0926rge]] *If Γ is loopless and k-regular, then the largest eigenvalue of A is k, with multiplicity at least $c(\Gamma)$.*

The actual multiplicity is exactly $c(\Gamma)$, but I won't prove it. (It follows from the rank of the incidence matrix. I will provide a more general proof in Section II.M.)

Proof. Notice that HH^{T} is a *Gram matrix* (which is defined as a matrix G of inner products of vectors in $\mathbb{R}^n$, i.e., where $g_{i,j} = v_i \cdot v_j$, the dot product of vectors v_i, v_j). This is positive semidefinite, which means that it is symmetric and $\forall x \in Y$, $Ax \cdot x \geq 0$. So all eigenvalues are greater than or equal to zero.

Let x be an eigenvector of A with eigenvalue λ. Then $Ax = \lambda x$. And $\mathrm{HH}^{\mathrm{T}}x = kIx - Ax = (k-\lambda)x$. This implies that x is an eigenvector of HH^{T} with eigenvalue $k - \lambda$.

To show k is an eigenvalue with multiplicity greater than or equal to $c(\Gamma)$, suppose the components have vertex sets $V_1 = \{v_1, \dots, v_{n_1}\}$, $V_2 = \{v_{n_1+1}, \dots, v_{n_1+n_2}\}$, ... So $\pi(\Gamma) = \{V_1, V_2, \dots, V_{c(\Gamma)}\}$. Let $x_i \in \mathbb{R}^n$ be the vector which is 0 except for being 1 on every vertex of V_i. It is easy to see that $Ax_i = kx_i$. Therefore we have at least $c(\Gamma)$ independent eigenvectors, hence k has multiplicity at least $c(\Gamma)$. □

Now we look at $B^{\mathrm{T}}B$, which is an $E \times E$ matrix. In this matrix the entry $x_{i,j}$ is the number of edges between the vertices v_i, v_j, and $x_{i,i}$ is the degree of vertex v_i. It is clear that $B^{\mathrm{T}}B = A(L) + 2I$, where $L = L(\Gamma)$. Since $\mathrm{H}^{\mathrm{T}}\mathrm{H}$ is positive semidefinite, the eigenvalues are greater than or equal to zero.

Theorem I.2. [[LABEL T:0926lge]] *The eigenvalues of a line graph are greater than or equal to* -2*.*

Proof. Let λ be an eigenvalue of $A(L)$ with eigenvector x. Then $A(L)x = \lambda x$. Now

$$B^{\mathrm{T}}Bx = (A(L) + 2I)x = (\lambda + 2)x.$$

This implies that $\lambda + 2$ is an eigenvalue of $B^{\mathrm{T}}B$. So $\lambda \geq -2$. □

J. Cycles, Cuts and their Spaces

[[LABEL 1.cyclecut]]

J.1. The binary cycle and cut spaces. [[LABEL 1.binarycyclecut]]

The binary edge space.

Oct 3: Peter Cohen, TZ

In $\mathcal{P}(E)$ there is the operation of symmetric difference, or set addition, written $\oplus$. Under set summation $\mathcal{P}(E)$ is a binary vector space, that is, a vector space over the two-element field $\mathbb{F}_2$, indeed $\mathcal{P}(E) \cong \mathbb{F}_2^E$ in a natural way.

Definition J.1. The *characteristic function* of an edge set $S \subseteq E$ is $1_S : E \to \{0, 1\}$, where $1_S(e) = 1$ if the edge e is contained in S and 0 otherwise.

The correspondence $S \leftrightarrow 1_S$ is the natural isomorphism of $\mathcal{P}(E)$ with $\mathbb{F}_2^E$. In view of this correspondence we may, and do, regard any subspace of $\mathcal{P}(E)$ as a subspace of $\mathbb{F}_2^E$ and vice versa. This kind of switching back and forth between different viewpoints (in this case, sets vs. functions) is a powerful tool in all of mathematics. Still, one should not forget that it is two different kinds of objects that are being treated as equivalent.

In the vector space $\mathcal{P}(E)$ there is an inner product $S \cdot T := |S \cap T| \pmod 2$. It corresponds to the dot product in $\mathbb{F}_2^E$, defined by $x \cdot y := \sum_{i \in E} x_i y_i \in \mathbb{F}_2$. By that I mean $1_{S \cdot T} = 1_S \cdot 1_T$ in $\mathbb{F}_2$.

The binary cycle space.

Sept 15: Yash Lodha

The first essential subspace of the binary edge space is the cycle space.

Definition J.2. [[LABEL D:0926z1f2]] The *binary cycle space* is the subspace spanned by all circles (if regarded as lying in $\mathcal{P}(E)$) or all characteristic functions of circles (if in $\mathbb{F}_2^E$). A *binary cycle* is any element of the binary cycle space. We denote the binary cycle space by $Z_1(\Gamma; \mathbb{F}_2)$.

Proposition J.1. [[LABEL P:1003z1even]] *The binary cycles are the even-degree subsets of E.*

In fact, the real proposition is stronger; it has Proposition J.1 is an immediate corollary.

Proposition J.2. [[LABEL P:0926evencircles]] *Any even-degree edge set is the disjoint union of circles.*

Proof. I'll sketch the proof. In one direction, it's easy to see that a sum of any number of circles (disjoint or not), or any other sets each of which has even degree, will itself have even degree. In the other direction, one has to prove that any even-degree edge set S that is not empty contains a circle. This is a standard lemma of introductory graph theory. Deducting the circle C from S leaves a smaller even-degree edge set, disjoint from C, so the proposition follows by induction. □

I said at the beginning of the course that the nullity or cyclomatic number of Γ equals the number of independent circles. It is time to explain the exact meaning of that statement.

Definition J.3. [[LABEL D:0915fundcircles]] Given a maximal forest T of Γ, if we add another edge e we obtain a circle. This circle is called the *fundamental circle* associated with e, written $C_T(e)$. The entire set $\{C_T(e) \mid e \notin T\}$ is called the *fundamental system of circles* associated with T.

Proposition J.3. [[LABEL P:0915fundbasis]] *Given T, every circle is a set sum of fundamental circles in a unique way.*

That is, a fundamental system of circles is a basis of the binary cycle space. (Not every basis has this form.)

Proof. **[NEEDS PROOF]**

□

The binary cut space.

The second essential subspace of $\mathbb{F}_2^E$ is the binary cut space. It is dual to the binary cycle space. There are many kinds of duality in graph theory; this one is orthogonal duality in the binary edge space $\mathcal{P}(E)$, but it also corresponds to planar duality, although we won't treat that in these notes. Sept 26: Yash Lodha

Definition J.4. [[LABEL D:0926cutset]] A *cut* or *cutset* is the set of edges between a vertex set $X \subseteq V$ and its complement X^c (if this set is nonempty). A *bond* is a minimal cut.

For $X, Y \subseteq V$ we define $E(X, Y)$ to be the set of those edges with one endpoint in X and the other in Y. Thus, a cutset is any nonempty set $E(X, X^c)$. A particular case is when X is a singleton:

Definition J.5. [[LABEL D:1003vertexcut]] A *vertex cut* is the set of all edges incident to a vertex, i.e., $E(\{v\}, V \setminus v)$. A *vertex bond* is a vertex cut that is also a bond.

[WERE THERE EXAMPLES? They would be good here, to illustrate the possibilities.]

Definition J.6. [[LABEL D:1003binarycutspace]] The *binary cut space* of Γ, written $B^1(\Gamma; \mathbb{F}_2)$, is the set $\{\text{cuts}\} \cup \varnothing$ in $\mathcal{P}(E) \cong \mathbb{F}_2^E$.

Proposition J.4. [[LABEL P:1003cutbonds]] *Every cut is a disjoint union of bonds in a unique way.*

This is remarkably similar to its dual, Proposition J.2, but the uniqueness is a difference; it fails to hold true in Proposition J.2.

Proof. Consider the vertex sets X and X^c. Let $E(X, X^c)$ be the cutset defined above. For $e \in E(X, X^c)$, let $v_1 \in X$ and $v_2 \in X^c$ be the vertices incident on e. Then consider the component C_1 of the subgraph induced by X which contains v_1 and the component C_2 of the subgraph induced by X^c which contains v_2. Let $E(C_1, C_2)$ be the edge set with one endpoint in C_1 and the other in C_2. Now it is clear that $E(C_1, C_2)$ is a bond, since C_1 and C_2 are connected, removing a proper subset of $E(C_1, C_2)$ will leave $C_1 \cup C_2$ connected, and hence not increase the number of components of our graph.

From here it follows that $E(X, X^c)$ is the unique disjoint union of edge sets (which are bonds) connecting a pair of components of X and X^c. □

Properties of the binary spaces.

Oct 3 (draft): Peter Cohen

Here is a list of the principal properties of the binary cycle and cut spaces, other than those already mentioned.

Theorem J.5. [[LABEL T:1003binarycutscycles]]

(1) The binary cycle space $Z_1(\Gamma; \mathbb{F}_2)$ is a subspace of $\mathcal{P}(E)$.
(2) The binary cut space $B^1(\Gamma; \mathbb{F}_2)$ is a subspace of $\mathcal{P}(E)$.
(3) $B^1(\Gamma; \mathbb{F}_2)$ is orthogonal to $Z_1(\Gamma; \mathbb{F}_2)$.
(4) The binary cycle space is the span of the set of circles. Any fundamental system of circles is a basis of $Z_1(\Gamma; \mathbb{F}_2)$.
(5) The binary cut space is spanned by the bonds. In fact, it is spanned by the vertex bonds, hence by the vertex cuts.
(6) $B^1(\Gamma; \mathbb{F}_2)$ and $Z_1(\Gamma; \mathbb{F}_2)$ are orthogonal complements in $\mathbb{F}_2^E$.
(7) The sum of dimensions $\dim B^1 + \dim Z_1 = |E|$, the number of edges.
(8) $\dim Z_1 = |E| - n + c(\Gamma)$ = the cyclomatic number.

The proof is a homework exercise, or rather, a series of exercises. For instance, two of the key parts of the proof are to show that

(a) the intersection of a circle and a cut always has even cardinality, and
(b) the set sum of two different cuts, $E(X_1, X_1^c) \oplus (E(X_2, X_2^c)$, is a cut.

J.2. **The cycle and cut spaces over a field.** [[LABEL 1.fieldcyclecut]]
Now we let F be any field.

Definition J.7. A *directed circle* is a circle given a direction. This is separate from any orientation of the edges of the circle.

The indicator vector of a circle is a kind of signed characteristic function.

Definition J.8. The *indicator vector* (or *indicator function*), $I_C : E \to F$, is defined for a directed circle C. If an edge $e \in C$ has the same orientation as the circle, $I_C(e) = 1$. If $e \in C$ is oriented oppositely to the circle's direction, $I_C(e) = 1$. If $e \notin C$, then $I_C(e) = 0$.

Reversing the direction of the circle negates the indicator vector. Thus, for our purposes, it doesn't matter which direction C has; the important point is to distinguish oppositely oriented edges within C.

Definition J.9. The *F-cycle space* $Z_1(\Gamma; F)$ is the span of the indicator vectors of the circles in the graph Γ.

Definition J.10. A *directed cut*, denoted by $\vec{E}(X, Y)$, is a cut with a direction from X to Y.

Lemma J.6 (2)**.** [[LABEL L:1003 2]] *$I_C \cdot I_D = 0$ for any directed circle C and directed cut D.*

Proof. [**NEEDS PROOF**]

□

Theorem J.7. [[LABEL T:1003cyclescuts]]

(1) $B^1(\Gamma; F)$ = span of $\{I_B : B$ is a bond$\}$ = span of $\{I_B : B$ is a vertex cut$\}$ = span of $\{I_B : B$ is a vertex bond$\}$.
(2) $B^1(\Gamma; F)$ and $Z_1(\Gamma; F)$ are orthogonal complements in F^E.
(3) $\dim B^1 + \dim Z_1 = |E|$.
(4) $\dim B^1 = n - c(\Gamma)$.
(5) $\dim Z_1 = |E| - n + c(\Gamma)$.
(6) $B^1(\Gamma, F_2) = \mathrm{Row}\, \mathrm{H}(\Gamma)$, and $Z_1(\Gamma, F_2) = \mathrm{Nul}\, \mathrm{H}(\Gamma)$.

Proof. [**NEEDS PROOF**]

□

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.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 $\mathbb{Z}_2$ as the sign group, or $\{+1, -1\}$. 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 a signed graph with signature $\sigma' = \sigma_{E'}$. An edge subset $S \subseteq E$ makes a natural signed subgraph, (V, S, σ_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 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.

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

A.2. **Examples of Signed Graphs.** [[LABEL 2.examples]]

Several different signed graphs can be constructed from an ordinary graph Γ. Oct 22: Jackie Kaminski

- $+\Gamma = (\Gamma, +)$, where all edges are positive. We call it the all-positive Γ.
- $-\Gamma = (\Gamma, -)$, where all edges are negative. Unsurprisingly, we call it the all-negative Γ
- $\pm\Gamma = (+\Gamma) \cup (-\Gamma)$, 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 $E(\pm\Gamma) = +(E(\Gamma)) \cup -(E(\Gamma))$. 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.
- Σ° is Σ with a negative loop adjoined to every vertex (that doesn't have one)
- We can also define $+\Gamma^\circ$, $-\Gamma^\circ$, $\pm\Gamma^\circ$, where we think of doing the $+$, $-$, or $\pm$ before the $^\circ$ so that the final result has one negative loop at each vertex, not a positive loop or two loops.

Definition A.4. [[LABEL Df:1022full]] We say Σ is *full* if every vertex supports a half edge or a negative loop.

In other words, referring ahead to Section A.4, 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.

A.3. **Walk and circle signs.** [[LABEL 2.walksigns]]

The signs of walks and (especially) circles are fundamental in the subject of signed graphs. Oct 6: Jackie Kaminski

Definition A.5. [[LABEL D:1006 Walk Sign]] For any walk $W = e_1e_2\cdots e_l$, the sign of W is

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

If $\sigma(W) = +$ we call the walk *positive*; otherwise we call it *negative*. In particular, a circle is positive if its sign is $+$; otherwise it is 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 of $\mathcal{C}(|\Sigma|)$ is $\mathcal{B}^c(\Sigma) := \{\text{negative circles of } \Sigma\}$.

A.4. **Balance.** [[LABEL 2.balance]]

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

Definition A.7. [[LABEL D:1006 Balanced Graph]] We say Σ is *balanced* if

(1) every circle is positive and
(2) Σ has no half edges.

A subgraph is balanced if it is a balanced signed graph. We similarly define a balanced edge set; specifically:

Definition A.8. [[LABEL D:1006 Balanced Edge Set]] For $S \subseteq E$, we say S is *balanced* if $\Sigma|S$ is balanced, where $\Sigma|S := (V, S, \sigma|_S)$.

Note that the existence of loose edges has no effect on the state of balance of either a graph or an edge set.

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

Definition A.9. [[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.1 (Balance Theorem (Harary 1953a))**.** [[LABEL T:1006 Harary]] Σ *is balanced* $\iff$ *there is a bipartition* $V = V_1 \cupdot V_2$ *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.5). 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.2. [[LABEL L:1006 balanced blocks]] Σ *is balanced if and only if every block (maximal 2-connected subgraph) is balanced.*

We recall that a graph is 2-connected 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 being balanced means *every* circle in Σ is balanced, so any circles in a particular block are certainly balanced.

For the reverse direction, assume every block of Σ is balanced, and let C be a circle in Σ. It is well known that every circle is contained within a single block (since the blocks are *maximal* 2-connected subgraphs), so C is balanced since it is a circle in a balanced subgraph. Therefore Σ is balanced. $\square$

A.5. **Switching.** [[LABEL 2.switching]]

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

Definition A.10. 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.3. [[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.11. [[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.4 (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.3, 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 observe that every circle in Γ is the set sum (symmetric difference) of the fundamental circles of the non-tree edges of C, in other words $C = \bigcup_{e \in C \setminus T} C_T(e)$. (This is closely related to Proposition J.3.)

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.3, $\sigma_1^{\zeta_{n-1}}(C_T(f)) = \sigma_2(C_T(f))$, 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. □

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

Corollary A.5. [[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.4) $V = X_1 \cupdot X_2$ 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. □

(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.4.)

Oct 8: Simon Joyce

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.

Lemma A.6. [[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.*

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.7. [[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.8. [[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.* □

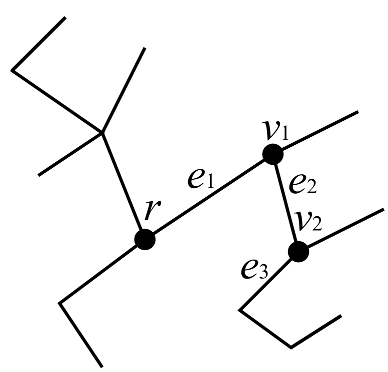

Figure A.1. F:1008 We want $\sigma^\zeta(e_i) = \tau(e_i)$.

Switching equivalence and switching isomorphism.

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

Definition A.12. [[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.9. [[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.1. 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.3.

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.9), 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. □

Oct 13:
Nate Reff

B. Characterizing Signed Graphs

[[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 $\sigma(S) = \sum_{e \in S} \sigma(e)$.

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

Oct 15 (draft): Peter Cohen and Thomas Zaslavsky

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

Proof. By induction, each part of the theorem implies another part.

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 $\bar\sigma = \tau|_e$. In this case, $U = nul\tau$, so $\bar\sigma^{-1}(0) = C \cap nul\tau$. Since $nul\tau$ is a subspace of Z, 2 implies 3.

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

2-1) Given that $\bar\sigma = \tau|_e$ **[Incomplete.]**

[WRITE A PROOF] □

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* (formerly called *circle additivity*) 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 $C_1 \oplus C_2$ of two circles with smaller values of ν.*

Proof. NEDS PROOF, NED.

□

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 $\mathcal{B}$. (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 $\bar\sigma(C) = \bar\sigma(C_1) + \bar\sigma(C_2)$.

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 oddity 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. □

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, $P = e_1 e_2 \ldots e_k$ (not containing any half edges), has a sign $\sigma(P) = \sigma(e_1)\sigma(e_2)\ldots\sigma(e_k)$. 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

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

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 $\{B_1, B_2, \ldots, B_k\}$ of disjoint, nonempty sets B_i such that $B_1 \cup B_2 \cup \ldots \cup B_k = V$. A *partial partition* of V is a partition of a subset of V; its *support* $\operatorname{supp}(\pi) := \bigcup \pi$ is that subset. (One should not 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 $\pi \cup \{\{0, 1, \ldots, n\} \setminus \operatorname{supp}(\pi)\}$ of $\{0, 1, \ldots, n\}$. (The block $\{0, 1, \ldots, n\} \setminus \operatorname{supp}(\pi)$ 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 that, given a maximal forest T of an 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 $Z_1(\Gamma; \mathbb{F}_2)$. In fact we can say that $C = \bigoplus_{e \in C \setminus E(T)} C_T(e)$ since we can rearrange the sum to correspond to a theta graph.

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 $Q_C = Q_{C_1} \cupdot Q_{C_2}$. 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 $Q_{C_1} \cupdot Q_{C_2} = Q_C$ 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.

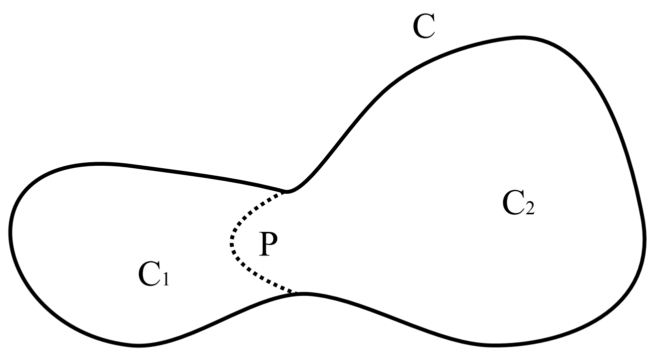

FIGURE C.1. F:1013Figure1

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 $Q_C = Q_{C_1} \cupdot Q_{C_2}$. 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 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]]

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.

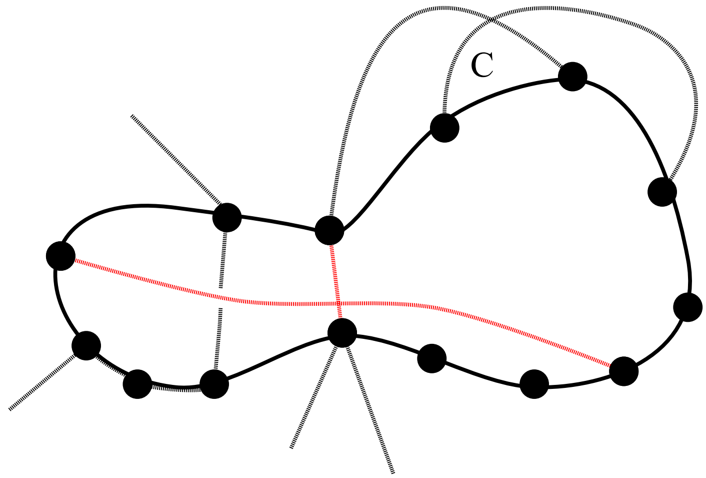

Figure C.2. F:1013Figure2

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.

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. $\square$

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. **Things that 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.

Dec 5b: Peter Cohen and Thomas Zaslavsky

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.

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

Proposition D.1. [[LABEL P:1205baledge]] *There are three types of partial balancing edge:*

(1) An edge e that is an isthmus between two components of $\Sigma \setminus e$, of which one (at least) is balanced. In this case, deleting e creates either one balanced component where previously there was an unbalanced component, or two balanced components where there was only one before. [[LABEL P:1205baledge isthmus]]

(2) A negative loop or half edge e in an otherwise balanced component. Then e is what prevents the component from being balanced. [[LABEL P:1205baledge loop]]

(3) A link $e{:}vw$ added to a component that is otherwise balanced, in such a way that a vw-path in the component has sign opposite to that of e. For example, a negative link in an otherwise all-positive component. [[LABEL P:1205baledge link]]

The proof is a good homework problem.

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.2. [[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.

Dec 8: Yash Lodha

Proposition D.3 (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. $\square$

Oct 22: Jackie Kaminski

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

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 D.3 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 Section D.3.

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 M. **[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], 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 (despite some early consideration in the psychology literature), Tomescu (1976a) and Popescu and Tomescu (1996a et al.) found things to say about it for signed complete graphs. By the way, evaluating the normalized (6) seems (to me) more interesting though (obviously) harder than (6) itself.

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.*

Proof. $\Sigma \setminus S$ 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.6 we can switch Σ so T is all positive. Then every edge not in S is positive, because its fundamental circuit 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. □

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]]

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 Σ = $\min_\zeta |E^-(\Sigma^\zeta)|$, 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 $d^-(v) \le \frac{1}{2}d(v)$ 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*

$$l(-\Gamma) = |E(\Gamma)| - |\textit{maximum cut of } \Gamma|.$$

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 $-\Gamma$.

Now notice that $\max_{X \subseteq V} |E(X, X^c)|$ will certainly 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 proves the proposition. □

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

Proof. The maximum-cut problem is already NP-hard, and "Is the max cut size $\le k$?" is NP-complete. (See any book on algorithmic complexity.) □

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.

Definition D.4. [[LABEL D:1022 D]] $l_{\max}(\Gamma) := \max_{\sigma: E \to \{+,-\}} l(\Gamma, \sigma)$, 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.13 (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 $-K_n$.

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 $i > \frac{n}{2}$, $\max_{i=0,1,\ldots,n} i(n-i) = \lfloor\frac{n}{2}\rfloor(n-\lfloor\frac{n}{2}\rfloor)$. 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 $\Sigma^- = K_{n/2} \,\dot\cup\, K_{n/2}$ 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(any signed forest) $= 0$, since it has no circles, and therefore $l_{\max}$(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.14. [[LABEL T: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 and tight bounds for all s, given a fixed value of r.

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

Problem D.1. [[LABEL Pr:maxfr]] Find necessary, sufficient, or necessary and sufficient conditions on a graph Γ for $l(-\Gamma)$ to equal the maximum frustration $l_{\max}(\Gamma)$.

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.15 (Slilaty (2007a))**.** [[LABEL T:1022 Slilaty]] Σ *has no vertex-disjoint negative circles if and only if one or more of the following are 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 contracting an edge set in a subgraph, so before we can discuss minors we must define contraction.

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.4 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, you 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 we have $V(\Sigma/e) = V(\Sigma) \setminus \{v\}$, 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 regular graphs. The process now differs for the balanced and unbalanced components of S. The edge set and vertex set of Σ/S will be as follows:

$$\begin{aligned} E(\Sigma/S) &:= E(\Sigma \setminus S), \\ V(\Sigma/S) &:= \{\text{vertex sets of balanced components of } (V,S) = \Sigma|S\} \\ &= \boldsymbol{\pi}_{\mathrm{b}}(S). \end{aligned}$$

To contract we first apply a switching function ζ so the balanced components of S are all positive. Lemma A.6 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.

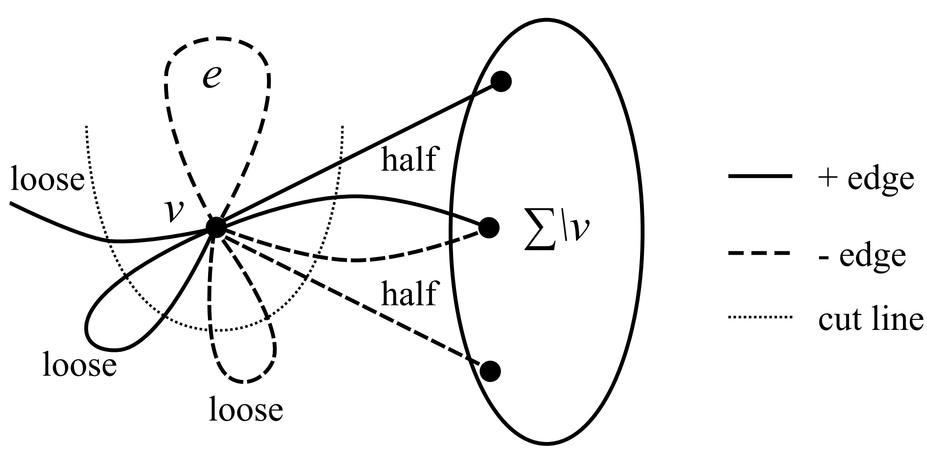

FIGURE E.1. Cutting out v leaves half and loose edges. [[LABEL F:1024Figure1]]

An example of how the contraction process 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. All the edges with an endpoint at the new contracted vertex become half edges and all the edges with endpoint in the unbalanced component become loose edges.

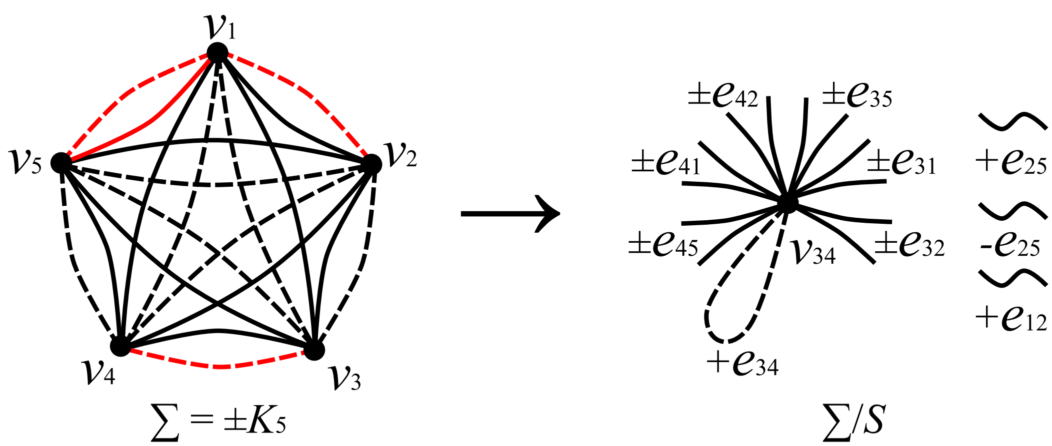

FIGURE E.2. A contraction of $\pm K_5$.
[[LABEL F:1024Figure2]]

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

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

(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.*

Proof. By theorem A.4, 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.7. 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 $\mathcal{B}(\Sigma^{\zeta_2}/S) = \mathcal{B}(\Sigma^{\zeta_2}/S)$.

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 an extension $\hat{\zeta}$ that is a switching function of Σ such that $\hat{\zeta}(v) = +$ for any $v \in V(S)$. 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.

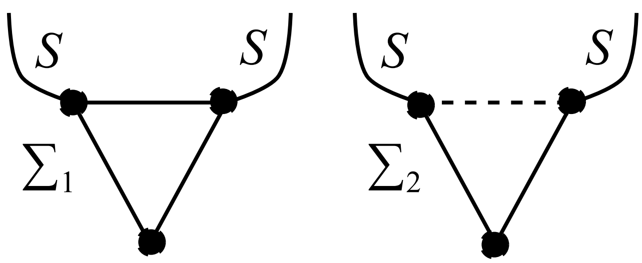

FIGURE E.3. Part (b) of Lemma E.2 fails for unbalanced S. [[LABEL F:1024Figure3]]

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

To summarise, by definition a minor of a signed graph Σ can be constructed as follows. First, delete all edges that are supposed to be deleted. Now all vertices to be deleted become isolated; delete these vertices. Finally, contract all edges that are supposed to be contracted. In short, a minor of a signed graph is defined as a contraction of a subgraph.

Theorem E.3. [[LABEL T:1024 minors are minors]] *Given a signed graph Σ, the result of any sequence of deletions and contractions of edge and vertex sets of Σ is a minor of Σ. In other words, a minor of a minor is a minor.*

[Proof of this is in the following day's notes.]

Oct 27: Nate Reff

Proof. $(\Sigma/S)/T = \Sigma/(S \cup T)$ where $S \cap T = \varnothing$. (Note that we have equality here in the sense of edges and not just an isomorphism; however in the sense of vertices we do not get this nice equality.) Notice that $V(\Sigma/S) = \boldsymbol{\pi}_{\mathrm{b}}(V, S)$. Also notice the following:

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

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

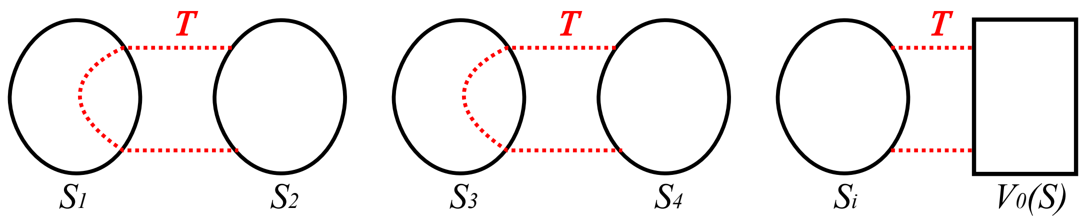

FIGURE E.4. If we throw in T what happens to the unbalanced and balanced components? [[LABEL F:1027Figure1]]

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.

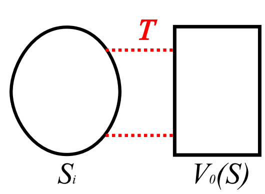

FIGURE E.5

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

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]]

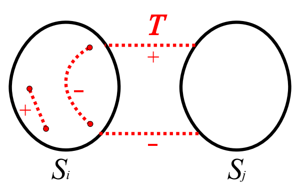

Figure E.6

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.*

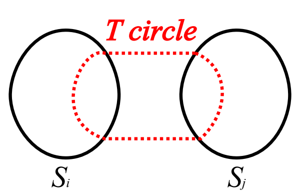

Figure E.7

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.

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 D.2 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

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

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

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

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

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

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

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 $\operatorname{bcl}(T) = \operatorname{clos}_{\Sigma^+}(T)$, 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 $\operatorname{clos}_\Gamma$ in Definition D.1, 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 $\operatorname{bcl} T$. That means $\operatorname{bcl}_\Sigma T = \operatorname{bcl}_{+\Sigma^+} T = \operatorname{clos}_{\Sigma^+} T$. □

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

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 $\boldsymbol{\pi}(S) = \{B_1, \dots, B_k, C_1, \dots, C_l\}$, 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 $\boldsymbol{\pi}_{\mathrm{b}}(S) = \{B_1, \dots, B_k\}$, 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,

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

Oct 29: 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]]***[FORMERLY P:1029closure]*** *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, $e \in E{:}V_0 \subseteq \operatorname{clos} S$. □

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

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

(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 $\operatorname{Lat} \Sigma$ is a lattice.

Proposition F.7. [[LABEL P:1029lattices]] $\operatorname{Lat} \Sigma$ *is a lattice with* $S \wedge T = S \cap T$*, and* $S \vee T = \operatorname{clos}(S \cup T)$.

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

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

Proof. □

In $\operatorname{Lat} \Sigma$ 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.

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* $\operatorname{supp} \pi$ is the union of the blocks. The set of all partial partitions of V is written Π_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 D.1) clearly agrees with the refinement ordering of partial partitions.

The set of partial partitions of $[n]$ is denoted by Π_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 Π_n° a poset. This poset is a geometric lattice. In fact:

Proposition F.8. [[LABEL P:1031ppartition]] $\Pi_n^{\circ} \cong \Pi_{n+1}$.

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

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

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

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

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 $\Pi_V^\circ(\{+,-\})$, or for short, $\Pi_V^\circ(\pm)$.

Signed partial partitions are partially ordered in the following way: $\theta \le \theta'$ if $\pi(\theta) \le \pi(\theta')$ 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 $\Pi_V^\circ(\{+,-\})$, or for short, $\Pi_V^\circ(\pm)$. In particular, the set of signed partial partitions of $[n]$ is written $\Pi_n^\circ(\pm)$. 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 $\{[B_1,\tau_1],\dots,[B_k,\tau_k],B_0\}$, where $\{B_1,\dots,B_k,B_0\}$ 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 $\Theta_\mathrm{b} : \operatorname{Lat}(\Sigma) \to \Pi_V^\circ(\pm)$, 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 IV [**GAINS CHAPTER**].) 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 $\Theta_\mathrm{b} : \mathcal{P}(E) \to \Pi_v^\circ(\pm)$, but it will not be an injection.

Now define

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

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

Theorem F.9. [[LABEL T:1031lattices]] $\Theta_\mathrm{b} : \operatorname{Lat}(\Sigma) \to \Pi^\circ(\Sigma)$ *is a poset isomorphism.*

Lemma F.10. [[LABEL L:1031ppartition]] $\Theta_\mathrm{b}(S) = \Theta_\mathrm{b}(\operatorname{clos}(S))$.

Proof. The partition $\pi(\Theta_{\mathrm{b}}(S))$ is unchanged by taking the closure: $\pi(\Theta_{\mathrm{b}}(S)) = \pi(\Theta_{\mathrm{b}}(\mathrm{clos}(S)))$ since $\boldsymbol{\pi}_{\mathrm{b}}(S) = \boldsymbol{\pi}_{\mathrm{b}}(\mathrm{clos}(S))$ 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 $\Theta_{\mathrm{b}}(S) = \Theta_{\mathrm{b}}(\mathrm{clos}(S))$. Hence the lemma is proved. ☐

Proof of Theorem F.9. The theorem follows easily from the lemma. ☐

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

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

To prove this we need to look at the flats of $\mathrm{Lat}(\pm K_n^{\circ})$. 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 Θ_{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 $\Pi_v^{\circ}(\pm)$ 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 $E(B{:}\,\mathrm{supp}(D_j))$ contains $E(A{:}\,\mathrm{supp}(C_i))$ so the signs associated with the vertices $\mathrm{supp}(D_j)$ must be switching equivalent to the signs associated with $\mathrm{supp}(D_j)$ restricted to $\mathrm{supp}(C_i)$. 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. ☐

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 A.8), that the balance-closure of S is

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

and the closure of S is

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

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

(1) $\pm K_n^\circ$ (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.

(2) $\pm\Gamma^\circ$ (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.

(3) 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.)

(4) $\pm K_n$ (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.)

(5) $\pm\Gamma$ (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.

(6) $+\Gamma$ (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.

(7) $+\Gamma^\circ$ (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 $+\Gamma^\circ$.

(8) $-\Gamma$ (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 $\Gamma{:}W$.

(9) $-K_n$ (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.

(10) $-\Gamma^\circ$ (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.

(11) $-K_n^\circ$ (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!

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 $\mathrm{H}(\Sigma) = (\eta_{ve})$, whose column indexed by e is:

$$\begin{array}{c} \\ \\ \\ v \\ \\ \\ \\ w \\ \\ \\ \\ \end{array} \begin{pmatrix} 0 \\ \vdots \\ 0 \\ \pm 1 \\ 0 \\ \vdots \\ 0 \\ \mp\sigma(e) \\ 0 \\ \vdots \\ 0 \end{pmatrix} \quad \text{for } e{:}vw \text{ a link,} \qquad \begin{array}{c} \\ \\ \\ v \\ \\ \\ \\ \end{array} \begin{pmatrix} 0 \\ \vdots \\ 0 \\ \pm 1 \mp \sigma(e) \\ 0 \\ \vdots \\ 0 \end{pmatrix} \quad \text{for } e{:}vv \text{ a loop,} \qquad \begin{array}{c} \\ \\ \\ v \\ \\ \\ \\ \end{array} \begin{pmatrix} 0 \\ \vdots \\ 0 \\ \pm 1 \\ 0 \\ \vdots \\ 0 \end{pmatrix} \quad \text{for a half edge } e{:}v,$$

and a zero column for a loose edge. Thus a link has two nonzero elements in 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 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.

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: Nov 5: Nate Reff

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

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 either of the following types of subgraph:

I. A positive circle or a loose edge.
II. 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.

The characteristic of a field $\mathbf{K}$ is denoted by $\operatorname{char}\mathbf{K}$.

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

(1) If $\operatorname{char}\mathbf{K} \neq 2$*, the columns corresponding to S are linearly dependent $\iff$ S contains a frame circuit.*
(2) If $\operatorname{char}\mathbf{K} = 2$*, the columns corresponding to S are linearly dependent $\iff$ S contains a circle or a loose edge.*

From a matroid perspective, this means the frame circuits are circuits of a matroid on a ground set E, and the incidence matrix represents the matroid. This is the *frame matroid* of Σ, sometimes called the *signed-graphic matroid* [SG] (and formerly called the *bias matroid* [BG2]).

Proof of sufficiency ($\Longleftarrow$). For (1) it suffices to prove that a frame circuit, is dependent. For (2) it suffices to prove that a circle or loose edge is dependent.

We write $\mathbf{c}_e :=$ the column of e in $\mathrm{H}(\Sigma)$ and $\mathbf{b}_i$ for the ith coordinate unit vector of $\mathbf{K}^n$.

Case I: $\{e\}$ is a loose edge. Then $\mathbf{c}_e = \mathbf{0}$ which is linearly dependent.

Case II: A circle $C = v_0e_1v_1e_2v_2\dots e_lv_l$ where $v_0 = v_l$. Then the submatrix of this is:

$$\begin{array}{cc}
 & \begin{array}{ccccccc} e_1 & e_2 & e_3 & e_4 & \dots & e_{l-1} & e_l \end{array} \\
\begin{array}{c} v_1 \\ v_2 \\ v_3 \\ \vdots \\ v_{l-1} \\ v_l \\ \vdots \\ \vdots \\ \vdots \end{array} &
\begin{bmatrix}
\pm1 & \pm1 & 0 & 0 & \dots & 0 & 0 \\
0 & \pm1 & \pm1 & 0 & \dots & 0 & 0 \\
0 & 0 & \pm1 & \pm1 & \dots & 0 & 0 \\
\vdots & 0 & \ddots & \ddots & \ddots & \vdots & \vdots \\
0 & 0 & \dots & \dots & 0 & \pm1 & \pm1 \\
\pm1 & 0 & \dots & \dots & 0 & 0 & \pm1 \\
0 & 0 & 0 & 0 & \dots & 0 & 0 \\
\vdots & 0 & \ddots & \ddots & \ddots & \vdots & \vdots \\
0 & 0 & 0 & 0 & \dots & 0 & 0
\end{bmatrix}
\end{array},$$

Look at the first two entries, $\eta_{v_1e_1}$ and $\eta_{v_1e_2}$. By negating columns we can make these opposite in sign. In other words, by negating columns we can ensure that $\eta_{ii} = -1$ for $i \in \{1, 2, \dots, l\}$ and $\eta_{i-1,i} = -\sigma(e_i)\eta_{ii} = \sigma(e_i)$ by the definition of the incidence matrix.

If we look at our circle as seen in Figure G.2 we get the following sum:

$$\begin{aligned}
&\mathbf{c}_{e_l} + \sigma(e_l)\mathbf{c}_{e_{l-1}} + \sigma(e_l)\sigma(e_{l-1})\mathbf{c}_{e_{l-2}} + \sigma(e_{l-2}e_{l-1}e_l)\mathbf{c}_{e_{l-3}} + \cdots \\
&+ \sigma(e_le_{l-1}e_{l-2}\dots e_{l-(i-1)})\mathbf{c}_{e_{l-i}} + \sigma(e_le_{l-1}e_{l-2}\dots e_{l-i})\mathbf{c}_{e_{l-i-1}} + \cdots .
\end{aligned}$$

By our definitions we notice the following patterns: **[This isn't very clear. Can you add explanation to the formulas?]**

$$\underbrace{\mathbf{c}_{e_l} + \sigma(e_l)\mathbf{c}_{e_{l-1}}}_{\text{cancels out row of } v_{l-1}} + \sigma(e_l)\sigma(e_{l-1})\mathbf{c}_{e_{l-2}} + \sigma(e_{l-2}e_{l-1}e_l)\mathbf{c}_{e_{l-3}} + \cdots$$

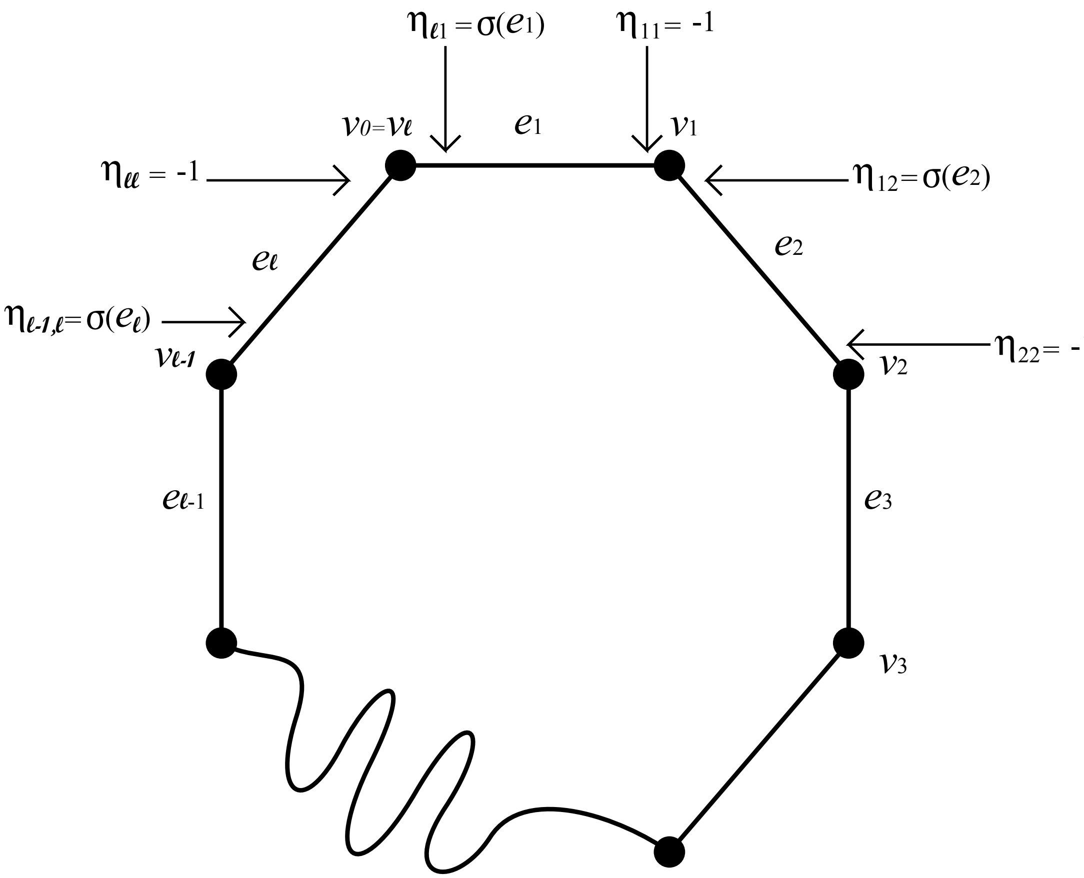

Figure G.1. The circle C.

$$+\,\sigma(e_l e_{l-1} e_{l-2} \dots e_{l-(i-1)})\mathbf{c}_{e_{l-i}} + \sigma(e_l e_{l-1} e_{l-2} \dots e_{l-i})\mathbf{c}_{e_{l-i-1}} + \cdots$$

and

$$\mathbf{c}_{e_l} + \underbrace{\sigma(e_l)\mathbf{c}_{e_{l-1}} + \sigma(e_l)\sigma(e_{l-1})\mathbf{c}_{e_{l-2}}}_{\text{cancels out row of } v_{l-2}} + \sigma(e_{l-2}e_{l-1}e_l)\mathbf{c}_{e_{l-3}} + \dots$$
$$+\,\sigma(e_l e_{l-1} e_{l-2} \dots e_{l-(i-1)})\mathbf{c}_{e_{l-i}} + \sigma(e_l e_{l-1} e_{l-2} \dots e_{l-i})\mathbf{c}_{e_{l-i-1}} + \cdots$$

And in general we see that at vertex v_{l-i-1}, $\mathbf{c}_{e_{l-i}}$ contributes $\sigma(e_l e_{l-1} \dots e_{l-i+1})\sigma(e_{l-i})$, and $\mathbf{c}_{e_{l-i-1}}$ contributes $\sigma(e_l e_{l-1} \dots e_{l-i})(-1)$. Together these sum to 0 in row $l-i-1$. This is valid for $l-1 \geq l-i-1 \geq 1$ where $0 \leq i \leq l-2$. So for rows 1 through $l-1$ we get a trivial sum [**WHAT is a trivial sum? I don't think it means zero!**], and for row l, $\mathbf{c}_{e_l}$ contributes -1 and $\mathbf{c}_{e_1}$ contributes $\sigma(e_l e_{l-1} \dots e_2)\sigma(e_2)$ which sum to $\sigma(C)-1$. Hence the vectors are linearly dependent if C is positive. They generate $2\mathbf{b}_l$ if C is negative and hence they are linearly dependent if $\operatorname{char}\mathbf{K} = 2$.

Case III: C is a loop. The conclusion is the same as in case II. ☐

Rephrasing the conclusions of Case II, if we have a closed walk $W = e_1 e_2 \dots e_l$ from v_k to v_k, then a suitable linear combination of vectors $\mathbf{c}_{e_1}, \mathbf{c}_{e_2}, \dots, \mathbf{c}_{e_l}$ equals $(\sigma(W)-1)\mathbf{b}_k$. The precise formula is

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

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

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

[This needs explanation/proof!]

Nov 7 (draft): Simon Joyce

Lemma G.3. [[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.*

[Known Proposition?? YES] A connected graph with no theta subgraph is a cactus, i.e., it is K_1 or every block is an isthmus or a circle.

Call a signed graph Σ *contrabalanced* if it contains no loose edges or balanced circles.

[this needs some attention] A *1-tree* is a tree with 1 extra edge on the same vertices. (Therefore either a half edge or framing a circle).

Lemma G.4. [[LABEL L:1107 circle]] *If S is connected and has minimum degree 2 or more and has cyclomatic number 2 or less, then it is a circle.*

Lemma G.5. [[LABEL L:1107 negcircle]] *A negative circle is independent.*

Theorem G.6. [[LABEL T:1107 dep rk clos]] *Given a signed graph Σ and $S \subseteq E(\Sigma)$. S may mean the columns of S in $\mathrm{H}(\Sigma)$.*

(1) S is linearly dependent if and only if it contains a frame circuit.
(2) S is linearly independent if and only if it is a tree or a negative 1-tree. (ie the circle, if any is negative.)
(3) S is intersection of column set $\{x_e : e \in E(Sigma)\}$ with a flat of $\mathbb{K}^n$ if and only if S is a closed edge set.
*(4) **[this line needs attention]*** $\mathrm{rk}(S)[:= n - b(S)] = \dim\{x_e : e \in S\} := \dim\langle x_e : e \in S\rangle$.
(5) $\mathrm{clos}(S) = \{e \in E(\Sigma) : x_e \in \langle x_f : f \in S\rangle\}$.

Corollary G.7. [[LABEL C:1107 matrix rank]] *The rank of $\mathrm{H}(\Sigma)$ is $n - b(\Sigma)$. Its nullity is $|E| - n + b(\Sigma)$ and that of $\mathrm{H}(\Sigma)^T$ is $b(\Sigma)$.*

Proof. **[NEEDS PROOF.]** □

There are two other important corollaries, which a reader who is not involved with matroids may ignore. Let us define 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 smalles such space is the zero space, generated by the empty set of columns, and the column space $\mathrm{Col}(M)$is the largest such space. It's well known that $\mathrm{Lat}(M)$ is a geometric lattice (in fact, that's where the name comes from).

Corollary G.8. [[LABEL C:1107 matriod 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,* $\Pi^\circ(\Sigma)$ *is a geometric lattice with rank function* $n - |\pi(\theta)|$*.* ***[notation??]***

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

[NEEDS MORE PROOF.]

□

Corollary G.9. [[LABEL C:1107 cor of cor]] *The set* $\Pi^\circ(\pm) := \{$*signed partial partions of* $[n]\}$ *is a geometric lattice.*

Proof. **[NEEDS PROOF.]**

□

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

We will now discuss the adjacency matrix $A(\Sigma)$ of a signed graph Σ. $A(\Sigma) = (a_{ij})_{n\times n}$, where

a_{ij} = (number of positive v_iv_j edges) - (number of negative v_iv_j edges) $(i \neq j)$.
a_{ii} = 2(number of negative loops) + (number of half edges).

$$A + \mathrm{H}\mathrm{H}^T = D(\Sigma), \text{ the degree matrix}$$

$$\begin{aligned} \text{Diagonal and } d_{ii} &:= \text{net degree}(v_i) = d^{\pm}(v_i) \\ &:= (\text{number of half edges}) + 2(\text{number positive loops}) \\ &\quad - 2(\text{number of negative loops}) + (\text{number of positive links}) \\ &\quad - (\text{number of negative links}). \end{aligned}$$

[WHY DO WE HAVE A'? I don't remember. – TZ] Nov 10 (draft): Peter Cohen

The definition of $A'(\Sigma)$ is a'_{ij} = (number of positive links between vertices i and j) $-$ (number of negative links). When v_i and v_j are the same vertex, a'_{ij} = 2(number of positive loops $-$ number of negative loops), as each loop has two orientations because in our formal definition we can distinguish its two ends from each other (although not by their incident vertex).

The degree matrix, denoted by $D(|\Sigma|)$, is the diagonal matrix with $d_{ii} = d_{|\Sigma|}(v_i)$. Note that a loop counts twice.

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. **[WE NEED A BETTER FORMULATED PROOF.]** $\mathrm{H}\mathrm{H}^{\mathrm{T}} = D(|\Sigma|) - A$ for each (i,j)-entry where $i \neq j$. $\mathrm{H}\mathrm{H}^{\mathrm{T}}$ will yield a matrix of dot products, specifically, row i with row j, where the kth column of H represents e_k. We will denote row i by η_{ik}, and row j by η_{jk}. $\eta_{ik}\eta_{jk} \neq 0$ if and only if e_k is an edge connecting v_i and v_j; it is $-\sigma(e_k)$, the sign of the edge. □

Adding up all the signs of the links, $+1$ for a positive link and -1 for a negative link, gives a'_{ij}.

For the case where $i = j$, each edge e_k contributes $+1$ if $\eta_{ik} = \pm 1$, that is, e_k is a link or half edge. e_k will contribute 4 if $\eta_{ik} = \pm 2$, that is, e_k is a negative loop. e_k contributes 0 if it is a positive loop.

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, $D - A = \mathrm{H}\mathrm{H}^{\mathrm{T}}$ 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 **??** [**TZ: THERE'S A MISSING THEOREM!**], 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 $D - A = \mathrm{H}\mathrm{H}^{\mathrm{T}}$, $k - \lambda \geq 0$ for every eigenvalue λ and the eigenvalue $\lambda = k$, corresponding to the eigenvalue 0 of $D - A$, has multiplicity $\operatorname{nul} \mathrm{H}^{\mathrm{T}} = b(\Sigma)$. ☐

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 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** $\mathrm{H}(B)$**.]**

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 $\mathrm{H}(\mathrm{B}) = (\eta_{ik})$ 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 $e = v_i v_j$ has i-th entry τ_{ik} and j-th entry τ_{jk}, and the remaining entries are zero. For a loop $v_i v_i$ 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,

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

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

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.*

I. Equations from Edges, and Signed Graphic Hyperplane Arrangements

[[LABEL 2.equations]]

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

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

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

For a signed edge $x_i = \sigma(e)x_j$ because from $\mathbf{x} \cdot \mathbf{c}_e = 0$ 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 $h_0 = \mathbb{R}^n$, 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 Corollary G.8. The exact formula is a matter of tracing the correspondences. ☐

11/14: Kaminski

Nov 17 (draft): Simon Joyce

I.1. **Binary, affine additive representations.** For this section we will be working over $\mathbb{F}_2$, whose additive group is $\cong \{+.-\}$.

Definition I.1. [[LABEL D:1117 AGIM]] Given a signed **[bidirected?]** graph Σ define the *augmented graphic incidence matrix*, $M(\Sigma)$ to be as follows.

$$M = \begin{array}{c} \\ \\ v_1 \\ \vdots \\ v_n \end{array} \left(\begin{array}{c} \text{edges} \rightarrow \\ \text{edge signs} 0, 1 \\ \text{H}(|\Sigma|) \\ \text{Incidence matrix} \\ \text{of} |\Sigma| \end{array} \right) \begin{array}{c} \\ x_0 \\ x_1 \\ \vdots \\ x_n \end{array}$$

[This looks awful. Still haven't figured out a way to make it look nice.]

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

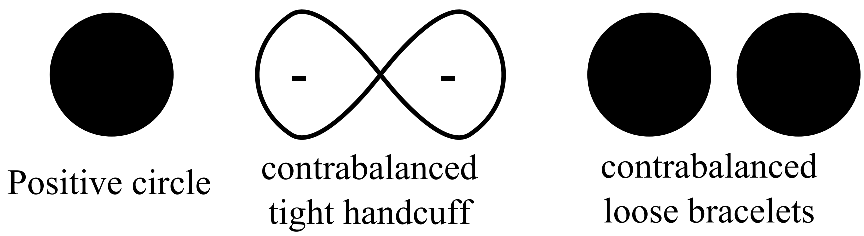

FIGURE I.1. The different kind of lift circuits.

[[LABEL F:1117liftcircuits]]

Theorem I.2. [[LABEL T:1117 lift circuits]] *Given M be the augmented graphic incidence matrix of a graph Σ, a set of columns of M is linear dependent if and only if the corresponding edge set contains a lift circuit.*

Lemma I.3. [[LABEL :1117 lemma1]] *The columns of M corresponding to $S \subseteq E$ generate b_0 if and only if S is unbalanced.* ***[not sure what this is saying]***

Lemma I.4. [[LABEL L:1117 lemma2]] *Given M be the augmented graphic incidence matrix of a graph Σ, row 0 of M is a linear combination of the other rows if and only if Sigma is balanced.*

For a signed graph Σ, denote by $\mathcal{A}[\Sigma]$ the *affinographic hyperplane arrangement* of Σ over $\mathbb{F}_2$.

$$e : vw \leftrightarrow \begin{bmatrix} \sigma(e) \\ 0 \\ 1 \\ 0 \\ 1 \\ 0 \end{bmatrix} \begin{array}{l} \leftrightarrow \text{Equation } x_j - x_i = \sigma(e)x_0 \text{ in } \mathbb{F}_2^n + 1. \\ \leftrightarrow \text{Linear hyperplane } \overline{h}_{ij}^{\sigma(e)} \text{ in } \mathbb{F}_2^{n+1}. \\ \leftrightarrow \text{Affine hyperplane } h_{ij}^{\sigma(e)} \text{ in } \mathbb{A}^n(\mathbb{F}_2). \end{array}$$

[an example goes here]

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

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

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

$$\mathcal{L}(\mathcal{A}[\Sigma]) \cong lat^b\Sigma$$

J. Chromatic Functions

Given a signed graph Σ with vertex set V, a k-coloration is a map γ where $\gamma : V \to \Lambda_k^*$ or $\gamma : V \to \Lambda_k$ where $\Lambda_k^* = \{\pm 1, \pm 2, \ldots, \pm k\}$ and $\Lambda_k = \Lambda_k^* \cup \{0\}$.

We call an edge $e : vw$ in Γ *proper* if $\gamma(v) \neq 0$ and$\gamma(w) \not= \sigma(e)\gamma(v)$. Call an edge $e : vw$ *improper* if $\gamma(w) = \sigma(e)\gamma(v)$. We also consider loose edges to be improper. [**the wording of this should be looked at**].

Nov 19: Nate Reff

K. Chromatic Functions

[[LABEL 2.chromatic]]

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

K.1. **Coloring a signed graph.** [[LABEL 2.coloring]]

Suppose that $\Sigma = (V, E, \sigma)$ is a signed graph.

Definition K.1. [[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

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

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. We write $\chi(\Sigma) :=$ minimum k such that there exists a proper coloration, and $\chi^*(\Sigma) :=$ minimum k such that there exists a zero-free proper coloration. If $\chi(\Sigma) = \infty$ (or $\chi^*(\Sigma) = \infty$) then there does not exist a proper coloration (zero-free coloration) at all.

Consider the example of a signed graph Σ in Figure K.1. There clearly does not exist a proper 0-coloration. There is, however, a proper 1-coloration as seen in Figure K.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 K.1. Therefore $\chi^*(\Sigma) = 2$.

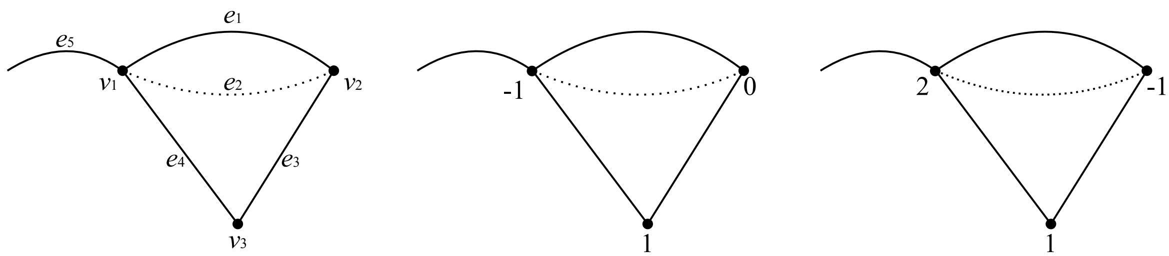

FIGURE K.1. Signed graph Σ, a proper 1-coloration of Σ, a proper zero-free 2-coloration of Σ.

[[LABEL 1119image1]]

K.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,

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

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 (K.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 K.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 $\Sigma \iff \gamma(w) \neq \sigma(e)\gamma(v) \iff \zeta(v)\zeta(v)\zeta(w)\gamma(w) \neq \zeta(v)\zeta(v)\zeta(w)\sigma(e)\gamma(v) \iff \zeta(w)\gamma(w) \neq \zeta(v)\sigma(e)\zeta(w)\zeta(v)\gamma(v) \iff \gamma^\zeta(w) \neq \sigma^\zeta(e)\gamma^\zeta(v) \iff e$ 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 $\Sigma \iff \gamma(v) \neq 0 \iff \zeta(v)\gamma(v) \neq 0 \iff \gamma^\zeta(v) \neq 0 \iff e$ is proper in Σ^ζ with γ^ζ. $\square$

Proposition K.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 K.1. □

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

Unitarity:

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

Multiplicativity:

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

and

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

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

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

Switching Invariance:

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

for every switching function ζ.

Deletion-Contraction:

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

and

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

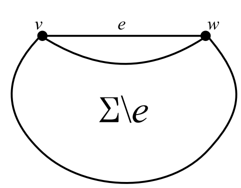

FIGURE K.2
[[LABEL 1119image2]]

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

The hard part is to prove the deletion-contraction property. To prove $\chi_\Sigma(2k+1) = \chi_{\Sigma\backslash e}(2k+1) - \chi_{\Sigma/e}(2k+1)$, we start by coloring $\Sigma\backslash e$ 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 $\Sigma\backslash e$ where $\gamma(v) = \gamma(w)$. So the number of proper colorations of $\Sigma\backslash e$ equals the sum of the number of proper colorations of Σ and Σ/e. Therefore our formula follows. ☐

Nov 21 (draft): Peter Cohen and Thomas Zaslavsky

In coloring a signed graph, a question arises about the case of the color 0. Intuitively, 0 would represent a blank or neutral color, and is treated as any other color in unsigned coloration. However, the 0 color can't be signed, and while it can be included in signed coloration, it has only 1 possible coloration (neutral), and can limit the graph in that respect. Some colorations will include the 0 color, while other colorations, which we will call "zero-free", do not use the color 0. **[CLARIFY]**

Definition K.3. [[LABEL T:1121full]] A graph is *full* if every vertex supports at least one unbalanced edge. We denote a signed graph, Σ, to be full by $\Sigma^\bullet$. **[FIX. Also, should this appear earlier? Does it?]**

Theorem K.4. [[LABEL T:1121dczero-free]] $\chi^*_\Sigma(\lambda) = \chi^*_{\Sigma\backslash e}(\lambda) - \chi^*_{\Sigma/e}(\lambda)$, *where* λ *is the number of colors.* ***[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 coloring of Σ. Vertex v has color $= 0 \iff$ it has a proper coloring of $\Sigma \neq \Sigma \setminus e$. **[FIX?]** ☐

In the zero-free case,

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

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

Lemma K.6. [[LABEL T:1121fulldelete]] $\Sigma^\bullet \setminus e$ *is full* $\iff$ e *is not an unbalanced edge.*

Proof. If e is an unbalanced edge, deleting e would make $\Sigma^\bullet$ no longer full. In the case where e is the only balanced edge in an otherwise full graph, deleting e would result in the graph being full. ☐

Therefore, if e is not an unbalanced edge,

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

Theorem K.7 (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 K.8 (Subset Expansion)**.** [[LABEL T:1121chromaticsubset]] *The chromatic polynomials have the subset expansions*

$$\chi_\Sigma^{[*]}(\lambda) = \sum_{\substack{S\subseteq E \\ \text{[balanced]}}} (-1)^{|S|} \lambda^{b(S)}.$$

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 find it to cite —— good point – TZ]

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 or negative loops, and so Σ 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;

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

□

Definition K.4. [[LABEL T:1121unbalcount]] The number of unbalanced components of a graph (or subgraph) is $u(\Gamma)$. This equals the number of components minus the number of balanced components; $u(\Gamma) = c(\Gamma) - b(\Gamma)$.

Thus, the chromatic polynomial of a signed graph is

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

and the zero-free chromatic polynomial is

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

We can define a comprehensive chromatic polynomial, which I call the *total chromatic polynomial*, as

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

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

Nov 24: Jackie Kaminski

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

We now take up the generalization to signed graphs of Stanley's theorem, Theorem H.19 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 H.17 and H.18) to a sesquijection between acyclic orientations of Σ and those of $\Sigma \setminus e$ and Σ/e.

Definition K.5. 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 K.9. [[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.[2] 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.

[2]If the circuit is a handcuff circuit, then this walk will simply repeat the circuit path.

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 $\alpha \cup e{:}\vec{vw}$ and $\alpha \cup e{:}\vec{wv}$, but that α "extended" to Σ/e contains an oriented cycle.

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.[3] 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

[3] 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 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

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 K.10. [[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 $\vec{\Sigma/e}$.*

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]]

□

11/26: Joyce

Dec 1: Nate Reff

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 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 K.3 is an example.

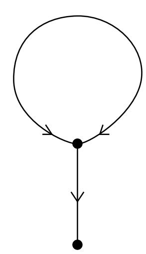

FIGURE K.3
[[LABEL F:1201image1]]

Lemma K.11. [[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.*

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 K.4. (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.

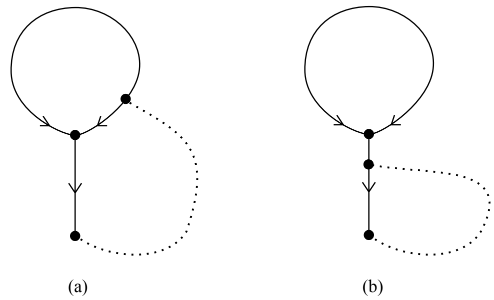

Figure K.4
[[LABEL F:1201image2]]

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 K.5. By Q_w we mean

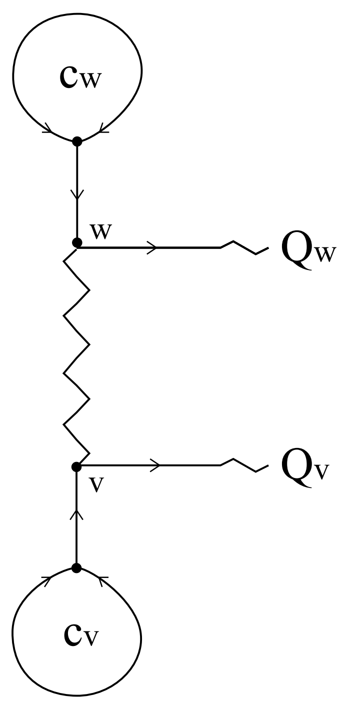

Figure K.5
[[LABEL F:1201image3]]

the part of Q we can get to by backtracking coherently along Q from w. If Q_w hits P_w, then lemma K.11 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.

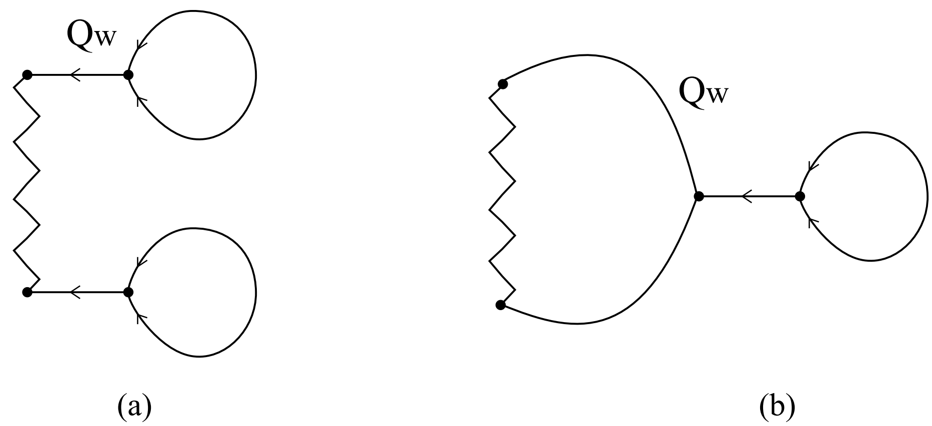

Figure K.6
[[LABEL F:1201image4]]

This means that Q must be one of the shapes seen in Figure K.6.
If Q_w does not hit P_w, then $P_w \cup Q_w$ is a cycle.

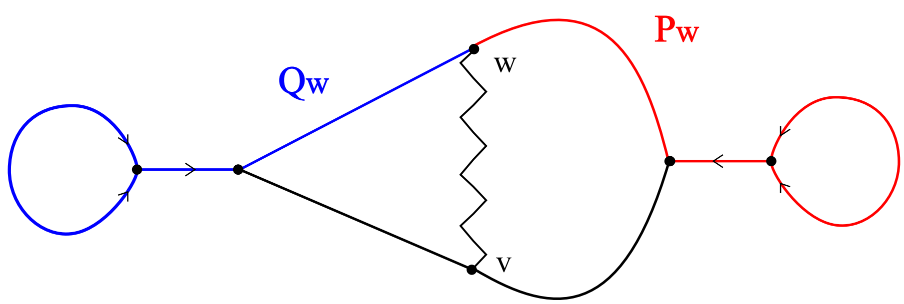

Figure K.7
[[LABEL F:1201image5]]

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 K.7, 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 K.11. 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.

Corollary K.12. [[LABEL C:1201CorollaryToCaseII]] *The union of two coherent balloons, not necessarily internally disjoint, that are joined coherently at their tips, contains a cycle.*

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]]

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 K.8).

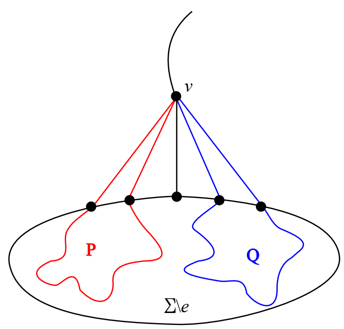

FIGURE K.8
[[LABEL F:1201image6]]

Therefore, P and Q are balloons and meet coherently. Apply Corollary K.12. Thus $\Sigma \setminus e$ has a cycle at v.

This concludes the proof of the Sesquijection Lemma. $\square$

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 K.9 defined in the following table.

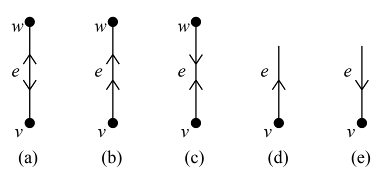

FIGURE K.9. An edge e which is (a) extraverted, (b) positive, (c) introverted, (d) an introverted half edge, (e) an extraverted half edge.
[[LABEL F:1201image7]]

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 H.19) 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 K.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 K.13. [[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

Dec 5 (draft): Peter Cohen and Thomas Zaslavsky

K.5. **The dichromatic and corank-nullity polynomials.** [[LABEL 2.dichromatic]]

The algebraic form of the chromatic polynomials, i.e., the subset expansion in Theorem K.8, 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 K.6. [[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 K.8).

Theorem K.14 (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
(K.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,

(K.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 (K.3), we have

$$\begin{aligned} Q_\Sigma(u,v,z) - Q_{\Sigma\setminus e}(u,v,z) &= v^{-n}\sum_{S\subseteq E: e\in S} (uv)^{b_\Sigma(S)} v^{|S|} z^{u_\Sigma(S)} \\ &= v^{-n}\sum_{T\subseteq E\setminus e} (uv)^{b_\Sigma(T)} v^{|T|+1} z^{u_\Sigma(T)} \\ &= v\cdot v^{-n}\sum_{T\subseteq E\setminus e} (uv)^{b_{\Sigma\setminus e}(T)} v^{|T|} z^{u_{\Sigma\setminus e}(T)} \\ &= vQ_{\Sigma\setminus e}(u,v,z) = vQ_{\Sigma/e}(u,v,z). \end{aligned}$$

Therefore, $Q_\Sigma = Q_{\Sigma\setminus e} + vQ_{\Sigma/e}$ 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 (K.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 (K.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 K.7. [[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 K.15 (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 K.14) and Proposition D.3.

[We need details here! WHAT IS THE PROP?]

□

Theorem K.16 (Theorem QRM). [[LABEL T:1208QRM]] $Q_\Sigma(u, v, z)$ *and* $R_\Sigma(u, v, z)$ *satisfy the following identities.*

(M) *Multiplicativity:*
$Q_{\Sigma_1 \,\dot\cup\, \Sigma_2} = Q_{\Sigma_1} Q_{\Sigma_2}$,
$R_{\Sigma_1 \,\dot\cup\, \Sigma_2} = R_{\Sigma_1} R_{\Sigma_2}$,
$Q_{\Sigma_1 \cup_v \Sigma_2} = Q_{\Sigma_1} Q_{\Sigma_2}$.

(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:*
$\Sigma_1 \cong \Sigma_2 \implies Q_{\Sigma_1} = Q_{\Sigma_2}$ *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.H.5 for guidance. ☐

K.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 K.17. [[LABEL T:1208allcolorations]]

$$X^{[*]}_\Sigma(k, w) = (-1)^{b(\Sigma)} (w-1)^n Q^{[*]}_\Sigma\Big(\frac{-\lambda}{w-1}, w-1\Big)$$

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 K.18 (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 K.19 (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

Dec 12: Nate Reff

L. 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.

L.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 L.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.1: 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).
(3) Assign l distinct labels from $[k]$ to the B_i's (there are $(k)_l = l!\binom{k}{l}$ ways to do this).

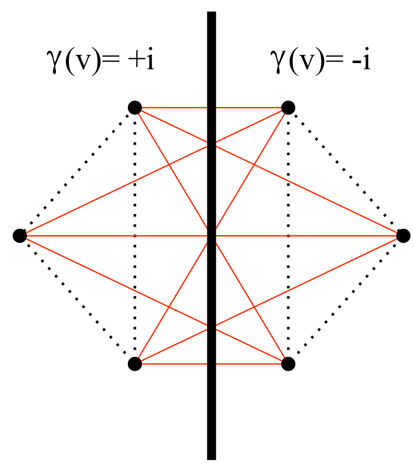

Figure L.1. Assigning signs to the vertices of $\gamma^{-1}(\pm i)$ in a signed K_n. The diagram shows the case in which there are 6 vertices colored $\pm i$. The positive edges (in red) are complete bipartite.[**NATE, NOTE: Figure uses too much white space!**]

[[LABEL 1212image1]]

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 L.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 L.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 L.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 L.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. □

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.

L.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:

$$\begin{aligned}\mathcal{T}(K_\Gamma) &:= \{ \text{ vertex sets of negative triangles } \} \\ &= \{ \text{ triples of vertices that support an odd number of edges in } \Gamma \text{ } \}.\end{aligned}$$

Lemma L.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 L.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 L.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 L.8. [[LABEL L:1212tga]] *Any adjacency matrix* A *of a two-graph* $\mathcal{T}$ *satisfies*

$$[[\text{LABEL E}:1212\text{tga}]]A^2 = (n-1)I + (n-2)A - 2(\sigma_{ij}t_{ij})_{ij}, \tag{L.1}$$

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 (L.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 (L.1). □

Proposition L.9. [[LABEL P:1212rtga]] *Any adjacency matrix A of a regular two-graph with t triples on each pair of vertices satisfies*

(L.2) $$[[\text{LABEL E : 1212rtga}]]A^2 = (n-1)I + (n-2-2t)A.$$

Conversely, if some adjacency matrix of a two-graph $\mathcal{T}$ satisfies a quadratic equation, then it satisfies (L.2) *and $\mathcal{T}$ is regular with t triples on each vertex pair.*

Proof. The first part is direct from Lemma L.8. The second part follows from comparing the presumed quadratic equation $A^2 = \beta I + \alpha A$ with (L.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 (L.2), this constant is t. □

Theorem L.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 (L.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 L.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) $\mu_1\rho_1 + \mu_2\rho_2 = \mathrm{tr}(A) = 0$.

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

$$\mu_1, \mu_2 = \frac{n}{2}\Big(1 \mp \frac{n-2-2t}{\sqrt{(n-2t)^2+8t}}\Big) = \frac{n}{2\sqrt{\Delta}}(\sqrt{\Delta} \mp \alpha).$$

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:

$$\text{[[LABEL E : 1212q]]}\, q(t-r) = r^2 - 2rt + 2t + t^2 = (t-r)^2 + 2t. \tag{L.3}$$

We conclude that either $t = r$ or

$$q = t - r + \frac{2t}{t-r}.$$

If $t = r$ then (L.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,

$$t = r + s,$$

$$q = s + 2 + \frac{2r}{s},$$

$$n = s + 2 + 2r + \frac{2r}{s}.$$

The eigenvalues are

$$\rho_1, \rho_2 = \frac{\alpha \pm q}{2} = \frac{1}{2}\left[\frac{2r - s^2}{s} \pm \Big(s + 2 + \frac{2r}{s}\Big)\right] = \begin{cases} 1 + \dfrac{2r}{s}, \\ -(s+1) \end{cases}$$

(the upper value is ρ_1, the lower is ρ_2) and the multiplicities are

$$\mu_1, \mu_2 = \frac{n}{2q}(q \pm \alpha) = \frac{n}{2}\ \frac{q \mp (n-2t-2)}{q} = \begin{cases} (s+1)\Big(1 + \dfrac{2rs}{s^2+2(r+s)}\Big), \\ 1 + \dfrac{2r}{s} + \dfrac{2r(s+2r)}{s^2+2(r+s)}. \end{cases}$$

We can therefore express n and t (the *parameters* of $\mathcal{T}$) 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.

2008 Dec 12: Zaslavsky

L.3. **Strongly regular graphs.** [[LABEL 2.twographs.srg]]

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 SRG(n, k, λ, μ), 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 $\mathcal{T}$ 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 L.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 L.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 $\Gamma = K_1 \mathbin{\dot\cup} C_5$, in other words, the pentagon (naturally) with an extra isolated vertex. It's clear from Proposition L.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 (L.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 L.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 L.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 L.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 L.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 L.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

(L.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 (L.4) not a polynomial in A. We use Equation (L.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 L.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

M. 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.

M.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 M.1 shows a graph Γ and its line graph $L(\Gamma)$.

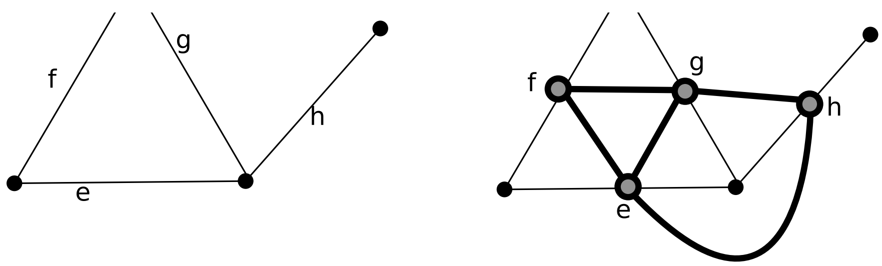

FIGURE M.1. A simple graph Γ, and Γ with its line graph $L(\Gamma)$ superimposed in heavy lines.

[[LABEL F:0129 Line Graph]]

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 M.2 which illustrates that if e, f are two parallel edges in Γ, then they are adjacent twice, which is reflected in $L(\Gamma)$ as the two edges between vertices e, f. Loops make things very messy, which is why we are restricting our attention to link graphs.

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 I.2, 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.

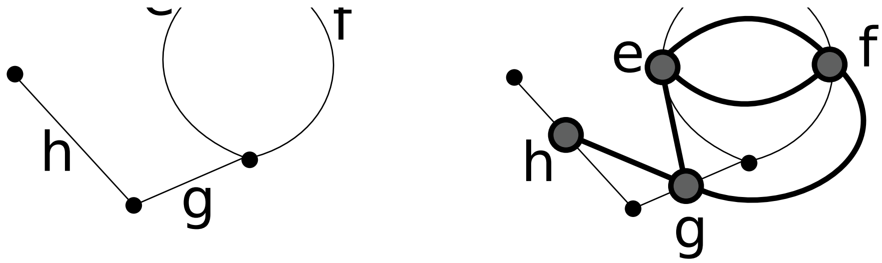

FIGURE M.2. A graph Γ with parallel edges, and Γ with its line graph $L(\Gamma)$ in heavy lines.
[[LABEL F:0129 Line Graph Parallel]]

M.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[4] (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 M.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 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.

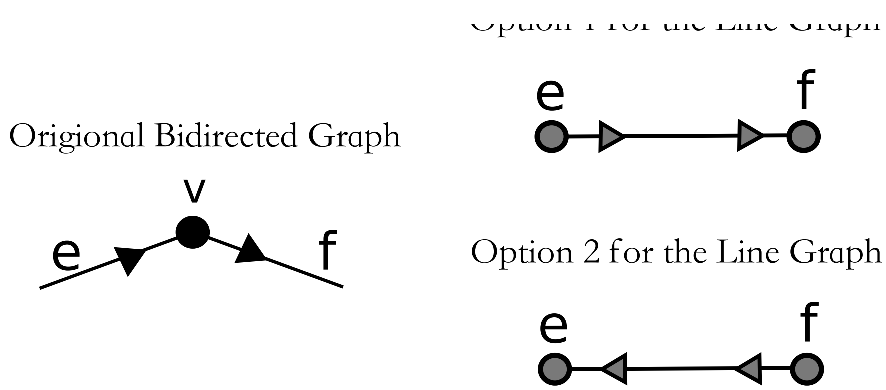

Figure M.3. Creating a bidirected line graph from a bidirected graph.
[[LABEL F:0129 Line Graph 2]]

[4]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 $\vec{\Sigma}$.

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 M.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.

2/3: Lodha

2009 Feb 5: Nate Reff

N. CIRCUITS, COCIRCUITS, AND THEIR SPACES

[[LABEL 2.cyclescircuitsspaces]]

N.1. **Unsigned graphs.** [[LABEL 2.]]

N.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).

N.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).

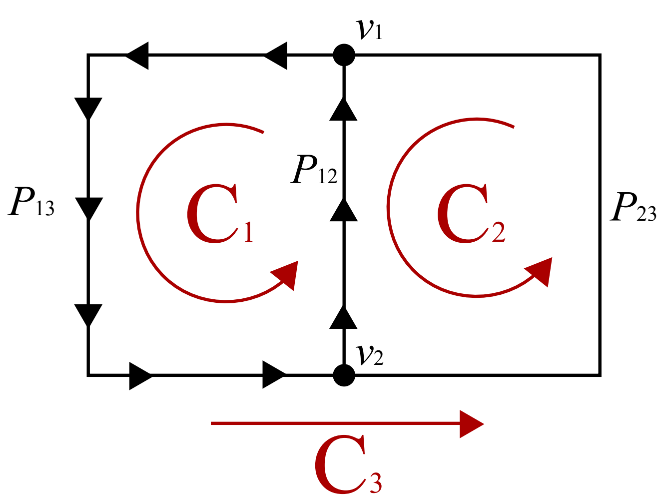

FIGURE N.1
[[LABEL 0205image1]]

Consider the theta graph in figure N.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

$$\begin{aligned} I_{\vec{C}_1} &= I_{P_{13}} - I_{P_{12}}, \\ I_{\vec{C}_2} &= I_{P_{12}} - I_{P_{23}}, \\ I_{\vec{C}_3} &= I_{P_{13}} - I_{P_{23}}. \end{aligned}$$

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.) $\square$

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 $\mathbf{K}$.

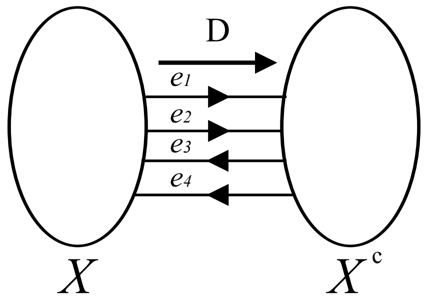

FIGURE N.2
[[LABEL 0205image2]]

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 N.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 N.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.

N.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

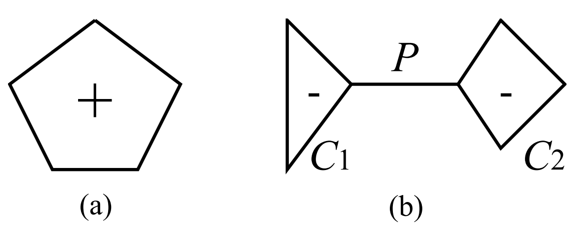

FIGURE N.3. (a) Positive circle, or Type I circuit, (b) Handcuffs, or Types II and III circuits.

[[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:

$$I_{\vec{C}}(e) = \begin{cases} \pm 1 & \text{if } e \in \vec{C} \text{ and } e \text{ is not in a connecting path } \vec{C}, \\ \pm 2 & \text{if } e \in \vec{C} \text{ and } e \text{ is in a connecting path } \vec{C}, \\ 0 & \text{if } e \notin \vec{C}. \end{cases}$$

2009 Feb 10 (draft): Simon Joyce

CYCLES AND CUTS (CONTINUED)

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 $\vec{W}$.

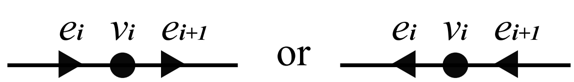

FIGURE N.4. F:0210 The two kinds of coherent edges you could have at $v_i \in \vec{W}$.

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.

$\tau_{\vec{\Sigma}} = \tau$ orients Σ where τ is a bidirection. We can think of τ as a map where τ : {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 $\tau_{\vec{W}} = -\tau_{\overleftarrow{W}}$.

N.3. **Indicator vector of $\vec{W}$.**

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 N.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 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.

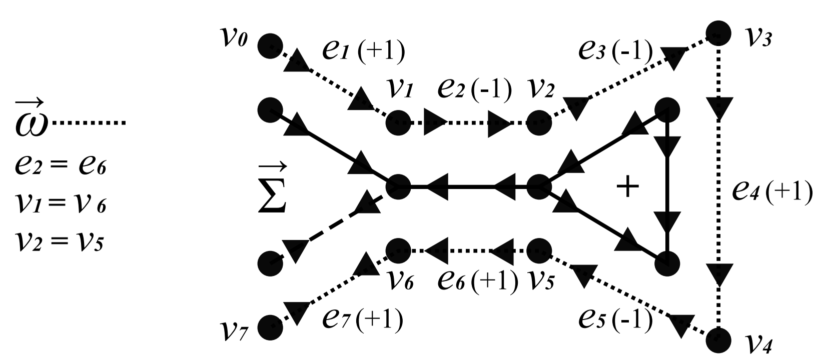

FIGURE N.5. F:0210 Edge signs of $\vec{W}$ for $\vec{\Sigma}$.

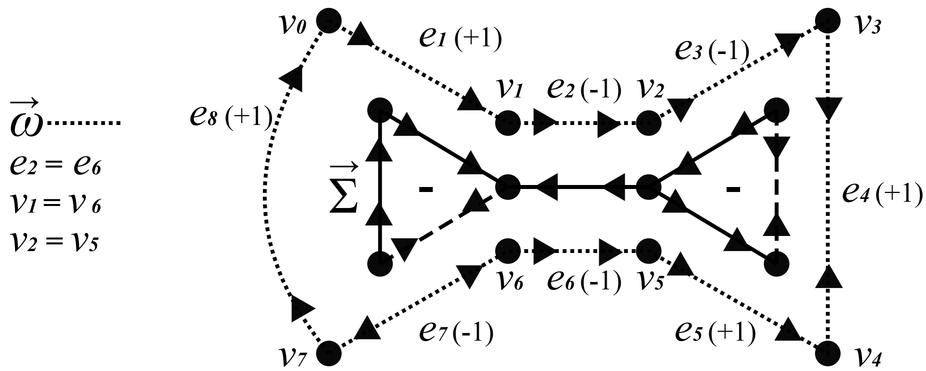

FIGURE N.6. F:0210 Indicator vectors on the graph $\vec{\Sigma}$.

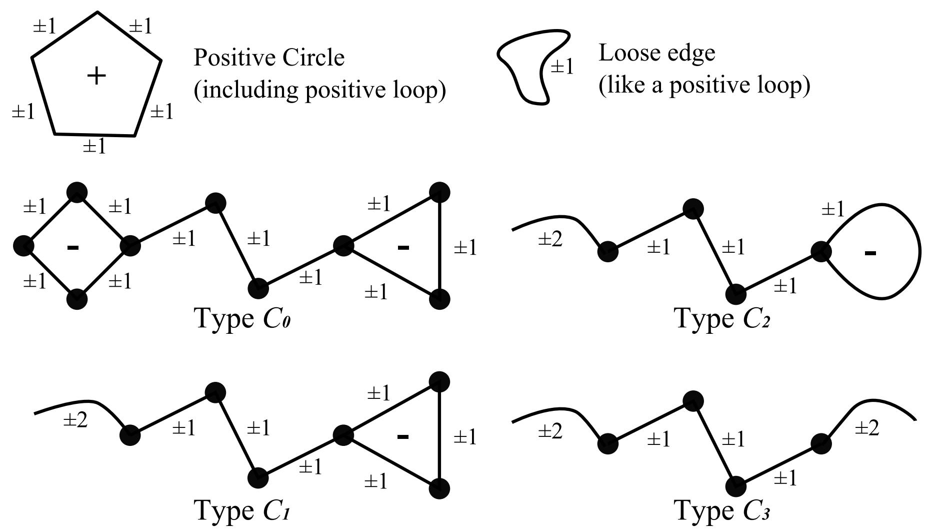

FIGURE N.7. F:0210 Indicator vectors on frame circuits.

2009 Feb 12: Jackie Kaminski

N.4. **Flows and Cycles.** [[LABEL 2.cyclespaces]]

N.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 N.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 N.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 N.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.

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.

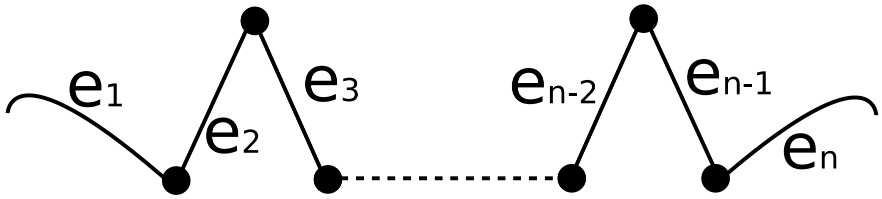

Figure N.8. A circuit in $\vec{\Sigma}$, with orientations omitted.
[[LABEL F:0212 Circuit with Half Edges]]

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 N.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 N.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.

N.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 N.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 $\Sigma{:}X$.

In Figure N.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

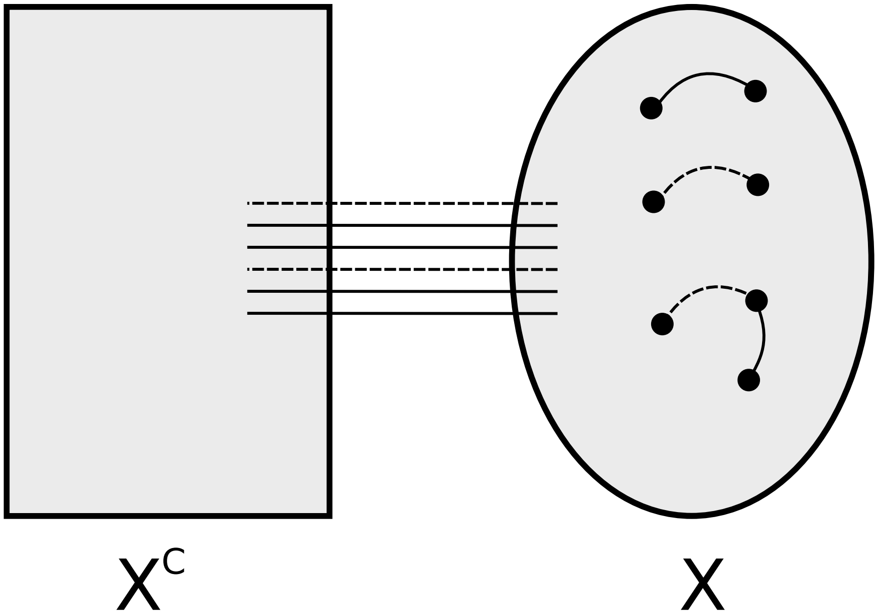

FIGURE N.9. A Cut is Σ
[[LABEL F:0212 cut]]

the cut, because it changes which set we need to balance, and consequently where the edges in U_X are taken from.

Definition N.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 N.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.

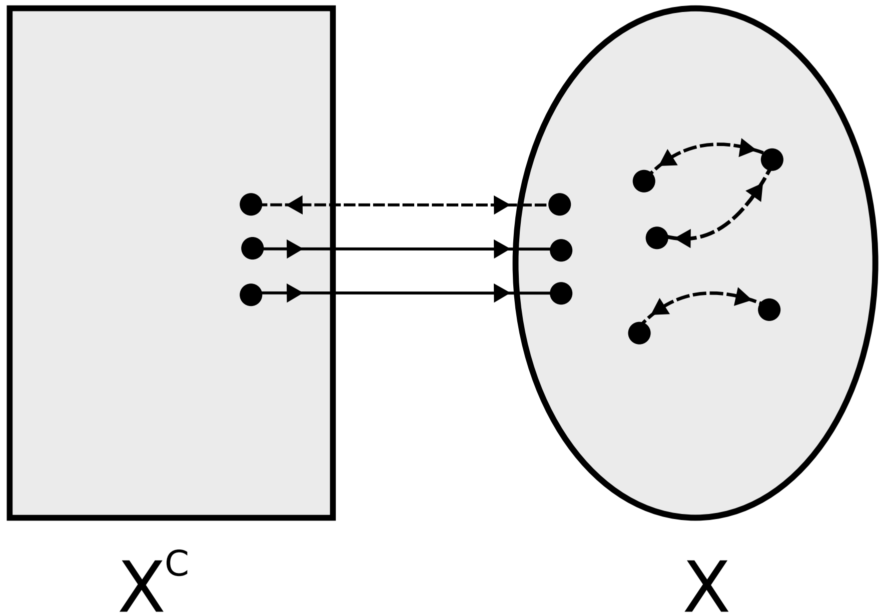

Figure N.10. A directed cut in Σ; notice that $(\Sigma{:}X) \setminus U_X$ is balanced [[LABEL F:0212 directed cut]]

Figure N.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 N.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 $I_{\vec{U}}(e)$.

2009 Feb 19: Nate Reff

N.6. **The three types of cut.**

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 N.11 for an illustration of the general form of a cut.

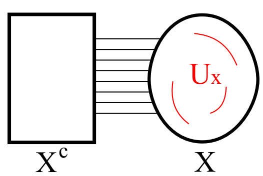

FIGURE N.11. A typical cut in Σ
[[LABEL F:0219image1]]

Here is an easy but important lemma.

Lemma N.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.**] □

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 N.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 N.13.

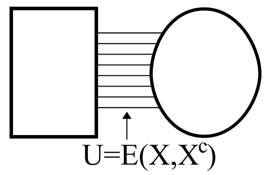

FIGURE N.12. A cut of Type I in Σ.
[[LABEL F:0219image2]]

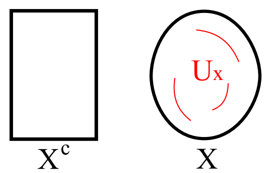

FIGURE N.13. A cut of Type II in Σ.
[[LABEL F:0219image3]]

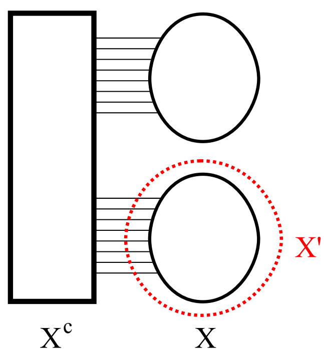

FIGURE N.14. Type I, not a bond.
[[LABEL F:0219image4]]

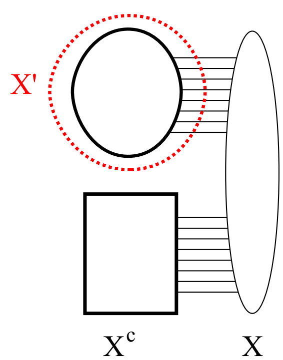

FIGURE N.15. Type I, not a bond.
[[LABEL F:0219image5]]

In Figures N.14, N.15, N.16 and N.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**

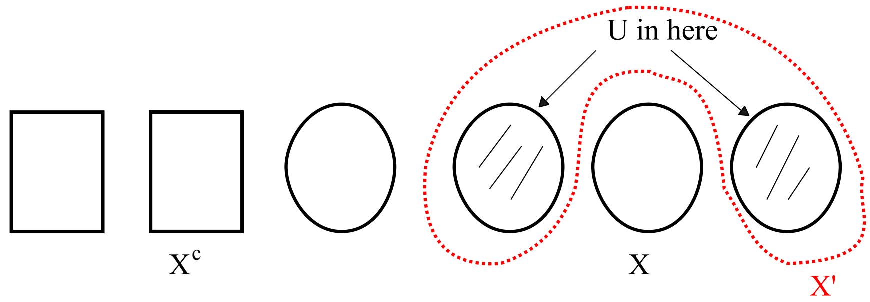

FIGURE N.16. Type II.
[[LABEL F:0219image6]]

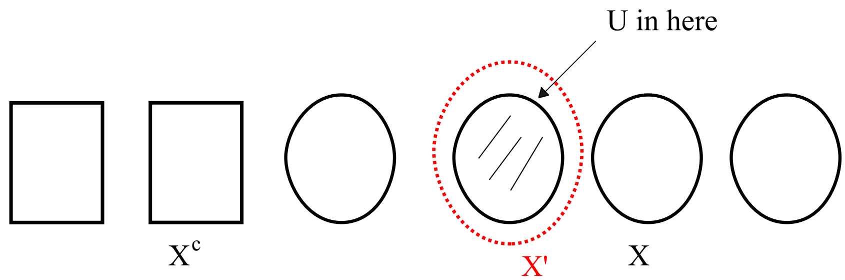

FIGURE N.17. Type II.
[[LABEL F:0219image7]]

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 N.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 N.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 N.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 N.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 N.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

N.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 N.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 N.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 N.10. [[LABEL P:0224]] $Z_1(\Sigma; \mathbf{K}) = \operatorname{Nul} \mathrm{H}(\Sigma)$.

Proof. By Lemma N.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 N.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 N.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:

$$\begin{gathered}
I_{C_1} \text{ is such that } I_{C_1}(e_1) \neq 0,\ I_{C_1}(e_2) = 0, \ldots, \\
I_{C_2} \text{ is such that } I_{C_2}(e_2) \neq 0,\ I_{C_2}(e_1) = 0, \ldots, \\
I_{C_3} \text{ is such that } I_{C_3}(e_3) \neq 0,\ I_{C_3}(e_1) = 0, \ldots, \\
\cdots \\
I_{C_m} \text{ is such that } I_{C_m}(e_m) \neq 0,\ I_{C_m}(e_1) = 0, \ldots.
\end{gathered}$$

These vectors are linearly independent and they span $\operatorname{Nul} \mathrm{H}(\Sigma)$. Therefore, $\operatorname{Nul} \mathrm{H}(\Sigma) \supseteq Z(\Sigma; \mathbf{K})$. □

N.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 N.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 N.13. [[LABEL L:0224BinB1]] $B \subseteq B^1$.

Proof. **[PROOF?]** □

Theorem N.14. [[LABEL T:0224B1]] $B = B^1$ *over a field* $\mathbf{K}$.

Proof. Exercise. Possibly a dimension argument. **[PROOF?]** □

Theorem N.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. $B^1 = \mathrm{Row}\,\mathrm{H}$, $Z_1 = \mathrm{Nul}\,\mathrm{H}$. □

Cuts and minimal cuts.

Now here are some contrasting facts.

For graph cuts: The set sum of graph cuts is a graph cut (or $\varnothing$). 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 N.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.)

Chapter III. Geometry

A. Geometrical Fundamentals

2009 Feb 26 (draft): Simon Joyce

[[LABEL 3.geometry]]

For the geometry of signed graphs, and later for gain graphs in Chapter IV, we have to understand vector (or linear), affine, and projective spaces themselves, especially those related to Euclidean space, and we need to know about point sets, positive span in real space, and real and complex hyperplane arrangements.

A.1. **Linear, affine, and projective spaces.** [[LABEL 3.spaces]]

Three kinds of geometry will be most important to us:

(1) $\mathbf{K}^n$: Linear/Vector spaces over a field (or division ring) $\mathbf{K}$.
(2) $\mathbb{A}^n(\mathbf{K})$: Affine spaces.
(3) $\mathbb{P}^n(\mathbf{K})$: Projective spaces.

Although one can do our kinds of geometry over a division ring, this leads to some annoying complexities, so we will always assume in this chapter that $\mathbf{K}$ is a field. Linear spaces are familiar to everyone, but affine and projective spaces are less so. There is one such space of each dimension for each field (or division ring); we write them $\mathbb{A}^n(\mathbf{K})$ and $\mathbb{P}^n(\mathbf{K})$. These will sometimes be shortened to just $\mathbb{A}^n$ and $\mathbb{P}^n$ where $\mathbf{K}$ is clear from context.

There are also other affine and projective geometries, which are defined axiomatically, without the use of coordinates. The coordinatizable geometries $\mathbb{A}^n$ and $\mathbb{P}^n$ (for which I reserve the word "spaces") are special cases—though not so terribly special, as non-coordinatizable geometries exist only in dimensions 1 and 2!

I'll begin with the geometry of a vector space coordinatized by a field; then I'll show one way to construct affine space from it. The next step is to construct projective space from affine space; but this is purely synthetic—independent of coordinates—so it applies to all projective and affine geometries.

The linear space $\mathbf{K}^n$. [[LABEL 3.linear]]

In the n-dimensional vector space $\mathbf{K}^n$, all subspaces contain the point $\mathbf{0}$. (They are called *homogeneous subspaces* because their equations have no constant term.) Lines are 1-dimensional subspaces, planes are 2-dimensional subspaces and hyperplanes are $n-1$-dimensional subspaces, or codimension 1 subspaces. In the lattice of flats of $\mathbf{K}^n$, all lattice points are subspaces. This lattice is graded by dimension. If a subspace has height k in the lattice—that is, dimension k—we call it a flat of *rank* k.

Construction of $\mathbb{A}^n(\mathbf{K})$ *from* $\mathbf{K}^n$. [[LABEL 3.lintoaff]]

The affine n-space $\mathbb{A}^n(\mathbf{K})$ has the same set of points but allows more flats than $\mathbf{K}^n$. Here all translates of linear subspaces are flats, but rank (in the lattice of flats) is now equal to dimension less 1. The points (0-dimensional subspaces, flats of rank 1) are the translates of $\mathbf{0}$. Two distinct points generate a line, which has rank 2. (A line is said to have dimension 1 because a point on a line has one degree of freedom. It has rank 2 because it takes 2 points to determine a line.) Any three distinct points that are not collinear generate a plane, which has dimension 2 and rank 3. $\varnothing$ is a flat in $\mathbb{A}^n(\mathbf{K})$. It has rank 0 or dimension -1.

Construction of $\mathbb{P}^n(\mathbf{K})$ *from* $\mathbb{A}^n(\mathbf{K})$. [[LABEL 3.afftoproj]]

This construction, unlike the two preceding ones, is synthetic. That means it can be carried out starting with any abstract affine geometry. (Since I haven't defined such things, we won't be using them, but it's worth remembering that coordinates are not necessary for this part of geometry.)

In the affine space $\mathbb{A}^n(\mathbf{K})$ we have parallel classes of lines in the familiar Euclidean way. We can define parallelism of lines in affine space of any dimension: affine lines l_1 and l_2 are called *parallel* if $l_1 \cap l_2 = \varnothing$ and there is an affine plane (that is, a translate of a 2-dimensional linear subspace) that contains both lines. Now, for each parallel class P, create a point $P^\infty \notin \mathbb{A}^n(\mathbf{K})$, and adjoin P^∞ to every line in P; the resulting set, $l_\mathbb{P} := l \cup P^\infty$, is called the *projective line* generated by the affine line l. These points are called *points at infinity* or *ideal points*.

The number of ideal points, when $\mathbf{K}$ is finite, is $\frac{|\mathbf{K}^n|-1}{|\mathbf{K}|-1}$. This is equal to the number of lines through a fixed point, since each parallel class has exactly one representative through each point of the affine geometry.

For a flat $f \subseteq \mathbb{A}^n(\mathbf{K})$, define

$$f^\infty := \{P^\infty : \exists\, l \subseteq f \text{ such that } l \in P\}$$

and

$$f_\mathbb{P} := \bigcup_{l \subseteq f} l_\mathbb{P} = f \cup f^\infty.$$

(According to our notation, then, l^∞ and P^∞, where $l \in P$, are two names for the same thing: the infinite point on the projective line l, which is also the infinite point of every line parallel to l.)

The points of $\mathbb{P}^n$ are the points of $\mathbb{A}^n$ and all the ideal points P^∞ of all parallel classes P. The flats of $\mathbb{P}^n$:

- $\varnothing_\mathbb{P} = \varnothing$ (rank 0);
- the rank-1 flats (points) are $p_\mathbb{P} = p$ for each affine point p and P^∞ for each parallel class P;
- the rank-2 flats (lines) are $l_\mathbb{P}$, one for each affine line, and for each affine plane (2-dimensional, rank 3) π the set π^∞ of all infinite points of parallel classes that contain a line in π;
- the rank-k flats for any k are the sets $f_\mathbb{P}$ where f is a rank-k flat in $\mathbb{A}^n$ and the sets g^∞ where g is an affine flat of rank $k+1$.

The new points constitute a new hyperplane, the set of all ideal points: $h_\infty := \{P^\infty : P \text{ is a parallel class of lines in } \mathbb{A}^n\}$. For any flat f we have $\dim(f_\mathbb{P} \cap h_\infty) = \dim(f) - 1$. Not only is h_∞ a hyperplane of $\mathbb{A}^n(\mathbf{K})$, but if f is any flat such that $\dim f \geq 0$, then $f^\mathbb{P} \cap h_\infty$ is a hyperplane of $f^\mathbb{P}$.

Conversely, given any hyperplane h, $\mathbb{P}^n \backslash h$ is an affine geometry. (This fact doesn't depend on coordinates.) Then adding back h_∞ (constructed as before) to $\mathbb{P}^n \backslash h$ gives $\mathbb{P}^n$ with $h_\infty = h$. Not necessarily true: $\mathbb{P}^n \backslash h$ is independent of choice of h up to isomorphism. [**that sentence needs attention**] It is true if we have a projective space because all $\mathbb{P}^n(\mathbf{K}) \setminus h \cong \mathbb{A}^n(\mathbf{K})$.

Construction of $\mathbb{P}^n(\mathbf{K})$ *from* $\mathbf{K}^{n+1}$. [[LABEL 3.lintoproj]]

This construction is analytic (it uses coordinates in $\mathbf{K}$).

The points of $\mathbb{P}^n(\mathbf{K})$ are defined to be the lines in $\mathbf{K}^{n+1}$. The projective flats are defined to be the sets

$$\bar{s} = \{l : l \text{ is a point of } \mathbb{P}^n(\mathbf{K}) \text{ and } l \subseteq s\}$$

for each subspace s of $\mathbf{K}^{n+1}$. For example, $\overline{\{\mathbf{0}\}} = \varnothing$, and $\bar{l} = \{l\}$ for each homogeneous line. The dimension of a flat is defined to be $\dim_{\mathbb{P}}(\bar{s}) := \dim(s) - 1$, so that a point is a 0-dimensional flat (in contrast to its linear dimension, which is 1), but the codimension is the same since $\operatorname{codim}_{\mathbb{P}}(\bar{s}) = \operatorname{codim}(s)$ (the latter being in $\mathbf{K}^{n+1}$).

The points of $\mathbb{P}^n(\mathbf{K})$ under this construction have what are called *homogeneous coordinates.* The homogeneous coordinates of the point $\{l\}$ are $[x_1, \ldots, x_n, x_{n+1}]$ where $(x_1, \ldots, x_n, x_{n+1})$ is any nonzero vector in l. Thus, $[x_1, \ldots, x_n, x_{n+1}]$ and $[cx_1, \ldots, cx_n, cx_{n+1}]$ are the same projective point, for any scalar $c \neq 0$.

[**The following stuff does need clarification**]

We could regard any hyperplane $\bar{h}$ as a linear hyperplane h given by $h = \{\mathbf{x} \in \mathbf{K}^{n+1} : \mathbf{a} \cdot \mathbf{x} = 0\}$, where $\mathbf{a}$ is some fixed non-zero vector. But the simplest way is to let $\bar{h}_{n+1}$ be h_∞ in $\mathbb{P}^n(\mathbf{K})$, where h_{n+1} is the coordinate hyperplane with $(n+1)$st coordinate $x_{n+1} = 0$.

Let $A = \{\mathbf{x} \in \mathbf{K}^{n+1} \mid x_{n+1} = 1\}$. Every line l in $\mathbf{K}^{n+1}$ meets A in one point or is parallel to A. Denote by $\bar{h}_{n+1}$ the set of homogeneous lines parallel to A. Let A_0 be the translate of A that goes through $\mathbf{0}$. The set A_0 as points, with the lines in $\bar{h}_{n+1}$, is $\mathbb{A}^n(\mathbf{K})$ in the original definition.

A.2. **Vector sets and hyperplane arrangements.** [[LABEL 3.repns]]

Real space and positive span. [[LABEL 3.realpos]]

Over $\mathbb{R}$, every homogeneous line, since it has two directions, is made up of two rays emanating from the origin, and every hyperplane has two sides. [**clarification**]

For a set $S \subseteq \mathbb{R}$, we define $\operatorname{pos}(S)$, the *positive span* of S, as

$$\operatorname{pos}(S) := \Big\{ \sum_{i=1}^{n} \alpha_i \mathbf{x}_i \mid \mathbf{x}_i \in S,\ \alpha_i \geq 0 \Big\}.$$

Thus, it is the 'Minkowski sum' of the closed rays generated by the vectors in S. (A closed ray includes the origin. The Minkowski sum is the set of all sums of vectors taken one from each ray.) If $S = \{\mathbf{x}_1, \ldots, \mathbf{x}_m\}$ (where no $\mathbf{x}_i = \mathbf{0}$), we define

$$\begin{aligned} h_i &:= \{\mathbf{x} \in \mathbb{R}^n \mid \mathbf{x}_i \cdot \mathbf{x} = 0\}, \\ h_i^+ &:= \{\mathbf{x} \in \mathbb{R}^n \mid \mathbf{x}_i \cdot \mathbf{x} > 0\}, \\ h_i^- &:= \{\mathbf{x} \in \mathbb{R}^n \mid \mathbf{x}_i \cdot \mathbf{x} < 0\}. \end{aligned}$$

This shows there is a duality between rays and oriented hyperplanes of $\mathbb{R}^n$. The duality extends to one between positive span and half-space intersection.

Lemma A.1. [[LABEL L:0226 lemma1]] $\mathbf{x} \in \operatorname{pos}(\mathbf{x}_1, \ldots, \mathbf{x}_m) \iff \bigcap_{i=1}^m h_i^+ \subseteq h_{\mathbf{x}}^+$.

I omit the proof, which is standard real geometry.

Let's look further into the properties of positive span. Assume the vectors $\mathbf{x}_i$ span $\mathbb{R}^m$. Then the union of all $\operatorname{pos}(\pm\mathbf{x}_1, \ldots, \pm\mathbf{x}_m)$, over all choices of signs for each generating vector, is $\mathbb{R}^n$. That is,

$$\bigcup_{(\varepsilon_1, \ldots, \varepsilon_m) \in \{\pm 1\}^n} \operatorname{pos}(\varepsilon_1 \mathbf{x}_1, \ldots, \varepsilon_m \mathbf{x}_m) = \mathbb{R}^m,$$

because every vector is a linear combination of $\{\mathbf{x}_1, \ldots, \mathbf{x}_m\}$, hence a nonnegative combination after choosing the signs ε_i. On the other hand, the intersection of all $\mathrm{pos}(\pm\mathbf{x}_1, \ldots, \pm\mathbf{x}_m)$ is $\{\mathbf{0}\}$, because $\mathrm{pos}(\varepsilon_1\mathbf{x}_1, \ldots, \pm\mathbf{x}_m) \subseteq (h_1^\varepsilon \cup h_1)$, with a similar formula for each other $i \in [m]$, and then

$$(h_1^+ \cup h_1) \cap (h_1^- \cup h_1) = h_1,$$
$$h_1 \cap \cdots \cap h_m = \{0\}.$$

Homogeneous, affine, and projective hyperplane arrangements. [[LABEL 3.hyps]]

An *arrangement of hyperplanes* $\mathcal{A}$ is a finite set of hyperplanes in $\mathbf{K}^n$. If the hyperplanes are linear subspaces (i.e., they are given by homogeneous equations), this is called a linear or homogeneous arrangement. An arrangement of hyperplanes in $\mathbb{A}^N(\mathbf{K})$ is called an affine arrangement and one in $\mathbb{P}^n(\mathbf{K})$ is called a projective arrangement.

An affine hyperplane has two sides, but a projective hyperplane has only one side. Two projective planes determine two 'sides'—I mean, regions.

Comparing a vector set $S \subseteq \mathbb{R}^n$ and its dual hyperplane arrangement, $\mathcal{A} = \{h_1, \ldots, h_m\}$, a region of $\mathcal{A}$ corresponds to a choice of side for each h_i. The region will be the intersection of these half spaces when this intersection is non-empty. This corresponds to a choice of $\pm\mathbf{x}_i$ in calculating the positive span, or equivalently a ray in the line generated by each $\mathbf{x}_i$. For example,

$$R = h_1^+ \cap h_2^+ \cap h_3^- \leftrightarrow \mathrm{pos}(\mathbf{x}_1, \mathbf{x}_2, -\mathbf{x}_3).$$

For a hyperplane arrangement $\mathcal{A}$ and $h \in \mathcal{A}$, the arrangement *induced* by $\mathcal{A}$ in h is

$$\mathcal{A}^h = \begin{cases} \{h \cap h' \ : \ h' \in \mathcal{A}\} & \text{if linear or projective,} \\ \{h \cap h' \ : \ h' \in \mathcal{A},\ h \cap h' \neq \varnothing\} & \text{if affine.} \end{cases}$$

For a real arrangement, $\mathcal{A}$ in $\mathbb{R}^n$, $\mathbb{R}^n \setminus \bigcup \mathcal{A}$ is divided into connected components, called regions, which are open polyhedra. Define $r(\mathcal{A})$ to be the number of regions.

Theorem A.2. [[LABEL T:0226 hyp del regions]] *Let $\mathcal{A}$ be a hyperplane arrangement in $\mathbb{R}^n$, $\mathbb{A}^n(\mathbb{R})$, or $\mathbb{P}^n(\mathbb{R})$. Let $h \in \mathcal{A}$. The number r of regions satisfies*

$$r(\mathcal{A}) = r(\mathcal{A} \setminus \{h\}) + r(\mathcal{A}^h).$$

Proof. I'll give a proof for the affine case. Each region R of $\mathcal{A} \setminus \{h\}$ is either disjoint from h, or bisected by h. In the first case R is a region of $\mathcal{A}$. In the second case, h splits R into two regions of $\mathcal{A}$ and one of $\mathcal{A}^h$, namely $R \cap h$. Thus, $r(\mathcal{A}) - r(\mathcal{A} \setminus \{h\})$, the number of new regions created by adjoining h, is equal to the number of regions cut out in h by the induced arrangement $\mathcal{A}^h$. ☐

The *characteristic polynomial* of an arrangement $\mathcal{A}$ in $\mathbf{K}^n$ or $\mathbb{P}^n(\mathbf{K})$ is

$$p_{\mathcal{A}}(\lambda) := \sum_{\mathcal{S} \subseteq \mathcal{A}} (-1)^{\dim(\bigcap \mathcal{S})}.$$

(This happens to depend only on the intersection lattice

$$\mathcal{L}(\mathcal{A}) := \{\bigcup \mathcal{S} : \mathcal{S} \subseteq \mathcal{A}\}$$

since it is the same as the characteristic polynomial of the matroid of $\mathcal{A}$ except for an extra factor of $\lambda^{\dim \bigcap \mathcal{A}}$; but I won't discuss that aspect here). The characteristic polynomial of an affine arrangement is just slightly different: the sets $\mathcal{S}$ whose intersection is empty must be omitted from the summation.

Theorem A.3 (Real Case [FUTA])**.** [[LABEL T:0303regions]] *The number of regions of a real hyperplane arrangement is* $r(\mathcal{A}) = (-1)^d p_{\mathcal{A}}(-1)$.

The proof by induction on the number of hyperplanes is an easy consequence of Theorem A.2 and a similar recursive property of the characteristic polynomial (which I omit, regretfully).

2009 Mar 3: Nate Reff

Complex hyperplane arrangements. [[LABEL 3.complexhyps]]

Write $\mathcal{A}$ for an arrangement of hyperplanes (that is, again, a finite set of hyperplanes) in $\mathbb{C}^d$.

Complex hyperplanes have real codimension equal to 2 and therefore do not disconnect the space $\mathbb{C}^d$ when they are deleted. The *complement* of $\mathcal{A}$, $\mathcal{M} := \mathbb{C}^d \setminus \bigcup \mathcal{A}$, consequently is connected. However, this space has nontrivial cohomology. Look at $H^i(\mathcal{M}; G)$, where $G = \mathbb{Z}$ or $\mathbb{Q}$. The rank $\operatorname{rk}(H^i(\mathcal{M}; \mathbb{Z})) = \beta_i(\mathcal{M})$, the i^{th} Betti number of the complement. The Poincaré polynomial, denoted by $P_{\mathcal{M}}(t) := \sum_{i \geq 0} \beta_i t^i$, is the generating function of the Betti numbers. Note that $\beta_i = 0$ if $i > d$.

Theorem A.4 (Complex Case: Orlik–Solomon Theorem [AH])**.** [[LABEL T:0303complexcoho]] $P_{\mathcal{M}}(t) = (-t)^d p_{\mathcal{A}}(-1/t)$.

This is the substantially more subtle complex analog of the real Theorem A.3. I will sorrowfully omit the proof because the theorem is too far off the subject of signed graphs.

The complex hyperplanes that come up in dealing with signed graphs are $(h^{\varepsilon}_{ij})_{\mathbb{C}} := \{z \in \mathbb{C}^n : z_j = \varepsilon z_i\}$. Thus, they have the same equations as the real hyperplanes h^{ε}_{ij}. A complex hyperplane arrangement whose equations have real coefficients corresponds to a real arrangement $\mathcal{A}$ with the same equations; it is called the *complexification* of $\mathcal{A}$. That is the kind of complex arrangement we get from a signed graph. We shall see other kinds of complex arrangements when we get to gain graphs in Chapter IV.

A.3. **Signed graphs, vectors, and hyperplanes.** [[LABEL 3.sgrepn]]

Now we turn our attention to vector representations and hyperplane representations of signed graphs. These are dual to each other.

Vector representations. [[LABEL 3.sgvector]]

Here we have vectors $e \mapsto x_e \in \mathbf{K}^n$, where x_e is the column of e in $\mathrm{H}(\Sigma)$. We allow $\mathbf{K}$ to be any field or division ring, but when we want orientations we use $\mathbb{R}$. Remember that for signed graphs we must have $\operatorname{char} \mathbf{K} \neq 2$ in order to distinguish positive from negative edges. For ordinary graphs we allow $\operatorname{char} \mathbf{K} = 2$.

By a *(vector) representation* over $\mathbf{K}$ we mean the dependent sets of vectors x_e, which we know by Theorem **??** are the vector sets that correspond to edge sets that contain a frame circuit. We are really representing the relationships between the edges. What we get as vectors are the following:

$$\pm\underset{+\text{links}}{\begin{pmatrix}0\\\vdots\\+1\\0\\\vdots\\-1\\0\end{pmatrix}},\quad \pm\underset{-\text{links}}{\begin{pmatrix}0\\\vdots\\+1\\0\\\vdots\\+1\\0\end{pmatrix}},\quad \underset{+\text{loops, loose edges}}{\begin{pmatrix}0\\\vdots\\0\\0\\\vdots\\0\\0\end{pmatrix}},\quad \pm\underset{\text{half edges}}{\begin{pmatrix}0\\\vdots\\+1\\0\\\vdots\\0\\0\end{pmatrix}},\quad \pm\underset{-\text{loops}}{\begin{pmatrix}0\\\vdots\\+2\\0\\\vdots\\0\\0\end{pmatrix}}.$$

Now, consider a finite group $\mathfrak{W}$ generated by reflections in $\mathbb{R}^n$. To explain reflections: We have a reflecting hyperplane h, and the reflection $\rho_h : \mathbb{R}^n \to \mathbb{R}^n$ with fixed point set $\mathrm{Fix}(\rho_h) = h$. We can define our hyperplane as $h = \{x \in \mathbb{R}^n \mid \langle \alpha, x\rangle = 0\}$ where $\alpha \in \mathbb{R}^n \setminus \{\mathbf{0}\}$. Note that here, as usual, the inner product is the normal Euclidean dot product.

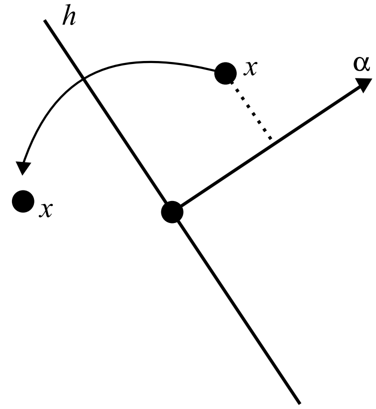

Figure A.1. Reflection across a hyperplane h.
[[LABEL F:0303image1]]

Looking at Figure A.1 we want to find the projection of x onto α and reverse it with respect to our hyperplane h. This gives us the following formula for the reflection:

$$\begin{aligned}\rho_\alpha(x) = \rho_h(x) &= x - 2\mathrm{proj}_\alpha x\\ &= x - 2\langle x, \hat{\alpha}\rangle\hat{\alpha}\\ &= x - 2\frac{\langle x, \alpha\rangle\alpha}{||\alpha||^2}.\end{aligned}$$

A *root system* connected with $\mathfrak{W}$ is a set $R_{\mathfrak{W}}$ of vectors (called the *roots*) such that:

(1) The roots span $\mathbb{R}^n$.
(2) The only scalar multiples of a root $x \in R_{\mathfrak{W}}$ that belong to $R_{\mathfrak{W}}$ are x and $-x$.
(3) If $x, y \in R_{\mathfrak{W}}$ then $\rho_x(y) \in R_{\mathfrak{W}}$. In other words $\rho_x(R_{\mathfrak{W}}) = R_{\mathfrak{W}}$.
(4) If $x, y \in R_{\mathfrak{W}}$ then $2\frac{\langle x,y\rangle}{||y||^2} \in \mathbb{Z}$.

If we are given $\mathfrak{W}$, $R_{\mathfrak{W}}$ is a set of normal vectors to the reflecting hyperplanes. We write $R_1 \perp R_2$ if $x_1 \in R_1, x_2 \in R_2 \implies \langle x_1, x_2\rangle = 0$ and $R_1, R_2 \neq \varnothing$. Suppose R is a root system. A *decomposition* of R is $R = R_1 \cup R_2$ where $R_1 \perp R_2$.

Obvious fact: if R_1 is a root system in $\mathbb{R}^{d_1}$, and R_2 is a root system in $\mathbb{R}^{d_2}$, then $R_1 \uplus R_2$ is a root system in $\mathbb{R}^{d_1+d_2}$. (The technically precise definition of $R_1 \uplus R_2$ is $\{(x,0) : x \in R_1\} \cup \{(0,y) : y \in R_2\}$.)

Therefore, to classify all root systems we can limit ourselves to those that are indecomposable (also known as *irreducible*).

Theorem A.5 (Killing)**.** *Let $\{b_1, b_2, \dots, b_n\}$ denote the standard orthonormal basis of $\mathbb{R}^n$. The only irreducible root systems are the following:*

Classical Root Systems

$$A_{n-1} = \{b_i - b_j\}_{i \neq j},$$
$$D_n = A_{n-1} \cup \{\pm(b_i + b_j)\}_{i \neq j},$$
$$B_n = D_n \cup \{\pm b_i\},$$
$$C_n = D_n \cup \{\pm 2b_i\}$$

Exceptional Root Systems

$$E_6, E_7, E_8,$$
$$F_4,$$
$$G_2$$

What we are interested in from a signed graph perspective are the classical root systems because of the following correspondences:

$$K_n \leftrightarrow A_{n-1},$$
$$\pm K_n \leftrightarrow D_n,$$
$$\pm K_n' \leftrightarrow B_n,$$
$$\pm K_n^\circ \leftrightarrow C_n.$$

Positive span, reflected in the graph. [[LABEL 3.posspangraphs]]

Let's examine how positive span of real vectors, denoted by pos, translates into transitive closure in graphs and signed graphs. Let $b_1, \dots, b_n$ be the standard unit basis vectors of $\mathbb{R}^n$.

If we look at $b_5 - b_2 = (b_5 - b_4) + (b_4 - b_3) + (b_3 - b_2)$,

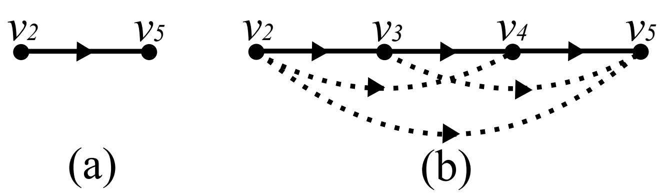

FIGURE A.2. (a)$b_5 - b_2$,(b)$(b_5 - b_4) + (b_4 - b_3) + (b_3 - b_2)$
[[LABEL F:0303image2]]

In Figure A.2 we indicate the edges of the transitive closure with dashed lines. So $b_5 - b_2 \in \mathrm{pos}(b_5 - b_4, b_4 - b_3, b_3 - b_2)$.

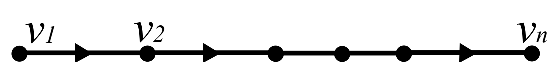

FIGURE A.3
[[LABEL F:0303image3]]

The transitive closure of the path in Figure A.3 is all of K_n oriented low-to-high and corresponds to $\mathrm{pos}(b_i - b_{i-1} : 2 \le i \le n) = A_{n-1}$.

This leads to a more general statement. Let R be a classical root system and let $E_{\mathbb{R}} \subseteq R$ be a subset corresponding to the edges of a signed graph Σ. Then

$$\mathrm{pos}(R) \leftrightarrow \text{ transitive closure in } \vec{\Sigma}.$$

The general definition of transitive closure in a signed graph such as $+K_n$, $\pm K_n$, which corresponds to a classical root system, is based on the following steps. Consider the pair of edges shown in Figure A.4.

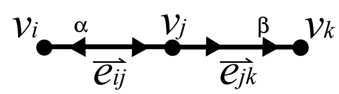

FIGURE A.4
[[LABEL F:0303image4]]

So $\sigma(\vec{e}_{ij}) = -\alpha$ and $\sigma(\vec{e}_{jk}) = \beta$. Then the following edge $\vec{e}_{ik}$, with sign $\sigma(\vec{e}_{ik}) = -\alpha\beta$ is

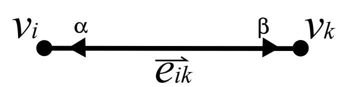

FIGURE A.5
[[LABEL F:0303image5]]

in the transitive closure. This leads us to define the transitive closure of S in an oriented signed graph $\vec{\Sigma}$ to be

$$\begin{aligned} \mathrm{trans}(S) := S \cup \{\vec{e}{:}vw \in \vec{E} \mid \exists \text{ a coherent walk } W = e_1e_2\ldots e_l \text{ from } v \text{ to } w \\ \text{with sign } \sigma(e) = \sigma(W) \text{ and orientation} \\ \tau(v,e_1) = \tau(v,e),\ \tau(w,e_l) = \tau(w,e)\}. \end{aligned}$$

The reason the definition requires a walk rather than a path is situations like that in Figure A.6, where $V = \{v_1, v_2, v_3\}$, $E = \{e^+_{12}, e^-_{12}, e^+_{13}, e^-_{13}\}$, oriented so e^+_{12} goes from v_1 to v_2, e^-_{12} is introverted, and e^+_{13} goes from v_3 to v_1; we leave the orientation of e^-_{13} unspecified for the moment. Let $S := \{e^+_{12}, e^-_{12}, e^+_{13}\}$ and $W := e^+_{12}, e^-_{12}, e^+_{13}$. Then $e^-_{13} \in \mathrm{trans}(S)$ if it is oriented to be introverted but not if it is extraverted. If the definition specified a path from v to w, then no orientation would put e^-_{13} into $\mathrm{trans}(S)$. But we know it should be in the transitive closure, because the vector $-b_1 - b_3 \in \mathrm{pos}\{b_2 - b_1, -b_2 - b_1, b_1 - b_3\}$.

FIGURE A.6. Transitive closure
[[LABEL F:0303image6]]

B. Day of Everything That's so Awful We've Been Putting It Off.

B.1. **Lift matroid.** [[LABEL 3.liftmatrixrepn]]

2009 Mar 5: Jackie Kaminski

So far in this course we've been neglecting lift circuits and the lift matroid, so now we return to them. Recall from Definition I.2 that a lift circuit is either a positive circle, two negative circles that share exactly one vertex (and no edges), or two vertex-disjoint negative circles (with the usual note that half edges act like negative loops). We also defined the augmented incidence matrix $\bar{\mathrm{H}}(\Sigma)$, over a field that contains $\mathbb{F}_2$ as a subfield, as

$$\begin{array}{c} \\ \\ v_1 \\ \vdots \\ v_n \end{array} \begin{array}{c} \begin{array}{ccc} e_1 & \dots & e_m \end{array} \\ \left(\begin{array}{ccc} \bar\sigma(e_1) & \dots & \bar\sigma(e_m) \\ \hline & & \\ & \mathrm{H}(|\Sigma|) & \\ & & \end{array} \right) \end{array},$$

where $\bar\sigma(e) = 0$ if $\sigma(e) = +$ and 1 if $\sigma(e) = -$.

Theorem B.1. *The minimal dependent sets of columns of* $\bar{\mathrm{H}}(\Sigma)$ *correspond to the lift circuits of* Σ*.*

The proof of this is an excellent exercise for the reader. We point to Theorem G.8 if you are in need of inspiration.

This theorem gives us a *lift representation* of Σ by vectors, namely the columns of $\bar{\mathrm{H}}(\Sigma)$, so named because the minimal dependent sets are the lift circuits. Notice that this representation *cannot* be done over the reals, as we need a field in which the sign group is an additive subgroup of order 2. Otherwise, the dependent columns don't correspond to lift circuits.

This leads us to mention that Σ has a frame representation only over fields of characteristic $\neq 2$, where we can represent the sign group as a multiplicative subgroup of the field, and a lift representation only over fields of characteristic 2, where we can represent the sign group as an additive subgroup of the field.

B.2. **Hyperplane representation (of the frame matroid).** [[LABEL 3.framehyprepn]]

March 5: Jackie Kaminski

Recall that for a link $e{:}v_iv_j$ with sign ε, the associated hyperplane (denoted by $h^{\varepsilon}_{ij}, h^{\varepsilon}_{e}, h_{ij}$, or h_e, depending on our mood, and whether the simpler notations will introduce ambiguity) is $\{x \in \mathbb{R}^n \mid x_j = \varepsilon \cdot x_i\}$. See Figure B.1. Similarly, for a half edge $f{:}v_i$, the associated hyperplane (denoted by h_i or h_f) is $\{x \in \mathbb{R}^n \mid x_i = 0\}$. Similarly for a negative loop at v_i, the associated hyperplane is $\{x \in \mathbb{R}^n \mid x_i = -x_i\} = \{x \in \mathbb{R}^n \mid x_i = 0\}$ (since $x_i = -x_i \implies 2x_i = 0$). For a positive loop ($g{:}v_i$ in Figure B.1), the associated hyperplane is the degenerate hyperplane, $\{x \in \mathbb{R}^n \mid x_i = +x_i\} = \mathbb{R}^n$. As before, the hyperplanes associated with a signed graph Σ are, when taken together, an arrangement of hyperplanes, $\mathcal{H}[\Sigma]$, in $\mathbf{K}^n = \mathbf{K}^{|V|}$, which in turn gives a representation of Σ by the intersection lattice $\mathcal{L}(\mathcal{H}[\Sigma]) \cong \operatorname{Lat}\Sigma$, where the isomorphism is the natural isomorphism induced by $h_e \leftrightarrow e$. Note that although $\mathbf{K}$ can be any field of characteristic $\neq 2$, we usually work in $\mathbb{R}$. In that case, $\mathcal{H}[\Sigma]$ divides up $\mathbb{R}^n$ into *regions*, which are the connected components of $\mathbb{R}^n \setminus \bigcup \mathcal{H}[\Sigma]$. Define $r(\mathcal{H}[\Sigma])$ to be the number of regions of $\mathbb{R}^n \setminus \bigcup \mathcal{H}[\Sigma]$. Recall that $\chi_\Sigma(\lambda)$ is the chromatic polynomial of Σ. (See Definition K.2 for more information.)

Lemma B.2. [[LABEL L0305: counting hyperplane regions]] *The number of regions of* $\mathcal{H}[\sigma]$) *in* $\mathbb{R}^n$ *satisfies* $r(\mathcal{H}[\Sigma]) = (-1)^n \chi_\Sigma(-1) = |\chi_\Sigma(-1)|$.

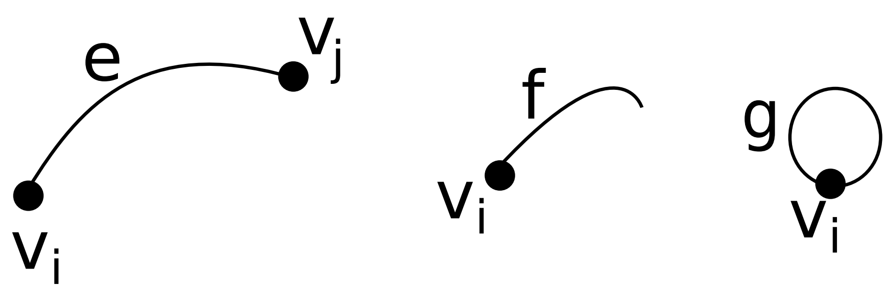

Figure B.1. Several edges, their associated hyperplanes are described above.
[[LABEL F0305: Edges]]

Proof. Notice that the first equals sign is the thing we need to prove. The second equals sign is obvious since χ_Σ actually counts something (namely the number of colorings).

We prove this using induction on $|E|$ via deletion/contraction of Σ. That corresponds to the formula $r(\mathcal{H}[\Sigma]) = r(\mathcal{H}[\Sigma] \setminus e) + r(\mathcal{H}[\Sigma]^{h_e})$, where $\mathcal{H}[\Sigma]^{h_e}$ means the hyperplane arrangement induced in the hyperplane h_e. This formula will be proved as a sublemma. The formula holds only when h_e isn't the degenerate hyperplane (that is, e is not a loose edge or a positive loop), and that for deletion/contraction of the chromatic polynomial (see Thm K.3) we will need to exclude e's being a negative loop or half edge.

So, let's handle all the 'troublesome' cases first, namely any graph without a link e on which we can use deletion/contraction and induction.

Special Case A: Σ contains a positive loop or loose edge.

There are no proper colorings of Σ from any set of colors, so χ_Σ is identically 0. On the other hand, if Σ has a positive loop or loose edge, then the associated hyperplane is the degenerate hyperplane $\mathbb{R}^n$, so $\mathbb{R}^n \setminus \bigcup \mathcal{H}[\Sigma] \subseteq \mathbb{R}^n \setminus \mathbb{R}^n = \varnothing$; therefore there are no regions. Thus,

$$r(\mathcal{H}[\Sigma]) = 0 = (-1)^n \chi_\Sigma(-1),$$

so our result holds for any signed graph containing at least one positive loop or loose edge.

Special Case B: Σ contains only negative loops and half edges.

When e is a negative loop or half edge, h_e is the coordinate hyperplane $h_e = \{x \in \mathbb{R}^n \mid x_i = 0\}$, where e is incident to v_i. Any set of k distinct coordinate hyperplanes divides $\mathbb{R}^n$ into 2^k distinct regions.

A graph whose only edges are negative loops and half edges is the disjoint union of components, each with one vertex (we assume that there are no multiple edges in our graph; we leave the generalization with multiple edges incident to the same vertex to the reader). From Definition K.4 we can determine the chromatic polynomial in this case. Recall that $b(S)$ is the number of balanced components of $\Sigma{:}S$ and $u(S)$ is the number of unbalanced components of $\Sigma{:}S$. Then,

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

For each subset S, $(-1)^{|S|}\lambda^{b(S)} = (-1)^{|S|+b(S)} = (-1)^{u(S)+b(S)}$, and since $u(S) = c(S) - b(S)$, and in our case $u(S) = n - b(S)$, our chromatic polynomial is actually of the form

$$\chi_\Sigma(-1) = \sum_{S\subseteq E}(-1)^{|S|}(-1)^{b(S)} = \sum_{S\subseteq E}(-1)^{u(S)+b(S)} = \sum_{S\subseteq E}(-1)^{n} = (-1)^n \cdot 2^{|E|}.$$

Therefore, in this case $r(\mathcal{H}[\Sigma]) = 2^{|E|} = (-1)^n \cdot 2^{|E|} \cdot (-1)^n = (-1)^n\chi_\Sigma(-1)$, which proves our theorem.

Before the inductive step, we prove two sublemmas we will need in the proof.

[**The following belongs in the geometry section III.1.)**] Recall from Section **??** that $\mathcal{H}^h$ denotes the hyperplane arrangement induced by $\mathcal{H}$ in the hyperplane h.

Sublemma B.3. [[LABEL L:0305 hyperplane regions]] *If h_e is not the degenerate hyperplane, then $r(\mathcal{H}[\Sigma]) = r(\mathcal{H}[\Sigma \setminus e)] + r(\mathcal{H}[\Sigma]^{h_e})$.*

Proof. [**This proof is not right.**]

Let R be a region in $\mathcal{H}[\Sigma]$. Either the topological closure $\bar{R}$ intersects h_e in more than just the origin, or it doesn't. If $\bar{R}$ intersects h_e in more than just the origin, let R' denote the region in $\mathcal{H}[\Sigma]$ that is the reflection of R about h_e. Then $R \cap h_e$ is a region in $\mathcal{H}[\Sigma]^{h_e}$, and $\bar{R}' \cap h_e$ is the same region in $\mathcal{H}[\Sigma]^{h_e}$. Furthermore, R and R' form a single region in $\mathcal{H}[\Sigma \setminus e]$. Therefore the two regions R, R' contribute 2 to the left side and also 2 to the right side.

If, on the other hand, $R \cap h_e = \mathbf{0}$, then R corresponds to a single region in $\mathcal{H}[\Sigma \setminus e]$ and does not correspond to a region in $\mathcal{H}[\Sigma]^{h_e}$. Therefore R contributes 1 to both the left and the right sides. This completes our proof. $\square$

Sublemma B.4. [[LABEL L:0305 hyperplane arrangement isomorphism]] *For any edge e, $\mathcal{H}[\Sigma]^{h_e} \cong \mathcal{H}[\Sigma/e]$. Furthermore, this isomorphism is equality, if we coordinatize $\mathbf{K}^{n-1}$ correctly.*

Proof. A hyperplane in $\mathcal{H}[\Sigma]^{h_e}$ has the form $h_f \cap h_e$, for f an edge in Σ. Under the natural isomorphism $h_f \cap h_e \mapsto$ the hyperplane of f in Σ/e. For convince of notation, we will write h''_f for the hyperplane of f in Σ/e. [5]Similarly we use Σ'' for Σ/e and h'' for hyperplanes in $\mathcal{H}[\Sigma'']$ and $\mathcal{H}''$ for $\mathcal{H}[\Sigma/e]$.

Furthermore, if we think of a hyperplane arrangement as a multiset, then this mapping is a one to one correspondence between hyperplanes in $\mathcal{H}[\Sigma]^{h_e}$ and $\mathcal{H}[\Sigma/e]$. Now we notice that this equality is trivial if h_e is the degenerate hyperplane. Now if h_e is a coordinate hyperplane (say $x_i = 0$) then a hyperplane in $\mathcal{H}[\Sigma]^{h_e}$ has the format

$$\{(x_1, \ldots, x_{i-1}, 0, x_{i+1}, \ldots, x_n) \mid (x_1, \ldots, x_{i-1}, 0, x_{i+1}, \ldots, x_n) \in h_f, \text{ for } h_f \in \mathcal{H}[\Sigma], f \neq e\},$$

now for h_e a coordinate hyperplane, $e \int v_i$ is a negative loop or half edge. Therefore the vertex v_i of Σ isn't in Σ/e. Therefore if we think about h''_f in $\mathbf{K}^{n-1}$ indexed by the vertices $v_1, \ldots, v_{i-1}, v_{i+1}, \ldots, v_n$, we notice that if f was not incident to v_i in Σ, then the defining equation of h_f doesn't involve x_i and

$$h''_f = \{(x_1, \ldots, x_{i-1}, 0, x_{i+1}, \ldots, x_n) \mid (x_1, \ldots, x_{i-1}, 0, x_{i+1}, \ldots, x_n) \in h_f, \text{ for } h_f \in \mathcal{H}[\Sigma], f \neq e\}.$$

If f was incident to v_i in Σ, then in Σ/e f is a half edge (since e was a negative loop or half edge). Furthermore, h_f satisfied $x_j = \sigma(f)x_i$ (where j may equal i). Then h''_f is a coordinate

[5]This notation is borrowed from Tutte, who wrote $\Gamma' := \Gamma \setminus e$ and $\Gamma'' := \Gamma/e$. That's all right when there is no confusion about which edge we are deleting and contracting.

hyperplane with $x_j = 0$ (unless f was half edge or loop, in which case h''_f is the degenerate hyperplane). In either case $h_f = h''_f$ with a 0 inserted as the i^{th} coordinate. Therefore the result is proved in the coordinate hyperplane case.

The case where h_f is not a coordinate hyperplane is similar, and left as an exercise to the reader. **[We need a proof. It's not quite so similar, given that we want to identify the lattices.]**

$\square$

We now resume our proof of Lemma B.2.

Main Case: Σ contains at least one link edge e.

Here we will assume that Σ is connected, since separate components can be treated separately, as is noted often. We will proceed by induction on the number of link edges in Σ. Notice that our base case is a graph with 0 links. If said graph contains loops, loose edges, or half edges, it is taken care of above. So the only part of the base case we are left with is a graph with a single vertex, and no edges. Note that $\chi_{K_i}(\lambda) = \lambda$, and since $\mathbb{R}^1$ cut by no hyperplanes has 1 region we notice that

$$r(\mathcal{H}[\Sigma]) = 1 = (-1)^1 \cdot (-1) = (-1)^n \chi_\Sigma(-1),$$

and so our result holds here, which completes the base now.

Now we assume that $r(\mathcal{H}[\Sigma]) = (-1)^n \chi_\Sigma(-1)$ holds for any graph with less than k links, and let Σ be a graph with k links, and let e be a particular link.

So for the inductive step, we assume that $r(\mathcal{H}[\Sigma]) = (-1)^n \chi_\Sigma(-1) = |\chi_\Sigma(-1)|$ holds for any graph on $k-1$ or less edges, and of course, since we are in the main case we are considering only graphs with at least one link e. Let Σ be a graph with k edges. By Lemma B.3, $r(\mathcal{H}[\Sigma]) = r(\mathcal{H}[\Sigma \setminus e]) + r(\mathcal{H}[\Sigma])^{h_e}$ (since e is a link). Furthermore, by the inductive hypothesis, $r(\mathcal{H}[\Sigma \setminus e)] = (-1)^n \chi_{\Sigma\setminus e}(-1)$, since $\mathcal{H}[\Sigma \setminus e]$ has $k-1$ edges.

Now we notice, by Lemma B.4, $\mathcal{H}[\Sigma]^{h_e} \cong \mathcal{H}[\Sigma/e]$, and therefore $r(\mathcal{H}[\Sigma]^{h_e}) = r(\mathcal{H}[\Sigma/e])$. And, since Σ/e has $k-1$ edges, we know that $r(\mathcal{H}[\Sigma/e]) = (-1)^n \chi_{\Sigma/e}(-1)$. Now we see that

$$\begin{aligned} r(\mathcal{H}[\Sigma]) &= r(\mathcal{H}[\Sigma \setminus e)] + r(\mathcal{H}[\Sigma])^{h_e} \\ &= (-1)^n \chi_{\Sigma\setminus e}(-1) + (-1)^n \chi_{\Sigma/e}(-1) \\ &= (-1)^n \big(\chi_{\Sigma\setminus e}(-1) + \chi_{\Sigma/e}(-1)\big) \\ &= (-1)^n \chi_\Sigma(-1). \end{aligned}$$

The last step is by deletion/contraction of χ_Σ, since e is a link. This proves our result in the main case, and therefore this concludes our proof of Lemma B.2. $\square$

Proposition B.5. *Let $S \subseteq E$, corresponding to $\mathcal{S} \subseteq \mathcal{H}[\Sigma]$. The flat $\bigcap \mathcal{S}$ of $\mathcal{H}[\Sigma]$ is the subspace $\{x \in \mathbb{R}^n \mid x_i = 0, v_i \in V_0(S)\} \cap \{x \mid \zeta(v_j)x_j = \zeta(v_i)x_i$ for each edge v_iv_j in any balanced component of S. $\}$. Here ζ is a switching function such that $S{:}[V \setminus V_0(S)]$ is all positive.*

We leave the proof as an exercise for the reader. To this end recall the notation that the balanced components of $\Sigma|S$ are $S{:}B_1, \dots, S{:}B_k$ and the unbalanced components are the components of $\Sigma{:}V_0(S)$, where $V_0(S) = V \setminus (B_1 \cup \cdots \cup B_k)$.

March 5: Zaslavsky

Isomorphism of hyperplane arrangements.

[MOVE to the appropriate place in Section III.A.2.]

The definition of isomorphism of arrangements was incomplete. There are several notions of isomorphism. In class today I only proved that $\mathcal{H}[\Sigma]^{h_e}$ and $\mathcal{H}[\Sigma/e]$ are lattice-isomorphic. This is good enough for the theorem about the number of regions.

There are several kinds of isomorphism between hyperplane arrangements. In every case one assumes they are over the same field. Otherwise, there are significant differences.

Definition B.1. [[LABEL D:realarrcombiso]] A *combinatorial isomorphism* of real hyperplane arrangements $\mathcal{H}_1$ and $\mathcal{H}_2$ is a bijection $\psi : \mathcal{H}_1 \to \mathcal{H}_2$ that induces an isomorphism of the intersection (semi)lattices and also an isomorphism of the face posets. (This implies that the bijection can be oriented, i.e., one can give h and $\psi(h)$ positive sides in a way that leads to the face-poset isomorphism.)

Definition B.2. [[LABEL D:realarriso]] An *isomorphism* of real hyperplane arrangements is a combinatorial isomorphism that preserves dimension. **[Probably this has no use.]**

Definition B.3. [[LABEL D:arrlatiso]] A *lattice isomorphism* of hyperplane arrangements over the same field is a bijection $\psi : \mathcal{H}_1 \to \mathcal{H}_2$ that induces an isomorphism of the intersection (semi)lattices. An equivalent definition is that ψ preserves codimension, i.e., $\operatorname{codim} \bigcap \psi(\mathcal{S}) = \operatorname{codim} \bigcap \mathcal{S}$ for every $\mathcal{S} \subseteq \mathcal{H}_1$, and it preserves the property of a subarrangement's having empty or nonempty intersection, i.e., $\bigcap \psi(\mathcal{S}) = \varnothing \iff \bigcap \mathcal{S} = \varnothing$ (these can be proved easily enough).

Definition B.4. [[LABEL D:0503 Lattice Isomorphism]] A *lattice isomorphism* $\mathcal{H}_1 \cong \mathcal{H}_2$ is a bijection $\theta : \mathcal{H}_1 \to \mathcal{H}_2$ such that for all $\mathcal{S} \subseteq \mathcal{H}_1$, either $\operatorname{codim}(\bigcap \theta(\mathcal{S})) = \operatorname{codim} \bigcap \mathcal{S}$ or both intersections are empty.

This is different from the usual isomorphism of arrangements. Two lattice-isomorphic hyperplane arrangements will not necessarily have exactly the same characteristic polynomial, whereas two isomorphic hyperplane arrangements will. The difference depends on the dimension of the ambient spaces, say $\mathbb{R}^{d_1}$ and $\mathbb{R}^{d_2}$, respectively; we have $p_{\mathcal{H}_2}(\lambda) = \lambda^{n_2 - n_1} p_{\mathcal{H}_1}(\lambda)$.

[MERGE the last two defs. Find a def. of isomorphism of arrangements over any field?]

3/10: Lodha

2009 Mar 12: Nate Reff

C. Geometric Representations of Signed Graphs

[[LABEL 3.vertexrepn]]

Suppose we have a signed simple graph Σ. Define $\sigma(vw) := \sigma(e_{vw})$ if there is an edge e_{vw}, or 0 if there is not.

C.1. **Gramian representations.** [[LABEL 3.gramianrepn]]

Definition C.1. [[LABEL D:0312gramian repn]] A *Gramian representation* of Σ is a mapping $\psi : V \to \mathbb{R}^\infty$ such that

$$\psi(v) \cdot \psi(w) = \sigma(vw) \text{ for } v \neq w.$$

This means that $A(\Sigma)$, except for the diagonal elements, is the Gram matrix of a set of vectors, $\{\psi(v) \mid v \in V\}$. (Remember that the *Gram matrix* G of vectors $v_1, \ldots, v_n$ is the matrix of dot products, $g_{ij} = v_i \cdot v_j$.) By $\mathbb{R}^\infty$ we really mean $\mathbb{R}^d$ of sufficiently high dimension that there is always plenty of room for any finite set of vectors.

In a Gramian representation, if you scale the whole representation by a constant, in other words $\psi'(v) := \alpha\psi(v)$ for $\alpha \neq \pm 1$, you preserve the Gramian property.

C.2. **Angle representations.** [[LABEL 3.anglerepn]]

Definition C.2. [[LABEL D:0312angle repn]] An *angle representation* is a Gramian representation such that

$$\cos\measuredangle\big(\psi(v), \psi(w)\big) = \hat{\psi}(v) \cdot \hat{\psi}(w) = \frac{\sigma(vw)}{\gamma} \text{ for all distinct } v, w \in V,$$

where γ is a fixed positive number. Equivalently, it is a Gramian representation such that $\|\psi(v)\|\|\psi(w)\| = \gamma$ for every adjacent pair of vertices. We call γ the *denominator invariant* of the representation.

An angle representation is *equinormal* if all vectors have the same length, i.e., all $\|\psi(v)\|$ are the same number, which clearly must be $\sqrt{\gamma}$ (if Σ has an edge).

A consequence of this definition is the following: In a walk $v_0v_1 \ldots v_l$, $\|\psi(v_1)\| = \gamma/\|\psi(v_0)\|$ and $\|\psi(v_2)\| = \gamma/\|\psi(v_1)\| = \|\psi(v_0)\|$. Therefore $\|\psi(v_{2i})\| = \|\psi(v_0)\|$ and $\|\psi(v_{2i+1})\| = \gamma/\|\psi(v_0)\|$. Therefore in an odd circle all $\|\psi(v_i)\| = \sqrt{\gamma}$, and consequently $\|\psi(v)\| = \sqrt{\gamma}$ for every vertex in the component of the odd circle. We have proved:

Proposition C.1. [[LABEL P:0312anglerepnnorm]] *Assume Σ is connected.*

(a) If Σ is not bipartite, then $\|\psi(v)\| = \sqrt{\gamma}$ for all $v \in V$.
(b) If Σ is bipartite, say with bipartition $V = V_1 \cup V_2$, then there are positive numbers γ_1 and γ_2 with $\gamma_1\gamma_2 = \gamma$ such that $\|\psi(v_1)\| = \gamma_1$ for all $v_1 \in V_1$ and $\|\psi(v_2)\| = \gamma_2$ for all $v_2 \in V_2$. The vectors $\psi(v_1)$ for $v_1 \in V_1$ are mutually orthogonal, and the vectors $\psi(v_{2i})$ for $v_{2i} \in V_2$ are mutually orthogonal.

Proof. Since $\cos\measuredangle(v_1, v_2) = \sigma(v_1v_2)/\gamma$, we have $v_1 \cdot v_2 = \sigma(v_1v_2)$. (Remember that $\sigma(vw) = 0$ if v and w are not adjacent.) The result follows. $\square$

In the bipartite case we get another angle representation if we replace all the vectors v_1 by scalar multiples αv_1 and all the vectors v_2 by $\alpha^{-1}v_2$ for any $\alpha \neq 0$. By choosing α so that $\alpha\gamma_1 = \alpha^{-1}\gamma_2$ we normalize so all vectors have the same length $\alpha = \sqrt{\gamma_2/\gamma_1}$. That is, every angle representation of a connected graph can be equinormalized.

Define two angle representations to be *similar* if they differ by an orthogonal transformation of $\mathbb{R}^\infty$ or by scaling individual $\psi(v)$. We show that γ is an invariant of similarity.

Proposition C.2. [[LABEL P:0312equivalence]] *Suppose ψ and ψ' are two equivalent angle representations of Σ with denominator invariants γ and γ', respectively. Then $\gamma = \gamma'$.*

Proof. The definition of an angle representation implies that $\gamma = \sigma(vw)/(\hat{\psi}(v) \cdot \hat{\psi}(w))$. The unit vectors $\hat{\psi}(v)$ and $\hat{\psi}(w)$ are independent of scaling individual vectors by positive numbers. Reversing a vector v is switching in Σ, so it negates both $\hat{\psi}(v) \cdot \hat{\psi}(w)$ and $\sigma(vw)$, therefore leaving γ constant. An orthogonal transformation, of course, also holds γ constant. $\square$

Proposition C.3. [[LABEL P:0312eigenvalues]] *Suppose* Σ *has an angle representation with denominator invariant* γ*. Then* $-\gamma \leq \lambda_{\min}(A(\Sigma))$*. The representation vectors are linearly dependent if and only if* $-\gamma$ *is the least eigenvalue of* $A(\Sigma)$*.*

Proof. Any angle representation can be equinormalized without changing γ. In an equinormal angle representation $\psi(v) \cdot \psi(v) = \gamma$, so the Gram matrix of the representation is $A(\Sigma) + \gamma I$. Therefore, the eigenvalues of A are those of G translated down by γ. Remember that G is positive semidefinite so its eigenvalues are non-negative. Therefore, the eigenvalues of A are all greater than or equal to $-\gamma$, with $-\gamma$ showing up if and only if $\operatorname{rk} G < n$, equivalently $\operatorname{rk}\{\psi(v_1), \psi(v_2), \ldots, \psi(v_n)\} < n$, or in other words the representation vectors are linearly dependent. $\square$

Corollary C.4. [[LABEL C:0312mineigenvalue]] *If* $\lambda_{\min}$ *is the smallest eigenvalue of* $A(\Sigma)$*, then* $\gamma \geq -\lambda_{\min}$ *for every angle representation of* Σ*.*

I think a better way to state Proposition C.3 would be in terms of $A = A(-\Sigma)$. Then $\gamma \geq \lambda_{\max}(A)$ and the vectors are dependent if and only if the representation has denominator invariant $\gamma = \lambda_{\max}(A)$. This suggests that a representation of Σ should have been called a representation of $-\Sigma$. A second, strong reason for thinking this will appear when we get to angle representations of line graphs. However, the established terminology is that it is Σ, not $-\Sigma$, which is represented by ψ, so I adhere to that.

C.3. **Root representations.** [[LABEL 3.rootrepn]]

Definition C.3. [[LABEL D:0312root repn]] A *root representation* is an angle representation with denominator invariant $\gamma = 2$; that is, where every vector has norm $\|\psi(v)\| = \sqrt{2}$. Equivalently, all angles are 60°, 90° and 120°, and therefore $\lambda_{\min}(A(\Sigma)) \geq 2$.

Example C.1. [[LABEL X:0312lg]] An example of a signed simple graph Σ with this eigenvalue property is $-\bar{\Lambda}(\Sigma_0)$, where Σ_0 be any simply signed graph whose edges are links. (We proved that $\lambda_{\max}(\bar{\Lambda}(\Sigma_0)) \leq 2$ in Theorem **??**.) Does $-\bar{\Lambda}(\Sigma_0)$ have an obvious root representation? The answer is yes. We use the columns of $\mathrm{H}(\Sigma_0)$, with $\Sigma_0 \subseteq \pm K_n$, as vectors $\psi(e)$ where $e \in V(-\bar{\Lambda})$ (which $= V(\bar{\Lambda}) = V(|\bar{\Lambda}|)$).

We will prove that this example is the norm (pardon the pun).

Theorem C.5 (Cameron–Goethals–Seidel–Shult [CGSS], Vijayakumar [Dinf, Dfs, E8])**.** [[LABEL T:0312rootrepn]] *A signed simple graph* Σ *has a root representation* $\iff$ Σ *is represented by a subset of some* D_n *(these are reduced line graphs) or* E_8 *(these are not line graphs, reduced or otherwise).*

Corollary C.6. [[LABEL C:0312rootrepnlg]] *Every signed simple graph* Σ *with* $\lambda_{\max} \leq 2$ *is a reduced line graph of a simply signed link graph, with a finite number of exceptions.*

The theorem's proof is long! The corollary is an easy consequence, since there are only finitely many signed graphs represented within E_8, and all those represented by D_n are (reduced) line graphs (cf. Theorem **??** OR Section **??**). **[There should be a discussion of how the line graph's vectors are derived from the graph's vectors. In Ch. II? Ch. III? Is it in missing notes?]**

Recall that in $\mathbb{R}^8$ we had $D_8 = \{\pm b_i \pm b_j\}_{i \neq j}$. So there are $\binom{8}{2}$ pairs (i, j) and 4 ways to put on signs, giving a total of 112 vectors. Since $E_8 = D_8 \cup \{\frac{1}{2}(\sum_{i=1}^{\infty} \varepsilon_i b_i) \mid \varepsilon_i = \pm 1 \text{ and } \prod \varepsilon_i = 1\}$, E_8 gives $2^7 = 128$ more vectors, and therefore $|E_8| = 240$.

You can easily verify that $\|\psi(v)\| = \sqrt{2}$ for every element of E_8. Consequently, $E_8 = \psi(\Sigma(E_8))$ for some signed simple graph $\Sigma(E_8)$ of order 240, via $E_8 = \{\psi(v_1), \dots, \psi(v_{240})\}$. In $\Sigma(E_8)$ we have

$$\sigma(v_iv_j) = \frac{1}{2}\psi(v_i)\cdot\psi(v_j) = \begin{cases} +2 & \text{if } \measuredangle(v_i, v_j) = 0^\circ \text{ (that is, } i = j),\\ +1 & \text{if } \measuredangle(v_i, v_j) = 60^\circ,\\ 0 & \text{if } \measuredangle(v_i, v_j) = 90^\circ,\\ -1 & \text{if } \measuredangle(v_i, v_j) = 120^\circ. \end{cases}$$

Question. What signed graph is this? It must be one with interesting properties; what are they?

3/24: Kaminski

3/26: Joyce

2009 March 31: Yash Lodha

Recall that $\mathfrak{L}$ is a star-closed line system at angles 60° and 90°. We took S to be the set of vectors of norm $\sqrt{2}$ that generate the lines of $\mathfrak{L}$, i.e.,

$$S := \{x \mid \langle x\rangle \in \mathfrak{L} \text{ and } \|x\| = \sqrt{2}\}.$$

We assumed that S is irreducible, with dimension ≥ 2. We chose $a, b, c \in S$ such that $a + b + c = 0$ and $a \cdot b = a \cdot c = b \cdot c = -1$ and define

$$\begin{aligned} S_a &:= \{x \in S \mid x \perp a,\ x \cdot b = 1,\ x \cdot c = -1\},\\ S_b &:= \{x \in S \mid x \perp b,\ x \cdot c = 1,\ x \cdot a = -1\},\\ S_c &:= \{x \in S \mid x \perp c,\ x \cdot a = 1,\ x \cdot b = -1\},\\ S_0 &:= \{x \in S \mid x \perp a, b, c\}. \end{aligned}$$

Note that $x\cdot a + x\cdot b + x\cdot c = 0$ for any vector x. We found that $S_b = c + S_a$ and $S_c = -b + S_a$.

Lemma C.7 (Sublemma 6)**.** [[LABEL L:0331 6]] $S_0 = \{x - x' \mid x, x' \in S_a,\ x \neq x'\}$.

Proof. Let S_0' be the set on the right-hand side.

First we prove that $S_0 \supseteq S_0'$. Let $x, x' \in S_a$ with $x \neq x'$. Then $x \cdot a = x' \cdot a$, therefore $(x - x') \cdot a = 0$. Also, $x \cdot b = x' \cdot b$, therefore $(x - x') \cdot b = 0$.

Now we prove that $S_0 \subseteq S_0'$. Suppose $S_0 \not\supseteq S_0'$. Then there is $y \in S_0$ such that y doesn't have the form $x - x'$. $y \perp a, b, c$. Let $x \in S_a$. Then $x' + y \in S_a$ or $x' + y \notin S$ (because $x' + y$ satisfies the product rules for S_a). Could $x' + y \notin S$? Suppose $y \cdot x' = \pm 1$. Then if necessary replace y by $-y$ to get $y \cdot x' = -1$. By star closure, $-(y' + x) \in S$. Therefore $y' + x \in S$ and $x' + y$ equals some $x \in S$. So $y = x - x'$.

Suppose $y \perp$ every $x' \in S_a$. Then let $S_i' = \{a, b, c\} \cup S_a \cup S_b \cup S_c \cup S_0'$. $y \perp S_0'$ because any member of S_0' is $x - x'$, for $x, x' \in S_a$ and $y \perp S_a$. Therefore $(S_0 \setminus S_0') \perp S$.

This gives a decomposition of S. $S = \pm S_1$, where $S_1 = \ldots \cup S_0$ and since S is irreducible, this can't happen. $\square$

Let Γ_a be the graph with vertex set S_a and edges xy whenever $x \perp y$.

Lemma C.8 (Sublemma 7)**.** [[LABEL L:0331 7]] *Let $x, y \in S$. Then $x \perp y \implies z = b - c - x - y \in S_a$ and $z \perp x, y$ so if there is an edge xy then there are edges xz, yz.*

Proof. $x + c \in S_b$ and $y - b \in S_c$. Therefore, $(x + c) \cdot (y - b) = \pm 1$ or 0. Therefore, star closure implies $(x + c) - (y - b) \in S$.

To show $z \in S_a$, we calculate:

$$z \cdot a = b \cdot a - c \cdot a - x \cdot y - y \cdot 0 = -1 - (-1) - 0 - 0 = 0,$$
$$z \cdot b = b \cdot b - c \cdot b - x \cdot b - y \cdot b = 2 + 1 - 1 - 1 = 1.$$ □

Lemma C.9 (Sublemma 8)**.** [[LABEL L:0331 8]] *Some x, y, z. Let $w \in S_a \neq x, y, z$. Then $w \perp$ exactly one of x, y, z.*

Proof. $w \cdot (x + y + z) = w \cdot b - w \cdot c = 1 - (-1) = 2$.

Therefore, $w \cdot x + w \cdot y + w \cdot z = 2$ so we must have one of $w \cdot x,\ w \cdot y,\ w \cdot z$ equal to 0 and the others equal to $+1$. □

Corollary C.10 (Corollary 78)**.** [[LABEL C:78]] *(a) Every edge of Γ_a is in one and only one triangle.*
(b) If Γ_a is a triangle and w is a vertex not in the triangle, then w is adjacent to exactly one vertex of the triangle.

This leaves us with the **Essential Question:** *What is Γ_a?*

Suppose xy is an edge. Then there exists a unique vertex $f(xy)$, also written $f(x, y)$, that makes a triangle with x and y. Note that $f(u, \cdot)$ is a self-inverse function: $f(u, f(u, y)) = y$.

Lemma C.11 (Sublemma 9)**.** [[LABEL L:0331 9]] *If Γ_a has a vertex adjacent to all other vertices, it is a windmill.*

Proof. Let x be a vertex and $y_1, \ldots$ its neighbors. The edge xy_1 belongs to a triangle. Therefore, there is an edge $y_1 y_i$ to make the triangle. Since we know that xy is not in any other triangle, there does not exist an edge $y_1 y_j$ with $j \neq i$.

Now delete x. We have a graph $\Gamma_a{:}N(x)$ of degree 1. Therefore this is a k-edge matching M_k, for some k. So, $\Gamma_a{:}(x \cup N(x)) = M_k \vee x = W_k$, where $k = \frac{1}{2}\deg(x)$ and $\vee$ denotes the *join* operation, in which every vertex of the first graph is made adjacent to every vertex of the second graph.

We conclude that every closed neighborhood in Γ_a is a windmill. □

Lemma C.12 (Sublemma 10)**.** [[LABEL L:0331 10]] *If Γ_a is neither K_n^c nor a windmill, then it is one of three special graphs Γ_n (of order n) for $n = 9, 15, 27$, which are strongly regular graphs.*

Proof. Step 1: There are no isolated vertices. Consider a vertex v. Since there exists an edge xy in Γ_a (since it is not the complement of K_n), there is a triangle xyz. If $v \neq x, y, z$, then $v \sim$ one of x, y, z. Therefore $\deg(v) > 0$ for all $v \in S_a$.

Step 2: The first part of regularity. We know for the case $|S_a| = 1$, Γ_a is K_1^c. For $|S_a| = 2$ it is K_2^c and for $|S_a| = 3$ it is K_3^c or K_3. Hence $|S_a| \geq 4$, by our assumption that it is neither a K_n^c nor a windmill.

Notice that Γ_a is not complete, by Corollary 78(a).

Choose any two nonadjacent vertices, $u \not\sim v$. The neighborhoods of u and v have the form $N(u) = N_u \cup N_{uv}$ and $N(v) = N_v \cup N_{uv}$; let $N_0 := \{x \mid x \not\sim u, v \text{ and } x \neq u, v\}$.

Looking at an x in N_{uv}, an edge xy in N_{uv} cannot exist since would make triangles xyu and xyv, contradicting Corollary 78(a). Thus no edge of the matching in $N(u)$ can be in N_{uv}. Therefore N_{uv} matches into N_u. This means $f(u, \cdot)$ is an injection $N_{uv} \to N_u$.

Let $y \in N_u$. Then $f(u,y) \in N_u \cup N_{uv}$. For v such that $v \not\sim y$ and $v \not\sim u$, by Corollary 78(b) $v \sim f(u,y)$. Therefore $f(u,y) \in N_{uv}$. Therefore $f(u,\cdot) : N_u \to N_{uv}$ and $f(u,\cdot)$ is an injection. It follows that $|N_u| = |N_{uv}|$ and all edges in $N(u)$ are $N_u N_{uv}$ edges.

Similar reasoning shows that $|N_u| = |N_{uv}| = |N_v|$; let k be their common value. It follows that $\deg(u) = 2k = \deg(v)$. We've shown that $u \not\sim v$ in Γ_a implies that $\deg(u) = \deg(v)$.

Step 3: The rest of regularity. We will show that Γ_a^c is connected, whence Γ_a is regular.

Suppose to the contrary that Γ_a^c has $r > 1$ components.

If $r \geq 3$ and $|V_1| \geq 2$, choose $x \in V_2$, $y \in V_3$, and $u, v \in V_1$. Then we get the two triangles uxy and vxy on the one edge xy in Γ_a [**(shown in missing figure)**]. This is not allowed. Therefore, each $|V_i| = 1$ if $r \geq 3$. Therefore $\Gamma_a = K_n$, which is excluded. Consequently, $r < 3$.

Now we look at the 2-component case.

Since there are triangles in Γ_a, there is an edge in, say, V_1, call it xy. If $|V_2| \geq 2$ then we have over lapping triangles, which is not allowed.

Therefore $|V_1| = 1$ and we're in Lemma DQ **??** with Γ_a as a windmill.

Therefore $r = 1$, Γ_a^c is connected, so Γ_a is regular.

Conclusion: Γ_a is $2k$-regular, where k equals the number of common neighbors of two nonadjacent vertices.

Since any two adjacent vertices have exactly one common neighbor, Γ_a is a strongly regular graph with parameters $(n, 2k, 1, k)$.

Furthermore, edges in $N(u)$ are a perfect matching from $N(u) \setminus N_{uv}$ to N_{uv} if $v \not\sim u$.

Also, N_u is a coclique and N_{uv} is a coclique. $\square$

4/2: Joyce

2009 July 9: Jackie Kaminski

Root representations (continued).

C.3.1. *A review of the proof of Lemma 10 C.12.*

Recall that Lemma C.12 states that if Γ_a is not equal to $\bar{K}_l$ or a windmill, then it is one of three special graphs, Γ_n with $n = 9, 15, 27$, which are regular of order 2, 3, 5 respectively. Since (this) proof of Lemma 10 C.12 is so long, let's take a minute and review what we've done to this point.

We recall first that a root representation (Defn C.3) is $\psi : V \to \mathbb{R}^\infty$ such that $\|\psi(v)\|^2 = 2$ and $\psi(v) \cdot \psi(w) = 0, \pm 1$ when $v \neq w$, and $\psi(v) \cdot \psi(w) = \sigma(vw)$ when vw is an edge. Recall also the definitions of the particular root sytems D_l, E_8, and A_{l-1} from **??**:

$$\begin{aligned} D_l &:= \{\pm b_1 \pm b_j \mid i \neq j\} \subseteq \mathbb{R}^l, \\ E_8 &:= D_8 \cup \left\{\tfrac{1}{2}\textstyle\sum_{i=1}^8 \varepsilon_i b_i \mid \varepsilon_i = \pm 1 \text{ and } \prod_1^8 \varepsilon_i = +1\right\}, \\ A_{l-1} &:= \{\pm(b_i - b_j) \mid i \neq j\} \subseteq D_l, \end{aligned}$$

where the b_i are the standard unit basis vectors. And in particular all these vectors (sometimes called "roots") have norm $\sqrt{2}$, and pairs have dot product $0, \pm 1$—except of course for pairs of the same vector twice, or a vector and its negative, in which case the dot product is ± 2. Recall also that we defined the map $\hat{\psi}(v) := \frac{\psi(v)}{\|\psi(v)\|}$, whose main effect is that inner products are now $0,\ \pm 1/2,\ \pm 1$ instead of $0,\ \pm 1,\ \pm 2$.

To prove that $\operatorname{Im}\psi \subseteq D_\infty$ or E_8, we now ignore the signed graph, and focus on any set of vectors S in $\mathbb{R}^l$, satisfying the following conditions (note that we drop the x notation, and instead refer to the vectors as x):

(1) For $x \in S$, x is a unit vector.
(2) For $x, y \in S$, the inner product is 0, $\pm\frac{1}{2}$, or ± 1.
(3) S is maximal with respect to first 2 properties.
(4) S is indecomposable (irreducible). Meaning there does not exist a partition $S = S_1 \uplus S_2$ such that $S_1 \perp S_2$ and neither part is the empty set.

These properties guarantee that S is centrally symmetric. We also proved Lemma **??** which guaranteed that if $a, b \in S$ and $\measuredangle(a, b) = 120°$, then for $c = -(a+b)$, $c \in S$; see Figure C.1. Notice also that Figure C.1 cannot be a root system as is, at a minimum it needs the negatives of the vectors a, b, c.

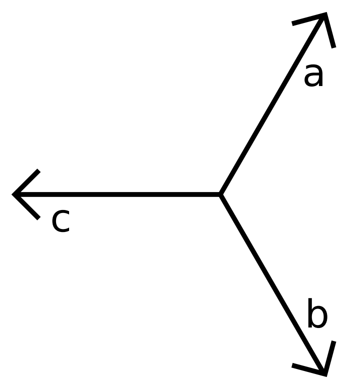

FIGURE C.1. Vectors a, b, c.
[[LABEL F0709: vectors]]

As an interesting side note, if there do not exist $a, b \in S$ such that $a \cdot b = 1/2$, then $S = S_1 \uplus S_2 \cdots \uplus \cdot S_k$, where every S_i one dimensional. In this case, every two vectors not negatives to each other are actually orthogonal.

Next we introduced some notation. Fix $a, b, c \in S$ such that $a \cdot b = \frac{-1}{2}$, and let c be as above, then we defined $S_a, S_b, S_c, S_0 \subseteq S$:

$$\begin{aligned} S_a &:= \left\{x \in S \mid x \perp a,\ x \cdot b = \tfrac{1}{2},\ x \cdot c = -\tfrac{1}{2}\ \right\}, \\ S_b &:= \left\{x \in S \mid x \perp b,\ x \cdot c = \tfrac{1}{2},\ x \cdot a = -\tfrac{1}{2}\ \right\}, \\ S_c &:= \left\{x \in S \mid x \perp c,\ x \cdot a = \tfrac{1}{2},\ x \cdot b = -\tfrac{1}{2}\ \right\}, \\ S_0 &:= \{x \in S \mid x \perp a, b, c\}. \end{aligned}$$

We have shown that these four sets are all pairwise disjoint. When we defined $S' := \{a, b, c\} \uplus S_a \uplus S_b \uplus S_c \uplus S_0$, we were able to prove several lemmas about the structure of S and these subsets, listed here to refresh your memory:

Lemma C.13 (Lemma **??**: Sublemma 2)**.** [[LABEL L:0709 Lemma 2]] $S = S' \cup (-S')$.

This is not a disjoint union because $S_0 = -S_0$ (which we show later).

Lemma C.14 (Lemmas **??**: Sublemmas 3–4)**.** [[LABEL L:0709 Lemma 3/4]] $S_b = c + S_a$ *and* $S_c = -b + S_a$*, by isometries.*

In other words, when seeking to understand the structure of S_a, S_b, and S_c, it suffices to understand the structure of S_a only.

Lemma C.15 (Lemma C.7: Sublemma 6). [[LABEL L:0709 Lemma 6]] *$S_0 = \{x - x' \mid x, x' \in S_a,\ x \not\perp x',\ x \neq x'\}$. Therefore $S_0 = -S_0$.*

What these lemmas combine to give us, is that to study all of S, it suffices to pick $a, b \in S$ such that $a \cdot b = \frac{1}{2}$. Then a, b, and S_a (once you understand its structure) give you all the information about S.

The next lemmas provide help understanding the geometry of S_a, and we're building toward being able to describe and understand S_a through a graph, Γ_a, that we will construct.

Lemma C.16 (Lemma **??**: Sublemma 5). [[LABEL L:0709 Lemma 5]] *In S_a, all dot products are ≥ 0.*

Lemma C.17 (Lemma C.8: Sublemma 7). [[LABEL L:0709 Lemma 7]] *For $x, y \in S_a$, $x \perp y \implies z := b - c - x - y \in S_a$ and $z \perp x, y$.*

This lemma rules out anything 2-dimensional and indecomposable.

Lemma C.18 (Lemma C.9: Sublemma 8). [[LABEL L:0709 Lemma 8]] *If $w \in S_a$ and x, y, z as above with $w \neq x, y, z$ then $w \perp$ exactly one of x, y, z.*

At this point we begin a new part of the proof. We move on to classifying the indecomposable root systems by graph theory, after reinterpreting the last few lemmas graphically. Information on graph structure will allow us to explore all the possibilities for S, which will turn out to be two infinite families and a handful of special examples.

We begin the second part of the proof of Lemma 10 C.12, by reviewing the definition of the graph Γ_a.

Definition C.4. [Definition **??**] [[LABEL D:0709 Gamma]] The *orthogonality graph* of S_a is Γ_a, which has vertex set S_a and has an edge $xy \in E(\Gamma_a) \iff x \perp y$.

Note that Γ_a is an ordinary, simple, unsigned graph. Lemma C.16 that implies it is an unsigned graph. There are clearly no loops, since $x \not\perp x$ for any vector $x \in S_a \subset S$, and by definition there are no multiple edges. Additionally, from Lemmas C.17 and C.18 we derived the following corollary:

Corollary C.19 (Corollary "78" **??**). [[LABEL C:0709 Cor 78]]

(a) Every edge is in a unique triangle.
(b) Given a triangle and a vertex w not in the triangle, w is adjacent to exactly one vertex of that triangle.

Knowing that each edge is in a unique triangle, we were able to define the following function on edges xy in Γ_a: $f(x, y) := f(xy) :=$ the third vertex of the unique triangle on edge xy.

This brought us to the question of what the possibilities for Γ_a are. $\bar{K}_l$ is a possibility as the conditions are vacuously satisfied; $\Gamma_a = \bar{K}_l$ can arise from a root system like Figure C.1 (and its negative). Another possibility is that Γ_a may have one vertex that is adjacent to all others; in this case by, Lemma C.20 Γ_a is a windmill (see Figure C.2.

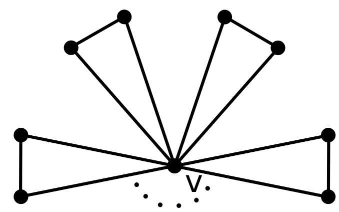

FIGURE C.2. A windmill at vertex v.
[[LABEL F:0709: windmill]]

Lemma C.20 (Lemma C.11: Sublemma 9)**.** [[LABEL L:0709 Lemma 9]] *If Γ_a has a vertex that is adjacent to all others, then Γ_a is a windmill.*

This left us at Lemma 10 C.12, where we examine (exhaustively) what the possibilities are for the orthogonality graph (Γ_a) if it is not edgeless or a windmill. The result is a graph which has a large number of edges (relative to $|V|$) and is most definitely not planar. Though we do our best to draw it as we examine its structure, bear in mind that the following diagrams often focus on what is most important, leaving out details that clutter the graphics (and our understanding).

So, now we have a graph that does have at least one edge, and there is no vertex that is adjacent to all others (and in particular every vertex has at least one other vertex that it is not adjacent to). Let u, v be a pair of non-adjacent vertices, and let $N_{u,v}$ be the set of vertices that are adjacent to both u, v, let N_u be the set of vertices adjacent to u and not adjacent to v, and similarly let N_v be the set of neighbors of v that are not adjacent to u. Finally, let N_0 be $V(\Gamma_a) \setminus (\{u, v\} \cup N_{uv} \cup N_u \cup N_v)$.

We conclude that we have as much of the picture of Γ_a as is in Figure C.3.

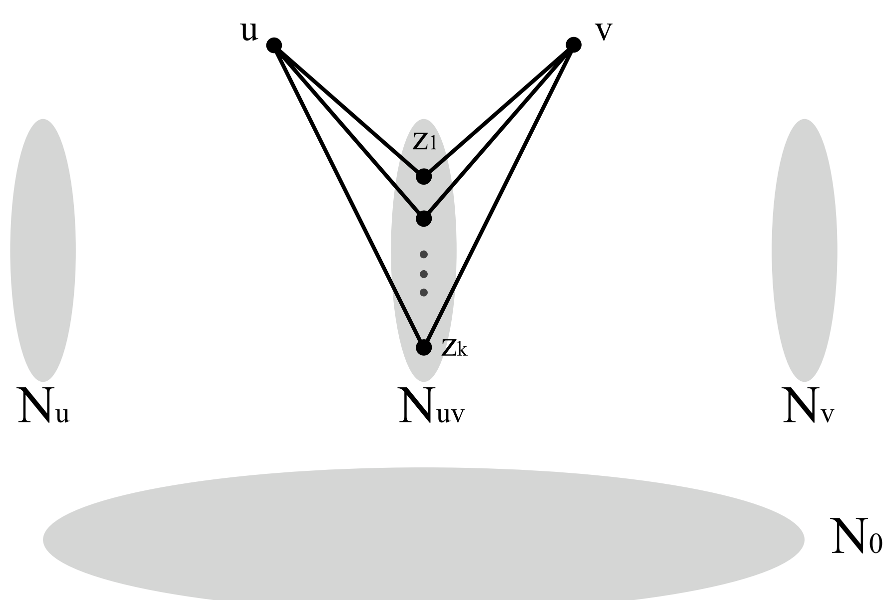

FIGURE C.3. Part of Γ_a.
[[LABEL F:0709: Gamma a first]]

Now, we know that each edge is in a (unique) triangle. Therefore the edge uz_i must be in a triangle, and $f(u, z_i) \neq z_j$, otherwise $z_i z_j$ would be in a triangle with u and also with v, a contradiction. Therefore $f(u, z_i) \notin N_{uv}$, but since the neighbors of u are $N_u \cup N_{uv}$ by construction, $f(u, z_i) \in N_u$, and by choice of notation, let's call $f(u, z_i) = x_i$. By a symmetric argument $f(v, z_i) = y_i$.

We then concluded that $|N_{uv}| = |N_u| = |N_v|$. Furthermore, we concluded that the only N_u to N_{uv} edges were of the form $x_i z_i$ [6], and similarly, the only N_u to N_{uv} edges were of the form $y_i z_i$. This affords us a more complete picture of Γ_a; refer to Figure C.4 to see what information we have thus far.

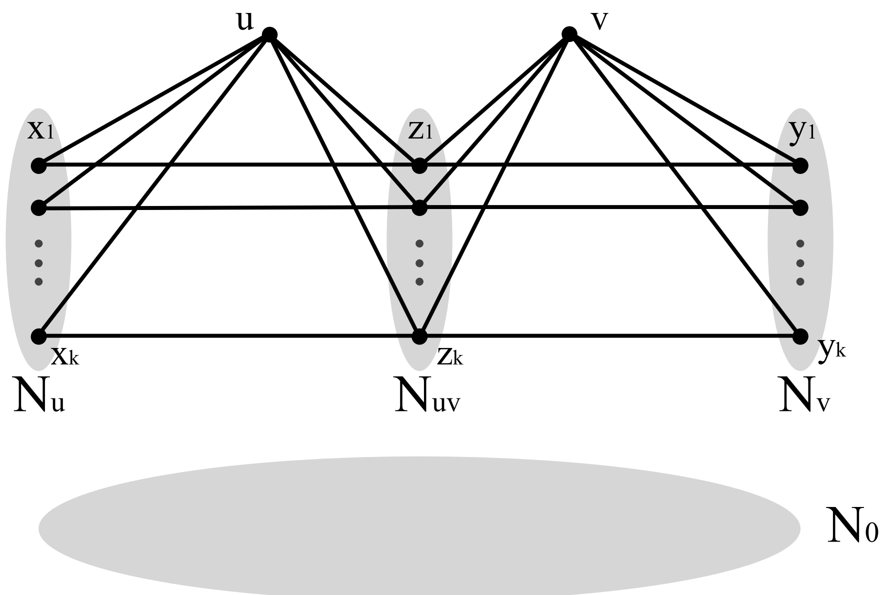

Figure C.4. Part of Γ_a, note the N_u to N_v edges are omitted.
[[LABEL F:0709: Gamma a second]]

At this point we had also proved that Γ_a is regular (of degree $2k$), that $\Gamma_a{:}N_u$, $\Gamma_a{:}N_v$, and $\Gamma_a{:}N_{uv}$ all have no edges, and that $\Gamma_a{:}(N_u \cup N_v)$ is a complete bipartite graph, minus the perfect matching $x_i \sim y_i$. Where this leaves us is that Figure C.4 is accurate, except for edges with at least one end point in N_0, and except for the $x_i \sim y_j$ for $i \neq j$ edges. We omit the N_u to N_v edges because they simply clutter up the diagram, and have yet to draw any conclusions about the edges with at least one endpoint in N_0.

From here we proved a few more things about the structure of Γ_a. We showed that Γ_a induced in the neighborhood of any vertex is a windmill (with k blades), which has proved to be a very useful fact. See Figure C.5 which focuses on the neighborhood of a vertex in N_{uv}. We also showed that Γ_a has $n = 6k - 3$ vertices total, and $m = k(6k - 3)$ edges total, and than, $\Gamma_a{:}N_0$ has $n_0 = 3k - 5$ vertices and $m_0 = k(k - 2)$ edges.

Next we recall a few sublemmas to Lemma 10 **??** to get a handle on the structure of N_0 (both within N_0, and how the vertices of N_0 are adjacent to the rest of Γ_a).

---

[6]If there were an $x_i z_j$ edge for $i \neq j$, then the edge uz_j would be in a triangle with both x_j and x_i, a contradiction.

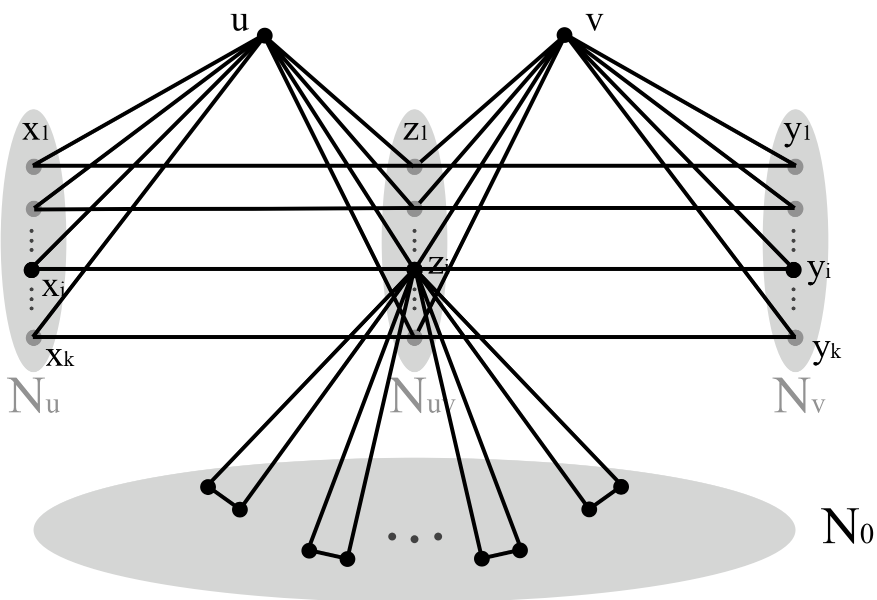

FIGURE C.5. Part of Γ_a, focusing on z_i.
[[LABEL F:0709: Gamma a third]]

Sublemma C.21. [[LABEL C:0709 sublemma 10.E]] *For $w \in N_0$ and for all $i \in [k]$, exactly one of the following is true:*

(1) $w \sim x_i$ and y_i.
(2) $w \sim z_i$.
(3) Neither (1) *nor* (2)*.*

In particular, (1) *and* (2) *cannot both be true.*

(This was originally Sublemma **??**.)

Furthermore, we established that for $w \in N_0$, either $w \sim x_i, y_i$ for either exactly 2 i's, or w is not adjacent to any vertices in N_u, N_v. Furthermore, $\deg_0(w) = \#$ of vertices $z_i \in N_{uv}$ that are adjacent to w, where $\deg_0(w)$ is the degree of $w \in \Gamma_a{:}N_0$. In particular, $\deg_0(w) = k$ if $w \sim x_i, y_i$ for exactly 2 i's, or $\deg_0(w) = k - 2$ if $w \not\sim x_i, y_i$ for any i. We then proved Sublemma 10F **??**.

Sublemma C.22. [[LABEL C:0709 sublemma 10.F]]

(1) $f(x_i, y_j) \neq f(x_i, y_k)$ for $k \neq j$.
(2) $f(x_i, y_j) = f(x_j, y_i)$ for $i \neq j$.

(This was originally Sublemma **??**.)

Part (2) of Sublemma C.22 is illustrated in Figure C.6. The sublemma says that given the solid edges (of the form $x_i y_j$ and $x_j y_i$) the two unique triangles on these two edges has the same third point (w in this figure), which will be in N_0.

The last part of the proof of Lemma 10 C.12 that we did last semester (and hence the last part in this review) is to look at the possibilities for Γ_a when $k = 2$. We can immediately start with $|N_{uv}| = |N_u| = |N_v| = 2$, and with those standard edges. Additionally, since $n = 6k - 3 = 9$, (or since $n_0 = 3k - 5 = 1$) we know that $|N_0| = 1$, we will name w the single

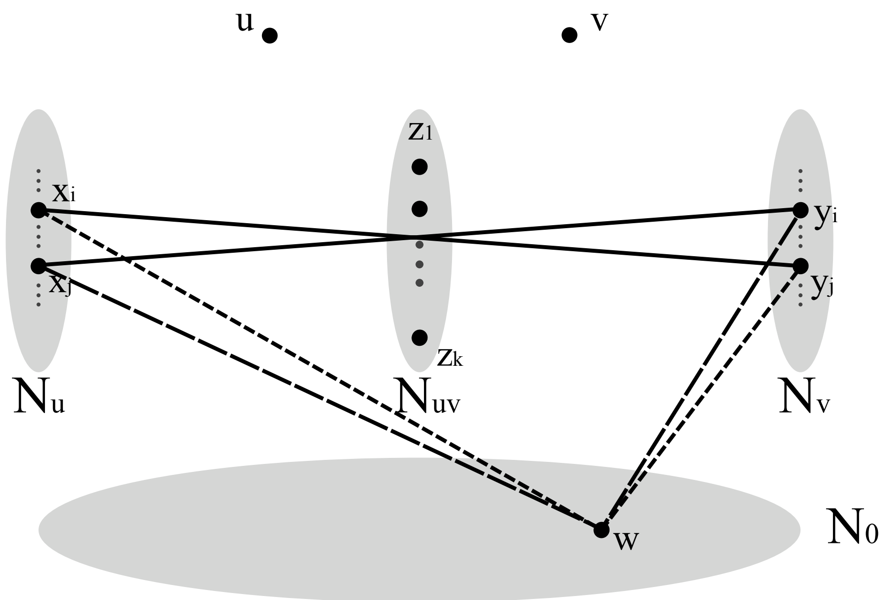

FIGURE C.6. Illustrating Part 2 of Lemma C.22 in Γ_a.
[[LABEL F:0709: Gamma a forth]]

vertex in N_0. We also know the graph should be $2k = 4$-regular. This leaves us with Figure C.7.

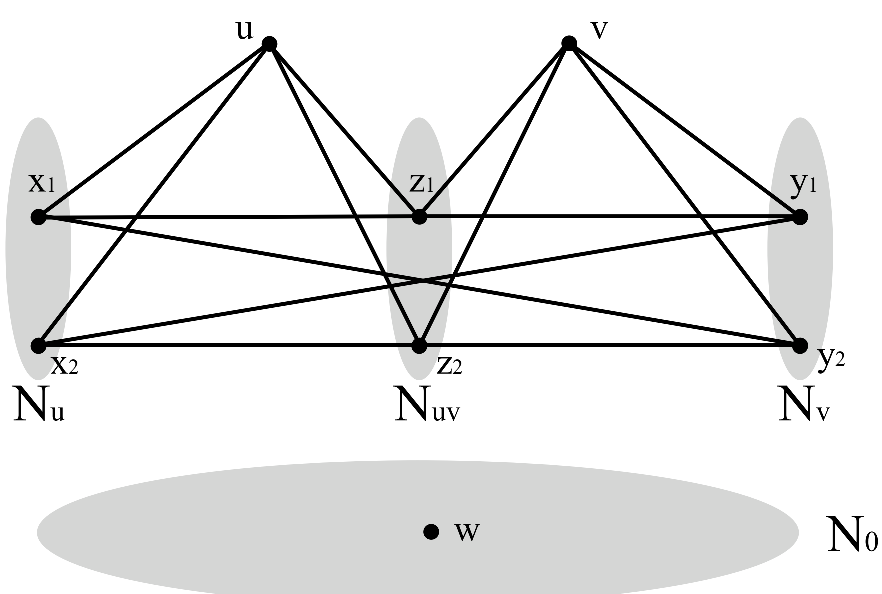

FIGURE C.7. The 'obvious' information about Γ_a when $k = 2$.
[[LABEL F:0709: k=2 first]]

This leaves us with very little information to add. Because u, v, z_1, z_2 all have degree $2k = 4$, there are no more edges incident to these four vertices. Furthermore, x_1, x_2, y_1, y_2 all need one edge to N_0, which in this case is the vertex w, and once we add those four edges, we are done. Note that we also could have considered the vertex w, and since $\deg_0(w)$ must be $0 = k-2$, w must be adjacent to x_i and y_i for exactly two i's in $\{2\}$—in other words, for every i. This leaves us with the unique Γ_a for $k = 2$, which is depicted in Figure C.8.

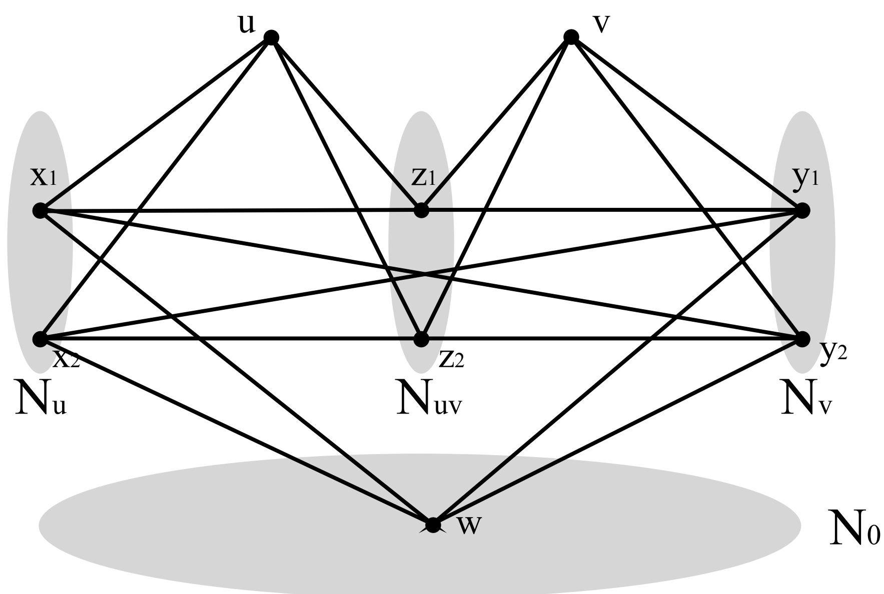

FIGURE C.8. Γ_a when $k = 2$.
[[LABEL F:0709: k=2 second]]

At this point what's left to do in the proof of Lemma 10 C.12 is to determine what the possible Γ_a's are for larger k (when Γ_a is not a windmill or $\bar{K}_l$).

(Incidentally, some of this proof about the structure of Γ_a might possibly have been simpler if we had first proved that in the closed neighborhood of any vertex Γ_a is a windmill.)

C.4. **Root systems (3/10–12).**

C.5. **Gramian representation. Gram matrix vs. adjacency matrix (3/10–12).**

C.6. **Angle representation. Bipartite and nonbipartite graphs (3/10–26).**

C.7. **Root representation vs. root systems (3/26–7/23).**

2009 July 14: Nathan Reff PARTIALLY EDITED

C.8. **Angle Representations.** [[LABEL 3.angle]]

C.8.1. *The proof of Lemma 10 C.12, so far.*
Recall the following data for Γ_a: $n = 6k-3 = 3(2k-1)$, $m = 3k(2k-1) = 6k^2 - 3k$, $d = 2k$. For $\Gamma_a{:}N_0$ we have: $n_0 = 3k-5$, $m_0 = k(k-2)$. Also, since Γ_a is $2k$-regular we can possibly find some upper bound on k.

We can construct the following table of information; we'll prove the rightmost column throughout the lecture:

k	n	d	n_0	m_0	Graph
2	9	4	1	0	$L(K_{3,3}) = P(9)$
3	15	6	4	3	$\overline{L(K_6)}$
4	21	8	7	8	(none)
5	27	10	10	15	Schläfli

where $P(9)$ is the Paley graph of order 9 [GR].

C.8.2. *Continuation of the proof of Lemma 10 C.12: The cases $k \geq 3$.*
See Figure C.9 for the basic picture in the case $k = 3$. There are two possibilities for the

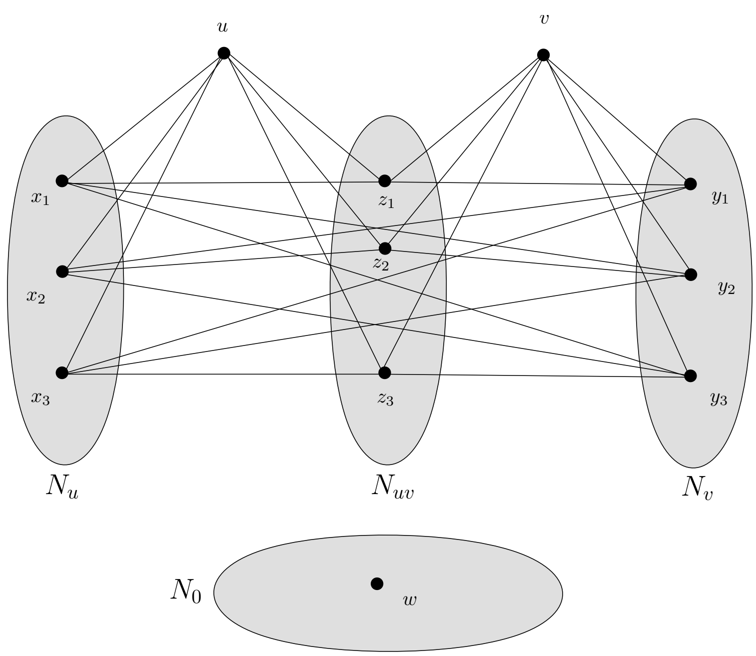

FIGURE C.9
[[LABEL F0714: k3case]]

valency of the w vertex in N_0. If w is adjacent to all the z vertices then $\deg_0(w) = k$. If w

is adjacent to all but two of the z vertices then $\deg_0(w) = k - 2$. Therefore the number of z vertices adjacent to w is $\deg_0(w)$.

In N_0 we are counting pairs of x vertices. Notice the different types of pairs in Figure C.10. For each pair of x vertices, there exists a unique w adjacent to both. Notationally, if

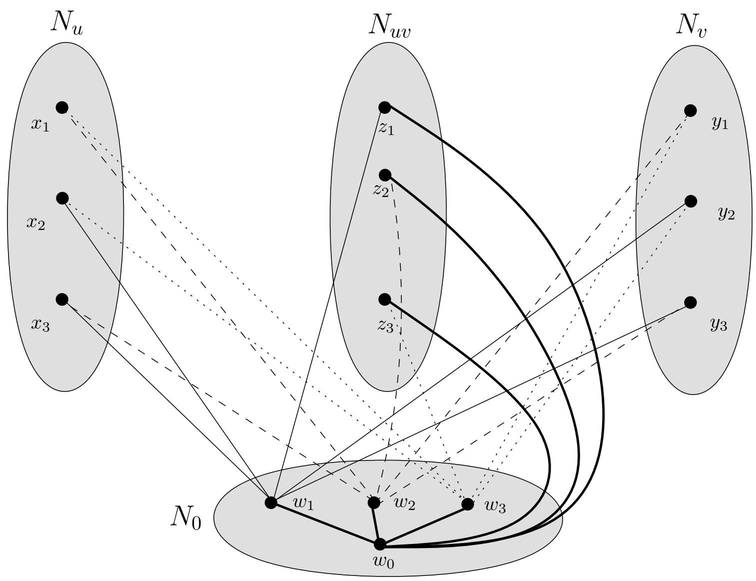

FIGURE C.10
[[LABEL F0714: pairs]]

$w_i \sim x_j, x_k$ then $w_i \sim z_i$ where $\{i, j, k\} = \{1, 2, 3\}$. This leaves the vertex w_0 in N_0 which is not adjacent to any of the x vertices. Hence w_0 is adjacent to z_1, z_2, and z_3. We also know that w_i and w_j are not adjacent for $i, j \in \{1, 2, 3\}$ because if they were adjacent, then the edge w_iw_j would be contained in one of the following two triangles: $w_iw_jx_k$ or $w_iw_jy_k$, which is not allowed. Therefore Γ_a exists and is unique.

Now we demonstrate that $\Gamma_a = \overline{L(K_6)}$. Let's draw $\overline{L(K_6)}$ by first looking at K_6 (see Figure C.11). So we can construct Γ_a in this case (see Figure C.12).

Now let's try to construct Γ_a in the case when $k = 4$ (see Figure C.13). Here we have $n_0 = 3(4) - 5 = 7$ vertices, $m_0 = 4(4-2) = 8$ edges, and $\binom{4}{2} = 6$ pairs of x vertices. Therefore six of the w vertices have $\deg_0(w) = k - 2 = 4 - 2 = 2$ and are adjacent to two z vertices and two xy's (meaning two of the x_1y_1, x_2y_2, x_3y_3, x_4y_4). So w_0 has $\deg_0(w) = k = 4$ and is adjacent to all four z vertices. However now suppose $w_1 \sim z_i, z_j$. Since $w_0 \sim z_i, z_j$ we have two possible triangles on edge w_0w_1 (see Figure C.13) which are $w_0w_1z_4$ and $w_0w_1z_2$. This is impossible, hence $k = 4$ is impossible.

Now we'll construct Γ_a in the case when $k = 5$ (see Figure C.14). Here we have $n_0 = 3(5) - 5 = 10$ vertices, $m_0 = 5(5-2) = 15$ edges, and $\binom{5}{2} = 10$ pairs of x vertices. So each w vertex is adjacent to either two or zero x vertices. Since there are 10 pairs of x vertices, each w is adjacent to exactly two x vertices. Let's say $w_{ij} \sim x_i, x_j, y_i, y_j$ and also w_{ij} is adjacent to all other z_h where $h \in \{1, 2, 3, 4, 5\} \setminus \{i, j\}$. Notice that w_{12} and w_{1k} have two common z neighbors z_l and z_m where $l, m \neq 1, 2, k$. Therefore there is no edge $w_{12}w_{1k}$. In general, there is no edge $w_{ij}w_{ik}$ since this would imply two triangles on the edge $w_{ij}w_{ik}$. Therefore the 15 edges are of the form $w_{ij}w_{kl}$. The Peterson graph P has $E = \{(ij, kl)\}$, and

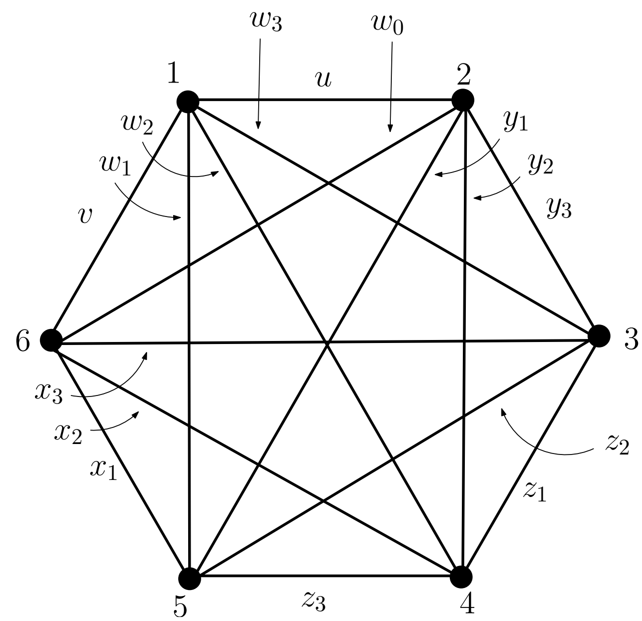

FIGURE C.11. K_6 with edges labelled as the vertices in Γ_a.
[[LABEL F0714: lk6]]

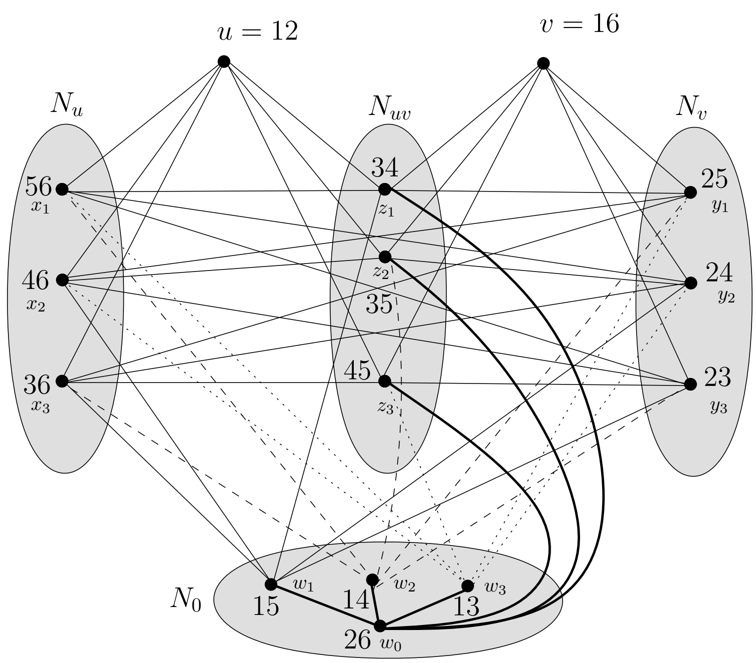

FIGURE C.12. $L(K_6)$ with vertices grouped as in Γ_a.
[[LABEL F0714: lk6const]]

therefore $\Gamma{:}N_0 \subseteq P$, but P has 15 edges and therefore $\Gamma{:}N_0 = P$ with standard labelling as $\overline{L(K_5)}$ (meaning we have labeled this exactly the way we want). Therefore Γ_a is completely determined. This completes the proof of Lemma 10 C.12. Hence, there exists Γ_a if $k = 2, 3, 5$ and there is no Γ_a for $k = 4$.

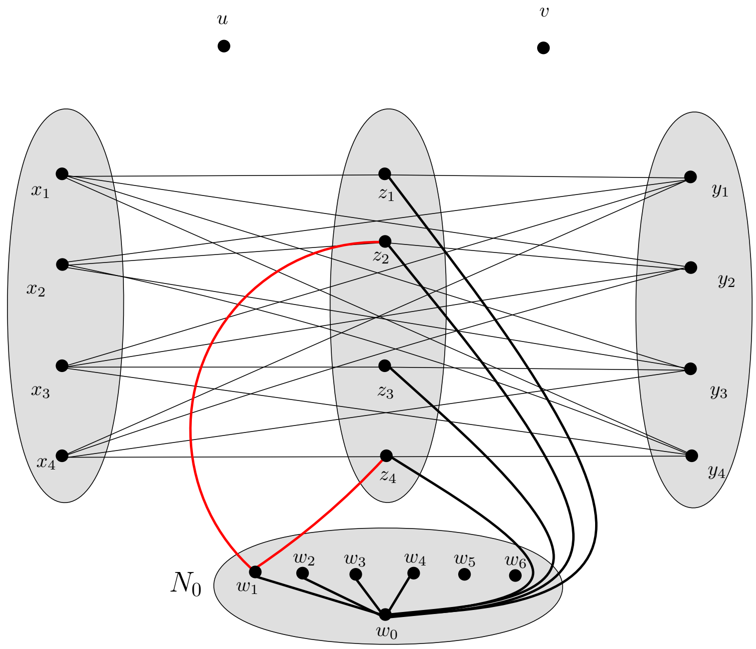

FIGURE C.13
[[LABEL F0714: k4case]]

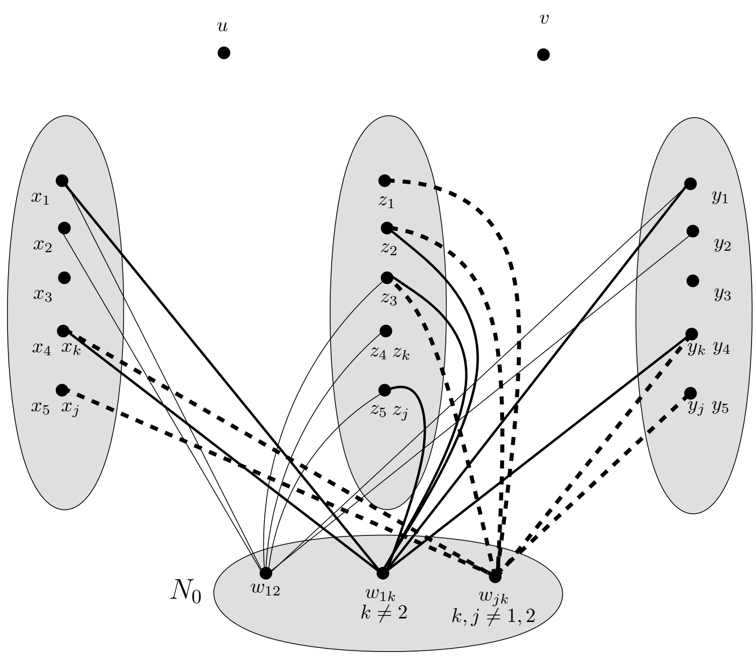

FIGURE C.14
[[LABEL F0714: k5case]]

Now let's do a much shorter proof of this Lemma 10 C.12. We make a new table and fill it in as we discuss strongly regular graphs.

k	n	d	λ	μ	graph	$\sqrt{\Delta}$	θ	τ	m_θ	m_τ
2	9	4	1	2	$L(K_{3,3}) = P(9)$	3	1	-2	4	4
3	15	6	1	3	$\overline{L(K_6)}$	4	1	-3	9	5
4	21	8	1	4	(none)	5	1	-4	fractional	fractional
5	27	10	1	5	Schläfli	6	1	-5	20	6

A *strongly regular graph* $\mathrm{SRG}(n, d, \lambda, \mu)$ is a d-regular simple graph of order n, such that any two adjacent vertices have λ common neighbors and any two nonadjacent vertices have μ common neighbors. (Note that the parameters here are given by different letters than in [GR, p. 218] where we would write $\mathrm{SRG}(n, k, a, c)$. Also see the notes on [**Dec 12 2008**] for a prior discussion of strongly regular graphs. [**Note: the parameters there are also different, we have written k for d, so perhaps we should change one of these to be consistent throughout. – Thanks Nate (TZ)**]

From Corollaries 7 and 8 [**REF ???**], every edge is contained in a unique triangle. We also proved that if u and v are nonadjacent vertices, then they must have k common neighbors and Γ_a is $2k$-regular (with $2 \le k \le 5$). This means that Γ_a is a $\mathrm{SRG}(6k-3, 2k, 1, k)$. We have the following result from the theory of strongly regular graphs.

Theorem C.23. *Suppose an* $\mathrm{SRG}(n, d, \lambda, \mu)$ *is connected. Then the eigenvalues of the adjacency matrix are* d *with multiplicity* 1, θ *with multiplicity* m_θ, *and* τ *with multiplicity* m_τ, *where*

$$\theta = \frac{\lambda - \mu + \sqrt{\Delta}}{2}, \qquad \tau = \frac{\lambda - \mu - \sqrt{\Delta}}{2},$$

$$m_\theta = \frac{1}{2}\left(n - 1 - \frac{2d + (n-1)(\lambda - \mu)}{\sqrt{\Delta}}\right), \quad m_\tau = \frac{1}{2}\left(n - 1 + \frac{2d + (n-1)(\lambda - \mu)}{\sqrt{\Delta}}\right),$$

$$\Delta = (\lambda - \mu)^2 + 4(d - \mu).$$

Proof. We find the discriminant of the quadratic polynomial in A (the adjacency matrix) for the strongly regular graph to be $\Delta = (1-k)^2 + 4k = (1+k)^2$ [**REF ??? I think showing the details of this was done on 12/12/2008 – NR**]. So we can write $\sqrt{\Delta} = 1 + k$ which is nice to work with. We have $\theta = \frac{1-k+k+1}{2} = 1$ and $\tau = \frac{1-k-k-1}{2} = -k$.

Notice that

$$\begin{aligned} \frac{2d + (n-1)(\lambda - \mu)}{\sqrt{\Delta}} &= \frac{4k + (6k-4)(1-k)}{k+1} \\ &= \frac{-2}{k+1}(3k^2 - 7k + 2) \\ &= -2(3k - 10) - \frac{24}{k+1}. \end{aligned}$$

Substituting into m_θ we obtain:

$$\begin{aligned} m_\theta &= \frac{1}{2}\left(n - 1 - \frac{2d + (n-1)(\lambda - \mu)}{\sqrt{\Delta}}\right) \\ &= \frac{1}{2}\left(6k - 4 + \left(6k - 20 + \frac{24}{k+1}\right)\right) \\ &= 6k - 12 + \frac{12}{k+1}. \end{aligned}$$

Since m_θ is a multiplicity, it must be an integer, which means that $k+1$ must divide 12. Compare the table we constructed with the information in Table 10.1 of [GR, p. 227]. □

7/16: Joyce

2009 July 21: Jackie Kaminski

C.9. **The signed graphs with root representations (continued).**

We need to finish showing the existence of root systems that have the signed graphs (and Γ_as) described above. We now turn our attention to two examples.

Example C.2. [[LABEL E:0721 empty graph]] To get $\Gamma_a = \bar{K}_l$, we take $S = A_{n-1}$ with $a = \mathbf{b}_2 - \mathbf{b}_1$ and $b = \mathbf{b}_3 - \mathbf{b}_2$. Then $\Gamma_a = \bar{K}_{n-3}$ (so $l = n-3$). The details of this example are from a previous day.

Example C.3. [[LABEL E:0721 windmill]] To get Γ_a to be a windmill with l blades, we take $S = D_n := \{\pm\mathbf{b}_i \pm \mathbf{b}_j \mid i \neq j \text{ and } i,j \leq n\}$ (where the $\mathbf{b}_i$'s are standard basis vectors). Then take $a = \mathbf{b}_2 - \mathbf{b}_1$ and $b = \mathbf{b}_3 - \mathbf{b}_2$, which makes $c = -(a+b) = \mathbf{b}_1 - \mathbf{b}_3$. Then $S_a = \{x \in D_n \mid x \perp a \text{ and } x \cdot b = -1[\ \mathbf{+1}\ ]\}$.

From the condition that $x \perp a$, we conclude that $i,j \geq 3$ or $x = \pm(\mathbf{b}_2 + \mathbf{b}_1)$. We will consider these two cases separately.

Case 1: $x = \pm\mathbf{b}_i \pm \mathbf{b}_j$ *with* $i,j \geq 3$.

From the condition that $x \cdot b = 1[\ \mathbf{+1}\ ]$, we get $(\pm\mathbf{b}_i \pm \mathbf{b}_j) \cdot (\mathbf{b}_3 - \mathbf{b}_2) = +1$ which implies $(\pm\mathbf{b}_i \cdot \mathbf{b}_3) + (\pm\mathbf{b}_j \cdot \mathbf{b}_3) + 0 + 0 = +1$ (since $i,j \geq 3$). The only way this can be true is if exactly one of $i,j = 3$ and the coefficient of $\mathbf{b}_3$ in x is $+1$. So in this case $x = \mathbf{b}_3 + \mathbf{b}_i$ for $i > 3$.

Case 2: $x = \pm(\mathbf{b}_2 + \mathbf{b}_1)$.

From $x \cdot b = +1$, so that $\pm(\mathbf{b}_2 + \mathbf{b}_1) \cdot (\mathbf{b}_3 - \mathbf{b}_2) = +1$, we see that $0 \mp 1 + 0 + 0 = +1$. So $x = -(\mathbf{b}_2 + \mathbf{b}_1)$ in this case.

Now we are close to determining Γ_a precisely. We know $S_a = \{-\mathbf{b}_2 - \mathbf{b}_1\} \cup \{\mathbf{b}_3 \pm \mathbf{b}_i \mid j > 3\}$. To find the orthogonality graph, Γ_a, we need to know which pairs of vectors in S_a are orthogonal.

Note that $-\mathbf{b}_2 - \mathbf{b}_1$ is orthogonal to all vectors $\mathbf{b}_3 \pm \mathbf{b}_i$ with $j > 3$. Since $(\mathbf{b}_3 + \mathbf{b}_i) \cdot (\mathbf{b}_3 - \mathbf{b}_i) = 0$, any pair $\mathbf{b}_3 + \mathbf{b}_i,\ \mathbf{b}_3 - \mathbf{b}_i$ are orthogonal. For $i,j > 3$ and $i \neq j$, we have $(\mathbf{b}_3 \pm \mathbf{b}_i) \cdot (\mathbf{b}_3 \pm \mathbf{b}_j) = 1 + 0 + 0 + 0 = 1$, so no two vectors in S_a of the form $\mathbf{b}_3 \pm \mathbf{b}_i,\ \mathbf{b}_3 \pm \mathbf{b}_j$ with $i \neq j$ are orthogonal. Therefore, Γ_a is the windmill in Figure C.15.

We conclude that, for D_n with $a = \mathbf{b}_2 - \mathbf{b}_1$ and $b = \mathbf{b}_3 - \mathbf{b}_2$, $S_a = \{-\mathbf{b}_2 - \mathbf{b}_1\} \cup \{\mathbf{b}_3 \pm \mathbf{b}_i \mid j > 3\}$ and Γ_a is a windmill with $n-3$ blades.

The signed graph represented by A_{n-1} is the all-positive line graph $+L(K_n)$ (Section I.I), which equals $-\Lambda(+K_n)$ (cf. Section II.M). We can see this from the fact that the vectors of A_{n-1} are the columns of an oriented incidence matrix of K_n (cf. Section I.**??**), or their negatives.

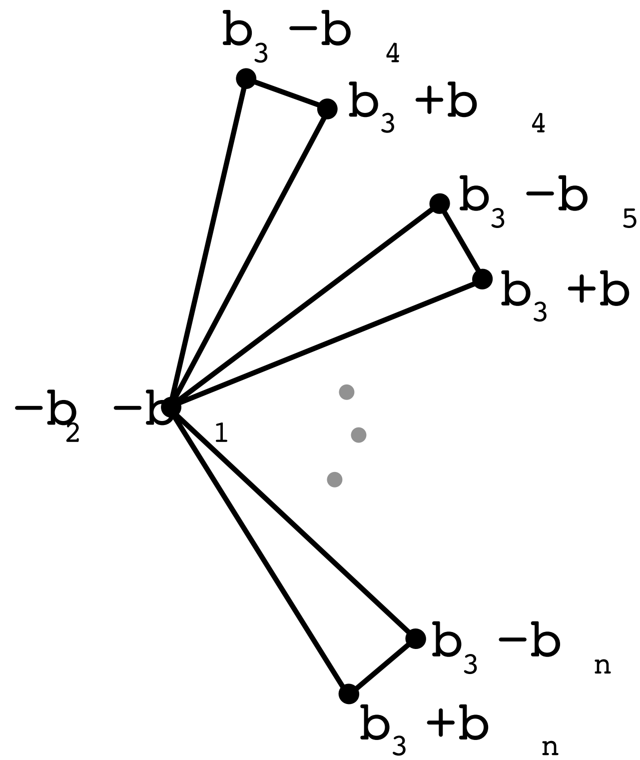

Figure C.15. Γ_a, the orthogonality graph for $S_a = \{-\mathbf{b}_2 - \mathbf{b}_1\} \cup \{\mathbf{b}_3 \pm \mathbf{b}_i \mid j > 3\}$
[[LABEL F:0721: windmill]]

But which signed graph is represented by D_n? The root system D_n is the set of all vectors of the form

$$\begin{pmatrix} \vdots \\ 0 \\ \pm 1 \\ 0 \\ \vdots \\ 0 \\ \pm 1 \\ 0 \\ \vdots \end{pmatrix},$$

where the $\pm$'s are independent of each other. For the signed graph representing D_n, remember that we are really talking about a switching class. For each pair of opposite vectors in D_n, $\mathbf{x}, -\mathbf{x}$, we have a vertex (denoted by x). The choice, across all vertices, between $\mathbf{x}$ and $-\mathbf{x}$ corresponds to a choice of a particular representative of the switching class. Once we have chosen a representative vector from each line $\langle \mathbf{x} \rangle$ (where $\mathbf{x} \in D_n$), we know the signed graph $\Sigma(D_n)$ has an edge vw with sign $\sigma(vw) \neq 0$ if and only if the vectors $\mathbf{v}$ and $\mathbf{w}$ have dot product $\sigma(vw)$.

Now let's order the vectors of D_n as the columns in a matrix M where

$$M = \left(\begin{array}{cc} \vdots & \\ 0 & \\ \pm 1 & \\ 0 & \text{with all} \\ \vdots & \text{possible} \\ 0 & \text{placements} \\ \pm 1 & \text{and signs} \\ 0 & \\ \vdots & \end{array} \right) = \big(\mathrm{H}(\pm K_n) \mid -\mathrm{H}(\pm K_n)\big)$$

In other words, the first half of the columns of M contains one choice of representatives from the pair $\mathbf{x}, -\mathbf{x}$, and the second half of the matrix contains the opposite choices in the same order. With the appropriate choice of labels, M is actually the oriented incidence matrix of $\pm K_n$,[7] augmented by the same incidence matrix from the opposite orientation of $\pm K_n$.

Now, $M^{\mathrm{T}}M$ is a $4\binom{n}{2} \times 4\binom{n}{2}$ matrix, indexed by the vectors in D_n, whose $(\mathbf{x}, \mathbf{y})$ entry is

$$\begin{cases} 0 & \text{if } \operatorname{supp}(\mathbf{x}) \cap \operatorname{supp}(\mathbf{y}) = \varnothing \text{ (hence } \mathbf{x} \perp \mathbf{y}), \\ 1 & \text{if } \operatorname{supp}(\mathbf{x}) \cap \operatorname{supp}(\mathbf{y}) = \{k\} \text{ and } \mathbf{x}_{(k)} = \mathbf{y}_{(k)}, \\ -1 & \text{if } \operatorname{supp}(\mathbf{x}) \cap \operatorname{supp}(\mathbf{y}) = \{k\} \text{ and } \mathbf{x}_{(k)} = -\mathbf{y}_{(k)}, \\ 0 & \text{if } \{\mathbf{x}, \mathbf{y}\} = \{\pm(\mathbf{b}_i - \mathbf{b}_j), \pm(\mathbf{b}_i + \mathbf{b}_j\}, \\ 2 & \text{if } \mathbf{x} = \mathbf{y}, \\ -2 & \text{if } \mathbf{x} = -\mathbf{y}. \end{cases}$$

(Here $\mathbf{x}_{(k)}$ denotes the kth component of $\mathbf{x} \in \mathbb{R}^n$.)

Going back to the structure of M as $\big(\mathrm{H}(\pm K_n) \mid -\mathrm{H}(\pm K_n)\big)$, we see that

$$M^{\mathrm{T}}M = \begin{pmatrix} \mathrm{H}(\pm K_n)^{\mathrm{T}}\mathrm{H}(\pm K_n) & -\mathrm{H}(\pm K_n)^{\mathrm{T}}\mathrm{H}(\pm K_n) \\ -\mathrm{H}(\pm K_n)^{\mathrm{T}}\mathrm{H}(\pm K_n) & \mathrm{H}(\pm K_n)^{\mathrm{T}}\mathrm{H}(\pm K_n) \end{pmatrix}$$

The upper left corner is $-A(\Lambda(\pm K_n)) = A(-\bar{\Lambda}(\pm K_n))$. Therefore, $\Sigma(D_n) = -\bar{\Lambda}(\pm K_n)$.

I point out here that representing a signed graph Σ by a set of vectors ($W \subseteq \mathbb{R}^\infty$) means we take one vector for each vertex. To represent Σ as a subset of D_n, we take at most one of each opposite pair $\pm x \in D_n$. Therefore we're choosing at most one orientation of the edges in $\pm K_n$ that corresponds to the vectors in D_n. Furthermore, for any root representation of Σ we have $\psi : V \to D_n$ where ψ is injective such that $\psi(v) \cdot \psi(w) = \sigma(v, w)$ (and $= 0$ if $v \not\sim w$). Obviously, this assumes that Σ is a signed simple graph, so that $\sigma(v, w)$ is well defined. This definition implies that we can't have $\psi(v) = -\psi(w)$ for any distinct vertices v, w, since $\psi(v) \cdot (-\psi(v)) = -2$. Therefore, for $x \in D_n$, either x or $-x$ or neither, but not both, is in $\operatorname{Im}(\psi)$; said another say, $\{x, -x\} \not\subseteq \operatorname{Im}(\psi)$.

C.9.1. *The matrix of a root representation in D_n.*

[7] Recall that $\pm K_n$ is the complete signed link graph; it is the signed graph on n vertices, whose edges are a positive and negative link between every pair of vertices.

Let's look more closely at the matrix representation of a simply signed link graph Σ in D_n. A Gramian representation of Σ, written as the columns of a matrix, has the form

$$G(\Sigma) = \begin{pmatrix} \psi(v_1) & \cdots & \psi(v_n) \end{pmatrix} \subseteq \mathrm{H}(\pm K_n)$$

for some orientation of K_n. Remember that this is the vertex matrix of the root representation of Σ; it is *not* the incidence or adjacency matrix of Σ!

Since the vectors of D_n have norm $\sqrt{2}$, $G(\Sigma)^{\mathrm{T}}G(\Sigma)$ is a $V \times V$ matrix with 2's along the diagonal, and $\sigma(vw)$ in the (v, w) entry otherwise.[8] So $G(\Sigma)^{\mathrm{T}}G(\Sigma) = A(\Sigma) + 2I$.

Examining the Gram matrix G further, we see that the largest possible Σ with a root representation ψ where $\mathrm{Im}(\psi) \subseteq D_n$ is when $\psi(V) \cup -\psi(V) = D_n$. In this case, $G(\Sigma) = (\pm K_n)$; therefore $A(\Sigma) + 2I = \mathrm{H}(\pm K_n)^{\mathrm{T}}\mathrm{H}(\pm K_n) = 2I - A(\bar{\Lambda}(\pm K_n))$. [9]Therefore the signed graph Σ represented by all of D_n is $-\overline{\Lambda}(\pm K_n)$.

C.9.2. *Signed graphs with angle representation in D_n.*

In other words, we have shown that D_n (which contains A_{n-1}) represents the negatives of reduced line graphs of simply signed link graphs, since all simply signed link graphs are $\subseteq \pm K_n$. (One could think the standard definition of a vector representation may have gotten the sign wrong. In other words, introducing an artificial negative into the definition of a representation would have led to the cleaner conclusion that D_n represents reduced line graphs (of simply signed link graphs). However, we'll continue to follow established tradition.) This leads us to the following theorem.

Theorem C.24. [[LABEL T: 0721 Signed Graphs with Root Rep's]] *The signed graphs that have root representations are the negatives of the reduced line graphs of simply signed link graphs, and a finite number of exceptions, where each exceptional graph has order $n \leq 120$. All the exceptions are subgraphs of $\Sigma(E_8)$*

Definition C.5. [[LABEL D: 0721 Sigma of a vector set]] Let $\Sigma(W) :=$ the largest signed graph which has a root representation in W, where $\Sigma(W)$ is understood as the switching class $[\Sigma(W)]$.

Recall that the choices of whether to label the vertices of $\Sigma(W)$ with $\mathbf{x}$ or $-\mathbf{x}$ corresponds to a choice of representative from $[\Sigma(W)]$.

Now that we have proved that all the negatives of reduced line graphs of simply signed link graphs all have root representations (as subsets of D_n), it is left to show that the only other signed graphs with root representations are the finite number of exceptions, all of which are subgraphs of $\Sigma(E_8)$.

C.9.3. *The mystery of the E_n signed graphs.*

I close with two small notes. As far as I know, no one knows much about the signed graphs $\Sigma(E_n)$. I remind you that signed line graphs were based on an orientation of the original signed graph (which is arbitrarily chosen in most cases), so signed line graphs are really about switching classes of signed graphs, not individual signed graphs.

2009 July 23: Nate Reff

[8]As usual, we interpret σ as the extended sign function, where $\sigma(vw) = 0$ if $v \not\sim w$.

[9]Recall that $\bar{\Lambda}$ means the reduced signed line graph, where pairs of oppositely signed parallel edges are canceled. In terms of notation, think of $\bar{\Lambda}(\Sigma)$ as $\overline{\Lambda(\Sigma)}$, since the operation 'take the reduced line graph' is made up of the operation 'take line graph' followed by the operation 'reduce.'

C.10. **Eigenvalues bounded below by -2.**

I begin today with a 'complete' version, so to speak, of the theorem of Cameron, Goethals, Seidel and Shult [CGSS], that every simple graph with all eigenvalues ≥ -2 is a Hoffman generalized line graph or one of a finite number of exceptions. I give it the form that seems more suitable to signed graphs, by negating all the signs.

Theorem C.25. [[LABEL T:20090723evalues2]] *With finitely many exceptions, every signed simple graph with all eigenvalues ≤ 2 is a reduced line graph of a simply signed link graph. The exceptions have ≤ 120 vertices and are, up to switching, subgraphs of one exceptional signed graph on 120 vertices.*

Proof. Suppose we have a signed simple graph Σ with root representation ψ. Then

$$\begin{aligned} \lambda_{\max}(\Sigma) \leq 2 &\iff 2I - A(\Sigma) \text{ is positive-semidefinite} \\ &\iff 2I - A(\Sigma) \text{ is a Gram matrix} \\ &\iff 2I + A(-\Sigma) \text{ is a Gram matrix} \\ &\iff \Sigma \subseteq \bar{\Lambda}(\pm K_n) \text{ or } \Sigma \subseteq \Sigma(E_8) \end{aligned}$$

since by our classification theorem from [CGSS], $-\Sigma \subseteq -\bar{\Lambda}(\pm K_n)$ or $\psi(\Sigma) \subseteq E_8$ and $\Sigma(E_8)$ is the representing graph of E_8,

$$\begin{aligned} &\iff \Sigma = \bar{\Lambda}(\Sigma_0) \text{ for some } \Sigma_0 \subseteq \pm K_n \\ &\qquad \text{or } \Sigma = \Sigma(S) \text{ where } S \subseteq E_8. \end{aligned}$$ □

[CGSS] would look at the theorem slightly differently, because to them the 'correct' sign is given by unsigned line graphs. Thus, instead of eigenvalues at most $+2$, they would look at eigenvalues at least -2, and they would express the theorem in these terms:

Corollary C.26. [[LABEL C: 20090723evalues-2]] *With finitely many exceptions, every signed simple graph with all eigenvalues ≥ -2 is the negative of a reduced line graph of a simply signed link graph. The exceptions have ≤ 120 vertices and are, up to switching, subgraphs of one exceptional signed graph on 120 vertices.*

The exceptional signed graph, of course, is $-\Sigma(E_8)$ in the theorem and $\Sigma(E_8)$ in the corollary.

I hope Cameron et al. forgive me for putting words in their mouths, since they didn't actually look at signed graphs. They did recognize that negating the graph gives the opposite theorem—i.e., with negated eigenvalue bound—and so they produced two theorems, one characterizing the unsigned graphs with eigenvalues ≥ -2, and the other characterizing those with eigenvalues ≤ 2. The latter was a much less interesting collection of graphs; I cannot say I fully understand why, but it may have to do with the way eigenvalues shift when a vertex is deleted; see [GR, ?] for more on that.

C.11. **The exceptional signed graphs.**

We need to know more about the exceptions in the theorem. Actually, not enough is known, but there are some elementary facts to look at closely.

C.11.1. *Verification of E_n.*

Let's look back at our original definitions for E_n where $n = 6$, 7, or 8. Recall that

$$\begin{aligned} E_8 &:= D_8 \cup \Big\{ \tfrac{1}{2} \sum_{i=1}^{8} \varepsilon_i b_i \;\Big|\; \varepsilon_i \in \{-1, 1\} \text{ and } \prod_{i=1}^{8} \varepsilon_i = +1 \Big\}, \\ E_7 &:= \{ x \in E_8 \mid x \perp w_0 \text{ for any one } w_0 \in E_8 \}, \\ E_6 &:= \{ x \in E_7 \mid x \perp w_1 \text{ for any } w_1 \in E_8 \setminus E_7 \text{ except } \pm w_0 \}. \end{aligned}$$

Alternatively, E_6 is the subset of E_8 formed by the set of vectors orthogonal to a fixed pair of vectors with inner product ± 1 [GR, p. 274].

Homework: (a) Verify that E_8 has the correct angles. (b) Why is E_8 "vertex transitive", meaning that $\mathrm{Aut}(E_8)$ is transitive on E_8's vectors?

C.11.2. *Construction of E_6 and E_7.*

Let's go through an explicit construction of E_7 and E_6. Choose $w_0 = \frac{1}{2} \sum_{i=1}^{8} b_i$ and define $y_S := \sum_{i \in S} b_i$, for $S \subseteq \{1, \dots, 8\} = [8]$.

Proposition C.27. [[LABEL P:20090723E7]] $E_7 = A_7 \cup \{ w_0 - y_S \mid S \in \mathcal{P}^{(4)}([8]) \}$.

Proof. Homework exercise. □

Let $w_1 = w_0 - (b_7 + b_8) = \frac{1}{2}(b_1 + \ldots + b_6 - b_7 - b_8)$

Proposition C.28. [[LABEL P:20090723E6]] $E_6 = A_5 \cup \{\pm(b_7 - b_8)\} \cup \{\pm(w_1 - b_i - b_j) \mid 1 \leq i < j \leq 6\}$.

C.11.3. *The exceptional signed graph/s.*

Suppose $\mathfrak{W} \subseteq \mathbb{R}^n$ and $\mathfrak{W}$ is a root system A_{n-1}, D_n, or E_n. We write $\Sigma(\mathfrak{W})$ for the signed graph of which $\mathfrak{W}$ is a root representation. [**WHAT IS $\mathfrak{W}$ doing here? I don't see the purpose.**] The graph Γ_a is the graph of orthogonality. Therefore the complement Γ_a^c is the graph of nonorthogonality so it must have all $+$ edge signs. [**NO; the fact that it has all + signs is proved.**] The subgraph of $\Sigma(E_n)$ (for $n = 6, 7, 8$) that is induced by $S_a := \{ x \in E_n \mid x \perp a,\ x \cdot b = 1 \}$, where $\langle a, b \rangle = -1$, is an identifiable all-positive signed graph, namely $+\Gamma_a^c$. Therefore $\Sigma(E_n)$ has an all-positive induced subgraph $+\Gamma_a^c$. (This may not be the largest all-positive induced subgraph. Homework exercise: Find the maximal all-positive induced subgraphs.)

Open Problem 1: Identify the signed-graph switching classes which are $[\Sigma(E_n)]$ or $[-\Sigma(E_n)]$, similarly to how we identified $[\Sigma(D_n)] = [-\bar{\Lambda}(\pm K_n)]$ and $[\Sigma(A_{n-1})] = [-\bar{\Lambda}(-K_n)]$.

Open Problem 2: Show how $+\Gamma_a^c$ fits into $\Sigma(E_n)$. We have a partial solution to this problem. We had to choose $\langle a, b \rangle = -1$, so that means $ab \in E^-(\Sigma(E_n))$. If $\langle x, a \rangle = 0$ we can say $x \notin N(a)$ in $\Sigma(E_n)$. If $\langle x, b \rangle = -1$ we can say $x \in N^-(b)$ in $\Sigma(E_n)$. Therefore $S_a \leftrightarrow N(a)^c \cap N^-(b)$ where a and b are any negative neighbors, and $\Gamma_a = \Sigma(E_n){:}S_a$.

Do we have to switch $\Sigma(E_n)$ in a particular way to get S_a as the alleged set of (non)neighbors? For example, do we have to assume $\Sigma(E_n)$ is switched so all edges at b are negative? This appears to imply that if we switch $\Sigma(E_n)$ so that all edges at a vertex v are negative, then $\Sigma(E_n){:}N(v)$ is all positive. Is that true?

We write $N[v]$ for the closed neighborhood of v, $N(v) \cup \{v\}$. Open Problem 2 is equivalent to the following: $\Sigma(E_n){:}N[v]$ is balanced for all v.

Semi-Open Problem: Identify the maximal induced balanced and antibalanced subgraphs of $\Sigma(E_n)$. The former are (or correspond to) maximal non-line **[CAN THAT BE RIGHT? maximal non-positive?]** graphs with eigenvalues ≥ -2, and the latter are (or correspond to) maximum non-negative line graphs with eigenvalues ≤ 2. This is possibly solved by [CGSS]. **[TZ SHOULD figure out what that means.]**

D. Equiangular and semi-equiangular lines

7/29: Joyce

2009 July 30: Jackie Kaminski **[Finishing the chapter!]**

D.1. **Equiangular Lines (continued).** [[LABEL equiangular]]

More on two-graphs.

In looking for 'large' sets of equiangular lines in 'small' dimensions we are lead to the problem of finding signed graphs $\Sigma = (K_n, \sigma)$ in which $\rho_{\min}(\Sigma)$ has high multiplicity, where $\rho_{\min}(\Sigma)$ denotes the minimum eigenvalue of $A(\Sigma)$. We would also like to acknowledge that we have made no effort to make precise what a 'large' sets of equiangular lines in 'small' dimensions is, this is mostly because there is no general theory yet about what pairs of 'large' and 'small' are possible or impossible. We now continue to explore this question from our point of view of switching and eigenvalues.

Recall from Section II.**??** that we associate to each signed complete graph $\Sigma = (K_n, \sigma)$ the two-graph $\mathcal{T} \subseteq \mathcal{P}^{(3)}(V)$ where $\{x, y, z\} \subseteq V$ is in $\mathcal{T}$ if and only if $\Sigma{:}\{x, y, z\}$ is a negative triangle.

Definition D.1. [[LABEL D:0730 two-graph of Sigma]] For Σ a signed graph, let $\mathcal{T}(\Sigma) :=$ $\{$triples of vertices that support negative triangles in $\Sigma\}$.

To see that this definition actually gives us a two-graph we need to show that every 4 vertices contain an even number of elements of $\mathcal{T}$. Note that if $\Sigma = +K_n$, then $\mathcal{T}$ is empty and the result holds vacuously. Now consider replacing any positive edge e with a negative edge. Any set of 4 vertices that does not containing e is unaffected, and for any set of 4 vertices containing e, the swap changes the sign of exactly two triangles on those 4 vertices[10], meaning there are still an even number of negative triangles. So inductively, we see that $\mathcal{T}(\Sigma)$ as defined above is actually a two-graph.

In general the map from signed graphs to two-graphs is not one-to-one, specifically, if Σ and Σ' are switching equivalent than $\mathcal{T}(\Sigma) = \mathcal{T}(\Sigma')$. However if we restrict our domain to switching equivalence classes then the map $\{[\Sigma]\} \to \{\mathcal{T}\}$ given by $[\Sigma] \mapsto \mathcal{T}(\Sigma)$ is one-to-one, and in fact a bijection.[11]

Recall from Chapter II (Theorem L.6) that Seidel did something similar with unsigned graphs Γ. The corresponding signed complete graph is simply K_Γ, where an edge $e \in E(K_n)$ is negative if and only if $e \in E(\Gamma)$.

Definition D.2. [[LABEL D:0730 regulartg]] $\mathcal{T}$ is *regular* is every vertex pair is in the same number of triples in $\mathcal{T}$.

[10] Since Σ is a signed complete graph.

[11] Where we interpret $\{[\Sigma]\}$ as the set of switching classes of signed complete graphs on V and $\{\mathcal{T}\}$ as the set of all two-graphs on V.

By the adjacency matrix of $\mathcal{T}$, we mean $A(\Sigma)$ for any signed graph Σ corresponding to $\mathcal{T}$. Note that for any subset $\mathcal{T} \subseteq \mathcal{P}^{(3)}(V)$, $\mathcal{T}$ is a two-graph if and only if $\mathcal{T}^c$ is a two-graph, and similarly $\mathcal{T}$ is regular $\iff$ $\mathcal{T}^c$ is regular.

Theorem D.1. [[LABEL T:0730 eigenvalues]] *A two-graph $\mathcal{T}$ is regular if and only if $A(\mathcal{T})$ has at most two distinct eigenvalues.*

Proof of ($\Longrightarrow$). For any $\mathcal{T}$, pick a $\Sigma \longleftrightarrow \mathcal{T}$, recalling that Σ is some signing of K_n. Since $\mathcal{T}$ is regular, every edge of Σ is in the same number of negative triangles, say t of them. Now choose a vertex v and switch Σ to Σ' where all the edges incident to v are positive.[12] See Figure D.1.

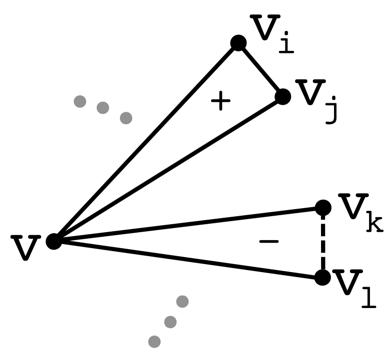

Figure D.1. A portion of a switched Σ.
[[LABEL F:0730: switched]]

Now, we can see by direct combinatorial calculation (this is Theorem II.??) that the (i, j) entry of $A(\Sigma)^2$ is the number of positive paths of length 2 between v_i and v_j less the number of negative paths of length 2 between them. This means that the diagonal entries are $n-1$, since a length-2 path from v_i to v_i in a signed K_n must leave v_i along an edge, then return along the same edge; regardless of the sign of that edge, the path is positive.

Now we consider any two different vertices v, w. If the edge vw is positive then any positive length-2 vw-path makes a positive triangle, and a negative length-2 vw-path makes a negative triangle. See Figure D.2. Since Σ is a regular two-graph, the edge vw is in t negative triangles; so there must be t negative length-2 vw-paths (and consequently $n-2-t$ positive ones). This means that the (v, w) entry of $A(\Sigma)^2$ is $(n-2-t)-t = n-2-2t$. Refer again to Figure D.2: if vw is a negative edge, then there must be t positive length-2 vw-paths (since Σ is a two-graph), and $n-2-t$ negative paths. In this case the (v, w) entry of $A(\Sigma)^2$ is $t-(n-2-t) = -n+2+2t$.

We can combine both cases elegantly by saying that the (v, w) entry of $A^(\Sigma)^2$ equals $\sigma(vw) \cdot (n-2-2t)$, for $v \neq w$. The consequence is that we can describe $A(\Sigma)^2$ completely as

$$A^2 = (n-1)I + (n-2-2t)A.$$

This means that the matrix $A(\Sigma)$ satisfies a quadratic equation, and therefore has at most two eigenvalues, and therefore $A(\mathcal{T})$ has at most two eigenvalues. $\square$

[12] Recall that every signed graph can be switched to have specified signs on a spanning tree.

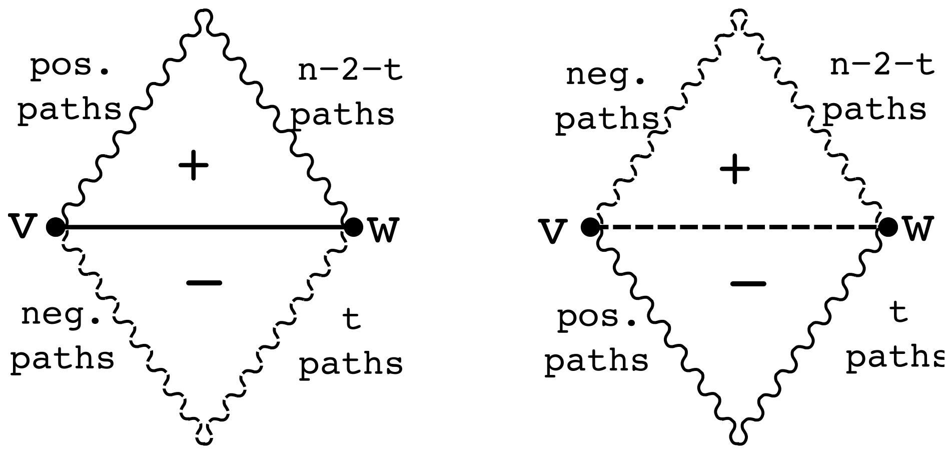

FIGURE D.2. Triangles on the edge vw in a signed complete graph Σ.
[[LABEL F:0730: triangles]]

One can calculate $A(\mathfrak{T})$ in general. Given any two-graph, not necessarily regular, the edge vw is in a certain number t_{vw} of negative triangles, and then

$$(A^2)_{vw} = (n-1)\delta_{vw} + \sigma(vw) \cdot (n-2-t_{vw}).[[\text{LABEL E}:0730\text{tga}]] \tag{D.1}$$

Proof of ($\Longleftarrow$). Now, we assume that $A(\mathfrak{T})$ has at most two eigenvalues. Thus, it satisfies a quadratic equation $A^2 - \alpha A - \beta I = 0$, or $A^2 = \alpha A + \beta I$. But by the calculation in the first part we see that $A^2_{vw} = (n-1)\delta_{vw} + \sigma(vw) \cdot (n-2-t_{vw})$, which means that β must equal $n-1$, and that $\sigma(vw)\alpha$ must equal $\sigma(vw)(n-2-t_{vw})$, and more importantly t_{vw} must equal some constant t for all v, w. Therefore each edge is in the same number of negative triangles, and therefore the two-graph is in fact regular. □

This mean that regularity of a two-graph is essentially an eigenvalue property. We would like to point out to the reader that in an earlier chapter we actually calculated the multiplicities of these eigenvalues.

[Tom: I have a note that 'this all started from people wanting a permutation representation of groups', but you weren't sure if that was right. So I don't really know what to do with that comment other than to remind you about it.]

Definition D.3. [[LABEL D:0730 Sigma v]] For $\Sigma = (K_n, \sigma)$, let Σ_v denote the unique switching of Σ where v has no negative neighbors.

Note that in some graphs there might be multiple switchings where a given vertex has no negative neighbors, but in a signed K_n, there is only one switching where this is true, hence the 'unique' is justified.

Theorem D.2. [[LABEL T:0730 G. R.]] $\mathfrak{T}(\Sigma)$ *is regular if and only if* $(\Sigma_v)^-$*, the spanning subgraph of the negative edges of* Σ_v*, is a* $SRG(n, k, \lambda, \mu)$*, with* $k = 2\mu$*.*

Readers with a further interest in the subject should see Godsil and Royle [GR], particularly Theorem 11.6.1, which is a stronger version of Therorem D.2.

Problem D.1. [[LABEL P:0730 SRG]] For Theorem D.2 , deduce k, λ, μ from n, t, to whatever extent possible. Are these independent of the choice of vertex v?

The end of equiangularity.

This treatment of regular two-graphs concludes our look at equiangular lines. Next, we look at a generalization suggested by general signed (simple) graphs.

D.2. **Semi-equiangular lines.**

The problem we have with equiangular lines is that root systems often give a system of lines with two angles, 60° and 90° for example. But we notice that 90° is a special angular, this leads to the definition of of semi-eqiangular lanes.

Definition D.4. [[LABEL D:0730 Semi-equi lines]] A set of lines is *semi-eqiangular* if every non-orthogonal pair of lines makes the same angle θ.

In other words, semi-equiangular lines allow any pair of distinct lines to make one of two angles, either a fixed angle $\theta < 90°$, or 90°. Examples of semi-eqiangular line sets include the root systems A_{n-1}, D_n, and E_8; simply think of each pair of opposite vectors $\mathbf{x}$, $-\mathbf{x}$ as defining a line.

Another important example is the root system G_2, a set of four lines in $\mathbb{R}^2$ at angles 45° and 90°. See Figure D.3.

[**Tom: You said to look up if G/R called this G_2, I couldn't find any info either way.**]

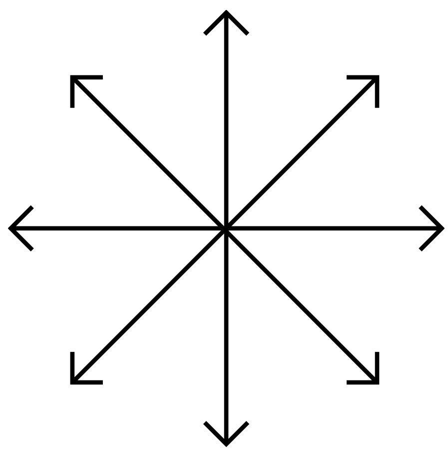

FIGURE D.3. A system of four semi-equiangular lines, generated by the root system G_2.

[[LABEL F:0730: 45]]

Now, for any system of semi-equianglar lines, we can choose one unit vector on each line. This will give us (a particular switching of) a signed graph, with the rule that $\sigma(vw) = \operatorname{sgn}(\cos \measuredangle(\hat{\mathbf{x}}_i, \hat{\mathbf{x}}_j)) = \operatorname{sgn}(\hat{\mathbf{x}}_i \cdot \hat{\mathbf{x}}_j)$. Since the angle between any two vectors is always $\leq \frac{\pi}{2}$, we may conclude that in a set of semi-eqiangular lines, $\theta \leq \frac{\pi}{2}$. And therefore $\hat{\mathbf{x}}_i \cdot \hat{\mathbf{x}}_j = \sigma(vw) \cdot \theta = \sigma(vw)/\gamma$ where $\gamma = 1/\cos\theta$, which we recognize from angle representations of signed graphs. Therefore we have $\mathbf{x}_i = \sqrt{\gamma} \cdot \hat{\mathbf{x}}_i$, so in an angle representation all vectors have norm $\sqrt{\gamma}$.

Conversely, if we start with a signed graph Σ, and create an angle representation, we will have $\|\mathbf{x}_i\| = \sqrt{\gamma}$ if $|\Sigma|$ is *not* bipartite. If $|\Sigma|$ is bipartite, we only know that $\|\mathbf{x}_i\| \cdot \|\mathbf{x}_j\| = \gamma$ for adjacent vertices v_i, v_j. We can then renormalize the vectors so all $\|\mathbf{x}_i\| = \sqrt{\gamma}$. In both cases, the angle between $\mathbf{x}_i$ and $\mathbf{x}_j$ is then $\cos^{-1}(\hat{\mathbf{x}}_i \cdot \hat{\mathbf{x}}_j) = \theta$ or $\pi - \theta$. (It is $\pi/2$ if $v_i \not\sim v_j$.)

Thus, $\theta = \arccos(1/\gamma)$ is the non-right angle of the semi-equiangular lines generated by the angle representation of Σ. Most of the elementary theorems about signed complete graphs go through, but we have no analog to two-graphs. This is a cop-out; we *could* develop more theory on the subject, but we won't. I do, however, leave the reader with the following conjecture (so named because I'm not certain whether it is true or not).

Conjecture D.1. [[LABEL C:0730 angle min]] Σ has an angle representation if and only if $\rho_{\min}(\Sigma) \leq 0$.

We leave this section by noting that one of the reasons we like regular two-graphs is that, unlike arbitrary two-graphs, they tend to correspond of systems of equiangular lines with a 'large' number of lines in 'small' dimension.

D.3. **Which signed graphs have angle representations?**

Recall that in a root system with $\gamma = 2$, the matrix $A(\Sigma) + 2I$ has eigenvalues ≥ 0, which means that $A(\Sigma)$ has e-values ≥ 2. Recall further that line graphs of signed graphs have eigenvalues ≥ -2. Therefore, if we had introduced an arbitrary negative into the definition of angle representations, we would have the same relationship.

Now notice that a signed graph Σ has an angle representation

$$\begin{aligned}
&\iff A(\Sigma) + \gamma I \text{ is positive semi-definite} \\
&\iff \rho_{\min}(A(\Sigma) + \gamma I) \geq 0 \\
&\iff \rho_{\min}(A(\Sigma)) + \gamma \geq 0 \\
&\iff \gamma \geq -\rho_{\min}(A(\Sigma)) = \rho_{\max}(A(-\Sigma)) \\
&\iff \gamma \geq \rho_{\max}(A(-\Sigma))
\end{aligned}$$

We are, of course, more interested in nontrivial angle representations.

Definition D.5. [[LABEL D:0730 nt angle rep]] An angle representation with n vectors is *nontrivial* if the vectors are in $\mathbb{R}^m$ and $m < n$.

An angle representation is nontrivial

$$\begin{aligned}
&\iff \operatorname{rk}(A(\Sigma + \gamma I)) < n \\
&\iff -\gamma \text{ is an eigenvalue of } A(\Sigma) \\
&\iff \gamma \text{ is an eigenvalue of } A(-\Sigma) \\
&\iff \gamma = \rho_{Max}(A(-\Sigma)), \text{ since } \gamma > 0 \text{ in the non-trivial angle representations.}
\end{aligned}$$

That proves the following theorem:

Theorem D.3. [[LABEL T:0730 nt angle rep]] *A signed graph Σ has a non-trivial angle representation $\iff \rho_{\max}(A(-\Sigma)) > 0$.* □

Notice that $\rho_{\max}(A(-\Sigma)) < 0 \implies \rho_{\max}(A(\Sigma)) > 0$. Therefore, if we think of signed graphs in pairs $\Sigma, -\Sigma$, than one of each pair of signed graphs has a nontrivial angle representation, unless all eigenvalues are 0, in which case $\Sigma = \overline{K_n}$.

We close this section with some questions about signed simple graphs in general, not necessarily complete.

(1) Q: What is the equivalent of a two-graph?
A: It's a switching class. (But what does that mean? Is there a more direct combinatorial description, analogous to that of complete switching classes by two-graphs?)
(2) Q: What's the equivalent of a *regular* two-graph?
A: We're not sure, but presumably it's a signed graph with ≤ 2 eigenvalues.
(3) Q: In view of the relevant theorem, what is a strongly regular signed graph? [**Does this mean the theorem about regular tg vs. srg? Do I have that?**]
A: Stick around a year or so.

Chapter IV. Gain Graphs and Biased Graphs

The monster in the cupboard throughout the theory of signed graphs has been the question: What about larger groups? It is time to answer the monster. Yes, we can have edge labels from groups with more than two elements—I call them 'gains'—but no, they are not just like signs, because for the theory to work, the gain of an edge has to depend on its orientation. Allowing for that difference, a great part of signed graph theory generalizes, not only to gains in any group, but even to a purely combinatorial abstraction, called a 'biased graph', that dispenses with groups altogether.

2010 Jan 28: Jackie Kaminski

A. Gain Graphs

[[LABEL 4.gg]]

We define our first subject.

A.1. **Basic definitions.** [[LABEL 4.ggdefs]]

Recall from Definition A.3 that a graph Γ is a triple (V, E, I) (but usually written (V, E)), where V and E are sets and I is an incidence multirelation between V and E in which each edge has incidences of total multiplicity at most 2. Consequently we have four types of edges: links, with two distinct endpoints, $I(e) = \{v, w\}$ with $v \neq w$; loops, which have two coinciding endpoints, $I(e) = \{v, v\}$ (these two kinds are *ordinary edges*); half edges, with one endpoint, $I(e) = \{v\}$; and loose edges, which have no endpoint, $I(e) = \varnothing$. See Figure A.1 for pictures of these four types.

Each type of edge, except a loose edge, has two possible orientations, as suggested by Figure A.2.

A.1.1. *What a gain graph is.* [[LABEL 4.ggis]]

We are now ready to define a gain graph. The set of oriented links and loops of a graph is $\vec{E}^*$. This set contains two copies of each ordinary edge, with opposite directions. If e is an edge in one direction, then e^{-1} denotes the same edge in the other direction.

Definition A.1. [[LABEL D:20100128: Gain Graph]] A *gain graph* is a graph whose links and loops are labelled invertibly by elements of a group. More precisely, a gain graph Φ consists of a graph $\|\Phi\|$, a group $\mathfrak{G}$, and a function $\varphi : \vec{E}^* \to \mathfrak{G}$ such that

$$[[\text{LABEL E}: 20100128\text{oriented}]]\varphi(e^{-1}) = \varphi(e)^{-1}, \tag{A.1}$$

where e^{-1} denotes the oriented edge e with its direction reversed.

We often write, for a gain graph, $\Phi = (\Gamma, \varphi)$, $(\Gamma, \varphi, \mathfrak{G})$, or (V, E, φ), etc., according to the needs of the situation; in particular, if we need to emphasize the gain group we call Φ a $\mathfrak{G}$*-gain graph*.

We call $\|\Phi\|$ the *underlying graph*, $\mathfrak{G}$ the *gain group*, and φ the *gain function*. The value $\varphi(e)$ is the *gain* of the (oriented) edge e.

By our definition, changing the gain group changes the gain graph. One can enlarge the group without changing the gains; that will give a new gain graph, though the difference is slight.

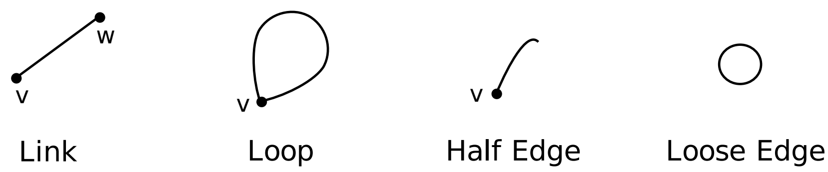

Figure A.1. The four types of edge in a graph
[[LABEL F:20100128: edges]]

I call a function *orientable*, or *invertible*, if it satisfies Equation (A.1). I am usually relaxed about how to write the gain function; for simplicity of notation I usually write $\varphi : E \to \mathfrak{G}$ or $\varphi : E^* \to \mathfrak{G}$; but φ is always defined on oriented ordinary edges and is always inverted by reversing the orientation; and it is never defined on half or loose edges.

Since φ is orientable, we must indicate (at least implicitly) the direction in which the gain is taken. I use several notations for the gain of e in the direction from v to w, depending on which is most clear in context; they all have the same meaning:

$$\varphi(e) = \varphi(e{:}vw) = \varphi(e_{vw}) = \varphi_{vw}(e) = \varphi(e{:}\overrightarrow{vw}).$$

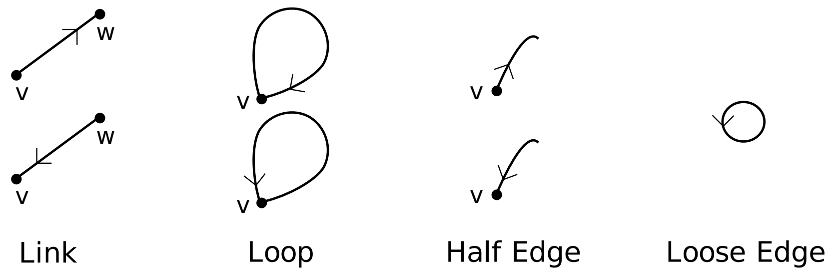

FIGURE A.2. The possible directions of an edge.
[[LABEL F:20100128: directed edges]]

The first of these is sufficient when we know the direction in which the gain is calculated, as for instance in computing the gain of a walk (see later). Several of the notations actually imply the direction of the edge, namely, $e{:}vw$, e_{vw}, and $e{:}\overrightarrow{vw}$. Notice that the notations are ambiguous for a loop; but that will almost never make any difficulty. (The reason for so many notations is partly that the one I like best, $e{:}vw$, is the least well known and the least immediately understandable, partly indecision, and partly that I have trouble remembering which one I'm using at the moment, so I may switch back and forth.) In whatever notation, we always have the inversion law (A.1), that is, $\varphi(e{:}wv) = \varphi(e{:}vw)^{-1}$.

A.1.2. *Groups.* [[LABEL 4.groups]]

We will write general gain groups multiplicatively. In the group $\mathfrak{G}$, $1 := 1_{\mathfrak{G}}$ is the identity element, unless (of course) the group is written additively. The conjugate of $\mathfrak{g}$ by $\mathfrak{h}$ is defined as

$$\mathfrak{g}^{\mathfrak{h}} := \mathfrak{h}^{-1}\mathfrak{g}\mathfrak{h} \quad \text{for } \mathfrak{g}, \mathfrak{h} \in \mathfrak{G}.$$

Some of the more important and common examples of gain groups are the two-element group $\mathfrak{G} = \{+, -\}$, in which case Φ is a signed graph, and the trivial group $\mathfrak{G} = \{1\}$. Other important gain groups are $\mathbb{Z}^+$, the additive group of integers, $\mathbb{R}^*$ (or $\mathbb{R}^\times$), the multiplicative group of reals, the multiplicative and additive groups F^* and F^+ of any field, and the finite cyclic groups $\mathbb{Z}_r$ or, isomorphically (but multiplicatively), the groups of complex r^{th} roots of unity.

A.1.3. *The free group of edges.* [[LABEL 4.freegroup]]

Another way to look upon a gain is as a homomorphism from a free group. Think of E^* as the generators of a free group, $\mathfrak{F}(E^*)$ (which for brevity I like to write $\mathfrak{F}_E$). We interpret a generator e as an edge in one orientation, and the free-group inverse e^{-1} as the same edge in the opposite direction. Then a gain function can be treated as a homomorphism $\varphi : \mathfrak{F}_E \to \mathfrak{G}$, since the values of φ on the generators, no matter what they happen to be, determined the homomorphism. Everything works out nicely; especially, the gain of a walk is the value $\varphi(W)$ of the homomorphism applied to the word in $\mathfrak{G}_E$ that expresses the walk.

This point of view is sometimes best, but usually I find the more direct interpretation of gains, as a function on edges, to be more suitable.

A.2. **Walk and circle gains.** [[LABEL 4.walkcirclegains]]

We now introduce more notation, in order to be able to talk about balance, which is a not-unexpected generalization of balance in a signed graph.

Definition A.2. [[LABEL D:20100128: Gain of a walk]] The *gain of a walk* is

$$\varphi(W) := \varphi(e_1)\varphi(e_2)\cdots\varphi(v_l),$$

where the walk is $W = v_0e_1v_1e_2v_2\cdots e_lv_l$, from v_0 to v_l.

According to this definition and the usual convention about an empty product, the gain of a walk of length 0 is $1_{\mathfrak{G}}$.

The gain of a walk is invertible. Since $W^{-1} = v_le_l^{-1}v_{l-1}\cdots e_1^{-1}v_0$,

$$\varphi(W^{-1}) = \varphi(e_l^{-1})\varphi(e_{l-1}^{-1})\cdots\varphi(e_1^{-1}) = \varphi(e_l)^{-1}\varphi(e_{l-1})^{-1}\cdots\varphi(e_1)^{-1} = \varphi(W)^{-1}.$$

Recall that a trail is a walk with no repeated edges and a path is a walk without repeated edges or vertices. A path or trail has a gain, as it is a walk. Recall also that a closed walk begins and ends at the same vertex, i.e., it has $v_0 = v_l$, where $l > 0$ (although $l = 0$ is allowed for a walk in general). The same applies to a closed trail or path (but remember that a closed path, since $v_0 = v_l$, is technically not a path!). Finally, recall that a circle is the edge set of a closed path.

Proposition A.1. [[LABEL P:20100128: Gain of a Circle]] *The gain of a circle is well defined up to conjugation and inversion.*

Proof. Let $W = v_0e_1v_1\cdots e_lv_l$ be a closed walk ($v_l = v_0$). Consider its gain, $\varphi(W)$. If we started and ended our walk at v_0, but traversed the circle in the opposite direction, the resulting gain would be $\varphi(W^{-1}) = \varphi(W)^{-1}$. If $l > 1$, we could also have started the walk at a different vertex. Choose $0 < k < l$, and define $W_k := v_k e_{k+1} v_{k+1} \cdots e_l v_0 e_1 \cdots e_k v_k$. Then

$$\begin{aligned}
\varphi(W_k) &= \varphi(e_{k+1})\varphi(e_{k+2})\cdots\varphi(e_l)\varphi(e_1)\cdots\varphi(e_k) \\
&= \big[\varphi(e_1)\cdots\varphi(e_k)\big]^{-1}\big[\varphi(e_1)\cdots\varphi(e_k)\big]\big[\varphi(e_{k+1})\varphi(e_{k+2})\cdots\varphi(e_l)\varphi(e_1)\cdots\varphi(e_k)\big] \\
&= \big[\varphi(e_1)\cdots\varphi(e_k)\big]^{-1}\big[\varphi(e_1)\cdots\varphi(e_k)\varphi(e_{k+1})\varphi(e_{k+2})\cdots\varphi(e_l)\big]\big[\varphi(e_1)\cdots\varphi(e_k)\big] \\
&= \big[\varphi(e_1)\cdots\varphi(e_k)\big]^{-1}\varphi(W)\big[\varphi(e_1)\cdots\varphi(e_k)\big] \\
&= \varphi(W)^{\varphi(e_1\cdots e_k)}.
\end{aligned}$$

Therefore the gain of a circle is only defined up to conjugation and inversion. □

While the gain of a circle, or any closed walk, isn't well defined, whether or not the gain equals $1_{\mathfrak{G}}$ is well defined, since the identity element is invariant under conjugation and inversion in $\mathfrak{G}$.

Definition A.3. [[LABEL D:20100128: Neutral]] We call a walk W *neutral* if $\varphi(W) = 1_{\mathfrak{G}}$.

Notice that the gain of a walk (ignoring direction) is well defined up to inversion, but only up to inversion, since it depends on the direction of the walk. The gain of a closed walk is well defined up to inversion and conjugation—conjugation, because the gain is conjugated if one changes the starting vertex; see Equation (A.2). Normally, though, a walk comes with direction and initial vertex, so these potential ambiguities do not arise.

A.3. **Balance.** [[LABEL 4.ggbal]]

Now we come to the fundamental notion of gain graph theory.

Definition A.4. [[LABEL D:20100128: Balance]] An edge set $S \subseteq E$ is *balanced* if every circle $C \subseteq S$ is neutral and S has no half edges. A subgraph is balanced if its edge set is balanced.

In particular, a circle is balanced if and only if it is neutral. We write $\mathcal{B}(\Phi)$ to denote the set of balanced circles of Φ.

There is a difference between balance and neutrality. The word "neutral" refers only to a walk, whereas the word "balanced" refers to an edge set (or subgraph). This is an important distinction. A walk can be neutral but not balanced, and vice versa. It is balanced if its underlying subgraph has only neutral circles, while it is neutral only when the product of edge gains in the order the edges appear in the walk is the group identity, regardless of whether its graph contains non-neutral circles.

The most obvious balanced gains are those which are identically 1. Write 1_E for that gain function; the gain graph is $(\Gamma, 1_E)$. This does not imply that the gain group is trivial; $\mathfrak{G}$ could be any group.

Definition A.5. [[LABEL D:20100128: Balanced Components]] Let $b(\Phi)$ denote the number of connected components of Φ that are balanced and contain at least one vertex. (That is, we exclude loose edges. A component that has a vertex is a *vertex component*; a loose edge is not a vertex component.)

Definition A.6. [[LABEL D:20100128: Contrabalance]] We call Φ *contrabalanced* if it contains no neutral circles.

Example A.1. [[LABEL X:20100128: Z3 Gain Graph]] This example is displayed in Figure A.3. Here Φ is a graph with gain group $\mathbb{Z}_3$, the integers modulo 3 under addition. The values of φ are given for each of the edges in the indicated direction. So $\varphi(a) = \varphi(g) = 2$ and $\varphi(b) = \varphi(e^{-1}) = 1$, etc. Let W denote the closed walk a, b, c, d, e, f, g (where vertex labels are omitted to streamline notatoin). Notice that,

$$\varphi(W) = \varphi(a) + \varphi(b) + \varphi(c) + \varphi(d) + \varphi(e) + \varphi(f) + \varphi(g) = 0 \in \mathbb{Z}_3.$$

So W is a neutral walk. But $\varphi(abcd) = 2$, so the circle *abcd* is not neutral. Therefore W is not balanced, and since $\varphi(efg) = 1$, W is in fact contrabalanced. However the set $S = \{a, b, f\}$ is balanced, since it contains no half edges and no circles (and consequently no unbalanced circles). Finally since $\|\Phi\|$ is connected, and Φ is contrabalanced, its only component is unbalanced, so $b(\Phi) = 0$.

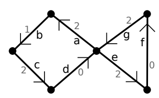

FIGURE A.3. A gain graph Φ with $\mathfrak{G} = \mathbb{Z}_3$. Gains are given for the indicated direction of each edge.
[[LABEL F:20100128: Z3 Gain Graph]]

Notice that for any $\|\Phi\|$, if φ is identically $1_{\mathfrak{G}}$, then Φ is in most ways just like $\|\Phi\|$. In particular $\mathcal{B}(\Phi) = \mathcal{C}(\|\Phi\|)$, where $\mathcal{C}$ is the set of all circles in the underlying graph.

The next result is the fundamental characteristic of balanced circles.

Proposition A.2. [[LABEL P:20100128: Theta Graph]] *No theta subgraph has exactly two balanced circles.*

Proof. Given a theta graph, let P_1, P_2, P_3 denote the three internally disjoint vw-paths. For a proof by contradiction, assume that two circles are neutral. By choice of notation we may assume that $\varphi(P_1P_2^{-1}) = \varphi(P_2P_3^{-1}) = 1_{\mathfrak{G}}$, and $\varphi(P_1P_3^{-1}) \neq 1_{\mathfrak{G}}$. Now notice that

$$\varphi(P_1P_2^{-1}) = \varphi(P_1)\varphi(P_2)^{-1} = 1_{\mathfrak{G}} \iff \varphi(P_1) = \varphi(P_2).$$

Similarly, $\varphi(P_2) = \varphi(P_3)$. Consequently, $\varphi(P_1) = \varphi(P_3)$. Therefore

$$\varphi(P_1P_3^{-1}) = \varphi(P_1)\varphi(P_3)^{-1} = \varphi(P_1)\varphi(P_1)^{-1} = 1_{\mathfrak{G}},$$

a contradiction. Thus, no theta subgraph can have two neutral circles unless the third circles is neutral as well. □

We'll see more properties of balance in Section A.5 after developing the theory of switching and potentials in Section A.4.

We end this section with two general notes. First, many interesting properties of gain graphs don't depend on the actual values of the gains; they only depend on the set of balanced circles. These properties tend to be closely related to Proposition A.2. Second, we don't have a Harary-type structure theorem for balanced gain graphs as we did for signed graphs. The nearest thing to such a theorem is (i) $\iff$ (iii) in Proposition A.6, in which the "structure" is a potential function which intrinsically depends upon the gain group.

2010 Feb 2: Nathan Reff

A.4. **Switching and potentials.** [[LABEL 4.sw]]

Let $\Phi = (\Gamma, \varphi, \mathfrak{G})$ be a gain graph with gain group $\mathfrak{G}$. A *switching function* (or *selector*) is any function $\zeta : V \to \mathfrak{G}$. *Switching* the graph Φ (or the gain function φ) means replacing φ by φ^ζ, defined by

$$\varphi^\zeta(e{:}vw) = \zeta(v)^{-1}\varphi(e{:}vw)\zeta(w),$$

giving us the gain graph $\Phi^\zeta = (\Gamma, \varphi^\zeta, \mathfrak{G})$. We reserve the notation $\bar\zeta$ for the inverse of values of ζ, i.e., $\bar\zeta(v) = \zeta(v)^{-1}$. We call Φ_1 and Φ_2 *switching equivalent*, written $\Phi_1 \sim \Phi_2$, when there exists a switching function ζ such that $\Phi_2 = \Phi_1^\zeta$. Notice that this entails equality—not isomorphism—of the underlying graphs (in other words, $\|\Phi_1\| = \|\Phi_2\|$). Notice also that $\Phi_2 = \Phi_1^{\bar\zeta}$.

Proposition A.3. [[LABEL P:20100202sw equiv]] *Switching equivalence is an equivalence relation.*

Proof. The proof is a good homework exercise. □

We say Φ_1 is *isomorphic* to Φ_2, written $\Phi_1 \cong \Phi_2$, if the gain groups are the same (not isomorphic) and there exists a graph isomorphism $\theta : \|\Phi_1\| \to \|\Phi_2\|$ which preserves gains—in other words, $\varphi_2(\theta(e)) = \varphi_1(e)$ (which we can write $\theta \circ \varphi_2 = \varphi_1$). The *switching class* of Φ, written $[\Phi]$ is the equivalence class under switching equivalence (note: not switching isomorphism).

We define Φ_1 and Φ_2 to be *switching isomorphic* if there exists a switching function ζ such that $\Phi_2 \cong \Phi_1^\zeta$.

Proposition A.4. [[LABEL P:20100202 Sw pres Balance]] *Switching preserves balance (and imbalance). In other words,* $\mathcal{B}(\Phi^\zeta) = \mathcal{B}(\Phi)$.

Proof. Let $C = v_0 e_1 v_1 \dots v_{l-1} e_l v_l$, where $v_l = v_0$, be a circle, considered as a walk (not as a subgraph or an edge set). Then $\varphi(C) = \prod_{i=1}^{l} \varphi(e_i)$ (remember that a walk has a direction; this gain is the one corresponding to the implied direction). If ζ is a switching function, then

$$\begin{aligned}
\varphi^\zeta(C) &= \varphi^\zeta(e_1)\varphi^\zeta(e_2)\cdots\varphi^\zeta(e_l) \\
&= [\zeta(v_0)^{-1}\varphi(e_1)\zeta(v_1)][\zeta(v_1)^{-1}\varphi(e_2)\zeta(v_2)]\cdots[\zeta(v_{l-1})^{-1}\varphi(e_l)\zeta(v_l)] \\
&= \zeta(v_0)^{-1}\varphi(e_1)\varphi(e_2)\cdots\varphi(e_l)\zeta(v_l) \\
&= \zeta(v_0)^{-1}\varphi(C)\zeta(v_0) \\
&= \varphi(C)^{\zeta(v_0)},
\end{aligned}$$

a conjugate of $\varphi(C)$. Thus, $\varphi^\zeta(C) = 1$ if and only if $\varphi(C) = 1$. □

What this means is that when we switch a gain graph we do not change anything at the level of balanced circles. (And you might notice that switching conjugates the gains of any closed walk W:

$$\text{[[LABEL E : 20100202walk]]}\varphi^\zeta(W) = \varphi(W)^{\zeta(v_0)}. \tag{A.2}$$

However, what the conjugating element is will depend not only on the order of edges in the walk but on the vertex v_0 where one chooses to start the walk.)

It is interesting that the gain remains unchanged under switching for any circle whose gain lies in the center, $Z(\mathfrak{G})$, and not just when it is the group identity. Furthermore, for such circles the gain, though it may depend on the direction, is independent of the choice of initial vertex. Naturally, that includes all circles if the gain group is abelian; you might guess (correctly) that sometimes abelian gains will be much more manageable than nonabelian ones.

I want to mention a special kind of switching. Any group element $\mathfrak{g}$ defines a constant switching function, whose value is $\mathfrak{g}$ on every vertex. Then $\varphi^{\mathfrak{g}}(e) = [\varphi(e)]^{\mathfrak{g}}$; that is, switching by $\mathfrak{g}$ truly is conjugation—every gain is conjugated by $\mathfrak{g}$. (This is why I say switching generalizes conjugation to gain graphs.) We can multiply a switching function by a constant group element. For $\mathfrak{g} \in \mathfrak{G}$ and a switching function ζ we have the identities

$$\varphi^{\zeta\mathfrak{g}} = (\varphi^\zeta)^{\mathfrak{g}} \quad \text{and} \quad \varphi^{\mathfrak{g}\zeta} = (\varphi^{\mathfrak{g}})^\zeta.$$

Recall that an *ordinary graph* has no half or loose edges, and that $\Gamma|S := (V, S)$ for any edge set S—this is the *restriction* of Γ to S. I will write $\Gamma|T$ for the restriction of Γ to the edge set of T when T is a forest; in fact, sometimes I do think of T as an edge set, although sometimes I have to think of it as a subgraph (and I hope the context will always make the intention clear).

Lemma A.5 (Unique Balanced Extension)**.** [[LABEL L:20100202uniquebal]] *Given an ordinary graph Γ, a maximal forest T, and a gain function $\psi : T \to \mathfrak{G}$ for $\Gamma|T$. (Here we think of T as an edge set.) Then there is a unique gain function $\hat\psi$ on Γ such that $\hat\psi|_T = \psi$ and $\hat\Psi$ is balanced.*

This means that, given any gains on a maximal forest, there is a unique balanced extension.

Proof. First we prove that a balanced extension of ψ to Γ, call it $\hat\psi$, is unique if it exists. MORE PROOF

To illustrate the proof see the following picture:

[FIGURE]

Here $C_T(e)$ is a fundamental circle we want to be balanced in our extension. In other words $\hat\Psi(C_T(e)) = 1 \Leftrightarrow \hat\Psi(eP) = 1 \Leftrightarrow \hat\Psi(e)\hat\Psi(P) = 1 \Leftrightarrow \hat\Psi(e) = \hat\Psi(P)^{-1}$. This tells us exactly what gain to assign for e to give $C_T(e)$ as balanced. By uniqueness of inverses we get the uniqueness of $\hat\psi$ for free.

Now we take a step back and define a new gain function, $\hat\psi'$, on Γ. Let $\Phi_1, \Phi_2, \ldots, \Phi_r$ be the components of Φ. Root each Φ_i at some vertex u_i. We define $\theta : V \to \mathfrak{G}$ by $\theta(v) := \psi(T_{u_iv})$ for $v \in V_i$, where T_{u_iv} is the path in T from u_i to v. Define $\hat\psi' := 1_E^\theta$. This gain function is balanced by Proposition A.4.

Next, we prove that $\hat\psi|_T$ extends ψ. Let $e{:}vw \in T$, and assume that v is closer to u_i than is w; in other words, $v \in V(T_{u_iw})$. Then $\hat\psi'(e) = \theta(v)^{-1}1\theta(w) = \psi(T_{u_iv})^{-1}\psi(T_{u_iw}) = \psi(T_{u_iv})^{-1}\psi(T_{u_iv})\psi(T_{vw}) = \psi(T_{vw}) = \psi(e) = \hat\psi(e)$. So, $\hat\psi'$ extends ψ.

Finally, we prove that $\hat\psi'$ is balanced. But this is immediate from Proposition A.4, because $\hat\psi \sim 1_E$, which is the most elementary balanced gain function.

Since $\hat\psi'$ is balanced and extends ψ, it must equal $\hat\psi$ by uniqueness. Therefore, $\hat\psi$ is a balanced extension of ψ. □

A.5. **Balance again.** [[LABEL 4.ggbalagain]]

Now we can state and prove a list of properties equivalent to balance, similar to some of those in Section II.A. A *potential function* for φ is a function $\theta : V \to \mathfrak{G}$ such that for every $e \in E$ we have $\varphi(e{:}vw) = \theta(v)^{-1}\theta(w)$. Equivalently, $\varphi = 1_E^\theta$, or $\varphi^{\bar\theta} = 1_E$.

Theorem A.6 (Equivalents of Balance)**.** [[LABEL T:20100202equivbal]] *Let Φ be an ordinary gain graph. The following properties are equivalent:*

(i) Φ *is balanced.* [[LABEL T:20100202equivbal Bal]]
(ii) $\Phi \sim (\|\Phi\|, 1)$*. In other words,* $\varphi \sim 1_E$*.* [[LABEL T:20100202equivbal Sw]]
(iii) φ *has a potential function.* [[LABEL T:20100202equivbal Pot]]
(iv) *For each v and w in V, every vw-path has the same gain.* [[LABEL T:20100202equivbal Path]]
(v) *For each v and w in V, every vw-walk has the same gain.* [[LABEL T:20100202equivbal Walk]]

Proof. The method that seems to work best is not the elegant circular implication but a chain of equivalences. I will prove that (i) $\iff$ (ii) $\iff$ (iii) $\iff$ (v) $\iff$ (iv).

Four implications are trivial or obvious. Trivial: (iv) $\implies$ (v), since every path is a walk. Obvious: (ii) $\implies$ (i), by Proposition A.4 and the fact (obvious) that 1_E is balanced. Also obvious: (ii) $\iff$ (iii), because $\varphi \sim 1_E \iff \varphi^\zeta = 1_E$ for some $\zeta \iff \varphi = 1_E^{\bar\zeta}$ for some $\bar\zeta$.

MORE TO COME

□

2010 Feb 4: Simon Joyce

Before the proof of Theorem A.6 we prove some necessary lemmas.

Lemma A.7. [[LABEL L:20100204 edge gain is tree path gain]] *Given a balanced gain graph Φ and a maximal forest T of Φ, then $\varphi(e{:}vw) = \varphi(T_{vw})$ for any edge.*

Proof. If $e \in T$, this is trivial, because e is the tree path T_{vw}. If $e \notin T$, then the edges of T_{vw} and e will form a circle in Φ. Let C be the closed path $T_{vw}e^{-1}$. Since Φ is balanced, $\varphi(C) = 1$. Therefore $\varphi(T_{vw})\varphi(e)^{-1} = 1$, or $\varphi(e) = \varphi(T_{vw})$. □

Lemma A.8. [[LABEL L:20100204 any desired tree gain]] *Given a gain graph Φ with gain group $\mathfrak{G}$, and a maximal forest T of Φ, then for any gain function $\psi : T \to \mathfrak{G}$ there is switching function ζ of Φ such that $\varphi^\zeta|_T = \psi$.*

Proof. Let Φ_i be the components of Φ and root each Φ_i at a vertex u_i. Define a switching function ζ of Φ by $\zeta(v) := \varphi(T_{u_iv})^{-1}\psi(T_{u_iv})$ for $v \in V_i$. Let $e \in E(T)$ with endpoints v and w. Then,

$$\begin{aligned}
\varphi^\zeta(e{:}vw) = \varphi^\zeta(T_{vw}) &= [\psi(T_{u_iv})^{-1}\varphi(T_{u_iv})]\varphi(T_{vw})[\varphi(T_{u_iw})^{-1}\psi(T_{u_iw})] \\
&= \psi(T_{u_iv})^{-1}[\varphi(T_{u_iv})\varphi(T_{vw})\varphi(T_{u_iw})^{-1}]\psi(T_{u_iw}) \\
&= \psi(T_{vu_i})[\varphi(T_{u_iv}T_{vw}T_{wu_i})]\psi(T_{u_iw}) \\
&= \psi(T_{vu_i})\psi(T_{u_iw}) \\
&= \psi(T_{vw}) = \psi(e).
\end{aligned}$$

□

Proof of Theorem A.6. It is obvious by Proposition A.4 that (ii) implies (i) since if $\varphi \sim 1_E$ then Φ is balanced. To see that (i) implies (ii), first switch Φ by a switching function ζ so that $\varphi^\zeta|_T = 1_T$ where T is a spanning forest of Φ. Such a ζ exists by Lemma A.8. Then $\varphi^\zeta = 1_E$ since $\varphi^\zeta(e) = 1$ for all $e \notin T$ by Lemma A.7.

That (ii) and (iii) are equivalent is obvious since $\varphi \sim 1_E$ if and only if $\varphi^\zeta = 1_E$ for some ζ, or $\varphi = 1_E^{\bar\zeta}$.

Trivially (v) implies (iv), since every path is a walk.

To see that (iv) implies (v), let P be a vw-path in Φ. We want to show that for any vw-walk W in Φ, $\varphi(W) = \varphi(P)$. The proof is by induction on the length of W.

First assume $v = w$. So W is a closed vv-walk. Then P has length 0 so $\varphi(P) = 1$. If W has length 0, then $\varphi(W) = 1 = \varphi(P)$. If W has length 1, it's a loop , so by assumption $\varphi(W) = 1$. If W has length 2 or more, then it has an internal vertex z (which could be repeated). Split W at one appearance of z into two subwalks, W_1 and W_2, so $W = W_1W_2$. Then $\varphi(W) = \varphi(W_1)\varphi(W_2) = \varphi(W_1)\varphi(W_2^{-1})^{-1}$. By induction $\varphi(W_1) = \varphi(W_2^{-1})$ since these are vz-walks and are shorter than W. Therefore $\varphi(W) = 1$.

Now assume $v \neq w$. If W has no repeated vertices it is a path, so $\varphi(W) = \varphi(P)$ by (iv). Now assume W does have some repeated vertex z. Split W at two appearances of z into three subwalks, W_{vz}, W_{zz}, and W_{zw}. Then $W = W_{vz}W_{zz}W_{zw}$ so $\varphi(W) = \varphi(W_{vz})\varphi(W_{zz})\varphi(W_{zw})$. But $\varphi(W_{zz}) = 1$ from the previous part, since W_{zz} is a closed zz-walk. Therefore $\varphi(W) = \varphi(W_{vz})\varphi(W_{zw}) = \varphi(W_{vz}W_{zw})$. Since $W_{vz}W_{zw}$ is a shorter vw-walk, we conclude by induction that $\varphi(W_{vz}W_{zw}) = \varphi(P)$ and we're done.

For (v) $\implies$ (iii), let Φ_i be the components of Φ and root each Φ_i at a vertex u_i. Define the switching function θ by the formula $\theta(v) = \varphi(W_{u_iv})$, where $v \in V_i$ and W_{u_iv} is any walk from the root u_i to v. This is well defined since all u_iv-walks have the same gain by (v). Now e_{vw} is a vw-walk, and consequently $\varphi(e) = \varphi(W_{vu_i}W_{u_iw})$ for any walks W_{vu_i} and W_{u_iw} by (v). But,

$$\begin{aligned}
\varphi(W_{vu_i}W_{u_iw}) = \varphi(W_{vu_i})\varphi(W_{u_iw}) &= \varphi(W_{u_iv})^{-1}\varphi(W_{u_iw}) \\
&= \theta(v)^{-1}\theta(w) = 1^\theta.
\end{aligned}$$

Therefore $\varphi = 1_E^\theta$, i.e., φ has the potential function θ.

The last thing to do is to show that (iii) implies (v). This is left to you as an (easy) homework exercise. □

[**An example of a gain graph for list of examples.**] A multiply signed graph has gain group $\mathfrak{G} = \{+,-\}^k$ (or $\mathbb{Z}_2^k$) for some integer $k \geq 1$. When $k = 1$ this is a signed graph, as we're used to. The nice property of this group for us is that it has exponent 2, that is, every element is self-inverse. Therefore, every edge e satisfies $\varphi(e^{-1}) = \varphi(e)$, so orientation of edges is unnecessary, as for signed graphs but not for other gain groups. Moreover, since $\mathfrak{G}$ is abelian, conjugation has no effect, so switching does not change the gains of closed walks.

A result we had for signed graphs was that for two signed graphs Σ_1 and Σ_2 with the same underlying graph Γ, then $\mathcal{B}(\Sigma_1) = \mathcal{B}(\Sigma_2)$ if and only if $\Sigma_1 \sim \Sigma_2$. However for any two gain graphs Φ_1 and Φ_2 on Γ, having $\mathcal{B}(\Phi_1) = \mathcal{B}(\Phi_2)$ does not imply that $\Phi_1 \sim \Phi_2$. A counterexample is given in Figure A.4. Here the gain group $\mathfrak{G} = \{+,-\}^2$ for both Φ_1 and Φ_2. Also $\mathcal{B}(\Phi_1) = \mathcal{B}(\Phi_2)$. Every element in $\mathfrak{G}$ is self inverse so the orientation does not matter and since the gain group is abelian in this case, switching does not change the gains of the circles. But the gains of the two unbalanced circles in both graphs are not the same so Φ_1 and Φ_2 cannot be switching equivalant.

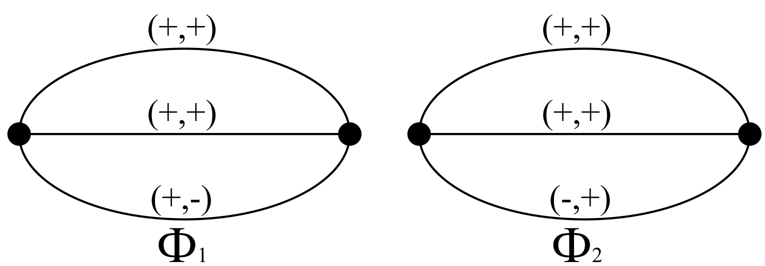

FIGURE A.4. $\mathcal{B}(\Phi_1) = \mathcal{B}(\Phi_2)$ but $\Phi_1 \not\sim \Phi_2$.
[[LABEL F:20100204swineq]]

A.6. **Two essential parameters.** [[LABEL 4.someparams]]

I want to mention here a crucial number and partition associated with an edge set in a gain graph, that generalize those of signed graphs (Section II.C.1). The right time to treat them, though, is not now, but in the next section where we perform the ultimate generalization; so this is a quick sketch.

Let $S \subseteq E$; then

$$b(S) := \text{ the number of components of } S \text{ that are balanced,}$$
$$\boldsymbol{\pi}_{\mathrm{b}}(S) := \text{ the family of vertex sets of components of } S \text{ that are balanced,}$$

where in both cases I mean to exclude loose edges; that is, I consider only components that have vertices. When I say S here, I mean the spanning subgraph (V, S). (I'll discuss this more carefully in the next section.) Obviously, $\boldsymbol{\pi}_{\mathrm{b}}(S)$ is a *partial partition* of V, i.e., a partition of a subset of V, and $b(S) = |\boldsymbol{\pi}_{\mathrm{b}}(S)|$. If S has no balanced components, $b(S) = 0$, then $\boldsymbol{\pi}_{\mathrm{b}}(S) = \varnothing$, which is (!) the unique partition of the empty set.

2010 Feb 9: Jackie Kaminski and Amanda Ruiz

B. Biased Graphs

[[LABEL 4.bg]]

Recall that for a gain graph Φ, the set $\mathcal{B}(\Phi)$ denotes the class of balanced circles. Important aspects of Φ, especially the number of balanced vertex components $b(\Phi)$, depend only on the class $\mathcal{B}$, not on the actual gains. This motivates us to study combinatorial properties of $\mathcal{B}$ that are independent of gains.

B.1. **Basic definitions.** [[LABEL 4.bgdef]]

Definition B.1. [[LABEL D:20100209 Biased Graph]] A *biased graph* is a pair $\Omega = (\Gamma, \mathcal{B})$, where Γ is a graph and $\mathcal{B} \subseteq \mathcal{C}(\Gamma)$ is such that no theta subgraph contains exactly two circles of $\mathcal{B}$. An edge set or subgraph of Γ is called *balanced* if every circle in it belongs to $\mathcal{B}$ and it has no half edges. In particular, a circle is balanced if and only if it belongs to $\mathcal{B}$; thus, we call $\mathcal{B}$ the *class of balanced circles* of Ω.

We call Γ the *underlying graph* of Ω. We call any class $\mathcal{B} \subseteq \mathcal{C}(\Gamma)$ that satisfies the theta-graph requirement a *linear class* of circles. If Ω is a biased graph, we denote its underlying graph by $\|\Omega\|$ and its balanced circle class by $\mathcal{B}(\Omega)$.

I like to think of the 'bias' as the departure from balance, so that the more balanced circles there are, the less 'biased' the graph is. Thus, a biased graph is completely unbiased when it is balanced, and completely biased when it is contrabalanced (Definition D.1).

Since we know from Proposition A.2 that the class $\mathcal{B}(\Phi)$ of balanced circles of any gain graph Φ is a linear class (it satisfies the requirements for a biased graph), from any gain graph Φ we obtain a biased graph $\langle\Phi\rangle := (\|\Phi\|, \mathcal{B}(\Phi))$, whose underlying graph and set of balanced circles are those of Φ. Biased graphs that arise from gain graphs are called *gainable*. A gainable biased graph may have gains in many different ways; we'll discuss this later on **[MAYBE]**, but you should notice that switching is a trivial way to change the gains without changing the bias, due to Proposition A.4:

Proposition B.1. [[LABEL P:20100209 switching]] *For any switching function ζ, $\langle\Phi^\zeta\rangle = \langle\Phi\rangle$.* □

All the concepts about connected components, balanced components, and their associated vertex partitions and partial partitions that we had for signed graphs extend to biased graphs, and since all gain graphs "are" biased graphs, they are included in these definitions. But first we need to clarify one definition. Recall that, for $S \subseteq E(\Gamma)$, $\Gamma|S$ is defined as $(V(\Gamma, S))$.

Definition B.2. [[LABEL D:20100209 restricted to S]]
For $S \subseteq E(\Omega)$, the *restriction* of Ω to S (or Ω restricted to S), denoted by $\Omega|S$, is the biased graph on $\Gamma|S$ with $\mathcal{B}(\Omega|S) := \{B \in \mathcal{B}(\Omega) : B \subseteq S\}$.

It's obvious that $\langle\Phi\rangle|S = \langle\Phi|S\rangle$.

B.2. **Balance. Balanced partial partition.** [[LABEL 4.balance]]

Recall that a component of a graph, more precisely a vertex component, means a connected component containing at least one vertex. So, loose edges do not count as components, but isolated vertices do.

Definition B.3. [[LABEL D:20100209 components]] For a biased graph Ω), $c(\Omega)$ denotes the number of (vertex) components of $\|\Omega\|$, and $b(\Omega)$ denotes the number of (vertex) components that are balanced in Ω.

Often, we want these numbers for $\Omega|S$ where S is any edge set in Ω. When Ω is clear from context, we write $c(S)$ as shorthand for $c(\Omega|S)$ and $b(S)$ as shorthand for $b(\Omega|S)$.

The numbers $c(S)$ and $b(S)$ are associated with partial vertex partitions.

Definition B.4. [[LABEL D:20100209 partition]] Let $S \subseteq E(\Omega)$. The partition of $V(\Omega)$ into vertex sets of components of $\Gamma|S$ is denoted by $\pi(\Omega|S) = \pi(S)$. The partial partition of $V(\Gamma)$ into the vertex sets of balanced components of $\Gamma|S$ is denoted by $\boldsymbol{\pi}_{\mathrm{b}}(\Omega|S) = \boldsymbol{\pi}_{\mathrm{b}}(S)$.

Obviously, $c(S) = |\pi(S)|$ and $b(S) = |\boldsymbol{\pi}_{\mathrm{b}}(S)|$. The partition $\pi(S)$ is always a partition of $V(\Omega)$, while $\boldsymbol{\pi}_{\mathrm{b}}(S)$ is a partition of V only when S is balanced; otherwise it partitions a proper subset of V. Do not forget the special case, which will often arise, where there are no balanced components; then $\boldsymbol{\pi}_{\mathrm{b}}(S) = \varnothing$, the empty partition of $\varnothing \subseteq V$.

The notations $c(\Omega), b(\Omega), \pi(\Omega)$, and $\boldsymbol{\pi}_{\mathrm{b}}(\Omega)$ mean the same as $c(E), b(E), \pi(E)$, and $\boldsymbol{\pi}_{\mathrm{b}}(E)$, respectively, and we'll use whichever is most convenient.

Definition B.5. [[LABEL D:20100209 edge induced]] For $S \subseteq E(\Omega)$, the *edge-induced subgraph*, $\Omega{:}S$, is $\Omega|S$ with isolated vertices deleted. We may write $V(S)$ for the vertex set of $\Omega{:}S$

We end this section with a small remark that is essentially about matroid theory. Suppose V is finite. Then $|V| - b(\Omega|S) = |V(S)| - b(\Omega{:}S)$. If V is infinite we definitely can't calculate the left-hand side, but we can calculate the right-hand side whenever $V(S)$ is finite. Readers familiar with matroids will like to know that this is the rank function of a matroid (the 'frame matroid').

2010 Feb 9: Thomas Zaslavsky

C. COMBINATORIALIZATION

[[LABEL 4.combinator]]

Let's talk philosophy for a moment.

Biased graphs are a 'combinatorialization' of gain graphs. We extracted from the properties of gain graphs a certain key feature that is independent of algebra (in this case, of groups) and made it into the foundation of a new structural definition. While every gain graph gives rise to a biased graph, we actually get more from biased graphs since there are biased graphs that are not gain graphs.

This procedure of combinatorial abstraction has given rise to an enormous range of combinatorial structures. One that's close to home (given the fact that there are matroids hidden in the background of graphs and biased graphs) is matroids. Matroids are a combinatorial abstraction of linear independence; but they are a much broader class of objects because there are matroids that are not the dependence matroid of any set of vectors.

Another example, a relatively simple one, is Latin squares. From a modern point of view (which is certainly not the one that led Euler to invent them), Latin squares are a combinatorial abstraction of the multiplication table of a group. Latin squares in turn suggest generalizing the algebraic definition of a group to loops and quasigroups. (That is roughly the historical course of development.)

I can't fail to mention a third example that has also generated a great deal of combinatorial mathematics (under the name "incidence geometry"), namely, projective planes. A projective plane (see Section III.A.1) exists with coordinates in any field and even any division ring—this is the algebraic or "analytic" definition, in the sense of analytic geometry.

By extracting the incidence properties of a coordinatized plane one can give a synthetic definition (as in synthetic geometry) which is essentially combinatorial and leads to additional structures, the "non-Desarguesian planes", that cannot be obtained algebraically. (That's my opinion. The algebra of division rings has been extended to "ternary rings" to cover these planes, but ternary rings are so ill-behaved and intractable that they barely deserve to be called algebraic. There is a tight connection amongst projective planes, Latin squares, and ternary rings—but I digress.)

Feb 9: Jackie Kaminski and Amanda Ruiz

D. Examples of Many Kinds

[[LABEL 4.x]]

There are all kinds of interesting biased graphs (if I do say so myself). Almost all of them are gainable—but some are most conspicuously not. Some of them naturally appear as gain graphs, but others are most natural as biased graphs. Let's begin with a few examples of the latter sort.

D.1. **Abstract graphical examples.** [[LABEL 4.xg]]

D.1.1. *Simple bias.*

A *simply biased graph* is a biased graph with no loose edge, no balanced loops or balanced digons, and no vertex supporting more than one edge (which must be a half edge or an unbalanced loop).

D.1.2. *A plain old graph.*

We can think of any graph Γ as a biased graph by defining $\mathcal{B}(\Gamma) := \mathcal{C}(\Gamma)$, the set of all circles. We write this bias of Γ as $\langle\Gamma\rangle := (\Gamma, \mathcal{C}(\Gamma))$, the biased graph where all circles are balanced. This is one biased graph that is gainable in a huge number of ways, in particular as $\langle +\Gamma\rangle$ or $\langle 1_{\mathfrak{G}}\Gamma\rangle$ for any group $\mathfrak{G}$, or indeed (Γ, φ) for any balanced gains in any group, since all those are ways to make all circles are balanced.

(It is gratifying, but not a coincidence, that in the terminology of 'bias', $\langle\Gamma\rangle$ is 'unbiased'. I think of bias as a departure from balance.)

Still, $\langle\Gamma\rangle$ is an unsatisfying example of multiple ways to "gain" a biased graph. We know from Theorem A.6(ii) that any gains that make (Γ, φ) balanced switch to identity gains, and since once the only gain value is the identity the rest of the group is superfluous, we may as well cut $\mathfrak{G}$ down to a single element. In other words, the differences in the ways to put gains on $\langle\Gamma\rangle$ are trivial. We'll see a more substantial example of multiple gain functions later, but here is one that is, at any rate, more substantial than a balanced biased graph.

D.1.3. *Contrabalance.*

Definition D.1. [[LABEL D:20100209 contrabalanced]] A biased graph is *contrabalanced* if it has no balanced circles.

The contrabalanced biased graph with underlying graph Γ is $(\Gamma, \varnothing)$. This also has a large number of gain functions, most of which are inequivalent under switching or any other simple group operation—I leave the proof of this to the reader. [**BUT CITE EXAMPLE OF GAIN GRAPHS WITH SAME BALANCE FROM SIMON.**]

Proposition D.1. [[LABEL P:20100209 contrabalanced]] *If Ω is contrabalanced, it has gains in some group. If E is finite, Ω has gains in a finite group, in fact, in any sufficiently large finite group.*

Proof. I leave the proof to the reader—it should be fun. A couple of hints: For the first two statements, think of the edges as generators of a group (that's the free-group homomorphism viewpoint on gains). The third statement is more challenging; if you get stuck, consult Gagola (1999a). □

D.1.4. *Fullness.*

Many of the constructions from earlier chapters still apply. For instance, by $\Phi^\bullet$ or $\Omega^\bullet$ we mean Φ or Ω with a half edge added to each vertex. Since half edges don't receive gains, there is no need to consider what gain to assign to the new half edges.

Definition D.2. [[LABEL D:20100209 full]] We say Ω is *full* if every vertex supports an unbalanced edge, that is, an unbalanced loop or a half edge.

By Ω° we mean Ω with an an unbalanced loop adjoined to every vertex that doesn't already have an unbalanced edge. This is different than $\Omega^\bullet$, where you add a half edge or unbalanced loop to every vertex.

D.2. **Group and group-subset expansions.** [[LABEL 4.xgg]]

Now we turn to some examples that are defined in terms of gains.

[THIS SECTION IS DUPLICATED LATER.]

Example D.1. [[LABEL X:20100209 pm gamma]] A nice example is $\langle \pm\Gamma \rangle$, the biased graph of $\pm\Gamma$. Recall that $\pm\Gamma := +\Gamma \cup -\Gamma$, i.e., Γ with each edge replaced by a pair of parallel edges, one with a $+$ sign and one with a $-$ sign (Section A.2).

A generalization of $\pm\Gamma$ to groups with more than two elements is the 'group expansion'.

Definition D.3. [[LABEL D:20100209 group expansion]] For a group $\mathfrak{G}$ and a graph Γ, the the *group expansion* of Γ, more precisely its $\mathfrak{G}$-*expansion*, is the gain graph $\mathfrak{G}\Gamma$ whose vertex set is $V(\Gamma)$ and in which each ordinary edge (loop or link) of Γ is replaced by $|\mathfrak{G}|$ edges having all possible gains in $\mathfrak{G}$. That is, for each ordinary edge e, each $\mathfrak{g} \in \mathfrak{G}$ is the gain of one of the edges replacing e (in a fixed direction; which direction doesn't matter). Formally, $E(\mathfrak{G}\Gamma) := \mathfrak{G} \times E(\Gamma)$; an edge of the expansion is written $(\mathfrak{g}, e)$ or more concisely $\mathfrak{g}e$, and the gain of this edge is $\mathfrak{g}$.

Since $\mathfrak{G}$ is closed under inversion the choice of direction of e is only due to the necessities of notation. By our convention about gains, $(\mathfrak{g}, e)$ is the same edge as $(\mathfrak{g}^{-1}, e^{-1})$. The fact that $\mathfrak{G}^{-1} = \mathfrak{G}$ means that we get the same edges no matter which orientation we give to e. We can perform the same construction with any self-inverse subset of the group.

Definition D.4. [[LABEL D:20100209 subset expansion]] Let $A \subseteq \mathfrak{G}$ such that $A^{-1} = A$. The A *expansion* (or *subset expansion*) of Γ is the gain graph $A\Gamma$ on vertex set $V(\Gamma)$ in which each ordinary edge of Γ is replaced by $|A|$ edges so that each $\mathfrak{g} \in A$ is the gain of one of the edges replacing e. Formally, $E(A\Gamma) := A \times E(\Gamma)$.

The most trivial example is $1_{\mathfrak{G}}\Gamma$, in which case $\langle 1_{\mathfrak{G}}\Gamma \rangle = \langle \Gamma \rangle = (\Gamma, \mathcal{C}(\Gamma))$.

We require $A = A^{-1}$ so we need not specify an orientation of Γ in order to know which edges appear with which gains in the expansion. We can generalize even further to an arbitrary subset of $\mathfrak{G}$ if we orient Γ first.

Definition D.5. [[LABEL D:20100209 arbitrary subset expansion]] For any $A \subseteq \mathfrak{G}$ and an arbitrary orientation $\vec{\Gamma}$ of Γ, the A *expansion* (or *subset expansion*) of $\vec{\Gamma}$ is the gain graph $A\vec{\Gamma}$ on vertex set $V(\Gamma)$ in which each edge $\vec{e} = e{:}\vec{vw}$ in $\vec{\Gamma}$ becomes the family $\{\mathfrak{g}\vec{e} : \mathfrak{g} \in A\}$ of edges in $A\vec{\Gamma}$ whose gains are $\varphi_{vw}(\mathfrak{g}\vec{e}) = \mathfrak{g}$ and $\varphi_{wv}(\mathfrak{g}\vec{e}) = \mathfrak{g}^{-1}$. Once again, formally $E(A\vec{\Gamma}) := A \times E(\vec{\Gamma})$.

If $V = [n]$ and we don't explicitly specify an orientation, we will assume $\vec{\Gamma}$ has edges oriented from lower to higher vertices, that is, $e{:}\vec{ij}$ where $i \leq j$.

Example D.2. [[LABEL X:20100209 Group]] Figure D.1 shows a group expansion of K_3. If the circle with the labeled gains is balanced then $\mathfrak{g}\mathfrak{h}\mathfrak{i}^{-1} = 1_{\mathfrak{G}}$, in other words $\mathfrak{g}\mathfrak{h} = \mathfrak{i}$. Notice that the directions on the edges make a difference here.

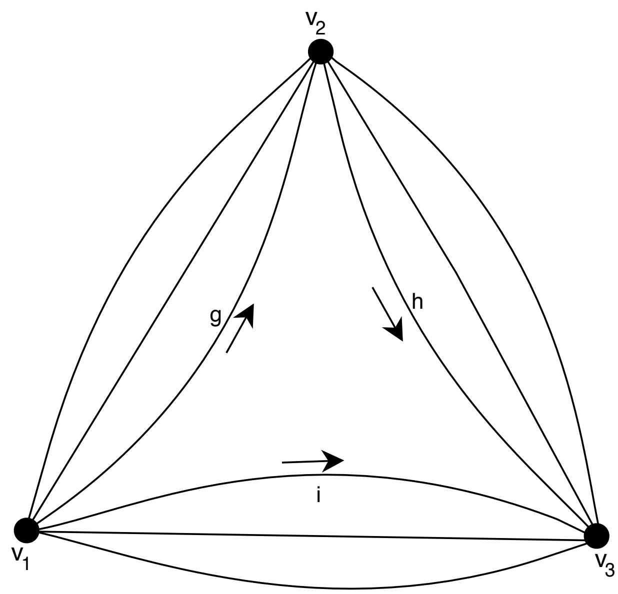

FIGURE D.1. A group expansion of K_3.
[[LABEL F:20100209 Group]]

Proposition D.2. [[LABEL P:20100209 determined group]] *Let $\mathfrak{G}$ be a group and $n \geq 3$. Then the biased graph $\langle \mathfrak{G}K_n \rangle$ determines $\mathfrak{G}$.*

Proof. I leave this as an exercise to the reader. It is not trivial to give a complete proof, but not hard if you persist. (Example D.2 contains a hint.) The key here is that for any path in $\mathfrak{G}K_n$ (of length ≥ 2) there exists a unique edge that closes the path into a balanced circle. □

The proposition certainly isn't true for any graph, but only for K_n, or any graph that contains K_3 as a subgraph. Notice that it doesn't say you can recover the gain graph $\mathfrak{G}K_n$ exactly, including the gains of all edges (which is obvious, since switching changes gains but not bias), just that you can recover the gain group $\mathfrak{G}$.

D.3. **Gain graphs from geometry.** [[LABEL 4.xgeom]]

A kind of gain graph that is popular in combinatorial geometry (although hardly anyone working with the geometry knows of gain graphs!) is an integral gain graph—whose gain group is $\mathfrak{G} = \mathbb{Z}^+$, the set of integers under addition—of a certain symmetrical form. Let $A = \{0, \pm 1\}$, or more generally, $A = [-l, l]_{\mathbb{Z}} := \{-l, -(l-1), \dots, 0, \dots, l-1, l\}$ where $l \geq 0$. This can also be generalized with $A = \{\pm 1\}$ or with $A = \pm[l] = \{-l, \dots, -1, 1, \dots, l\}$. The gain graph we want is AK_n, which is well defined since A is symmetric (i.e., closed under group inversion) so there is no need to specify directions on the edges K_n.

The *Catalan gain graph* $\{0, \pm 1\}K_n$ is associated with the *Catalan arrangement* of affine hyperplanes in $\mathbb{A}^n(\mathbb{R})$, defined by

$$\mathcal{C}_n := \{x_j = x_i,\ x_j = x_i \pm 1 : i \neq j\}.$$

(The edge $c\vec{ij}$ corresponds to the hyperplane $x_j = x_i + c$.) A related arrangement is the *hollow Catalan arrangement* defined by

$$\mathcal{C}_n^\circ := \{x_j = x_i \pm 1 : i \neq j\},$$

corresponding to the gain graph $\{\pm 1\}K_n$. As a reminder of former notation, observe that $\mathcal{C}_n = \mathcal{C}_n^\circ \cup \mathcal{H}[K_n]$. (The lattice of flats, $\mathcal{L}(\mathcal{C}_n)$, can be thought of as the semilattice of composed partitions of $[n]$. For readers unfamiliar with composed partitions—that is, all readers—it should be enough to think of them as partitions with some additional structure that we'll ignore here.)

Another important geometrical example is the *Shi gain graph* $\{0, 1\}\vec{K}_n$, which is associated with the *Shi arrangement* of hyperplanes, the set

$$\mathcal{S}_n := \{x_j = x_i, x_j = x_i \pm 1 : i < j\} = \mathcal{H}[K_n] \cup \mathcal{S}_n^\circ$$

of hyperplanes in $\mathbb{A}^n(\mathbb{R})$. A related arrangement is $\mathcal{S}_n^\circ$, the *hollow Shi arrangement*, given by

$$\mathcal{S}_n^\circ := \{x_j = x_i \pm 1 : i < j\},$$

whose gain graph is $\{1\}\vec{K}_n$.

Let's illustrate these arrangements and gain graphs for K_3.

Example D.3. [[LABEL X:20100209 K 3]] Figure D.2 shows $\{0, \pm 1\}K_3$, the graph for the Catalan arrangement $\mathcal{C}_3$. The indicated edge is $1e_{12}$, which we could also call $-1e_{21}$. (The difference is only in the reference orientation of the edge.) Figure D.3 shows $\{0, 1\}\vec{K}_3$, the gain graph of the Shi arrangement $\mathcal{S}_3$. In both cases the group operation is addition in $\mathbb{Z}$.

2010 Feb 11: Nathan Reff

D.4. **An example with no gains.** [[LABEL 4.nogains]]

Example D.4. [A biased graph that has no gains][[LABEL X:20100211 biasg with no gains]] In the graph of Figure D.4, let

$$\mathcal{B} := \big\{abcd \text{ and any quadrilateral that has exactly two edges from } \{a, b, c, d\}\big\}.$$

We'll see that this biased graph cannot be obtained from any gain function.

The first step is a simple proposition about any biased graph Ω.

Proposition D.3. [[LABEL P:20100211baldigons]] *If ef is a balanced digon in Ω, and if C is a circle such that $e \in C$, $f \notin C$, then C is balanced $\iff$ $(C \setminus e) \cup f$ is balanced.*

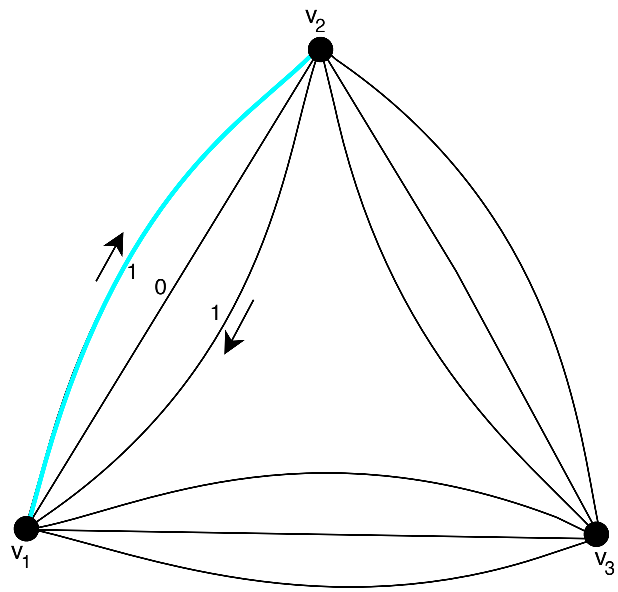

FIGURE D.2. The gain graph for the Catalan arrangement $\mathcal{C}_3$.
[[LABEL F:20100209 Catalan]]

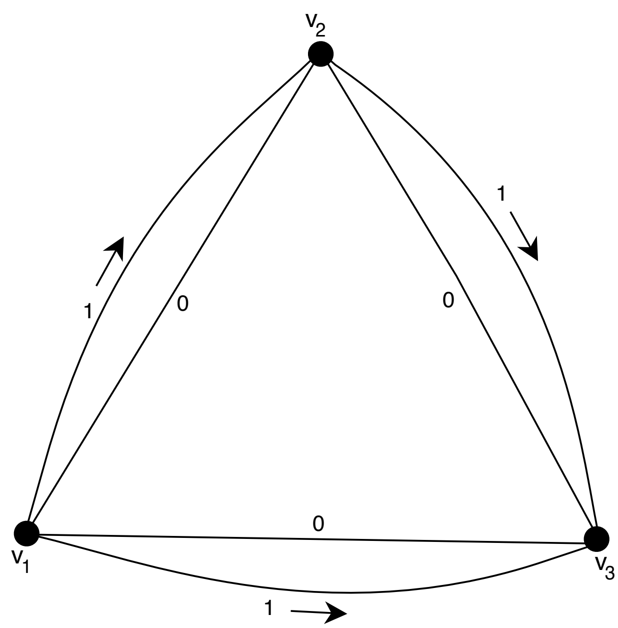

FIGURE D.3. The gain graph for the Shi arrangement $\mathcal{S}_3$.
[[LABEL F:20100209 Shi]]

Proof. This is another homework exercise. Hint: Think about the theta-graph property of a biased graph. $\square$

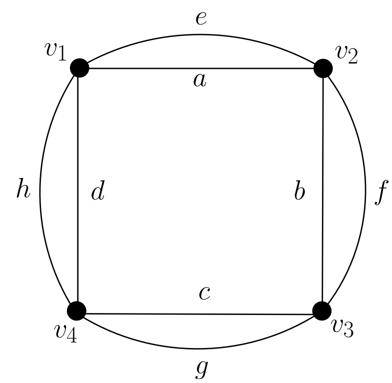

FIGURE D.4
[[LABEL F:20100211: biasedgnogains]]

This is why balanced digons are trivial and we don't need or want them in our example. In the example we can choose $\mathcal{B}$ to be any of the quadrilaterals such that no two quadrilaterals in $\mathcal{B}$ have three common edges. Then $(\Gamma, \mathcal{B})$ will be a biased graph.

Proof. Two quadrilaterals that have differences in at least two edges do not form a theta graph. This is why our example works. $\square$

If $\mathcal{B}$ contains only quadrilaterals which have an even intersection with $abcd$ then $\mathcal{B}$ is a linear class (it satisfies the theta condition). (Proof: Homework! This is not quite trivial.)

Suppose $(\Gamma, \mathcal{B})$ has gains. Switch the gains so a, b, c, d have gain 1. (For simplicity I'll use the same letter for the edge and its gain, the former in italic and the latter in Fraktur.) Since $abcd$ is balanced and $abgh$ is balanced, $\mathfrak{gh} = 1$. Because $efcd$ is also balanced, $\mathfrak{ef} = 1$. Therefore $\mathfrak{efgh} = 1$, so $efgh$ is balanced. However, we assume $efgh \notin \mathcal{B}$. The conclusion is that $(\Gamma, \mathcal{B})$ cannot have gains.

Thus, there are biased graphs not derived from gain graphs; biased graphs are a strict generalization of gain graphs.

This example works if $abcd$, $abgh$, and $efcd$ are all in $\mathcal{B}$ but $efgh \notin \mathcal{B}$, regardless of what other quadrilaterals may be balanced (subject to the theta-graph condition, of course). So it is really several examples of biased graphs without gains.

D.5. **Circuits.** [[LABEL 4.circuits]]

A kind of biased graph that is important not so much in itself but as a subgraph within a biased graph (including gain graphs) is a circuit. Actually, there are two distinct families of circuits: the *frame circuits* and the *lift circuits*. The two families overlap, but the difference is very important.

Example D.5. [[LABEL X:20100211 biased subgraphs]] Figure D.5 shows some important biased graphs. Note that half edges count as unbalanced loops.

The two kinds of circuit are very much like the frame and lift circuits of signed graphs, but there is the extra type, the contrabalanced theta, which cannot exist in a signed graph because the gain group is just too small. ("Too small" means order 2. Already order 3 is not too small!)

---

16 Feb
2010:
Simon Joyce

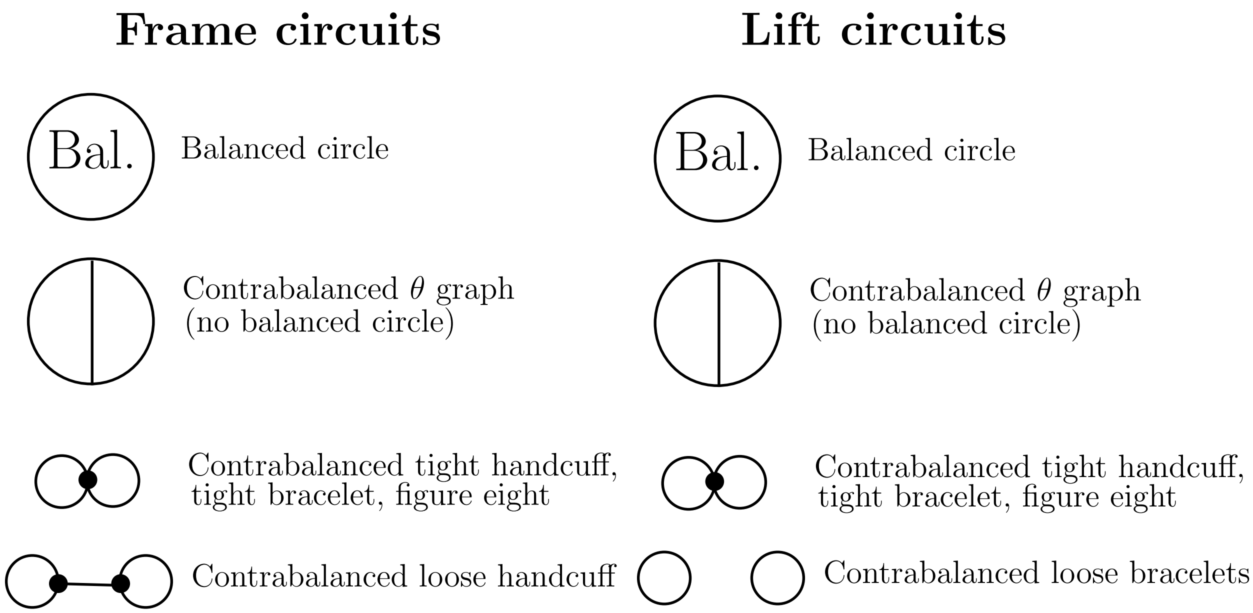

Figure D.5
[[LABEL F:20100211: frameandliftcircuits]]

D.6. **Balloons.** [[LABEL 4.balloons]]

An unbalanced frame circuit that is not a theta graph is made by pasting together two unbalanced circles (or half edges) with tails (of length not less than 0). These half handcuffs are significant structures in their own right, which appear in many constructions and proofs.

Definition D.6. [[LABEL D:20100216Balloon]] A balloon is a subgraph of a biased graph composed of an unbalanced circle or half edge and a path (called the string) such that the first vertex of the path is incident with a vertex in the circle of half edge (called the tie) and none of the edges in the path are contained in the unbalanced circle (or half edge, which is obvious). The last vertex in the path is called the *tip* of the balloon.

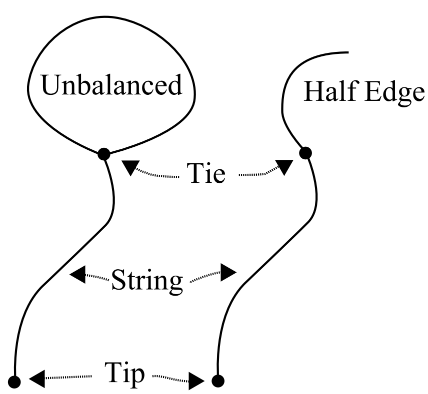

Figure D.6
[[LABEL F:20100216balloons1]]

Two examples of balloons are given in Figure D.6. Suppose you have a balloon in a biased graph with minimum degree 2. If you extend the string until you either return to a vertex in the balloon or the extended string or end the extended string in a half edge, then there is a frame circuit contained in this subgraph, since the extended balloon will contain a balanced circle, a contrabalanced theta graph, tight handcuffs, or loose handcuffs. If the extension returns to the original circle, other than at the tie, then you've formed a theta graph, which is contrabalanced or contains a balanced circle (see the red extension in Figure D.7). If the extension returns to the original string, possibly at the tie, you've formed a handcuff, which is contrabalanced or contains a balanced circle (see the blue extension in Figure D.7). If the extension returns to the extended string or ends with a half edge, you've also formed a handcuff (see the green or dashed extension in Figure D.7).

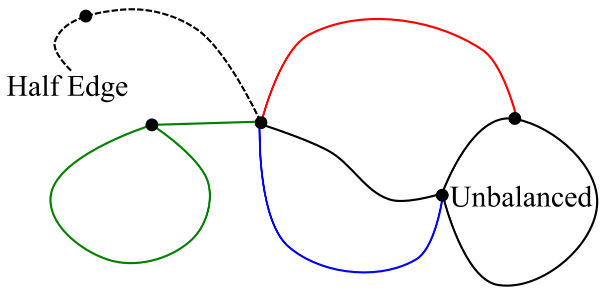

FIGURE D.7
[[LABEL F:20100216balloons2]]

---

11 Feb 2010: Nathan Reff

D.7. **Group expansions and biased expansions.** [[LABEL 4.expansions]]

Group and biased expansions are a complicated kind of example that have a rich theory of their own. Group expansions generalize the signed expansion of a graph, $\pm\Gamma$, to arbitrary groups; they were inspired by Dowling lattices (Dowling 1973b), which correspond to complete graphs. Biased expansions are the combinatorial abstraction of group expansions, just as biased graphs are the combinatorialization of gain graphs. Biased expansions are related to group expansions in the same way Latin squares are related to groups; indeed, biased expansions correspond to the higher-dimensional Latin squares called "Latin hypercubes" (see Denes and Keedwell, *Latin Squares and Their Applications* [**ADD DETAILS**]), but are even more general (see Zaslavsky 2006a).

D.7.1. *Group expansions.* [[LABEL 4.gx]]

[**DUPLICATES EARLIER SECTION in 20100209kaminskiruiz.tex.**]

Assume that Γ is a graph with no loose or half edges. We define the *expansion graph* $\mathfrak{G}\Gamma := (V, \mathfrak{G} \times E)$ and call Γ the *base graph*. The endpoints for the edge $(\mathfrak{g}, e)$ are just the endpoints of e. The gain of $(\mathfrak{g}, e{:}vw)$ is $\mathfrak{g} = \varphi_{\mathfrak{G}\Gamma}(\mathfrak{g}, e{:}vw)$.

Example D.6. [Group expansions of K_3][[LABEL X:20100211 Eg group expansions]] In Figure D.8 we have K_3 with two group expansions $\pm K_3$ and $\mathfrak{S}_3K_3$. Note that for $\mathfrak{S}_3K_3$ only four edges are actually labelled.

Example D.7. [$\mathfrak{S}_3C_4$][[LABEL X:20100211 bal circ property of ge]] In Figure D.9 we have the group expansion $\mathfrak{S}_3C_4$. Observe the following interesting property about group expansions and circles of the graph. Consider the path

$$W := v_1[(321)e_{12}]v_2[(13)e_{23}]v_3[1_{\mathfrak{S}_3}e_{34}]v_4$$

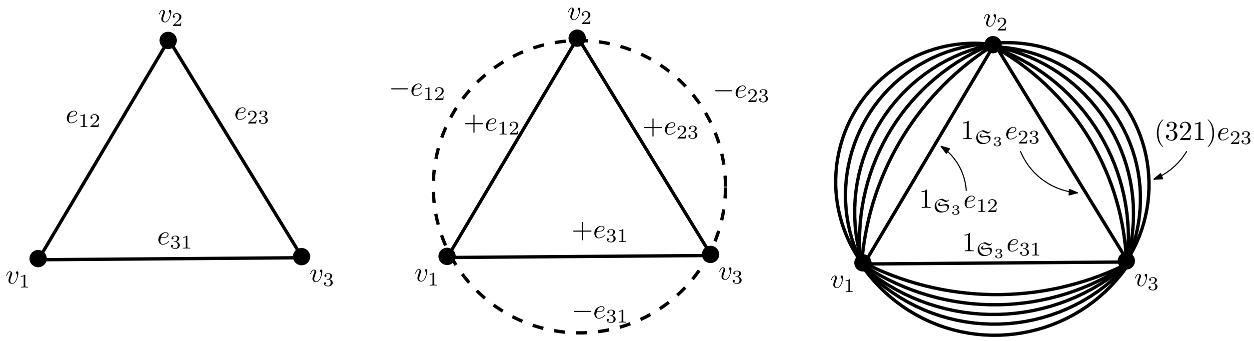

FIGURE D.8. (a) K_3, (b) $\pm K_3$, (c) $\mathfrak{S}_3 K_3$ (with only four edges labeled).
[[LABEL F:20100211: groupexpexamples]]

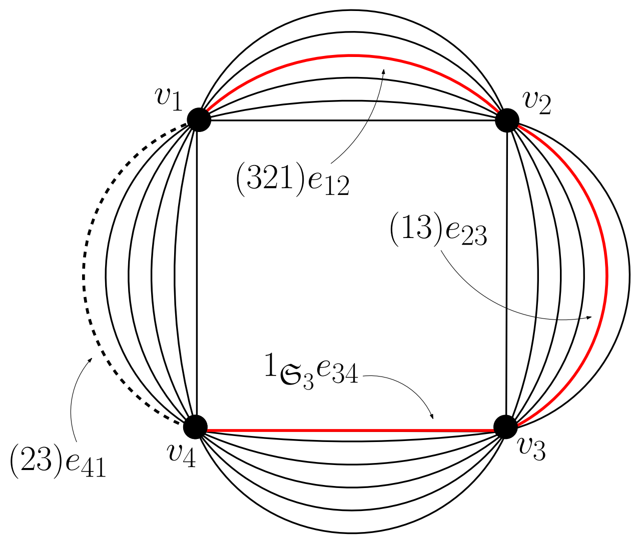

FIGURE D.9. $\mathfrak{S}_3 C_4$.
[[LABEL F:20100211: C4example]]

in the group expansion $\mathfrak{S}_3 C_4$ (colored red in Figure D.9). Since there is exactly one edge from v_4 to v_1, for each element of the gain group $\mathfrak{S}_3$ there must be a unique edge $\mathfrak{g}e_{41} \in \mathfrak{S}_3 \times E$ that completes a balanced circle in the group expansion $\mathfrak{S}_3 C_4$. In other words, $W \cup \mathfrak{g}e_{41}$ will be balanced. The edge in our example is $(23)e_{41}$, drawn as a dashed line in Figure D.9. Multiplying from left to right, the product is $(321)(13)1_{\mathfrak{S}_3}(23) = 1_{\mathfrak{S}_3}$, so $W \cup (23)e_{41}$ is a balanced circle.

It seems that with our construction there should be some sort of "projection" from the expansion graph to the base graph. This is true. In fact, it is still true if we consider something called a "biased expansion".

D.7.2. *Biased expansions.* [[LABEL 4.bx]]
Suppose Γ is an ordinary graph. A *biased expansion* of Γ is a biased graph Ω together with a graph homomorphism $p : \|\Omega\| \to \Gamma$, called the *projection*, that is surjective, is the identity function on V, and has the property that, for any cycle C in Γ, any $e \in C$, and any

path $\tilde{P}$ in Ω that projects onto $C \setminus e$, there exists a unique $\tilde{e} \in p^{-1}(e)$ such that $\tilde{P} \cup \tilde{e}$ is balanced. We call $p^{-1}(e)$ the *fiber over* e.

Because p is the identity on V and is one-to-one on $E(\tilde{P})$, p necessarily projects tP isomorphically onto $C \setminus e$. Thus, $\tilde{P}$ is a path simply because it contains only one edge for each edge of $C \setminus e$ and because $C \setminus e$ is a path.

Example D.8. [Projection for group expansions][[LABEL X:20100211 Eg group expansions]] In Figure D.10 we have an ordinary graph Γ and its group expansion $\mathfrak{S}_3\Gamma$. To

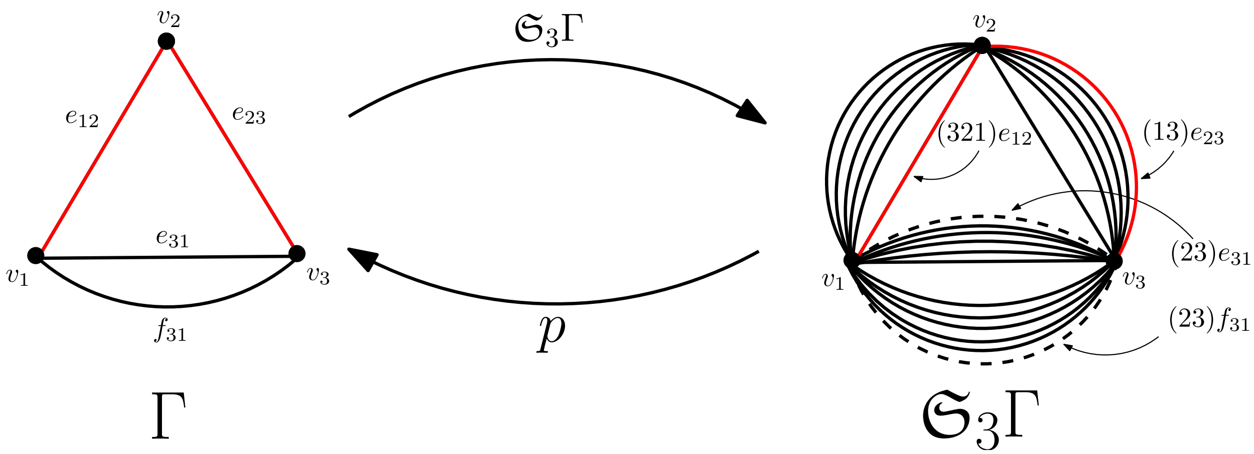

FIGURE D.10. Γ and its group expansion $\mathfrak{S}_3\Gamma$.
[[LABEL F:20100211: gammaexample]]

visualize how we defined the map p consider the following two circles in Γ:

$$C_1 = v_1 e_{12} v_2 e_{23} v_3 e_{31} v_1 \text{ and } C_1 = v_1 e_{12} v_2 e_{23} v_3 f_{31} v_1$$

And consider the path $P = C \setminus e_{31} = C \setminus f_{31}$ (drawn red in Γ of Figure D.10). Lifting P to the expansion graph $\mathfrak{S}_3\Gamma$ we pick some path $\tilde{P}$ that projects onto P. For example let $\tilde{P} = v_1[(321)e_{12}]v_2[(13)e_{23}]v_3$ (drawn red in $\mathfrak{S}_3\Gamma$ of Figure D.10). Now there exists a unique element $\tilde{e_{31}} \in p^{-1}(e_{31})$ such that $\tilde{P} \cup \tilde{e_{31}}$ is a balanced circle. This edge $\tilde{e_{31}} = (23)e_{31}$ is drawn in Figure D.10 as a dashed line. Similarly we have the unique element $\tilde{f_{31}} = (21)f_{31} \in p^{-1}(f_{31})$ which makes $\tilde{P} \cup \tilde{f_{31}}$ a balanced circle.

Proposition D.4. [[LABEL P:20100211fiber]] *If Γ is connected, every edge fiber $p^{-1}(e)$ has the same cardinality.*

Proof. A homework exercise. □

Proposition D.5. [[LABEL P:20100211bxdigons]] *If ef is a digon in the base graph Γ, then there exists a unique pairing (i.e., bijection) $\psi : p^{-1}(e) \to p^{-1}(f)$ such that $\tilde{e}\tilde{f}$ is balanced $\iff$ $\tilde{f} = \psi(\tilde{e})$.*

Proof. A homework exercise. □

Example D.9. [Biased expansion of K_3][[LABEL X:20100211 biasedexp of k3]] If $m = |p^{-1}(e)|$, sometimes we may use the notation $\Omega = mK_3$ to remind us that $\Gamma = K_3$. Also we may use the notation $\|\Omega\| = mK_3$ as the underlying graph of the biased expansion of K_3. In

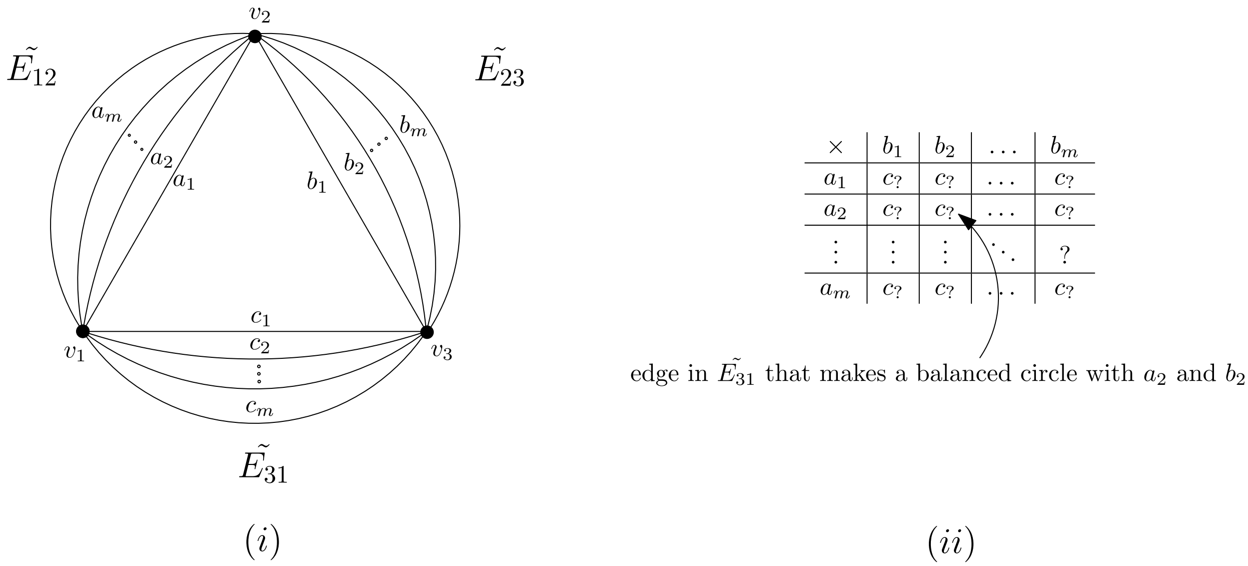

FIGURE D.11. (i) Biased expansion $\Omega = m \cdot K_3$, (ii) multiplication table of $m \cdot K_3$. [[LABEL F:20100211: biasexp latinsq]]

Figure D.11 we see the biased expansion $\Omega = mK_3$. We label the edges of from vertex v_1 to v_2 as $\tilde{E_{12}} = \{a_1, a_2, \dots, a_m\}$, the edges of from vertex v_2 to v_3 as $\tilde{E_{23}} = \{b_1, b_2, \dots, b_m\}$ and the edges of from vertex v_3 to v_1 as $\tilde{E_{31}} = \{c_1, c_2, \dots, c_m\}$. Now we can set up a multiplication table as follows (see Figure D.11). the product of $a_j \in \tilde{E_{12}}$ and $b_k \in \tilde{E_{23}}$ is some edge $c_? \in \tilde{E_{31}}$ which makes a balanced circle with a_j and b_k. From this construction we have the following theorem.

Theorem D.6. [[LABEL T:20100211bxk3]] *The multiplication table is a Latin square and is the multiplication table of mK_3 where m is the order of the Latin square.*

We say the Latin square is well defined (from mK_3) up to *parastrophe*, meaning it is well defined up to *isotopy* (permuting rows, columns, or names) and *conjugation* (changing roles). We can think of a biased expansion as a generalized latin square.

Because of this nice connection between biased expansions and Latin squares we can now say that biased graphs are indeed a strict generalization of gain graphs.

Corollary D.7. *There exist biased expansions that are not group expansions.*

Proof. There exist Latin squares that are not equivalent (technically, not isotopic, or more subtly, not parastrophic) to a group multiplication table (that is, a "Cayley table" of a group). □

Example D.10. [[LABEL X:20100211 dlattice]] The full group expansion $(\mathfrak{G}K_n)^\bullet$ gives the rank n Dowling lattice of the group $\mathfrak{G}$ (Dowling 1973b). Specifically, the Dowling lattice $Q_n(\mathfrak{G})$ consists of the edge sets of $fGK_n^\bullet$ that are closed (see Section IV.**??**), ordered by set inclusion.

E. MINORS

[[LABEL 4.minors]]

2010 Feb 16: Simon Joyce

A *minor*, as always, is a contraction of a subgraph. And, as usual, the difficult part of the concept is contraction.

E.1. **Subgraphs and deletion.** [[LABEL 4.deletion]]

We already discussed what a subgraph is of a gain graph or a biased graph: the subgraph simply inherits the structure of the original.

E.2. **Contraction.** [[LABEL 4.contraction]]

Technically, one can say the same for contraction. The complication is in the way the structure passes to the contraction.

E.2.1. *Contraction for gain graphs.* [[LABEL 4.ggcontraction]]

For a gain graph Φ we will define how to contract an arbitrary edge set $S \subseteq E(\Phi)$, rather than starting with the special case of contraction of a single edge as we did in Chapter II. As with signed graphs, how we contract balanced components will be different from how we contract unbalanced components.

Let B_i be the balanced components of S, let $V_i := V(B_i)$, and let

$$V_0(S) := \{\text{vertices of unbalanced components of } S\} = V(S) \setminus \bigcup V_i.$$

To contract S, first we switch so the edges of each balanced component have identity gain. Now contract each balanced component to a single vertex as you would normally. Then delete all vertices in V_0 and all remaining edges of S. This may create some number of half and loose edges.

Formally, $\Phi/S = (V(\Phi/S), E(\Phi/S), \varphi_{\Phi/S})$, where $V(\Phi/S) := \boldsymbol{\pi}_{\mathrm{b}}(S) = \{V_1, \dots, V_k\}$ (with $k = b(S)$) and $E(\Phi/S) := E \setminus S$. Each edge $e \in E \setminus S$ is incident to $V_i \in \Phi/S$ once for each endpoint of e that is in V_i in Φ. Contractions that give loops and half edges are shown in Figure E.1. Let ζ be the switching function we applied so we could contract the balanced components of S. We define $\varphi_{\Phi/S} = \varphi^{\zeta}(e)$ for $e \in E_*(\Phi/S) := \{\text{ordinary edges in } \Phi/S\}$. If the gain of $e{:}vw$ is from v to w in Φ, then its gain is from V_i to V_j in Φ/S where $v \in V_i$ and $w \in V_j$.

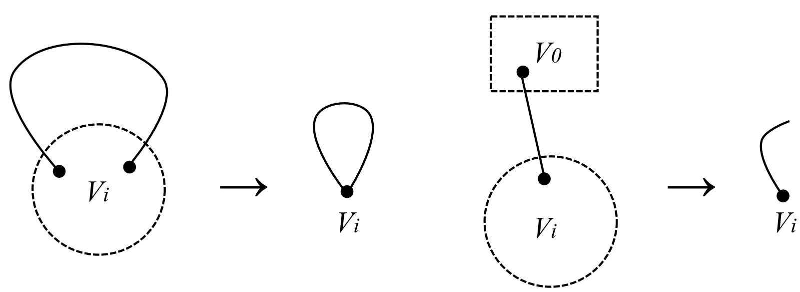

FIGURE E.1
[[LABEL F:20100216halfloop]]

An example of a gain graph contraction is shown in Figure E.2. The graph being contracted is Φ and the gain groups is $\mathfrak{S}_3$. Here the red edges are being contracted. The balanced components are circled in blue. The vertex switches are also given that make all the edge gains in each balanced component identity.

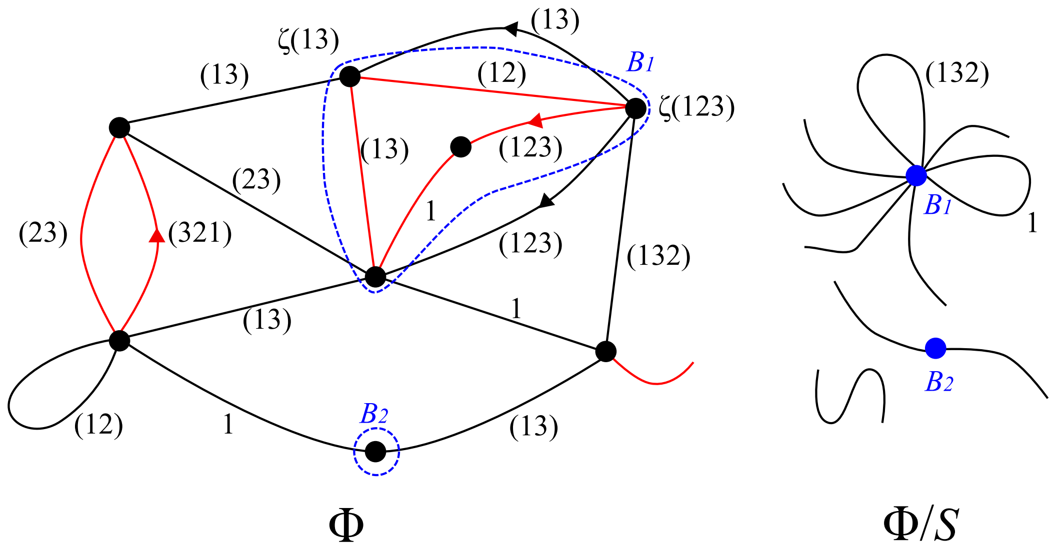

FIGURE E.2
[[LABEL F:20100216gaincontraction1]]

Is there a rational explanation of this complicated procedure? Here is the *Gestalt* of gain graph contraction (as Professor Rota would say[13]). Think of each vertex v_i as representing a variable x_i. Then an edge $e{:}v_iv_j$ with gain $\mathfrak{g}$ corresponds to the equation $x_j = x_i\mathfrak{g}$. A switching function ζ changes the variable x_i to $x_i\zeta(v_i)$. Switching $e{:}v_iv_j$ to have identity gain ought to make $x_i\zeta(v_i) = x_j\zeta(v_j)$. Since the switched edge gain is $\zeta(v_i)^{-1}\mathfrak{g}\zeta(v_j)$, we should have $x_i\zeta(v_i)\zeta(v_i)^{-1}\mathfrak{g}\zeta(v_j) = x_i\mathfrak{g}\zeta(v_j) = x_j\zeta(v_j)$ so this works out. If we can switch so all the edges in a component of some subset of the edges have identity gain, then it makes sense that we should be able to treat this component as a single vertex, or variable, since all the variables in such a component are necessarily equal.

Since the definition involves an arbitrary choice of switching function, we must—as we did for signed graphs—ask in what sense the contraction of a gain graph is well defined.

Proposition E.1. [[LABEL P:20101216 Gain contractions are switching equiv]] *Given a gain graph Φ and $S \subseteq E(\Phi)$:*

(i) *Any two contractions Φ/S are switching equivalent.*
(ii) *Any switching of Φ/S is a contraction of Φ.*

This proposition means: (i) the contraction Φ/S is well defined up to switching, since $[\Phi/S]$ is a well defined switching class, but (ii) there is no more refined characterization, since Φ/S could be any member of the switching class.

Proof. We may assume by switching that the edges of the balanced components of $\Phi|S$ have identity gain. Let V_i be the balanced vertex components of $\Phi|S$. These will be the vertices of Φ/S.

For Part (i), let ζ be a switching function of Φ such that all the edges of of the balanced components still have identity gain in Φ^ζ so that we can contract by S. Let $e{:}vw$ be an edge of S in a balanced component. Notice that $\zeta(v) = \zeta(w)$ since this edge still has identity

---

[13] *Combinatorics: The Rota Way*, Joseph P.S. Kung and Catherine Yan, Cambridge University Press.

gain after switching. It follows that for $v, w \in V_i$, a balanced vertex component of $\Phi|S$, then $\zeta(v) = \zeta(w)$ since the balanced components are connected.

Now define a switching function ζ' of Φ/S that is a restriction of ζ such that $\zeta'(V_i) = \zeta(v)$ for some $v \in V_i$. This is well defined by the previous observation. Then $(\Phi/S)^{\zeta'} = \Phi^{\zeta}/S$ since $\varphi_{\Phi^{\zeta}/S}(e) = \varphi_{(\Phi/S)^{\zeta'}}(e)$ for any edge in $E \setminus S$.

For Part (ii), let ζ be a switching function of Φ/S. Define a switching function $\hat{\zeta}$ on Φ that is an extension of ζ such that for v a vertex of Φ, $\hat{\zeta}(v) = \zeta(V_i)$ where v is in the balanced vertex component V_i of $\Phi|S$. Notice that the edge gains of each balanced component of $\Phi^{\hat{\zeta}}|S$ are still the identity. Then $\Phi^{\hat{\zeta}}/S = (\Phi/S)^{\zeta}$ since every edge in $E(\Phi^{\hat{\zeta}}) \setminus S$ has the same gain as in $(\Phi/S)^{\zeta}$. □

2010 Feb 18: Jackie Kaminski and Amanda Ruiz

In defining contraction of a gain graph we have some freedom of choice in how we switch to contract. However, given a particular edge set S, regardless of which switching function we use to give every balanced component of $\Phi|S$ all-identity gains, all contractions $\Phi \setminus S$ are switching equivalent. In other words, contraction is really defined on switching classes $[\Phi]$ of gain graphs rather than on individual gain graphs, and its result is really a switching class $[\Phi/S]$, not an individual gain graph. This observation is made precise in the following proposition.

Proposition E.2. [[LABEL P:20100218: contraction on switching classes]]

(i) *Any two contractions Φ/S are switching equivalent, so $[\Phi/S]$ is a well defined switching class.*

(ii) *Any switching of a contraction Φ/S is another contraction Φ/S. In other words Φ/S is well defined only up to switching classes and not any more narrowly.*

2010 Feb 16: Simon Joyce

E.2.2. *Contraction for biased graphs.* [[LABEL 4.bgcontraction]]

For a biased graph Ω and $S \subseteq E(\Omega)$, let B_i be the balanced components of S and $V_i = V(B_i)$. We define $\Omega/S := (V(\Omega/S),\ E(\Omega/S)$, and $\mathcal{B}(\Omega/S))$, where $V(\Omega/S) = \boldsymbol{\pi}_{\mathrm{b}}(S) = \{V_1, \ldots, V_k\}$, where $k = b(S)$ and $E(\Omega/S) = E \setminus S$. $e \in \Omega/S$ is incident to V_i once for each endpoint of e that's in V_i in Ω. $\mathcal{B}(\Omega/S) = \{C \in \mathcal{C}(\Omega/S) | C = C' \setminus S$ for some $C' \in \mathcal{B}(\Omega)\}$. Any C' of this form must be in $\Omega \setminus V_0(S)$.

Proposition E.3. *For a biased graph Ω and $S \subseteq E(\Omega)$, an equivalent formulation of the balanced circles in Ω/S is that*

$$\mathcal{B}(\Omega/S) = \{C \in \mathcal{C}(\Omega/S) \mid \forall\ C' \in \mathcal{C}(\Omega) \textit{ such that } C = C' \setminus S \textit{ and } C' \in \mathcal{B}(\Omega)\}.$$

The proof is a nontrivial exercise.

An example of biased graph contraction is given in Figure E.3. If C_1' is balanced in Ω, then C_1 is balanced in Ω/S. C_2' becomes two circles in Ω/S. The balance or unbalance of these circles does not depend on the sign of C_2', but on the signs of the other two circles in the theta subgraph formed by C_2' and the red path in B_4. The other circles are lost in the contraction since they have vertices in V_0.

Now suppose $\Omega = \langle\Phi\rangle$. We want to show $\langle\Phi/S\rangle = \Omega/S$. Clearly $\|\Phi/S\| = \|\Omega/S\|$, so we only need to show a circle is balanced in Ω/S if and only if it has identity gain in Φ/S. Before we prove this result we need the following lemma.

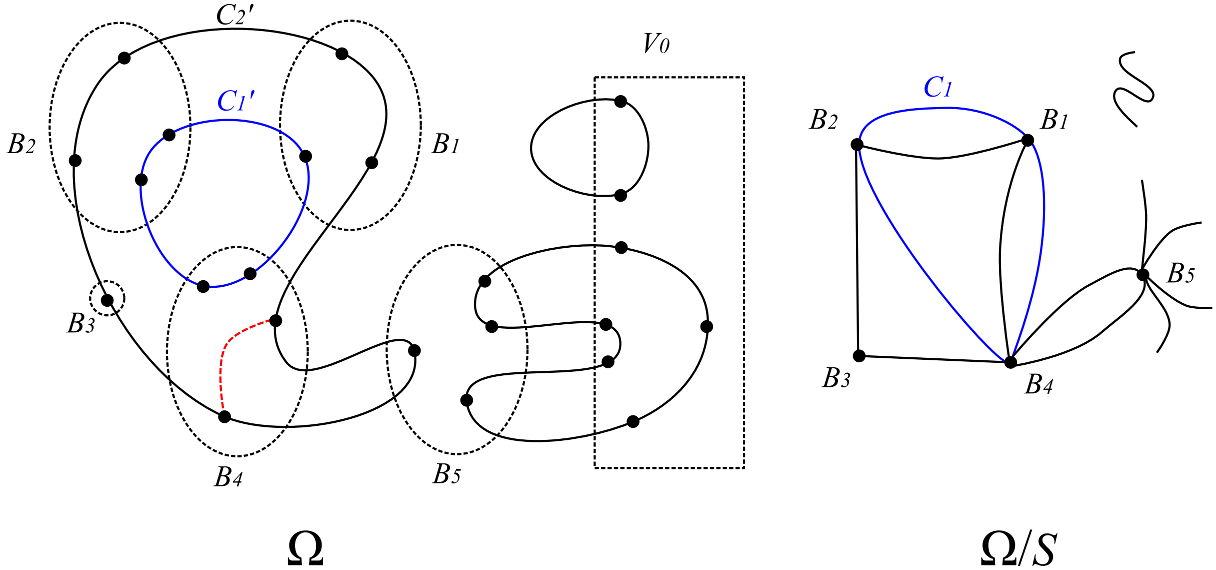

FIGURE E.3
[[LABEL F:20100216biasedcontraction]]

Lemma E.4. [[LABEL L:20100216 circles in gain graphs]] *Given a gain graph Φ and $S \subseteq E(\Phi)$. A circle C in $\|\Phi/S\|$ has the form $C = e_1 e_2 \cdots e_l$, where there is a circle $e_1 P_1 e_2 P_2 \cdots e_l P_l$ in $\|\Phi\|$, with each P_i a path in some balanced component B_i of $\Phi|S$.*

Proof. Since every vertex of Φ/S corresponds to a balanced vertex set V_i in $\Phi|S$, there will be a path $P_i \in B_i$ from the endpoints of e_i and e_{i+1} that are in V_i since B_i is a connected component as shown in Figure E.4. So $e_1 P_1 e_2 P_2 \cdots e_l P_l$ forms a circle in $\|\Phi\|$. □

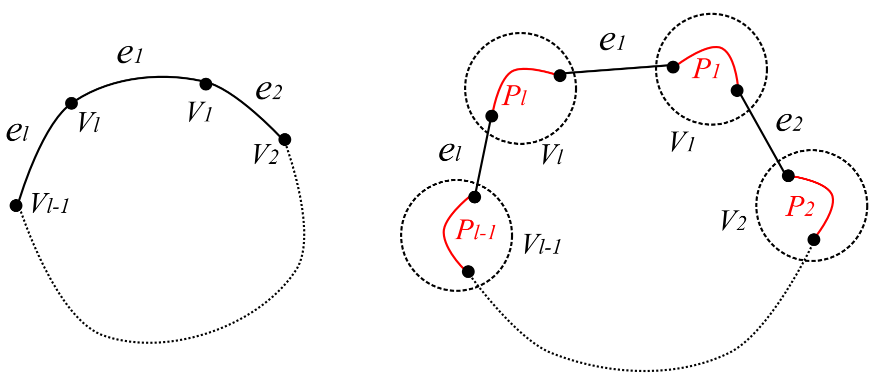

FIGURE E.4. C in Φ/S on the left and a circle with the edges of C in Φ on the right.
[[LABEL F:20100216circleblowup]]

Proposition E.5. [[LABEL P:20100216Gain and biased contraction agree]] *Given a gain graph Φ and $S \subseteq E(\Phi)$. A circle C in $\langle\Phi\rangle/S$ is balanced if and only if C is balanced in $\langle\Phi/S\rangle$.*

Proof. Let $\Omega := \langle\Phi\rangle$ and let C be a balanced circle in Ω/S. Let C' be a circle in $\|\Phi\|$ that exists by Lemma E.4. There is a switching function ζ of Φ such that the paths in C' have identity gain since each path is contained in a balanced component. Then,

$$\begin{aligned}\varphi_{\Phi/S}(C) &= \varphi_{\Phi/S}(e_1)\varphi_{\Phi/S}(e_2)\cdots\varphi_{\Phi/S}(e_l)\\ &= \varphi^{\zeta}(e_1)\varphi^{\zeta}(e_2)\cdots\varphi^{\zeta}(e_l)\\ &= \varphi^{\zeta}(e_1)\varphi^{\zeta}(P_1)\varphi^{\zeta}(e_2)\varphi^{\zeta}(P_2)\cdots\varphi^{\zeta}(e_l)\varphi^{\zeta}(P_l)\\ &= \varphi^{\zeta}(C').\end{aligned}$$

It follows that $\varphi(C) = 1 \iff \varphi(C') = 1$. So C is balanced in Ω/S if and only if C has identity gain in Φ/S. □

2010 Feb 18: Jackie Kaminski and Amanda Ruiz

E.3. **Minors.** [[LABEL 4.minorsminors]]

Now that we have defined subgraphs and contraction on gain graphs and biased graphs, we will spend some time proving an essential property of their relationship.

Definition E.1. [[LABEL D:20100218: minor]] A *minor* of a gain graph or biased graph is a contraction of a subgraph.

Theorem E.6. [[LABEL T:20100218: minor of minor]]

(i) *Any sequence of taking subgraphs and contracting edge sets in a gain graph Φ results in a minor of Φ.*
(ii) *Any sequence of taking subgraphs and contracting edge sets in a biased graph Ω results in a minor of Ω.*

More simply stated, *a minor of a minor is a minor.*

Usually we need this theorem for gain graphs, but it is true of all biased graphs; it would be a shame (and foolish) not to prove it for biased graphs in general. We do need separate proofs for biased graphs and gain graphs, since a proof for biased graphs does not give the information on the contracted gains that is needed to have the theorem for gain graphs, while on the other hand, since there are biased graphs that aren't gain graphs, it certainly won't suffice to prove the gain-graph version only.

Still, because the proof for biased graphs is more difficult, this is a fine example of how useful it is to have gains and switching. "Switching makes the proof go smoother!"

The essential part of the proof of Theorem E.6 is to prove that a contraction of a contraction is a contraction. This fact is expressed by the following two lemmas.

Lemma E.7. [[LABEL L:20100218: gain]] *If S, T are disjoint subsets of $E(\Phi)$, then $(\Phi/S)/T$ and $\Phi/(S \cup T)$ are the same, up to vertex labels and switching equivalence.*

Lemma E.8. [[LABEL L:20100218: biased]] *If S, T are disjoint subsets of $E(\Omega)$, then $(\Omega/S)/T$ and $\Omega/(S \cup T)$ are the same, up to vertex labels.*

To explain "up to vertex labels": The vertex labels in Φ (or Ω; the facts are the same for both) are different from those in Φ/S because a vertex in the latter is a set of vertices of the former. For instance, if S has one edge, a link $e_v w$, then $v, w \in V(\Phi)$ become the vertex $\{v, w\} \in V(\Phi/S)$. Even if v is an isolated vertex in $\Phi|S$, v is not is a vertex in the contraction; the vertex is $\{v\}$. That is why we have to say "up to vertex labels": the names of the vertices change under contraction. However, the change is not arbitrary. There is a *natural* correspondence of vertices of Φ and Φ/S: a vertex $v \in V(\Phi)$ which is not deleted in contraction corresponds to a unique vertex $V_i \in V(\Phi/S)$, namely, that $V_i \in \boldsymbol{\pi}_{\mathrm{b}}(S)$ such that $v \in V_i$. (It's hard to explain the naturality of this correspondence in a principled way in terms of sets. That's one reason I think set theory is not an adequate foundation for mathematics; or if you prefer, mathematics is not merely a higher development of set theory.)

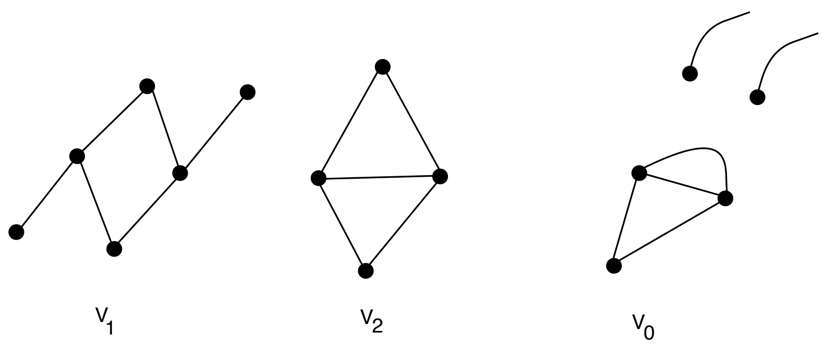

Φ/S

FIGURE E.5. In the gain graph Φ, the edge set S has balanced components V_1 and V_2 (more correctly, $S{:}V_1$ and $S{:}V_2$), and a union of unbalanced components V_0. Φ/S consists of two vertices V_1 and V_2, which are its components (as vertex graphs). The unbalanced vertices, those in V_0, are deleted in the process of contraction.

[[LABEL F:20100218: contract S]]

Before beginning the proof, I want to remind you that in a contraction Φ/S, the unbalanced components of $\Phi|S$ disappear (potentially turning some edges of S^c into half or loose edges), and a balanced component of $\Phi|S$ is contracted down to a single vertex. Figure E.5 shows $\Phi|S$ and then what happens to these graph elements in Φ/S.

Contraction[2] of a gain graph.

We now begin the proof of Theorem E.6 by establishing Lemma E.7.

Proof of Lemma E.7. We need to examine both $(\Phi/S)/T$ and $\Phi/(S \cup T)$ and compare them. Let's begin with the former. To consider $(\Phi/S)/T$, we must first consider Φ/S. To this end, choose a switching function ζ_S such that φ^{ζ_S} is identically $1_{\mathfrak{G}}$ on the balanced components of $\Phi|S$. In most cases there are many possible choices for ζ_S; fix one of them. As is often true in a proof by switching, we can simplify the details by adequate switching: according to Proposition E.2(ii), we may as well assume Φ has been switched by ζ_S before we begin; in other words, we may assume $\varphi|_S \equiv 1_{\mathfrak{G}}$.

For Φ/S, notice that $V(\Phi/S) = \boldsymbol{\pi}_{\mathrm{b}}(\Phi|S)$ by definition. So as to have a way to talk about these vertices, let $\boldsymbol{\pi}_{\mathrm{b}}(\Phi|S) =: \{V_1, \ldots, V_k\}$, let $V_{\mathrm{b}}(S) := V_1 \cup \cdots \cup V_k$, and let V_0 denote the set of vertices that are in unbalanced components of $\Phi|S$. So, $\{V_1,\ldots,V_k,V_0\}$ is a partition of $V(\Phi)$. Φ/S has gain function $\varphi_{\Phi/S}(e)=\varphi^{\zeta_S}(e)$, or more concisely, $\varphi_{\Phi/S}=\varphi^{\zeta_S}|_{S^c}$.

Consider as an example Figure E.6, which depicts $\Phi|S$. The circles represent the balanced vertex components (including isolated vertices) and the unbalanced components are represented by the rectangle. Φ is likely a much larger graph, with many edges outside of the V_i's, but all the vertices of Φ are contained this depiction of $\Phi|S$. In Φ^{ζ_S}, all gains within any $S{:}V_i$ (for $i>0$) are $1_{\mathfrak{G}}$, and we have essentially no information about the gains outside of the $S{:}V_i$'s. Finally, the vertices of Φ/S are $V_1,\ldots,V_k$. Figure E.7 shows $\Phi|(S\cup T)$, with T edges shown in red.

Moving on to $(\Phi/S)/T$, let ζ_T be a switching function on Φ/S such that all balanced components of $(\Phi/S)|T$ have identity gains. Again there may have been many choices for ζ_T, but we pick one. Now, $V(\Phi/S)/T)=\boldsymbol{\pi}_{\mathrm{b}}((\Phi/S)|T)=:\{W_1,\ldots,W_l\}$, and the gain function is $\varphi_{(\Phi/S)/T}=\varphi_{\Phi/S}^{\zeta_T}|_{(S\cup T)^c}$. The vertices of $(\Phi/S)/T$ are $W_1,\ldots,W_l$.

Let's consider for a moment the vertex set W_0, the vertex set of the unbalanced components of $(\Phi/S)|T$. There are four different ways a vertex V_i of Φ/S can wind up in W_0. First, in Φ, a vertex of V_i might have been adjacent to a vertex in V_0 through an edge $e\in T$ (see edge e in Figure E.7). In Φ/S, e becomes a half edge, so the component in $(\Phi/S)|T$ with vertex V_i is unbalanced. Secondly, there might be an edge $f\in T$ that is an unbalanced loop at V_i in Φ/S, which means that in Φ there is an unbalanced circle (possibly a loop) in $S\cup f$ which contains f and at least one vertex of V_i. Similarly, there might be an unbalanced circle of two or more edges, $C=gh\cdots$, in Φ/S whose vertices include V_i. This circle must come from an unbalanced circle C' in $[\Phi{:}V_{\mathrm{b}}(S)]|(S\cup C)$ such that $g,h,\ldots\in C'$ and at least one vertex in V_i is on C', since C must arise from a circle in $\Phi{:}V_{\mathrm{b}}(S)$ and a circle that is the contraction of a balanced circle is balanced. Finally, V_i can be connected, via a path in $(\Phi/S)|T$, to any of the above.

We would like to compare this to $\Phi/(S\cup T)$, so now we consider the structure of $\Phi/(S\cup T)$. To talk about this contracted graph we need a switching function ζ such that all the balanced components of $S\cup T$ have identity gains in Φ^ζ. (We'll pick a particular ζ shortly, in Lemma E.9.) We know $V(\Phi/(S\cup T))=\boldsymbol{\pi}_{\mathrm{b}}(\Phi|(S\cup T))=:\{X_1,\ldots,X_n\}$, say (to have a simple notation); let X_0 denote the vertex set of the union of all unbalanced components of $\Phi|(S\cup T)$; and recall that $\varphi_{\Phi/(S\cup T)}(e)=\varphi^\zeta(e)$ for $e\notin S\cup T$.

We wish to prove that the partition $\{X_1\ldots,X_m,X_0\}$ of $V(\Phi)$ is the same as the partition $\{\bar W_1,\ldots,\bar W_l,V_0\cup\bar W_0\}$, where $\bar W_j=\bigcup_{V_i\in W_j}V_i$ (so $\bar W_j\subseteq V(\Phi)$).

First we demonstrate that each $\bar W_j$ with $j>0$ is contained in a single X_p for some $p>0$. Any two vertices V_r and V_s in W_j are connected by a path $f_1f_2\cdots f_l$ of T-edges in Φ/S. The common vertex of f_{h-1} and f_h in Φ/S is a V_i. The second endpoint of f_{h-1} and the first endpoint of f_h are connected in $(\Phi|S){:}V_i$. Any vertex of V_r is connected to the first vertex of f_1 in $(\Phi|S){:}V_r$, and any vertex of V_s is connected to the last vertex of f_l in $(\Phi|S){:}V_s$. Therefore, the vertices in $\bar W_j$ are connected in $\Phi|(S\cup T)$. **[THE FOLLOWING NEEDS TO BE IMPROVED. POSSIBLY, A GENERAL GAIN-GRAPH CONTRACTION LEMMA ABOUT BALANCE?]** Moreover, each $(S\cup T){:}\bar W_j$ is balanced, because $S{:}\bar W_j$ is balanced, and $(\Phi/S)|T$ is balanced. (If you don't believe this, wait until Lemma E.9 which proves that the gains of $\Phi|(S\cup T)\setminus X_0$ can be switched to all identity gains.)

Next, we prove that $V_0\cup\bar W_0$ is contained in X_0. It is easy to see by looking at Figure E.7 that any vertex in $V_0\cup\bar W_0$ is in an unbalanced component of $\Phi|(S\cup T)$, but we have to explain the details. That $V_0\subseteq X_0$ is obvious. Consider a V_i that is connected to V_0 in

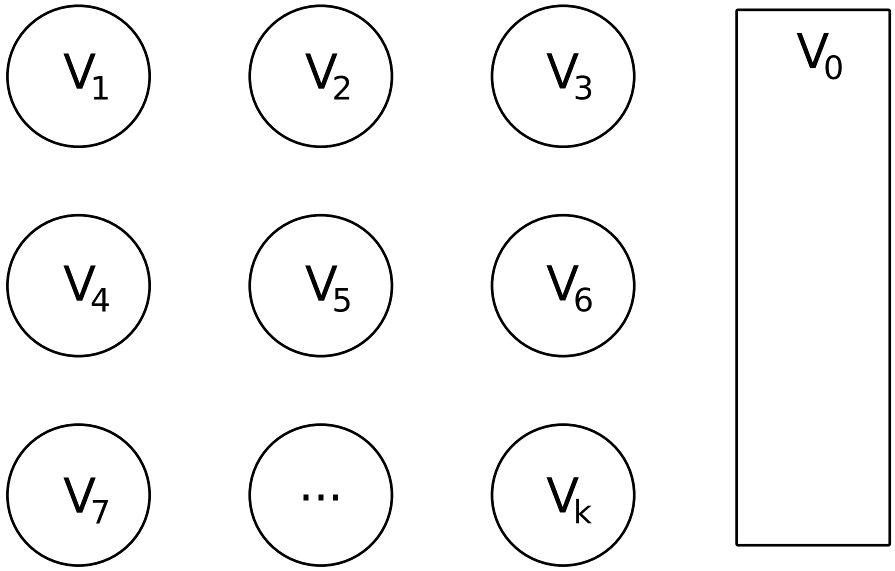

Figure E.6. $\Phi|S$. Balanced (vertex) components are represented by circles, all the unbalanced components are represented by the rectangle, and loose edges are omitted.

[[LABEL F:20100218: V]]

$\Phi|(S \cup T)$. It is part of an unbalanced component of $S \cup T$, so is contained in X_0. If V_i is adjacent to V_0 in Φ by a T-edge f, then in $(\Phi/S)|T$, V_i supports f as a half edge. If V_i

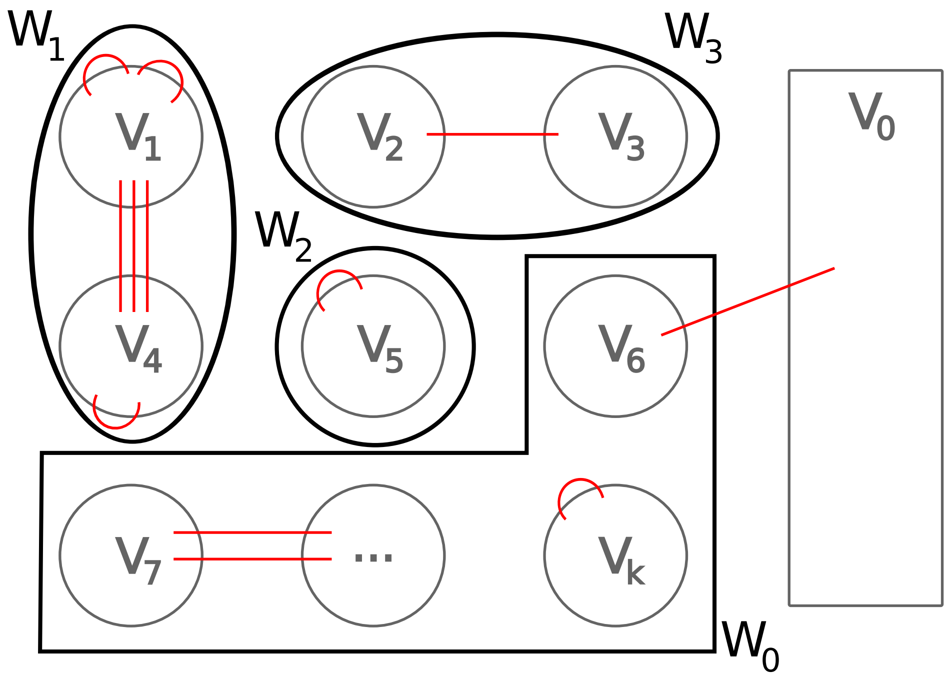

FIGURE E.7. $\Phi|(S \cup T)$. Balanced (vertex) components are represented by circles, all the unbalanced components are represented by the rectangle, and edges of T shown in red.

[[LABEL F:20100218: W]]

is connected to V_0 by a path $ff_1 \cdots f_l$ longer than one edge, whose edge incident to V_0 is f, then in $(\Phi/S)|T$, V_i is connected to the half edge f by the path $f_1 \cdots f_l$; consequently,

$V_i \in W_0$. Now, suppose $V_i \in W_0$. In $\Phi|(S \cup T)$, one possibility is that V_i is connected to V_0; then as we just showed, $V_i \subseteq X_0$. The other possibility is that V_i is in an unbalanced component $\Phi|(S \cup T){:}X$ that is not connected to V_0 by a path in $S \cup T$. Then $\Phi|(S \cup T){:}X$ contains an unbalanced circle C, or a half edge e. In the latter case, V_i is connected to the half edge e in Φ/S. In the former case, switching so the gains on S are all $1_{\mathfrak{G}}$, C contracts to a closed walk W in Φ/S whose gain is not the identity. [**NEEDS MORE—a lemma in a future day's notes?—to prove that W contains an unbalanced circle.**]

[**We will show the other containment on another day.**]

Let's return to the task of showing that $(\Phi/S)^{\zeta_T}/T$ and $\Phi^{\zeta}/(S \cup T)$ have the same gains on their edges, for a suitable switching function $\zeta : V(\Phi) \to \mathfrak{G}$. That will certainly suffice to show they are switching equivalent. Our ζ_T (on Φ/S) corresponds to a switching function ζ'_T on Φ, defined by $\zeta'_T(v) := \zeta_T(V_i)$ for any $V_i \in \boldsymbol{\pi}_{\mathrm{b}}(S)$ and any vertex $v \in V_i$, and letting $\zeta'_T(v)$ be arbitrary for $v \in V_0(S)$. Then $\zeta := \zeta'_T$ is the switching function we want.

Lemma E.9. [[LABEL L:20100218 gain functions]] $\Phi^{\zeta'_T}$ *has identity gains on the balanced components of* $\Phi|(S \cup T)$.

Proof. First, consider $e{:}vw \in S$ ($v = w$ is permissible), where e is in a balanced component of $\Phi|(S \cup T)$. Since every subset of a balanced subset is balanced, e is in a balanced component of $\Phi|S$, so $\varphi(e) = 1_{\mathfrak{G}}$. Therefore,

$$\varphi^{\zeta'_T}(e) = \zeta'_T(v)^{-1}\varphi(e)\zeta'_T(w) = \zeta'_T(v)^{-1}1_{\mathfrak{G}}\zeta'_T(w) = 1_{\mathfrak{G}},$$

where the last equality is true because, since v and w are in the same balanced component of $\Phi|S$, $\zeta'_T(v) = \zeta'_T(w)$.

Now consider $e{:}vw \notin S$, where $v \in V_i$ and $w \in V_j$. The calculation begins similarly but moves into the contracted graph:

$$\varphi^{\zeta}(e) = \zeta'_T(v)^{-1}\varphi(e)\zeta'_T(w) = \zeta_T(V_i)^{-1}\varphi_{\Phi/S}(e)\zeta_T(V_j) = (\varphi_{\Phi/S})^{\zeta_T}(e).$$

If $e \in T$, this is $1_{\mathfrak{G}}$ by the definition of ζ_T since e is in a balanced component of $(\Phi/S)|T$. Therefore $\varphi^{\zeta'_T}$ is identically $1_{\mathfrak{G}}$ on the balanced components of $\Phi|(S \cup T)$. □

A conclusion is that we can use $\varphi^{\zeta'_T}$ to contract by S and then, without further switching, by T. More importantly, any edges that are not in $S \cup T$ have the same gain in both Φ/S and $\Phi^{\zeta'_T}$ and consequently in $(\Phi/S)/T$ and $\Phi^{\zeta'_T}/(S \cup T)$. By Proposition E.2, $[(\Phi/S)/T] = [\Phi^{\zeta'_T}/(S \cup T)]$ (up to vertex labels). This observation concludes the proof of Lemma E.7.

[**Proof will be finished in the next n sets of notes [i.e., the missing step]**] □

2010 Feb 23: Nathan Reff

Contraction[2] *of a biased graph.*

The task now is to prove Lemma E.8, which we do by way of an explicit formula for the contraction of a contraction.

Lemma E.10. [[LABEL E:20100223 bgccecu]] $(\Omega/S)/T = \Omega/(S \cup T)$ *up to vertex labels.*

Proof. Here is an outline of the proof:

(1) Show the edge sets agree.
(2) Show the vertex components and loose edges agree.
(3) Show the balanced components agree.

The first step is easy, but the rest are difficult, more so than with gain graphs since, as we no longer have gains, we won't have the considerable convenience of switching.

Let's review some old and new definitions and notation. As usual, $\|\Omega\| = (V, E)$. The vertices of Ω/S are the vertex sets of the balanced components of the restriction to (V, S); the edges are those not in S:

$$\begin{aligned} \boldsymbol{\pi}_{\mathrm{b}}(S) &:= \boldsymbol{\pi}_{\mathrm{b}}(\Omega|S), \\ V(\Omega/S) &:= \boldsymbol{\pi}_{\mathrm{b}}(\Omega|S) = \{V_1, V_2, \ldots, V_k\}, \\ E(\Omega/S) &:= E \setminus S. \end{aligned}$$

Two handy notations, new and old (written in short, medium, and long forms, of which each has its use):

$$\begin{aligned} V_{\mathrm{b}} := V_{\mathrm{b}}(S) := V_{\mathrm{b}}(\Omega|S) &:= \bigcup \boldsymbol{\pi}_{\mathrm{b}}(S) = V_1 \cup V_2 \cup \cdots \cup V_k, \\ V_0 := V_0(S) := V_0(\Omega|S) &:= V_{\mathrm{b}}(S)^c = V \setminus V_{\mathrm{b}}(S). \end{aligned}$$

We'll also want notations for the vertex sets of $(\Omega/S)/T$ and $\Omega/(S \cup T)$; thus,

$$\begin{aligned} V(\Omega/(S \cup T)) &:= \boldsymbol{\pi}_{\mathrm{b}}(S \cup T) := \{X_1, X_2, \ldots, X_m\}, \\ X_{\mathrm{b}} &:= V_{\mathrm{b}}(\Omega|(S \cup T)) := \bigcup \boldsymbol{\pi}_{\mathrm{b}}(S \cup T) = X_1 \cup X_2 \cup \cdots \cup X_m, \\ X_0 &:= V_0(\Omega|(S \cup T)) := X_{\mathrm{b}}^c = V \setminus X_{\mathrm{b}}, \end{aligned}$$

and

$$\begin{aligned} V((\Omega/S)/T) &:= \boldsymbol{\pi}_{\mathrm{b}}((\Omega/S)|T) := \{W_1, W_2, \ldots, W_l\}, \\ W_{\mathrm{b}} &:= V_{\mathrm{b}}((\Omega/S)|T) := \bigcup \boldsymbol{\pi}_{\mathrm{b}}((\Omega/S)|T) = W_1 \cup W_2 \cup \cdots \cup W_l, \\ W_0 &:= V_0((\Omega/S)|T) := W_{\mathrm{b}}^c = V(\Omega/S) \setminus W_{\mathrm{b}} = \boldsymbol{\pi}_{\mathrm{b}}(S) \setminus W_{\mathrm{b}}. \end{aligned}$$

Now, compare the definitions in the two contractions we hope to be equal. For $\Omega/S/T$:

I. $V(\Omega/S/T) = \{W_1, W_2, \ldots, W_l\}$. This is a partial partition of $V(\Omega/S)$, so each $W_i \subseteq V(\Omega/S) \subseteq \mathcal{P}(V)$. Let's define $\bar{W}_j = \bigcup\{V_i \mid V_i \in W_j\}$.
II. $E(\Omega/S/T) = E \setminus (S \cup T)$.

On the other hand, in $\Omega/(S \cup T)$:

I′. $V(\Omega/(S \cup T)) = \{X_1, X_2, \ldots, X_m\}$.
II′. $E(\Omega/(S \cup T)) = (E \setminus S) \setminus T$.

We've proved (1) (with virtually no effort) since $E(\Omega/S/T) = E(\Omega/(S \cup T)$ by II and II′; but still we need to show that the endpoints of an edge $e \notin S \cup T$ correspond in the two contractions. Let's compare them. First,

$$V_{\Omega/(S\cup T)}(e) = \{X_h \in \boldsymbol{\pi}_{\mathrm{b}}(S \cup T) \mid X_h \ni v \text{ for } v \in V_\Omega(e)\}.$$

(This last is the multiset—not just a set—of all X_h which contain a vertex $v \in V_\Omega(e)$. If e is a loop in $\Omega/(S \cup T)$, and if its endpoints in Ω are v which is in X_h and w which is in $X_{h'}$, then $V_{\Omega/(S\cup T)}(e) = \{X_h, X_{h'}\}$. Each $V(e)$ in our analysis of endpoints of e must be treated as a multiset, since e may be a loop in any or all of the graphs.) Next,

$$V_{\Omega/S}(e) = \{V_i \in \boldsymbol{\pi}_{\mathrm{b}}(S) \mid V_i \ni v \text{ for } v \in V_\Omega(e)\},$$

so that

$$V_{(\Omega/S)/T)}(e) = \{W_j \in \boldsymbol{\pi}_{\mathrm{b}}(\Omega/S|T) \mid V_i \in W_j \text{ for } V_i \in \boldsymbol{\pi}_{\mathrm{b}}(S) \text{ such that } V_i \ni v \text{ for } v \in V(e)\}$$
$$= \{W_j \in \boldsymbol{\pi}_{\mathrm{b}}(\Omega/S|T) \mid \bar{W}_j \ni v \text{ for } v \in V(e)\}.$$

(This last is the multiset of all W_j such that there exists $v \in V_i \in W_j$ for $v \in V_\Omega(e)$.) This brings us to an important question. If W_2 is balanced, then what is $\bar{W}_2$? Is $\bar{W}_2$ a vertex in $\Omega/(S \cup T)$? To start with, is $(S \cup T){:}W_2$ balanced?

For an example see Figure E.8. The edges of S are red and the edges of T are blue. We distinguish balanced components by drawing them as circular or oval sets and unbalanced components as rectangular sets. The component $\bar{W}_2$ is drawn with a dashed line to indicate that we are unsure whether it is balanced or not.

[PROOF TO BE CONTINUED.] □

Now we have a crucial lemma about the effect of contraction on balance. We say "vertex components" instead of "components" because it's important that loose edges are not involved in this lemma.

2010 Mar 2: DRAFT Simon Joyce

Lemma E.11. [[LABEL L:20100302 balanced components of contractions are contractions of balanced components]] *Given a biased graph Ω and $S \subseteq E$. The balanced vertex components of Ω/S are the contractions of the balanced vertex components of Ω.*

Proof. We know $V(\Omega/S) = \pi_b(\Omega|S)$. Obviously we can treat each component of Ω separately so we may assume Ω is connected.

First assume Ω is balanced. Then S is balanced and Ω/S is balanced by definition. (In fact $\Omega/S = \langle\|\Omega\|/S\rangle$.)

Now assume Ω is unbalanced but S is balanced. Then $\boldsymbol{\pi}_{\mathrm{b}}(S) = \boldsymbol{\pi}(S)$, so $\|\Omega/S\| = \|\Omega\|/S$. A circle $C = e_1e_2\cdots e_l$ in $\|\Omega/S\|$ is a contraction of a circle $D = e_1P_1e_2P_2\cdots e_lP_l$ in $\|\Omega\|$, where each P_i is a path in $S{:}V_i$. **[This may follow from a previous lemma]** By definition, C is balanced if D is balanced for some choice of the P_i paths. We want to know if this is true for any choice of paths. Let $D' = e_1P_1'e_2P_2'\cdots e_lP_l'$ be another circle in $\|\Omega\|$ with paths $P_i' \in S{:}V_i$. We want to show that D is balanced if and only if D' is balanced.

First we suppose that only $P_1 \neq P_1'$. Then we'll use induction to get the following result.

Sublemma E.12. [[LABEL L:20100302 balanced in contraction means balanced in graph]] *Given a biased graph Ω and $S \subseteq E(\Omega)$, then if a circle C in Ω/S is the contraction of a circle D in Ω, then C is balanced if and only if D is balanced.*

Proof. We give two proofs for this sublemma, by two different methods. The first proof uses Tutte's Path Theorem.

Theorem E.13 (Tutte's Path Theorem [TLect])**.** [[LABEL T:20100302tpt]] *Given an inseparable graph Γ, a linear class $\mathcal{L}$ of circles, and two circles D and D', there is a path of circles, $D = D_0, D_1, \ldots, D_r = D'$ such that $D_i \notin \mathcal{L}$ except possibly when i is 0 or 1.*

A *path of circles* is a sequence of circles, $D_0, D_1, \ldots, D_r$, such that each $D_{i-1} \cup D_i$ is a theta graph.

We apply Tutte's Path theorem to the linear class

$$\mathcal{L}_e := \{\text{circles in } \Gamma \text{ that contain the edge } e\}.$$

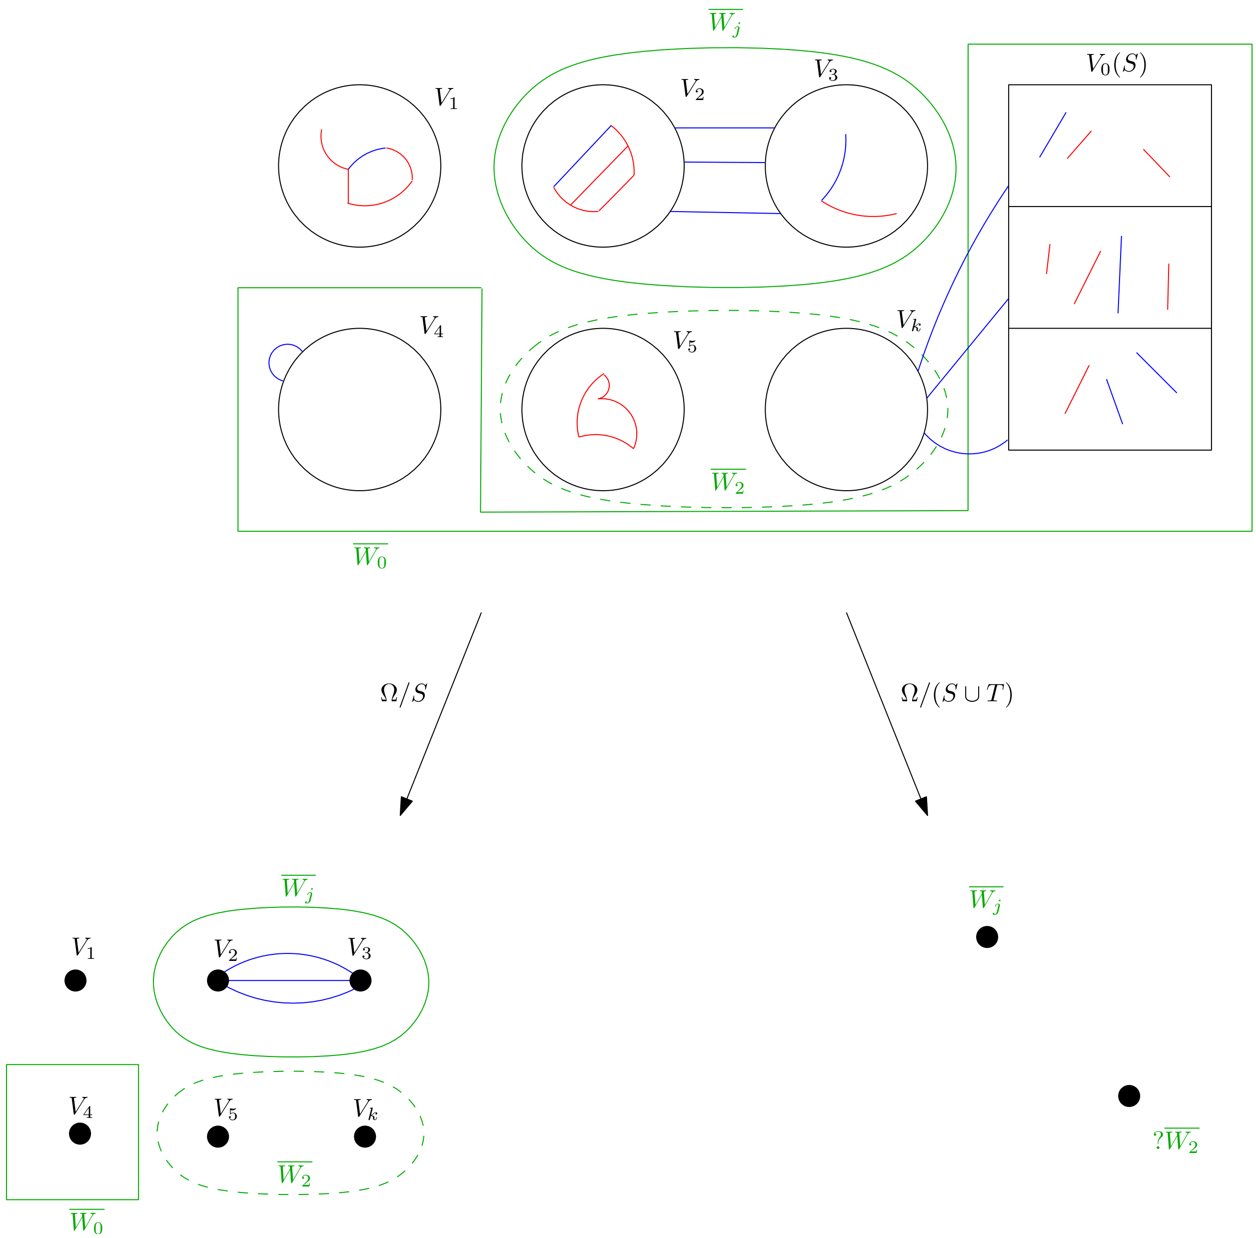

FIGURE E.8. Ω together with its contractions Ω/S and $\Omega/(S \cup T)$. [[LABEL F:20100223: omegacontractgraph]]

Corollary E.14. [[LABEL C:20100302 edge in path of circles]] *For e an edge of an inseparable graph Γ, any two circles D and D' that contain e are connected by a path of circles $D = D_0, D_1, \ldots, D_r - D'$, where each D_i contains e.*

Proof. **[NEEDS PROOF.]** □

More generally we can apply Tutte's path theorem to the linear class

$$\mathcal{L}_Q := \{\text{circles in } \Gamma \text{ containing the path } Q\},$$

where all the internal vertices of Q are divalent. This situation is that of Corollary E.14 with e subdivided into a path Q.

Corollary E.15. [[LABEL 20100302 path in path of circles]] *In an inseparable graph Γ, containing a path Q whose internal vertices are divalent, any two circles D and D' that contain Q are connected by a sequence $D = D_0, D_1, \ldots, D_r = D'$, where each $D_{i-1} \cup D_i$ is a theta graph and each D_i contains Q.* □

Now let $Q = e_2P_2 \cdots e_lP_le_1$, so that $D = PQ$ and $D' = P'Q$. Let Γ be the block of $Q \cup (S{:}V_1)$ that contains P_1 and P_1'. So we apply Corollary E.15 to D and D' to get a sequence $D = D_0, D_1, \ldots, D_r = D'$ with each D_i containing Q. Now $D_i \oplus D_{i+1}$ is a circle contained in $S{:}V_i$ and are therefore balanced since V_i is a balanced vertex component. Therefore by the theta property D_0 and D_1 have the same character, and D_1 and D_2 have the same character and so on. It follows that D_0 is balanced if and only if D_r is balanced. The result now follows by induction.

The second method of proof is by "hacking and hewing".[14] Let P_2 be the path we get by following P_1' past the first vertex where it first splits from P_1 to the vertex where it first intersects P_1. Then follow P_1 to its end. As before, $P_1Q \oplus P_2Q$ is a balanced circle; since $P_1Q \cup P_2Q$ is a theta graph, P_1Q and P_2Q have the same character. Now let P_3 be the path we get by following P_1' past the first vertex where it splits from P_2 to the vertex where it first intersects P_2. Then follow P_2 to its end. Again, P_2Q and P_3Q have the same character. We continue in this manner until eventually we get $P_r = P_1'$ for some r. We conclude that P_1Q and $P_1'Q$ have the same character and again the result follows by induction.

[**Does more need to be said about the induction step? ANSWER: I'm not sure what you're proving by induction here, i.e., what's "the result"? If you're trying to prove P_1Q and $P_1'Q$ have the same character, it's done. If not, then more should be said.**]

□

It follows that if S is balanced and Ω/S is unbalanced, then Ω is unbalanced. So if Ω is balanced then Ω/S is balanced.

Now we want to prove that if Ω is unbalanced, then Ω/S is unbalanced. We have shown that the definition of $\mathcal{B}(\Omega)$ is independent of the choice of how you obtain each circle in Ω/S from a circle in Ω.

So, we assume that S is unbalanced. Let $X := \bigcup\{V_i \mid V_i \text{ is not connected to } V_0 \text{ in } \Omega\}$. $\Omega{:}X$ is a union of components of Ω in which the corresponding part of S is balanced so these components are covered by the cases where we assumed S was balanced.

[**The proof is continued the next day.**] □

In order to finish proving our lemma from yesterday, we need to prove that if $S \subseteq E(\Omega)$ is balanced, Ω is connected, and Ω is unbalanced, then Ω/S is unbalanced. 2010 Mar 4: Jackie Kaminski and Amanda Ruiz

Suppose C is an unbalanced circle in Ω. Then C contracts to a closed walk W in $\|\Omega\|/S = \|\Omega/S\|$ (see Figures E.9 and E.10). Let $W = e_1e_2 \ldots e_l$, or written with vertices, $W = V_0e_1V_1e_2V_2 \ldots V_{l-1}e_lV_l$ where $V_l = V_0$ and the V_i are not necessarily distinct. Then in Ω, $C = e_1P_1e_2P_2 \ldots P_{l-1}e_lP_l$ where P_i is a path in $S{:}V_i$ as in Figure E.9. In particular, P_i is one of the components of $C{:}V_i$.

[14]Peter Cook, "Sitting on the Bench", in the American recording of *Beyond the Fringe*.

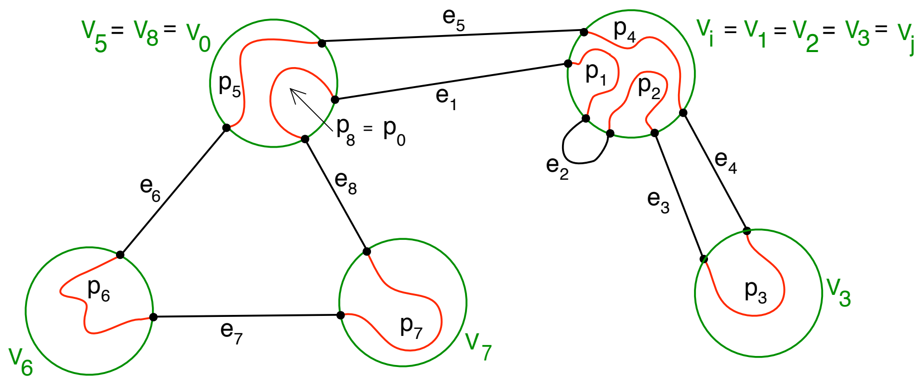

FIGURE E.9. C is an unbalanced circle in Ω. The V_i's are balanced components of $\Omega|S$. Since we think of S as a spanning subgraph, each e_i in C outside of S is an edge.
[[LABEL F: 20100304 circle]] **[The p should be capitalized. The v also (it's hard to tell).]**

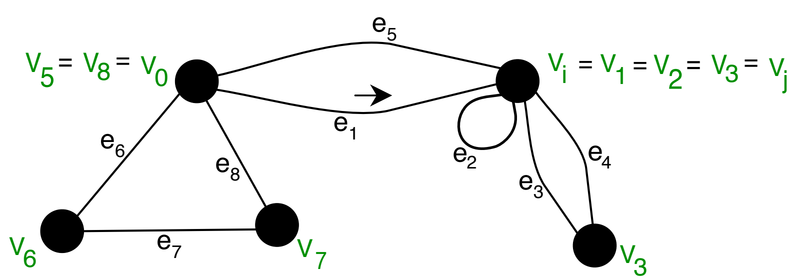

FIGURE E.10. C contracts to a closed walk in $\|\Omega/S\|$.
[[LABEL F: 20100304 walk]]

If W is a circle, then it is unbalanced [**(by some sublemma from before)**], since if $C_{\Omega/S} = C_\Omega/S$, then $C_{\Omega/S}$ is balanced if and only if C_Ω is balanced. In other words, it does not matter which choices we make in constructing C_Ω from $C_{\Omega/S}$.

In the case where W is not a circle, suppose we have a counterexample W', a walk in Ω whose contraction is balanced in Ω/S and which has the minimum number of edges outside of S over all such walks. Since W' is not a circle, it must repeat a vertex: there are i, j such that $V_i = V_j$. $S_i = S{:}V_i$, the component of $\Omega|S$ corresponding to $V_i = V_j$, is connected. Therefore, there exists a path Q from P_i to P_j, internally disjoint from $P_i \cup P_j$. [**I'M NOT FOLLOWING THIS EXACTLY. What are the hypothesis and conclusion?**]

There may be several such paths Q in various balanced components S_k of S. Choose Q to be a shortest such path. Then Q is internally disjoint from W, so $C \cup Q$ is a theta graph. Let C_1 and C_2 be the two circles in $C \cup Q$ that contain Q. So, in Ω we have constructed C_1 and C_2 such that $C_1 \cup C_2$ is a theta graph and $C_1 \oplus C_2 = C$. One of C_1 and C_2 must be unbalanced, and this circle contracts to a shorter walk in Ω/S than our "minimal" W. [**MORE NEEDED to finish up the argument.**]

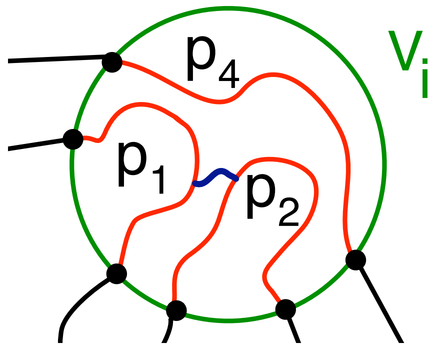

Figure E.11. This pic seems unnecessary
[[LABEL F: 20100304 Zoom S_i]]

In proving that Ω/S is unbalanced if $S \subseteq E(\Omega)$ is balanced and Ω is connected and unbalanced, we found ourselves working mostly with walks. We could generalize Lemma **??** to walks, with a similar proof.

Lemma E.16. [[LABEL L: 20100304 hypolemma]] *Suppose C is an unbalanced closed walk in Ω. Then C contracts to a closed walk W in $\|\Omega\|/S = \|\Omega/S\|$ which is unbalanced.*

Recall that a walk W is called balanced when its edge set is a balanced edge set.

Proof Outline. Although this follows as a corollary from Lemma E.16 it is natural to prove it without using that lemma by assuming a counterexample that has the fewest edges in S^c. We leave this proof as a healthy exercise for the reader. □

Our next step toward the proof of Lemma E.8 is the following lemma.

Lemma E.17. [[LABEL L: 20100304 biased contraction]] *Let Ω be a biased graph and let S, T be disjoint subsets of $E(\Omega)$. Then $(\Omega/S)/T = \Omega/(S \cup T)$*

The equals sign here means equal except for the vertex labels, which have a natural correspondence. Readers familiar with Tutte's approach to graph theory (see Tutte [GT]) will notice that this kind of very natural isomorphism is what he calls a vertex isomorphism, $=_V$.

Proof. Our first step is to prove that the vertex sets of $(\Omega/S)/T$ and $\Omega/(S \cup T)$ have a natural one-to-one correspondence. Recall that $V(\Omega|(S \cup T)) = \boldsymbol{\pi}_{\mathrm{b}}(\Omega|(S \cup T))$. Call the parts of this partition $\{X_1, X_2, \dots, X_m\}$ and let X_0 denote the vertex set of the unbalanced components of $\Omega|(S \cup T)$. For another bit of helpful notation, let $R_j = (S \cup T){:}X_j$. In order to compare these to the vertices of $(\Omega/S)/T$, we need to introduce similar notation, which must, of course, begin with Ω/S.

Recall that $V(\Omega/S) = \boldsymbol{\pi}_{\mathrm{b}}(\Omega|S)$; let $\{V_1, \dots, V_k\}$ denote this partial partition and let V_0 denote the union of the vertex sets of the unbalanced components of $\Omega|S$. Similarly, $V((\Omega/S)/T) = \boldsymbol{\pi}_{\mathrm{b}}((\Omega/S)|T)$; let $\{W_1, \dots, W_l\}$ denote this partition, and let W_0 denote the union of the vertex sets of the unbalanced components of $(\Omega/S)|T$. (See Figure E.7. This is the same argument as for the gain-graph version of this lemma.) Every component A of $(\Omega/S)|T$ has a vertex set which is some collection of V_i's, and $E(A) \subseteq T$ (in Ω/S). Now, if A is balanced **[if ?]** then $V(A)$ is a W_j for some j.

For ease of notation let $U = V(A)$, which is a set of vertices in Ω/S, and let $\bar{U} = \bigcup U = \bigcup_{V_i \in U} V_i$, a set of vertices in Ω.

Lemma E.18. [[LABEL L:20100304 U]] *$\bar{U}$ is connected in Ω; in fact it is connected by $S \cup T$ edges.*

Proof. Since $\bar{U} = \bigcup U = \bigcup_{V_i \in U} V_i$, for $x, y \in \bar{U}$, if $x, y \in V_i$ for some i, then there is a path in $S{:}V_i$ from x to y, since $S{:}V_i$ is connected. Therefore x and y are connected by a path of S-edges.

Now suppose $x, y \in \bar{U}$ with $x \in V_a$ and $y \in V_b$ where $a \neq b$. Since $x, y \in \bar{U}$, V_a and V_b are connected in $(\Omega/S)|T$. Fix a minimal **[shortest?]** V_aV_b-path P in T. Let e_a, e_b be the first and last edges of P, so $e_a, e_b \in T$. Since P is a minimal V_aV_b-path in $(\Omega/S)|T$ and $V_a \neq V_b$, then e_a and e_b are links in $(\Omega/S)|T$, and consequently are links in Ω. Furthermore, exactly one endpoint of e_a is in V_a (in Ω), call it v_a; similarly, let v_b denote the unique endpoint of e_b in V_b. Finally, there is an S-path from x to v_a and from v_b to y, in Ω. Furthermore, there is a T-path in Ω/S from v_a to v_b. If this T-path has length 1, we have the desired $(S \cup T)$-path in Ω from x to y. If the T-path is longer then its internal vertices (in Ω/S) are sets V_k, each of which is connected by S (in Ω). Therefore, for $x, y \in \bar{U}$ with $x \in V_a$, $y \in V_b$, and $a \neq b$, x and y are connected by edges of $S \cup T$. □

Lemma E.18 demonstrates that each $\bar{U}$ is contained in a single X_i for some i. So any balanced component W_j of $(\Omega/S)|T$ is contained in some balanced component X_i of $\Omega/(S \cup T)$.

We would like to show the other containment, and the same result for unbalanced components. Define $\overline{W}_j := \bigcup_{V_k \in W_j} V_k$ for each $W_j \in \boldsymbol{\pi}_{\mathrm{b}}((\Omega/S)|T)$

Lemma E.19. [[LABEL L:20100304 X]] *Each $X_i \in \boldsymbol{\pi}_{\mathrm{b}}(\Omega|(S \cup T))$ has the property that $X_i \subseteq \overline{W}_j$ for some j.*

Proof. Let R_i denote $(S \cup T){:}X_i$. R_i is balanced because $X_i \in \boldsymbol{\pi}_{\mathrm{b}}(\Omega|(S \cup T))$. Since $R_i \subseteq S \cup T$, $R_i/(S \cap R_i$ has $T{:}X_i'$ as its balanced components, where by X_i' we mean the set of vertices in Ω/S corresponding to $V_i \subseteq V(\Omega)$. $T{:}X_i'$ is connected and balanced.

The fact that $T \in X_i'$ is connected implies that it is contained in some W_j. If $j \neq 0$ we are done. It remains to to show that the connected component of $(\Omega/S)|T$ containing $T{:}X_i'$ cannot be W_0.

Suppose $T \in X_i' \subseteq W_0$, because $T \in X_i'$ is balanced, it must be connected to some unbalanced circle in $O = W_0$, counting half edges as unbalanced circles. But the only edges in O are T edges see Figure F: 20100304 O. If there is an unbalanced circle in $O \subseteq (\Omega/S)|T$, it must have come from an unbalanced circle in Ω. But, $(T \cup S){:}X_i = R_i$ is balanced by assumption, which is a contradiction.

This concludes the proof of Lemma E.19 □

The final step to proving that the vertices of $(\Omega/S)/T$ have a canonical correspondence with those of $\Omega/(S \cup T)$ is to show that the vertices in X_0 correspond to $\overline{W}_0$.

We are left to show that W_0‘ $='$ X_0 (another class).

[MORE TO COME!] □

March 9, 2010: Nathan Reff UNEDITED

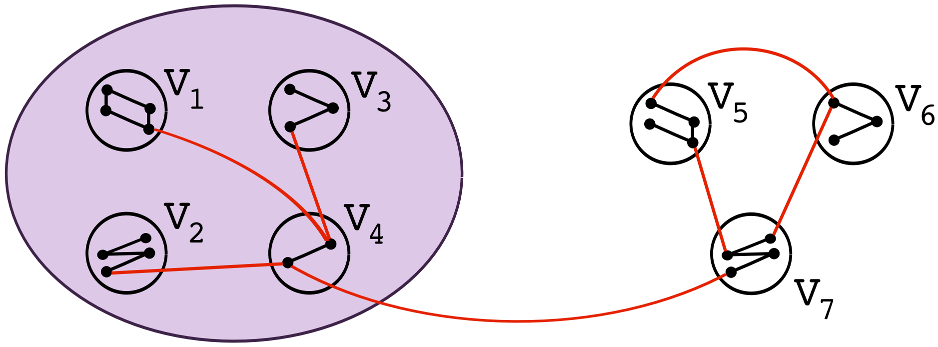

Figure E.12. The component O of $(\Omega/S)|T$. S edges are black, T edges are red, and $T{:}X_i'$ is indicated in purple.
[[LABEL F: 20100304 O]]

E.4. **Minors.** We know proper balanced components of $\Omega|(T \cup S)$ correspond to those of $\Omega/S/T$.

Corollary E.20. [[LABEL C:030910 bpp]] *Specifically: If the balanced patial partition* $\boldsymbol{\pi}_{\mathrm{b}}(\Omega|(S \cup T)) = \{X_1, \ldots, X_m\}$, $\boldsymbol{\pi}_{\mathrm{b}}(\Omega|S) = \{V_1, \ldots, V_k\}$ *and* $\boldsymbol{\pi}_{\mathrm{b}}(\Omega/S|T) = \{W_1, \ldots, W_l\}$ *then* $l = m$ *and each* $X_j = \bigcup\{V_i \mid V_i \in W_j\}$.

Explanation:

(1) A balanced component $(S \cup T){:}X_j$ of $\Omega|(S \cup T)$ contracts to a component of $(\Omega/S|T) = \Omega|(S \cup T)/S$ which is balanced. This component is $T{:}W_j = [(S \cup T)/S]{:}W_{j'}$ for some j. So each $X_j \mapsto W_{j'}$ for a unique j'. By suitable labeling $j' = j$. So we have each $X_j \mapsto W_j$.

(2) We want to show that any $T{:}W_j$ in $\Omega|(S \cup T)/S$ is the contraction of a balanced component $(S \cup T){:}X_{j'}$. Then $j' = j$ by (1).

From the proposition $T{:}W_j$ being a balanced component of $(S \cup T)/S$, must be the contraction of a balanced component of $S \cup T$, i.e., a $(S \cup T){:}X_{j'}$.

So we need to show an X_j gives a whole component in the contraction, and each whole component in the contraction comes from some X_j.

Lemma E.21 (Original Goal)**.** $\Omega/S/T$"="$\Omega/(S \cup T)$

Proof. We have a natural correspondence:

1. (Vertices) $V(\Omega/S/T) = \boldsymbol{\pi}_{\mathrm{b}}((\Omega/S)|T) \leftrightarrow V(\Omega/(S \cup T)) = \boldsymbol{\pi}_{\mathrm{b}}(\Omega|(S \cup T))$ so the verticies are the blocks in the balanced partial partitions.
2. (Edges) $E(\Omega/S/T) \leftrightarrow E(\Omega/(S \cup T))$
3. (Incidences) $I(\Omega/S/T) \leftrightarrow I(\Omega/(S \cup T))$

$$
\begin{aligned}
e \in (S \cup T)^c \text{ is incident with } v \in V &\iff e \text{ incident with } X_j \text{ such that } v \in X_j \\
&\iff e \text{ incident with } V_i \in \boldsymbol{\pi}_{\mathrm{b}}(\Omega|S) = V(\Omega/S) \text{ where } v \in V_i \\
&\iff e \text{ incident with } W_{j'} \in \boldsymbol{\pi}_{\mathrm{b}}((\Omega/S)|T) = V(\Omega/S/T) \\
&\qquad \text{such that } V_i \in W_{j'}
\end{aligned}
$$

We know from corollary E.20 if there exists X_j then there exists a W_j (with $j = j'$), and if there exists $W_{j'}$ then there exists a $X_{j'}$ (with $j = j'$). Also notice in the above equivalences there might not exist an X_j in which case we have a loose edge. So the vertices with $\Omega/(S \cup T)$ and $\Omega/S/T$ are naturally corresponding vertices.

In short, (1)–(3) give us:

$$\|\Omega/S/T\| = \|\Omega/(S \cup T)\|$$

(4) (Balance) $\mathcal{B}(\Omega/(S \cup T)) = \mathcal{B}((\Omega/S)/T)$

By (1)-(3) we have $\mathcal{C}(\Omega/(S \cup T)) = \mathcal{C}((\Omega/S)/T)$. Suppose $C \in \mathcal{B}(\Omega/S/T)$ then $C = C'/T$ where there exists $C' \in \mathcal{B}(\Omega/S)$ and there exists $C'' \in \mathcal{B}(\Omega)$ such that $C' = C''/S$. Then $C = C''/(S \cup T)$ (by "=" of $\mathcal{C}$'s since C'' is balanced). Therefore $C \in \mathcal{B}(\Omega/(S \cup T))$. Suppose now that $C \in \mathcal{B}(\Omega/(S \cup T))$. Then $C = C''/(S \cup T)$ where there exists $C'' \in \mathcal{B}(\Omega)$. This means $C''/S \in \mathcal{B}(\Omega/S)$ and $(C''/S)/T \in \mathcal{B}(\Omega/S/T)$. Therefore $\mathcal{B}(\Omega/S/T) = \mathcal{B}(\Omega/(S \cup T))$.

□

Theorem E.22. [[LABEL T:2010309]]

(1) Every result of a sequence of deletions and contractions of edge sets has the form $(\Omega \backslash S)/T$ for some disjoint $S, T \subseteq E$.

(2) Every minor of a minor is a minor.

Proof. (1) You can delete edges first, then you will have a sequence of contractions $(\Omega \setminus S)/T_1/T_2/\ldots/T_k = (\Omega \setminus S)/(T_1 \cup \ldots \cup T_k)$.

(2) If a vertex is deleted in a minor, say Ω_1/T_1 where $\Omega_1 \subseteq \Omega$, then we can delete the edges at v first (leaving it as an isolated vertex), then contract T_1 and delete v. Note that v is not affected by contracting T_1 since it is isolated. Suppose we have $\Omega_1 \subseteq \Omega$, Ω_1/T_1, $\Omega_2 \subseteq \Omega_1/T_1$, Ω_2/T_2, etc. Then we define

$$\Omega_1 := (\Omega \setminus S_1) \setminus Z_1$$

where $Z_1 \subseteq V = V(\Omega)$ and $S_1 \subseteq E$ is the set of all edges incident to Z_1's vertices.

$$\Omega_2 = [(\Omega_1/T_1) \setminus S_2] \setminus Z_2$$

where $Z_2 \subseteq V(\Omega_1/T_1)$ and $S_2 \subseteq E$.

Etc.

Note that all S_i's and T_j's are pairwise disjoint. Let $\bar{Z}_2 = \{v \in V \mid v \in V_{1i} \in \boldsymbol{\pi}_{\mathrm{b}}(\Omega_1/T_1), V_{1i} \in Z_2\}$. Notice in the definition of Ω_2 we have $Z_2 \subseteq V(\Omega_1/T_1) = \boldsymbol{\pi}_{\mathrm{b}}([(\Omega \setminus S_1) \setminus Z_1]|T_1) = [\boldsymbol{\pi}_{\mathrm{b}}([\Omega \setminus S_1])|T_1] \setminus Z_1 \subseteq \boldsymbol{\pi}_{\mathrm{b}}([\Omega \setminus S_1]|T_1)$ since every vertex in Z_1 forms a singleton block in $\boldsymbol{\pi}_{\mathrm{b}}(\Omega \setminus S_1)$ (the vertices of Z_1 were isolated). Notice also $\bar{Z}_2 \subseteq V(\Omega) \setminus Z_1$ and Ω_2/T_2.

Vertex deletion:
$\Omega_1' = \Omega \setminus S_1$ (so $\Omega_1 = \Omega_1' \setminus Z_1$)
$\Omega_2' = (\Omega \setminus S)/T_1 \setminus S$ (so $\Omega_2 = \Omega_2' \setminus Z_1 \setminus Z_2$)

$$\begin{aligned}
\Omega_2/T_2 &= (\Omega_2' \setminus Z_1 \setminus Z_2)/T_2 \text{ (where } Z_1 \text{ and } Z_2 \text{ are isolated vertices in } \Omega_2')\\
&= (\Omega_2'/T_2) \setminus Z_1 \setminus Z_2\\
&= ((((\Omega \setminus S_1)/T_1) \setminus S_2)/T_2) \setminus Z_1 \setminus Z_2\\
&= ((\Omega \setminus (S_1 \cup S_2))/(T_1 \cup T_2)) \setminus (Z_1 \cup Z_2) \text{ (where } Z_1 \text{ and } Z_2 \text{ are isolated vertices in } \Omega\ (S_1 \cup S_2))
\end{aligned}$$

So $\Omega_2/T_2 = ((\Omega \setminus (S_1 \cup S_2))/(T_1 \cup T_2)) \setminus (Z_1 \cup Z_2)$.

Let $S = \{$edges in $S_1 \cup S_2$ and edges in T_1 that are in $T_1{:}V_i$ where $V_i \in Z_2\}$.

Let $T = (T_1 \cup T_2) \setminus S$, then $(\Omega_2/T_2) \setminus Z_2 = ((\Omega \setminus S) \setminus (Z_1 \cup \bar{Z}_2))/T$. Take original vertices that became the Z_2 's and contract.

□

2010 Mar 11 DRAFT: Simon Joyce

F. Balancing Sets in a Biased Graph

Biased graphs will now receive a similar balancing set treatment as was given to signed graphs in Chapter II.

Let Ω be a biased with edge set E and $S \subseteq E$. If $\Omega\backslash S$ is balanced, then we call S a *total balancing set*. If $b(\Omega\backslash S) > b(\Omega)$, call S a *partial balancing set*. We call S a *strict balancing set* if S is a partial balancing set and $\boldsymbol{\pi}(\Omega\backslash S) = \boldsymbol{\pi}(\Omega)$, i.e. Ω and $\Omega\backslash S$ have the same number of components.

Now for some observations. "Frustration index", or "deletion index" $l(\Omega) = |S|$ where S is a minimal total balancing set. There's a definition by "change index" which gives the same number as deletion index. [**This is a potential prop or homework**] This is a generalization of D.9. A minimal partial balancing set is analogous to a bond in a graph or signed graph. It's also a copoint complement in the lattice of closed sets. The empty set $\varnothing$ is a total balancing set if and only if Ω is balanced. But $\varnothing$ is never a partial balancing set. A minimal partial balancing set of Ω is the same as a minimal total balancing set of a component of Ω. A minimal total balancing set of Ω is the same as the union of minimal total balancing sets in each component of Ω. A superset of a partial balancing set is a partial balancing set. A superset of a total balancing set is a total balancing set.

F.1. **A total or partial balancing edge in a biased graph.** An edge e in a biased graph Ω is called or total (partial) balancing edge if the set $\{e\}$ is a total (partial) balancing set. I would define a *balancing edge* as a total balancing edge such that Ω is unbalanced, i.e. $\{e\}$ is a minimal total balancing set.

Problem F.1.

(1) Is a partial balancing edge a total balancing edge of a component?
(2) Is a strict balancing edge a total balancing edge of a component?

We call a vertex v of Ω a *balancing vertex* if Ω is unbalanced and $\Omega\backslash v$ is balanced.

G. Closure and Closed Sets

[[LABEL 4.closure]]

Closure in a biased graph Ω is a function from $\mathcal{P}(E)$ to $\mathcal{P}(E)$. In Chapter II I used one definition of closure in a graph and proved it was equivalent to a second definition. Closure for a biased graph is very similar, but more complicated. Just for the heck of it, now I'll define closure by the second definition and prove it's equivalent to the first.

We denote by E_0 the set of loose edges in a biased graph Ω.

G.1. **Closure operators.** [[LABEL 4.closops]]

Definition G.1. [[LABEL D:20100311 biased clos and bcl]] Given S a subset of the edges of a biased graph Ω, we define $\operatorname{clos}(S) = S \cup \{e \notin S \mid e$ is contained in a frame circuit $C \subseteq S \cup e\}$. Also define $\operatorname{bcl}(S) = S \cup \{e \notin S \mid e$ is contained in a balanced circle $C \subseteq S \cup e\} \cup E_0$.

This definition is based on closure in the frame matroid, where the circuits are the frame circuits. **[To be added in Chapter V.]**

[The E_0 may no longer be necessary in a lot of the following stuff.]

Theorem G.1. [[LABEL T:20100311 clos is clos]]

(1) clos *is an abstract closure operator.*

(2) $\operatorname{clos}(S) = [E{:}V_0(S)] \cup \left[\bigcup_{B \in \boldsymbol{\pi}_{\mathrm{b}}(S)} \operatorname{bcl}(S{:}B) \right] \cup E_0.$

Proof. For gain graphs this can be proved using switching. This is left as an exercise for the reader. It is similar to the proof for biased graphs using more hard work.

Lemma G.2. [[LABEL L:20100311 L1]] *If S is balanced, then* $\operatorname{bcl}(S) = \operatorname{clos}(S)$

Proof. Proof is left as an exercise for the reader. □

Notice that if a biased graph is balanced, then biased closure is the same as regular graph closure in the underlying graph.

Lemma G.3. [[LABEL L:20100311 L2]]

(1) $S \subseteq \operatorname{bcl}(S)$

(2) If $S \subseteq T$, then $\operatorname{bcl}(S) \subseteq \operatorname{bcl}(T)$.

(3) There exists S such that $\operatorname{bcl}(S) \subset \operatorname{bcl}(\operatorname{bcl}(S))$.

Proof. Proof of the first two parts are left as exercises for the reader.

[third part needs an example] □

Lemma G.4. [[LABEL L:20100311 L3]] *If S is balanced, then* $\operatorname{bcl}(S)$ *is balanced.*

Proof. The proof for gain graphs uses switching.

For biased graphs, using induction we add one edge at a time since $\operatorname{bcl}(S) \subseteq \operatorname{bcl}(S \cup e)$ if $e \in \operatorname{bcl}(S) \backslash S$. We have to show $S \cup e$ is balanced. We know there is a balanced circle $C \subseteq S \cup e$ containing e. If $C' \subseteq S \cup e$ is another circle containing e we can show that C' is balanced using Tutte's Path Theorem or directly in a similar way as in the proof of Sublemma E.12. □

Lemma G.5. [[LABEL L:20100311 L4]] *If S is balanced, then* $\operatorname{bcl}(S) = \operatorname{bcl}(\operatorname{bcl}(S))$.

Proof. In a gain graph, use switching.

In a biased graph, if there is $f \in \operatorname{bcl}(\operatorname{bcl}(S)) \backslash \operatorname{bcl}(S)$ we look at a circle $C \subseteq (\operatorname{bcl}(S) \cup f)$ that contains f. There is at least one circle D containing f that is balanced. D can be factored as a product of edges and paths so that $D = fP_0e_1P_1e_2P_2 \cdots e_kP_k$ where each $e_i \in \operatorname{bcl}(S) \backslash S$ and each $P_i \subseteq S$. If $k = 0$, then $f \in \operatorname{bcl}(S)$. If $k > 0$, use induction on k. The details are left as a homework exercise. **[This needs to be checked. Something was not right in the notes.]** □

On a side note, the method of proof of Lemma G.5 works in the case of infinite graphs. An interesting research question arrises if you want infinite circles. What are infinite circles and what happens to closure in this case? [**Reference Bruce Richter et al for graphs**]

Lemma G.6. [[LABEL L:20100311 L5]]

$$\operatorname{clos} S \subseteq [E{:}V_0(S)] \cup \left[\bigcup_{B \in \boldsymbol{\pi}_{\mathrm{b}}(S)} \operatorname{bcl}(S{:}B) \right] \cup E_0.$$

Proof. The proof is left as an exercise. Just check the frame circuits. □

Lemma G.7. [[LABEL L:20100311 L6]]

$$\operatorname{clos} S \supseteq [E{:}V_0(S)] \cup \left[\bigcup_{B \in \boldsymbol{\pi}_{\mathrm{b}}(S)} \operatorname{bcl}(S{:}B) \right] \cup E_0.$$

Proof. Let $e \in \Big([E{:}V_0(S)] \cup \Big[\bigcup_{B \in \boldsymbol{\pi}_{\mathrm{b}}(S)} \operatorname{bcl}(S{:}B)\Big] \cup E_0\Big) \setminus \operatorname{clos}(S)$. Clearly $e \notin E_0$ since then $\{e\}$ would be a frame circuit. So first suppose $e \in \Big[\bigcup_{B \in \boldsymbol{\pi}_{\mathrm{b}}(S)} \operatorname{bcl}(S{:}B)\Big] \setminus S$. Then there is $B \in \boldsymbol{\pi}_{\mathrm{b}}(S)$ such that $e \in [\operatorname{bcl}(S{:}B)] \setminus S$. So there is a balanced circle $C \subseteq (S : B) \cup e$ containing e. But then C is a frame circuit containing e, so $e \in \operatorname{clos}(S)$.

Now suppose $e \in [E{:}V_0(S)]$. Let $S_1', S_2', \ldots, S_m'$ be the unbalanced components of S. Let the endpoints of e be v and w. First suppose the endpoints of e are in different components, say $v \in S_i$ and $w \in S_j$. S_i and S_j both contain an unbalanced circle or half edge. So in S_i there is a shortest path from v to an unbalanced circle or half edge. Similarly there is a shortest path in S_j from w to an unbalanced circle or half edge. This union of e, the unbalanced circles or half edges and the paths gives us a loose handcuff. Therefore e is contained in a frame circuit.

Now suppose v and w are both contained in the same unbalanced component S_i. First suppose S_i contains a half edge or unbalanced loop f with endpoint z. There is a shortest path P_1 from v to z in S_i and a shortest path P_2 from w to z. There will be a first vertex x where P_1 and P_2 intersect. This forms a circle C containing e. If C is balanced then e is contained in a frame circuit. If C is unbalanced then $C \cup P_1 \cup f$ is a contrabalanced handcuff containing e.

Now suppose S_i contains no half edges or unbalanced loops. Then S_i must contain an unbalanced circle D containing at least two vertices. Let P_1 be a path from v to D and P_2 a path from w to C. If P_1 and P_2 intersect, then a similar argument to that above shows e is contained in either a balanced circle or a contrabalanced handcuff so suppose P_1 and P_2 do not intersect. Then $e \cup P_1 \cup P_2 \cup D$ forms a theta graph. Let C_1 and C_2 be the two circles in this theta graph that contain e. If either is balanced then e is contained in a frame circuit. If not then both are unbalanced and then e is contained in a contrabalanced theta graph and we are done. □

Lemma G.6 and Lemma G.7 prove part (2) of Theorem G.1. To finish the proof we must show that clos is an abstract closure operator.

Lemma G.8. [[LABEL L:20100311 L7]]

$$\boldsymbol{\pi}_{\mathrm{b}}[\operatorname{clos}(S)] = \boldsymbol{\pi}_{\mathrm{b}}(S)$$

.

Proof. Obviously, $E{:}V_0$ has no balanced components. If $B \in \boldsymbol{\pi}_{\mathrm{b}}(S)$, then $\mathrm{bcl}(S{:}B)$ is balanced by Lemma G.4 and a component of $\mathrm{clos}(S)$ by Theorem G.1 part (2). □

Now we can finish the proof of G.1 part (1). $S \subset \mathrm{clos}(S)$ by definition. Clearly $V_0(\mathrm{clos}(S)) = V_0(S)$. For $B \in \boldsymbol{\pi}_{\mathrm{b}}(S)$, $\mathrm{clos}(S){:}B = \mathrm{bcl}(S{:}B)$ by part (2) of Theorem G.1. Using these facts and Lemmas G.5 and G.8, it follows that

$$\begin{aligned} \mathrm{clos}(\mathrm{clos}(S)) &= [E{:}V_0(\mathrm{clos}(S))] \cup \left[\bigcup_{B \in \boldsymbol{\pi}_{\mathrm{b}}(\mathrm{clos}(S))} \mathrm{bcl}(\mathrm{clos}(S){:}B) \right] \cup E_0 \\ &= [E{:}V_0(S)] \cup \left[\bigcup_{B \in \boldsymbol{\pi}_{\mathrm{b}}(S)} \mathrm{bcl}(\mathrm{bcl}(S{:}B)) \right] \cup E_0 \\ &= \mathrm{clos}(S). \end{aligned}$$

Finally if $S \subseteq T$, then $\mathrm{clos}(S) \subseteq \mathrm{clos}(T)$ by definition and we're done. □

There is a similar definition of closure using the lift circuits, but "frame closure" is the more natural definition of closure for several reasons, including geometry, chromatic polynomials, Tutte polynomials, etc., all of which will be developed in later sections.

G.2. **Closed sets.**

G.3. **Group-labelled partial partitions.**

G.4. **The two matroids. Rank.**

G.5. **Examples.**

H. Incidence Matrices of a Gain Graph

H.1. **Multiplicative gains.**

H.1.1. *Canonical and standard form.*

H.1.2. *Vector representation from multiplicative gains.*

H.2. **Additive gains.**

H.2.1. *Incidence matrix.*

H.2.2. *Vector representation from additive gains.*

I. Hyperplane Representations

I.1. **Linear representation from multiplicative gains.**

I.2. **Affine representation from additive gains.**

J. Coloring Gain Graphs

J.1. **Proper vs. improper. Set of improper edges.**

J.2. **Chromatic polynomials.**

K. Chromatic Functions of Gain and Biased Graphs

K.1. **Chromatic polynomials. Deletion-contraction and multiplicativity. Subset expansion.**

K.2. **Dichromatic polynomials. Deletion-contraction and multiplicative identities.**

K.3. **Tutte polynomial.**

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`

[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

[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).