Graphs, Gain Graphs, and Geometry

a.k.a.
Signed Graphs and their Friends

Course notes for

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

Math 581: Signed Graphs
Fall, 2014

Binghamton University (SUNY)

Version of November 3, 2014

Lectures by

Thomas Zaslavsky

Notes by

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

Redaction by

Thomas Zaslavsky

Contents

Chapter 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

Let's have a fast overview of graphs, signed graphs, and their equations and hyperplanes. Aug 25: Simon Joyce

So, what is a graph, exactly? There are all kinds of definition of a graph. We'll begin with a few, 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. 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. Also, it has a problem that I'll explain shortly.

The following is approximately what we will use (but simplified in a couple of ways, mainly by omitting half and loose edges. The full definition is in Section A.1).

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

The vertices incident to an edge e are its *endpoints*. The crucial difference between Definition A.3 and the previous two is that we treat edges as objects in themselves, not as pairs of endpoints, with or without multiplicity. This is valuable for several reasons. A big one is that different edges with the same endpoints may have different properties—such as different signs. Other reasons will appear in due course.

The most basic graph structure is connectivity. Briefly, 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, as long as we don't jump.

Now, our second level of mathematical structure.

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. Some signed graphs. Solid edges are positive, dashed ones are negative. [[LABEL F:0825sg]]

Signed graphs and their generalizations are what we are headed for in these notes, though only after setting the terms and properties of (unsigned) graphs.

And here is a third level of mathematics. Each edge of a graph implies an equation between two variables. The variables correspond to the n 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, $\mathcal{H}[\Gamma]$, is called the hyperplane arrangement generated by Γ. It divides up $\mathbb{R}^n$ into regions that have a remarkable combinatorial meaning to be revealed later. For a signed graph, a positive edge $+v_iv_j$ has hyperplane $x_i = x_j$ (so the edges of an unsigned graph behave like positive edges) and a negative edge $-v_iv_j$ has hyperplane $x_i = -x_j$. The set of hyperplanes corresponding to the edges of Σ is $\mathcal{H}[\Sigma]$. 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 to 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. **Definitions for a graph.** [[LABEL 1.defsgraph]]

We give a formal definition in terms of incidence between vertices and edges. It 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.

The essentials are that an edge is a thing in itself, not just a pair of vertices, and that an incidence between a vertex and an edge is not a relation, but is also itself an object. An incidence in our sense is probably better thought of as an edge end—and watch out; not every edge need have two ends!

Definition A.1. [[LABEL D:0827graph]] An *incidence multirelation* between sets V and E is a set I together with a mapping $\varepsilon : I \to V \times E$. If $\varepsilon(i) = (v, e)$ we say that v and e are *incident*, with incidence i, or that v and e participate in the incidence i.

A *graph* $\Gamma = (V, E, I)$ is an ordered triple consisting of disjoint sets V and E and an incidence multirelation I between them, such that each edge participates in at most two incidences. Usually, we don't mention I and ε explicitly; we simply write $\Gamma = (V, E)$. We often avoid confusion by writing $V(\Gamma)$, $E(\Gamma)$, and (if necessary) I_Γ and ε_Γ for V, E, et al. instead of $\Gamma = (V, E)$.

The elements of V are called the *vertices* of the graph Γ. The elements of E are called the *edges* of the graph Γ. The *elements* of a graph are its edges and vertices. We often call the incidences *edge ends* or the *ends* of the edge; think of each end as attached to a vertex. We usually write (v, e) for an incidence i between v and e; when there are two such incidences we may distinguish them by arbitrary subscripts, as in $(v, e)_1$ and $(v, e)_2$.

An example: In Figure A.1 edge q is incident to vertex v_4 twice, so there are two incidences called (v_4, q). This is consistent with our definition since we do not need edges to be incident to distinct vertices.

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

Definition A.2. These are the types of edge:

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 $e{:}v$, or e_v.

A *loose edge* 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

$$E^* := E^*(\Gamma) := \{ \text{ links and loops } \}.$$

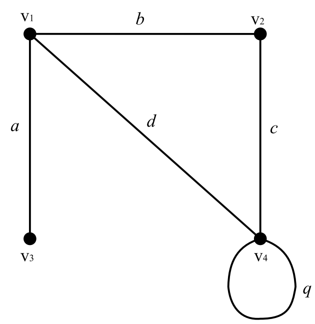

FIGURE A.1. A graph with a loop but no multiple edges. It is not simple, because of the loop.

[[LABEL F:0827graph]]

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 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).
- A *link graph* is a graph whose edges are links. A simple graph is a link graph, but not vice versa, obviously.
- An *ordinary graph* is a graph with no half edges or loose edges; that is, all edges are ordinary. A link graph is an ordinary graph.

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 confess that, in graph theory, a link graph is called a *multigraph* when it isn't simply called a "graph". Both names are also applied to ordinary graphs where loops are permitted. I won't confuse the reader by listing any other of the names applied to different kinds of graph, and to avoid all confusion I'm adopting the specific terms just defined.

A.3. **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 $\{x_1\}, \{v_1, v_2\}, \{w_1, w_2\}$.

A.4. **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.5. **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.3. [[LABEL D:0829degree]] The *degree* or *valency* of a vertex, denoted by $d(v)$, is the number of edge ends incident with v.

A vertex of valency 1, 2, or 3 is called respectively *monovalent, divalent*, or *trivalent.*

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

Definition A.5. [[LABEL D:0829regular]] A *k-regular* graph is a graph where every vertex has degree k.

A.6. **Types of subgraph.** [[LABEL 1.subgraphtypes]]

There are, of course, subobjects in graph theory; not only subgraphs in general, but also several special kinds. Here we assume $X \subseteq V$ and $S \subseteq E$. Aug 27a: Nate Reff

- A *subgraph* of Γ is Γ' such that $V' \subseteq V$, $E' \subseteq E$, has the same incidence multirelation 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 $\Gamma \setminus X$, 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.
- The *reduction* of Γ by a vertex set, denoted by $\Gamma \operatorname{red} X$, technically is not a subgraph. It is the graph $(V \setminus X, E)$ where for an edge $e \in E$, $V_{\Gamma \operatorname{red} X}(e) := V_\Gamma(e) \setminus X$. Reduction differs from deletion in that edges that lose endpoints when Γ is reduced do not disappear; they become loose or half edges.
  The reduction of Γ by a subgraph $\Delta = (X, S)$ is $(\Gamma \operatorname{red} X) \setminus S$.
- 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 $(X, S{:}X)$.

A.7. **Special vertex sets: stability and cliques.** [[LABEL 1.specialvsets]]

Aug 29c: Jackie Kaminski

Two types of vertex subset 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 $E{:}W = \varnothing$. The *stability number* or *independence number* $\alpha(\Gamma)$ is the size of a largest stable set. Note that a stable set may be maximal yet not maximum, i.e., it may have size less than $\alpha(\Gamma)$. In Figure C [**NEED GOOD FIGURE**], $\{x_1\}, \{v_1, v_2\}, \{w_1, w_2\}$ are five stable sets, all maximal. The stability number is 2.

A *clique* is a vertex set whose members are pairwise adjacent. The *clique number* $\omega(\Gamma)$ is the size of a largest clique.

These two types of set are complementary. For a simple graph Γ, a stable set of Γ is a clique of its complement Γ^c and vice versa.

A.8. **Equality and isomorphism.** [[LABEL 1.isomorphism]]

Aug 27 2014 Sep 27:

When are two graphs “the same”? There are two kinds of answer.

Equality. Strictly according to our definitions, Γ_1 and Γ_2 should be equal when they have the same vertex and edge set and the same set of incidences with the same vertex-edge pair associated to the same incidence in each graph. That is not what we really want: we want the incidences to be secondary to vertices and edges. Therefore, we define equality in a special way.

Definition A.6. [[LABEL D:equality]] Graphs Γ_1 and Γ_2 are *equal* if $V_1 = V_2$, $E_1 = E_2$, and there is a bijection $I_1 \leftrightarrow I_2$ such that $i_1 \leftrightarrow i_2 \implies \varepsilon_1(i_1) = \varepsilon_2(i_2)$.

This defines equality of labelled graphs, where we know the individual elements of the graph.

For simple graphs we can (as is usually done) regard edges as unordered pairs of vertices; then two graphs are equal if they have the same vertex and edge sets.

Isomorphism and automorphism.

Definition A.7. [[LABEL D:isom]] Graphs Γ_1 and Γ_2 are *isomorphic* if there are bijections $\alpha_V : V_1 \to V_2$ and $\alpha_E : E_1 \to E_2$ and a bijection $\alpha_I : I_1 \to I_2$ such that $\varepsilon_1(i_1) = (v_1, e_1) \implies \varepsilon_2(\alpha_I(i_1)) = (\alpha_V(v_1), \alpha_E(e_1))$. The pair (α_V, α_E) is called an *isomorphism* of Γ_1 with Γ_2. (When it is necessary to mention the incidence mappings ε_i, we say the isomorphism is the triple $(\alpha_V, \alpha_E, \alpha_I)$.)

An isomorphism of Γ with itself is an *automorphism* of Γ. The automorphisms, with functional composition, form the *automorphism group*, $\operatorname{Aut}\Gamma$.

The *isomorphism type* of a graph may be defined as the class of all graphs isomorphic to it. A simpler way to look at it is as an *unlabelled graph*, where we don't know the individual elements of the graph, but only their relationships to each other.

In the usual treatment of simple graphs, regarding edges as unordered pairs of vertices, two graphs are isomorphic if there is a bijection α of their vertex sets that preserves adjacency and nonadjacency, or preserves adjacency in both directions.

A.9. **Contraction of an edge.** [[LABEL 1.edgecontraction]]

Intuitively, contraction means shrinking an edge to a point. The two endpoints of a link therefore become one vertex; the two endpoints of a loop, being already identical, are not affected. 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.) Aug 27a: Nate Reff

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.

B. Basic Structures

[[LABEL 1.importantkinds]]

Aug 29a: Jackie Kaminski

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. To describe different ways of moving around a graph we use the following terms:

- A *walk* is a sequence $v_0e_1v_1\cdots e_lv_l$ 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; it is l in the preceding definition. A walk of length zero is just a vertex.
- A *closed walk* is a walk where $v_0 = v_l$ and $l \geq 1$. A walk is *open* if it is not closed.
- A *trail* is a walk with no repeated edges, but it may repeat vertices.
- 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. Thus, it must have positive length. (Oddly, a closed path is not a path.)
- A *circle* is the graph or edge set of a closed path, that is, it has no repeated vertices or edges except that the initial and the final vertex are the same. A circle differs from a closed path in that it is simply a graph or a set, while a (closed) path is ordered, with a beginning and an end.

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

Connection of vertices. [[LABEL 1.connectvert]]

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.

- A *connected component* (briefly, a *component*, or most precisely, a *vertex 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 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 $c(\Gamma) = 1$; that is, the relation of connection on V has exactly one equivalence class. and there are no loose edges.

  This means that a graph consisting of only a single loose edge is not "connected". That may seem strange but it is the most useful way to define connectedness. For loose edges we have a slightly different definition, which will appear very soon.
- 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 of edges. [[LABEL 1.connectedge]]

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 walk* is a sequence $x_0 x_1 \cdots x_k$ where the x_i's are alternately vertices and edges. x_0 may be a vertex or an edge, and the same for x_k. 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 and $x_0 x_1 = ev$ or $x_{k-1} x_k = ve$. Lastly, if x_i is a loose edge the walk is simply x_i.

  Similarly, there are generalized trails and paths.

  Be aware that a generalized walk is not necessarily a walk, 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 walk containing both.

  Since a generalized walk can consist of a single edge and no vertices, the relation of being connected is reflexive on $V \cup E$.

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

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

The proof is similar to that of Theorem B.1 so I omit it. There is also a generalization of Proposition B.2 which the reader can state and prove easily enough.

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 (generalized) relation of connection.

A *component* (or *vertex component*, or *connected component*) of Γ is the subgraph induced by an equivalence class of the connectedness relation on V. That is, a component is a topological component that has a vertex. Thus, a loose edge is not a component; all other topological components are components.

An *edge component* is the subgraph induced by an equivalence class of the connectedness relation on E. That is, an edge component is a topological component that has an edge. Thus, an isolated vertex is not an edge component; in fact, it is the only kind of topological component that is not an edge component.

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.

We say 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. We defined edge components specifically to prepare for line graphs (Section I).

Bridges, cutpoints, and blocks. [[LABEL 1.bridges]]

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 (and is not a loose edge). That is, the connecting walks (or generalized walks) may have elements of Δ as initial or final elements but no intermediate elements of those walks may be in Δ. (This implies that if an element of Δ is in a connecting walk of a bridge, it must be a vertex that is initial or final in the walk.)

An equivalent definition is that a bridge of Δ is the subgraph of Γ induced by a component of $\Gamma \operatorname{red} \Delta$, provided that component is not a loose edge.

If we wanted to allow loose edges as bridges we would define *topological bridges*, but we don't have a use for that concept.

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 (and does not contain any loose edges). A *block* (or *block graph*) is a graph that has only one block (and no loose edges).

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.

Ours is not the only existing definition. 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 the graph has no loops and half edges; but our definition is better because it has the important property given in Theorem B.4.

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

Aug 29b: Jackie Kaminski

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.

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

For example, a loop is a circle, as is Figure B in C.1 [**GOOD FIGURE NEEDED**]. We require the subgraph to have a vertex in order to exclude loose edges as circles.

As you can see, 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 an initial and final point. The graph of a closed path has neither.

Although there is real ambiguity in our use of the term 'circle', as sometimes we mean the edge set, sometimes the graph, the context should always make the meaning clear.

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 link 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). See Figure C in C.1 [**GOOD FIGURE NEEDED**].

  Equivalently, we can define a forest as a graph whose components are all trees. Or we can define a forest first and define a tree as a connected forest.

  An empty graph (no vertices or edges) is a forest, but not a tree since a connected graph must have exactly one connected component.
- A *spanning forest* is a spanning subgraph of Γ which is a forest. Any forest will do; for instance, the subgraph $(V, \varnothing)$ is a spanning forest in Γ.

  Similarly, a *spanning tree* is a spanning subgraph of Γ which is a tree. A graph has a spanning tree if and only if it is connected
- A *maximal forest* is a forest which is not properly contained in any other forest. A maximal forest is therefore spanning. Every graph has a maximal forest, which is connected (a spanning tree, in fact) if and only if the graph is connected. See Figure C in C.1 [**GOOD FIGURE NEEDED**].

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:

(a) Γ *contains a spanning tree* $\iff$ *it is connected.*
(b) *A maximal forest consists of a spanning tree of each component of* Γ*.*
(c) *Every tree except K_1 has a monovalent vertex. [*[LABEL T:0829forest mono]]

The proof is left to the reader—or, see a 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.2), 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]]

There are other tree-like graphs. Here I list some of them:

A *1-tree* is a tree with one extra edge (not a loose edge). In an ordinary graph, it is a connected graph with cyclomatic number 1. See Figure D in C.1 [**FIGURE NEEDED**].

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

[[LABEL 1.dcminors]]

Aug 29b: Jackie Kaminski

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.6. 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]]

We are now restricting ourselves to ordinary graphs again. [**In the following that is mostly true but the writing is confused.**]

- We already defined how to contract a link, loop, half edge, and loose edge.
  Refer to Figure D in C.1 [**GOOD FIGURE NEEDED**] 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 in C.1 [**GOOD FIGURE NEEDED**] for an example.
- For a graph Γ, let $\pi(S) :=$ the partition of V such that each block is the vertex set of a 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 in C.1 [**GOOD FIGURE NEEDED**].
- 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".

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

(1) [Kuratowski (mainly) and Wagner] Γ *is planar if and only if* Γ *does not contain either* K_5 *or* $K_{3,3}$ *as minors.*
(2) [Archdeacon, Glover, and Huneke] Γ *is projective planar if and only if* Γ *does not contain as a minor any of a list of* 35 *graphs.*

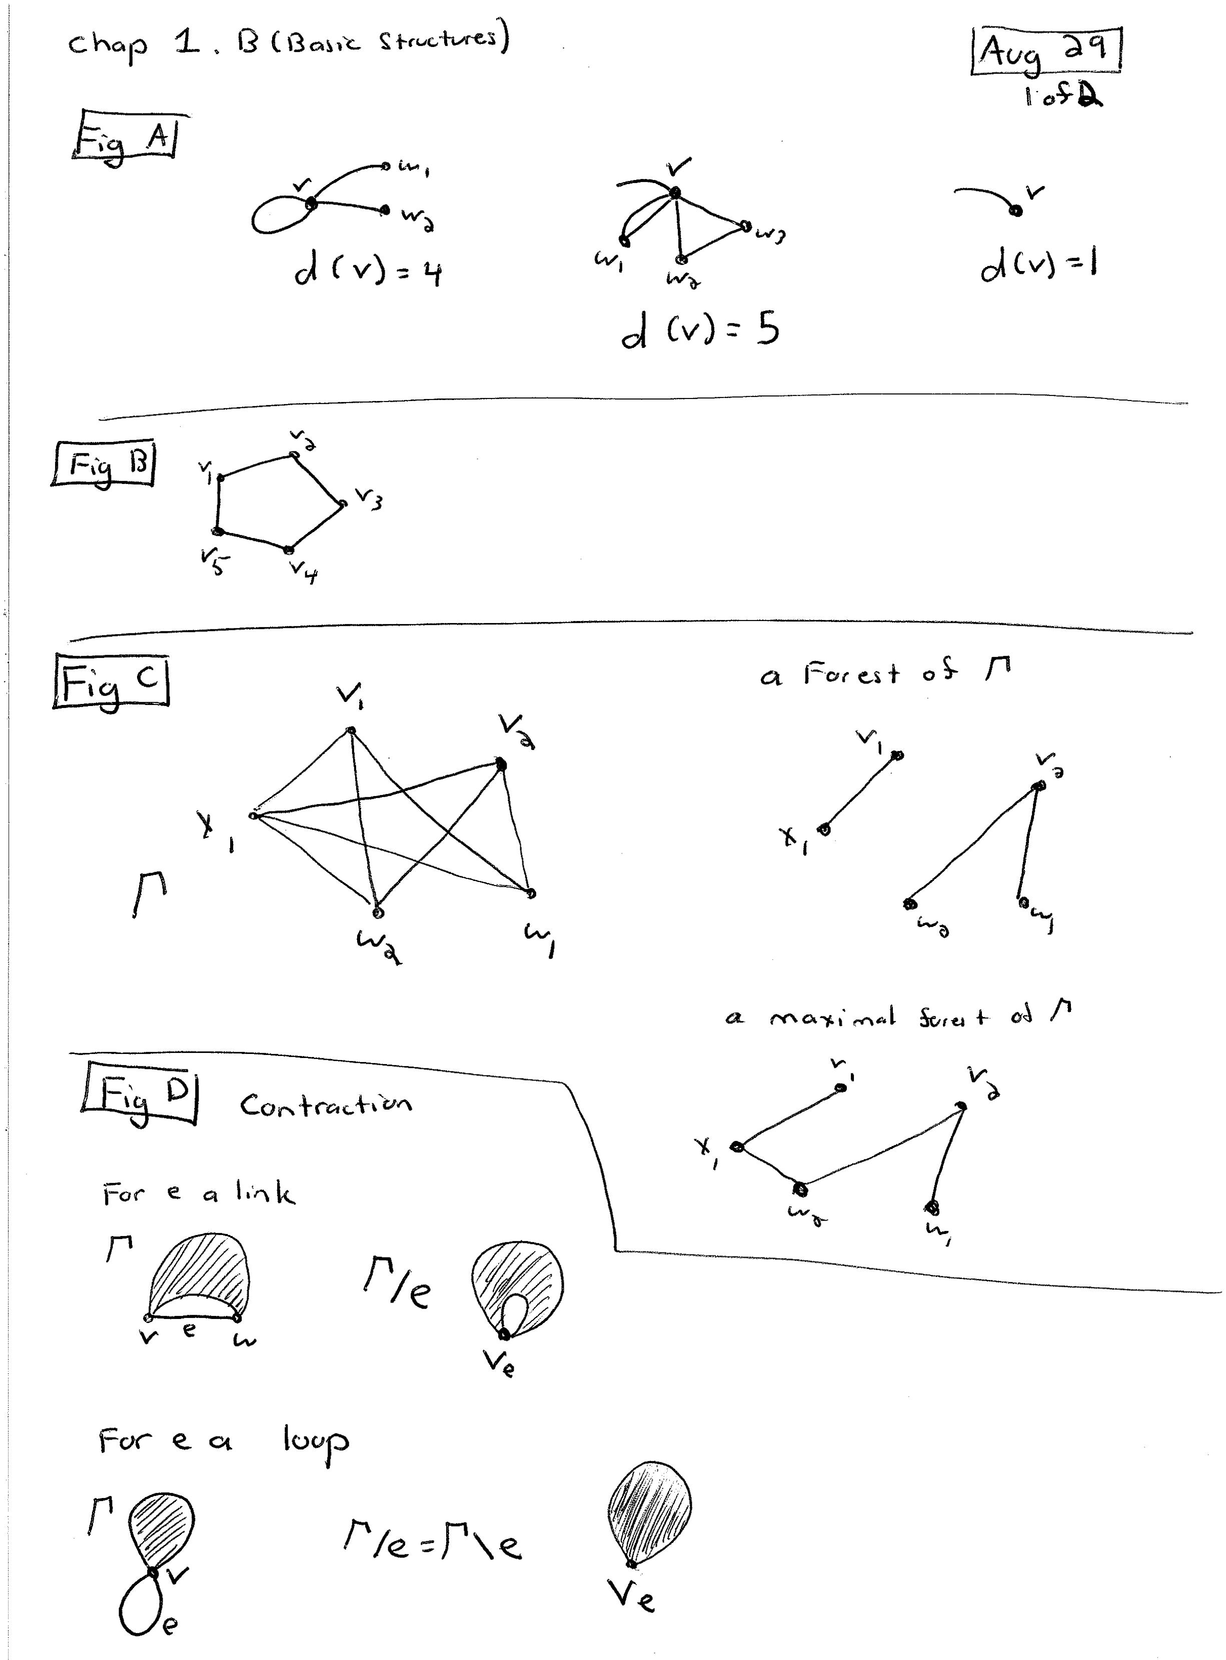

FIGURE C.1. Handmade Kaminski figure.

(3) [Robertson and Seymour; Bodendieck and Wagner; Glover and Huneke for orientable surfaces] **[VERIFY]** Γ *embeds in a surface* S *if and only if* Γ *does not contain as a minor any of a list of graphs, which depends on* S *but is always finite.*

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. Sept 3a: Yash Lodha

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 (to avoid confusion with the blocks of a graph) the *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, $|\pi|$. The only partition of the empty set, $V = \varnothing$, is $\pi = \{\}$. The only partition of the singleton set, $V = \{u\}$, is $\pi = \{V\}$. In all other cases $|\pi| \geq 2$. The partition $\hat{1} = \hat{1}_V := \{V\}$ with one block is called the *trivial partition*; the partition $\hat{0} = \hat{0}_V := \big\{\{x\} : x \in V\big\}$, in which every block is a singleton, 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

$$\Pi_V := \{\text{all partitions of } V\}, \qquad \text{and in particular,} \qquad \Pi_n := \Pi_{[n]}.$$

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

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 $\boldsymbol{\pi}(S) := \boldsymbol{\pi}_\Gamma(S) := \boldsymbol{\pi}(V,S) :=$ 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).

For instance, $\Pi(K_V) = \Pi_V$, but for an incomplete graph, $\Pi(\Gamma) \subset \Pi_V$. (Exercise: Prove both statements!)

Caution: Once we introduce half edges we need a more flexible and complicated extension of the concept of a connected partition. But for now we keep things relatively simple.

Connected partitions are intimately related to closed sets, the next topic.

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

Let's remind ourselves of the definition of an abstract closure operator.

Definition D.1. [[LABEL D:0903closure]] An *abstract closure operator* on a set E is a function $\mathcal{P}(E) \to \mathcal{P}(E)$, which we write $S \mapsto \overline{S}$, 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 whose meet and join operations are as follows: For closed subsets S, T of E, $S \wedge T = S \cap T$ and $S \vee T = \overline{S \cup T}$.*

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

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

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

Technically, 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.

[THIS DUPLICATES later stuff.]

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. $\square$

Theorem D.3. *[*[LABEL T:0903closedptns]] *The poset of closed edge sets of* Γ*, 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 several lemmas about $\boldsymbol{\pi}$.

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

We usually apply this definition 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.

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 $\operatorname{clos} S$.

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

Proof. Let $\boldsymbol{\pi}(S) = \{B_1, \dots, B_k\}$. 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. □

Lemma 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. □

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

Proposition 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* $\pi \in \Pi_V$*.* □

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

Whenever $S \subseteq S' \subseteq E$, each of the parts of $\boldsymbol{\pi}(S)$ is contained in a part of $\boldsymbol{\pi}(S')$, which is to say that $\boldsymbol{\pi}(S)$ is a refinement of $\boldsymbol{\pi}(S')$. Readers familiar with partitions of a set V will think of the last statement as; in symbolic terms, $\boldsymbol{\pi}(S) \leq \boldsymbol{\pi}(S')$ in the refinement ordering. It is well known that the set Π_V of all partitions of V with the refinement ordering forms a lattice. I leave it to the reader to check this (if necessary) and to verify that the set of connected partitions of a graph also forms a lattice, in which the join and meet operations have the properties in (d)–(f) of the following list.

Exercise D.1. [[LABEL X:0908ptnlatts]] [**MODIFY to be consistent with Theorems D.3 (earlier) and D.8 (later).**]

Verify these statements about partition lattices. V is a set and $\Gamma = (V, E)$ is an ordinary graph (both finite, as usual).

(a) Π_V is a poset.
(b) The partitions of V correspond bijectively to the equivalence relations on V. (I will write $\sim_\pi$ for the equivalence relation that corresponds to $\pi \in \Pi_V$.)
(c) Π_V is a lattice with lattice operations that satisfy

$$\pi \wedge \tau = \{B \cap C : B \in \pi,\ C \in \tau,\ B \cap C \neq \varnothing\},$$
$$\pi \vee \tau = \{D : D \text{ is an equivalence class of the transitive closure of } \sim_\pi \cup \sim_\tau\},$$
$$\pi \wedge \tau = \bigvee\{\rho \in \Pi_V : \rho \leq \pi, \tau\},$$
$$\pi \vee \tau = \bigwedge\{\rho \in \Pi_V : \rho \geq \pi, \tau\},$$

for $\pi, \tau \in \Pi_V$.
(d) $\Pi(\Gamma)$ is a poset.
(e) $\Pi(\Gamma)$ is a lattice with lattice operations that satisfy

$$\boldsymbol{\pi}(S) \wedge_\Gamma \boldsymbol{\pi}(T) = \boldsymbol{\pi}(S \cap T),$$
$$\boldsymbol{\pi}(S) \vee_\Gamma \boldsymbol{\pi}(T) = \boldsymbol{\pi}(S \cup T),$$
$$\pi \vee_\Gamma \tau = \pi \vee_V \tau,$$
$$\pi \wedge_\Gamma \tau = \bigvee_\Gamma\{\rho \in \Pi(\Gamma) : \rho \leq \pi, \tau\},$$
$$\pi \vee_\Gamma \tau = \bigwedge_\Gamma\{\rho \in \Pi(\Gamma) : \rho \geq \pi, \tau\},$$

for $\pi, \tau \in \Pi(\Gamma)$ and $S, T \subseteq E$. (I append the subscripts V and Γ to distinguish operations in Π_V and $\Pi(\Gamma)$. Later I will not use the subscripts.)
(f) $\boldsymbol{\pi}(S) \wedge_\Gamma \boldsymbol{\pi}(T)$ may not equal $\boldsymbol{\pi}(S) \wedge_V \boldsymbol{\pi}(T)$. Find a counterexample and also find necessary and sufficient conditions on the edge sets for equality to occur.

When τ, τ' are two partitions of V such that $\tau \leq \tau'$, then $E{:}\tau \subseteq E{:}\tau'$. This observation and the following definition lead 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]] *The set* $\operatorname{Lat}\Gamma$ *is a lattice; moreover, it is a join subsemilattice of* Π_V. *Furthermore,* $\operatorname{Lat}\Gamma \cong \Pi(\Gamma)$*; specifically,* $\boldsymbol{\pi} : \operatorname{Lat}\Gamma \to \Pi(\Gamma)$ *is an order isomorphism.*

[SAME AS THEOREM D.3?]

Lemma D.9. *[*[LABEL L:0908biject]] $\boldsymbol{\pi}$ *is a bijection from* $\operatorname{Lat}\Gamma$ *to* $\Pi(\Gamma)$.

Proof. We know from Lemma D.4 that $\boldsymbol{\pi}$ maps into $\Pi(\Gamma)$.

To see that $\boldsymbol{\pi}$ is injective, let S and S' be closed subsets of E and assume $\boldsymbol{\pi}(S) = \boldsymbol{\pi}(S')$. By Theorem D.2, $S = E{:}\boldsymbol{\pi}(S) = E{:}\boldsymbol{\pi}(S') = S'$. On the other hand, $\boldsymbol{\pi}$ is surjective because, for a connected partition τ of V, $E{:}\tau$ is closed (by Theorem D.2) and $\boldsymbol{\pi}(E{:}\tau) = \tau$ (by ??). □

Proof. Lemma D.9 shows that $\boldsymbol{\pi}$ is a bijection; we have to prove it respects the partial order. We already noted that $S \subseteq S' \subseteq E \implies \boldsymbol{\pi}(S) \leq \boldsymbol{\pi}(S')$ and that for $\tau, \tau' \in \Pi(\Gamma)$, $\tau \leq \tau' \implies E{:}\tau \subseteq E{:}\tau'$. It follows that $\boldsymbol{\pi}$ is order preserving.

Since $\operatorname{Lat}\Gamma$ is the lattice of the closure operator clos, $\Pi(\Gamma)$ is a lattice. To show that join in $\Pi(\Gamma)$ is the same as in Π_V, **[PROOF NEEDED]** $\square$

E. Incidence and Adjacency Matrices

[[LABEL 1.matrices]]

Incidence and adjacency matrices let graph theory benefit from the use of matrix theory. Sept 3b: Yash Lodha

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 0's 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.**??**. For our lines of interest, the right incidence matrix is (almost always) 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.4 and G.5. The fact, which the reader will have noticed, that the oriented incidence matrix is uniquely defined only up to negating columns, does not affect the columns' linear dependence.

Proposition E.1. [[LABEL P:0903incidrank]] *The oriented incidence matrix has rank $n - c(\Gamma)$, hence nullity $|E| - n + c(\Gamma) = \xi(\Gamma)$, the cyclomatic number.*

The rank of the unoriented incidence matrix is $n - b(\Gamma)$, where $b(\Gamma)$ is the number of bipartite components of Γ.

We shall prove both statements as special cases of a signed-graph theorem, Theorem **??**(4) in Section II.**??**.

E.2. **Adjacency, degree, and Kirchhoff matrices.** [[LABEL 1.adjmatrix]]

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, with a loop counting twice, once for each end.

[ADD PROPERTIES?: symmetric; 0 diagonal if no loops; A^l counts walks.]

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.2. *[*[LABEL T:0903incidence-adjacency]] *In a link graph, 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 $\mathrm{HH^T} = D - A$ we check the cases $i \neq j$ and $i = j$ separately when multipying the ith row of H with the jth column of $\mathrm{H^T}$. 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 **??** in Section **??**.

The matrix $D - A$ is itself important because it gives information (through its eigenvalues and eigenvectors) about graph structure different from that obtainable from the adjacency matrix. It is often called "the" *Laplacian matrix* of Γ. I prefer to call it the *Kirchhoff matrix* of Γ because it is only one of several kinds of Laplacian matrix of a graph, and because it was (as far as I remember—it was a long time ago) introduced by Kirchhoff for electrical network analysis.

The matrix $D + A$ has recently been discovered to be perhaps even more informative than the Kirchhoff matrix (see CITATION). It has been called the "signless Laplacian" matrix of Γ, but it is best understood as the Kirchhoff or Laplacian matrix of a signed graph; see Section II.**??**.

E.3. **Eigenvalues.** [[LABEL 1.evalues]]

[ADD eigenvalues and $\mathrm{HH^T}$, $BB^{\mathbf{T}}$ including the following.]

[MOVED FROM LINE GRAPHS]

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_{ij} for $i \neq j$, of BB^{T} is the number of edges v_iv_j and x_{ii} is the degree of v_i. So, $BB^{\mathrm{T}} = D + A$ where D is the degree matrix and A is the adjacency matrix. The entry x_{ij} for $i \neq j$ of $\mathrm{HH^T}$ is minus the number of v_iv_j-edges, and the entry x_{ii} of $\mathrm{HH^T}$ is the degree of the vertex v_i.

Theorem E.3. *[*[LABEL T:0926rge]] *If* Γ *is loopless and* k*-regular, then the largest eigenvalue of* A *is* k*, with multiplicity* $c(\Gamma)$*.*

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

Proof. Notice that $\mathrm{HH^T}$ 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 $\mathrm{HH^T}x = kIx - Ax = (k - \lambda)x$. This implies that x is an eigenvector of $\mathrm{HH^T}$ 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 $c(\Gamma)$. □

F. Orientation

[[LABEL 1.orientation]]

Sept 5: Peter Cohen and Yash Lodha

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.

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.

The notation for an oriented edge can be a bit tricky. We could write v_1v_2, or $v_1\vec{v}_2$,but this presents a problem with parallel edges. We will call an oriented edge $e_i{:}v_1\vec{v}_2$ and show it in drawings with an arrow on it, pointing from v_1 to v_2. Accordingly, we call v_1 (or the end of e at v_1) the *tail* and v_2 (or the end of e at v_2) the *head* of the oriented edge. If e is a loop, remember that it has two distinguishable ends, so it has two orientations!

I still have not formally defined an oriented edge. Formal notation is best explained in terms of the incidence matrix; see Section F.2.

[**MAKE SURE THERE'S A FORMAL DEFINITION SOMEWHERE.**]

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 walk (e.g., a path or circle) means giving the walk as a whole a direction. This is a completely separate property of the walk from directions on the edges.

Acyclic orientations.

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. I will call a linear ordering π of V *compatible* with an acyclic orientation $\vec{\Gamma}$ if π gives rise to the orientation $\vec{\Gamma}$.

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

Hence there is a relationship between acyclic orientations of Γ and linear orderings of V. We'll now study this relationship, beginning with a proof of Theorem F.1.

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. □

Proof of Theorem F.1. We perform induction on $|V|$. If $\vec{\Gamma}$ is acyclic, then it must have a source s. Then $\vec{\Gamma} \setminus s$ is acyclic and by our inductive hypothesis it has a compatible total ordering $v_2 < v_3 < \cdots < v_n$. The ordering for $\vec{\Gamma}$ is $s < v_2 < \cdots < v_n$. □

A variation is to let S be the set of all sources, say with k elements, and number them $v_1, \ldots, v_k$ in any order. Then delete S and apply the same method to $\Gamma \setminus S$, numbering its sources $v_{k+1}, \ldots$, and continue until all vertices are labelled. The resulting linear ordering is obviously compatible with the initial acyclic orientation.

We don't necessarily need a total ordering of V to construct an acyclic orientation. A partial ordering may be enough; in fact, usually it is.

Theorem F.3. *[*[LABEL T:0903posetao]] *For each acyclic orientation $\vec{\Gamma}$, 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.*

Proof. Define $v < w$ if there is an edge $e{:}\vec{vw}$ and extend by transitivity and reflexivity to a partial ordering of V. This is the required poset. □

We say this smallest partial ordering is *the partial ordering implied by* the acyclic orientation.

Exercise F.1. [[LABEL Ex:0903posetao]] Prove that the procedure described in the proof does indeed result in a partial ordering.

This discussion leads to a question. Given a graph Γ, can one describe all the posets on V that are implied by acyclic orientations of Γ? I don't know the answer—but see proposition below.

An important example is the complete graph.

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

Corollary F.4. *[[LABEL C:0903aoKn]] 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. □

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.

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 consists of the elements of the poset and we connect elements u and v with an edge if they are comparable in the partial ordering.

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

Comparability graphs are closely connected to acyclic orientations in general. Take an acyclic orientation, extend it to the induced partial ordering P, and compare Γ to both the comparability graph $C(P)$ and the Hasse diagram $H(P)$ of P, both considered as unoriented graphs. Then $H(P) \subseteq \Gamma \subseteq C(P)$. Conversely, suppose P is a partial ordering of V. Is it implied by an acyclic orientation of Γ?

Proposition F.5. *[[LABEL P:0903aocomparability]] A partial ordering P of $V(\Gamma)$ is implied by some acyclic orientation of Γ if and only if $H(P) \subseteq \Gamma \subseteq C(P)$.*

The proof is an exercise.

Totally cyclic orientations.

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.6. *[[LABEL P:0903orexist]] Γ has an acyclic orientation if and only if it has no loops. Γ has a totally cyclic orientation if and only if 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. □

The second part is an exercise for the reader.

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.

G. Equations and Inequalities from Edges

[[LABEL 1.equations]]

Sept 8b: Jackie Kaminski

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.

We can extend any of these correspondences to correspondences between subsets.

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 $\mathcal{A}$, $\mathcal{A}$ has no regions.

Now we define the *intersection lattice* of the arrangement. It is the poset

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

We partially order $\mathcal{L}(\mathcal{A})$ by reverse inclusion, so its top element is $\hat{1} = \bigcap \mathcal{A}$ and its bottom element is $\hat{0} = \mathbb{R}^n = \bigcap \varnothing$.

Proposition G.1. *[[LABEL P:0908interslattice]] The partially ordered set $\mathcal{L}(\mathcal{A})$ is a lattice.*

First Proof. The mapping $X \mapsto \bigcap\{h \in \mathcal{A} : h \supseteq X\}$ for $X \subseteq \mathbb{R}^n$ is a closure operator on $\mathbb{R}^n$. The closed sets are the sets in $\mathcal{L}(\mathcal{A})$. It follows from Proposition D.1 that $\mathcal{L}(\mathcal{A})$ ordered by inclusion is a lattice and therefore that the order dual, which is $\mathcal{L}(\mathcal{A})$ under reverse inclusion, is a lattice. $\square$

Second Proof (with some details elided). Define a mapping $\mathcal{L}(\mathcal{A}) \to \mathcal{P}(\mathcal{A})$ by $t \in \mathcal{L} \mapsto \mathcal{A}(t)$, where

$$\mathcal{A}(t) := \{h \in \mathcal{A} : h \supseteq t\}.$$

It is easy to see that this is an order-preserving injection, that $\bigcap \mathcal{A}(t) = t$ for $t \in \mathcal{L}$, and that if $t = \bigcap \mathcal{S}$ for a subarrangement $\mathcal{S} \subseteq \mathcal{A}$, then $\mathcal{A}(t) \supseteq \mathcal{S}$.

It is easy to check that $\mathcal{S} \mapsto \mathcal{A}(\bigcap \mathcal{S})$ is a closure operator on $\mathcal{A}$, the closed sets are the images of $\mathcal{L}(\mathcal{A})$, and the mapping $t \mapsto \mathcal{A}(t)$ is an order-preserving injection of $\mathcal{L}(\mathcal{A})$ into $\mathcal{P}(\mathcal{A})$. It follows from the first two observations that the image of $\mathcal{L}(\mathcal{A})$ is a lattice and from the third that $\mathcal{L}(\mathcal{A})$ itself is a lattice. □

Later, we'll 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 allow us to switch freely amongst the perspectives of geometry, lattices, and graphs.

[The following is duplicative of later results but has to be carefully edited to avoid bad cross-references.] Finally, we close with two lemmas that we will revisit later.

Lemma G.2 (= Lemma G.4)**.** *[*[LABEL L:0908closspan]] *For* $e \in \mathrm{clos}(S)$, $c_e \in \langle c_f : f \in S \rangle$.

Lemma G.3 (= Lemma G.7)**.** *[*[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.4. *[*[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.

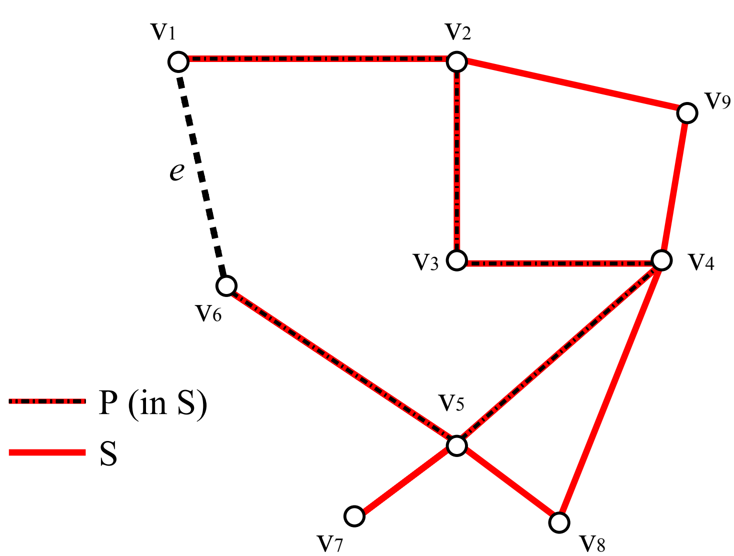

The red lines denote edges of $S \subseteq E$ in a graph $\Gamma = (V, E)$. If $e \in (\mathrm{clos}(S) \backslash S)$ as in the picture, then there exists a path $P \subseteq S$ such that there is a circle. We will show that $\langle c_f : f \in S \rangle$. Because $e \in (\mathrm{clos}(S) \backslash S)$ and thus $e \in \mathrm{clos}(S)$, 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.5. *[[LABEL L:0910lemma1b]] $c_e \in \langle c_f : f \in S \rangle \implies e \in \mathrm{clos}(S)$.*

Proof. Suppose $e \notin \mathrm{clos}(S)$. 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:

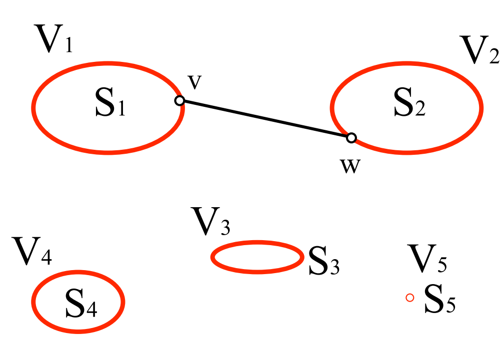

The incidence matrix $\mathrm{H}(\Gamma)$ looks like this, where O is a zero matrix, and $\mathbf{0}$ is a column vector of zeros:

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

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

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

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 $e \in \operatorname{clos}(S)$. □

Lemma G.6. *[[LABEL L:0910lemmaSubLemma]] For a hyperplane $H_e \in \mathcal{H}[\Gamma]$, $\bigcap \mathcal{H}[S] \subseteq H_e \iff e \in \operatorname{clos}(S)$.*

Lemma G.7. *[[LABEL L:0910lemma2]] $\bigcap \mathcal{H}[S] = \bigcap \mathcal{H}[\operatorname{clos}(S)]$.*

Proof. Use Lemma G.6, and dualize Lemmas G.4 and G.5. □

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

Proposition G.8. *S is independent $\iff$ S is a forest.*

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

Theorem G.9. *[[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.4 and G.5. □

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

Theorem G.11. *[[LABEL T:0910thm2]] There are natural isomorphisms $\Pi(\Gamma) \cong \operatorname{Lat}(\Gamma) \cong \{$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. □

[THERE SHOULD BE a corollary that rank in the oriented incidence matrix = rank of $\operatorname{clos} S$ in $\operatorname{Lat}\Gamma$. Cf. Section H.5 on R.]

Corollary G.12. *[[LABEL C:0910geomlattice]] $\operatorname{Lat}\Gamma$ 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 open half-space* of the hyperplane. If we orient $e{:}v_iv_j$ from v_i to v_j, the positive open half-space is the set $\{x \in \mathbb{R}^n : x_i < x_j\}$. The *positive closed half-space* is defined with weak instead of strict inequality. For each orientation, therefore, there is a family of (open and closed) positive half-spaces.

Lemma G.13. *[[LABEL L:0910lemma3]] A cyclic orientation of Γ gives an empty intersection of positive open half-spaces.*

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

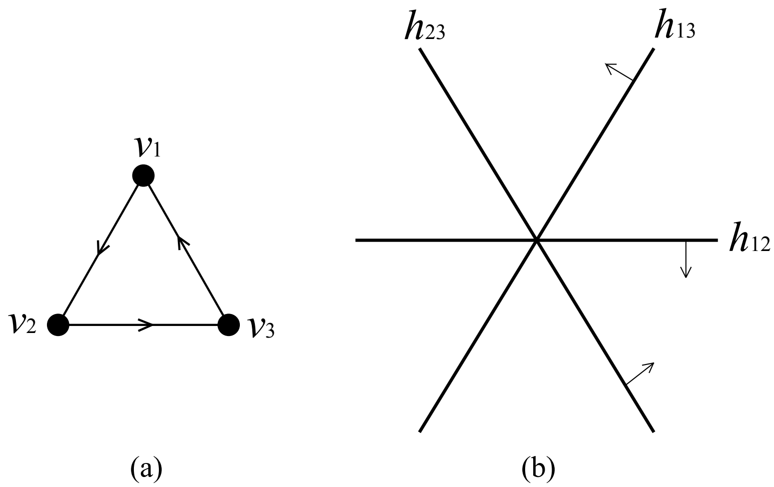

Figure G.1.

An example illustrates the proof. Suppose that the graph $\Gamma = K_3$ is oriented cyclically, as in Figure G.1(a). The corresponding orientation on each hyperplane is shown in (b). By the definition of the positive open half-space, the corresponding intersection of all the positive open 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 open half-spaces in a unique orientation of Γ, which is necessarily acyclic. Let $R(\alpha)$ be the intersection of positive open half-spaces of the acyclic orientation α. Write $\bar{R}(\alpha)$ for the topological closure of $R(\alpha)$; we call it a *closed region*. Then $\bar{R}(\alpha)$ is the intersection of the closed positive half-spaces.

Theorem G.14. *[*[LABEL T:0910thm3]] *The intersection of positive open 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. The correspondence $\alpha \mapsto R(\alpha)$ between acyclic orientations and regions is a bijection.*

Proof. In the cyclic case we just use Lemma G.13. 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 every positive open half-space. Therefore the intersection is nonempty, and in fact a region. $\square$

Exercise G.1. [[LABEL Ex:0910connected]] Fill in gaps in the proof of Theorem G.14 by proving the following properties.

(a) $R(\alpha)$ is connected for each acyclic orientation α of Γ.
(b) If $\mathbf{x} \in R(\alpha)$, then the entire region containing $\mathbf{x}$ is contained in $R(\alpha)$.
(c) Explain why (a) and (b) imply that $R(\alpha)$ is precisely a region of $\mathcal{H}[\Gamma]$ and that the correspondence in the theorem is a bijection.

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. Sept 12a: Simon Joyce & T.Z.

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, $\chi(\Gamma) = 2$. A graph with no edges has chromatic number 1, unless it is the empty graph, whose chromatic number is 0. At the opposite extreme, $\chi(\Gamma) = \infty$ if, and only if, Γ has a loop. We shall have little to say about the last two examples.

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

We now turn 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 nonnegative integer. Two trivial examples: $\chi_\varnothing(\lambda) = 1$ and $\chi_{K_1}(\lambda) = \lambda$. Obviously, the first nonnegative 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. □

A second valuable property is multiplicativity on components, which is usually stated this way:

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. □

The third valuable property is invariance, i.e., dependence only on the isomorphism type of the graph. Sept 12: T.Z.

Lemma H.3. *[*[LABEL L:0912 chrom invar]] *If $\Gamma \cong \Gamma'$, then $\chi_\Gamma(\lambda) = \chi_{\Gamma'}(\lambda)$.*

Proof. To the reader: Supply the proof! □

Theorem H.4. *[*[LABEL T:0912chromaticpoly]] *For an ordinary graph Γ with a loop, $\chi_\Gamma(\lambda) = 0$.*

If Γ has no loops, then $\chi_\Gamma(\lambda)$ is a monic, integral polynomial of degree n of the form

$$\chi_\Gamma(\lambda) = w_0\lambda^n + w_1\lambda^{n-1} + w_2\lambda^{n-2} + \cdots + w_{c(\Gamma)}\lambda^{c(\Gamma)} + \cdots + w_n\lambda^0,$$

where $w_i = 0$ if and only if $i > n - c(\Gamma)$ and for $i < n - c(\Gamma)$, $\operatorname{sgn} w_i = (-1)^i$. If Γ is simple, $w_1 = -|E|$.

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.4 we are entitled to call $\chi_\Gamma(\lambda)$ by its right name.

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

The numbers $w_i = w_i(\Gamma)$ are—if Γ has no loops, so the chromatic polynomial is not identically zero—called the *Whitney numbers of Γ of the first kind.* (Yes, *that* Whitney: Hassler, the famous topologist, was a great graph theorist before he moved up in dimension.)

Now that we know χ_Γ is a polynomial we may substitute any number for λ. The key identities of Lemmas H.1 and H.2 are now polynomial identities, because they are valid for infinitely many values $\lambda \in \mathbb{R}$, i.e., all positive integers. (Any polynomial equation that is valid for infinitely many real numbers is an identity.) Since the Whitney numbers are integers we can regard λ as an indeterminate over an arbitrary commutative ring with identity (in particular, any finite field), or it may take on any value in that ring.

Example H.1. [[LABEL X:0912kn cp]] A simple combinatorial argument shows that

$$\chi_{K_n}(\lambda) = (\lambda)_n := \lambda(\lambda-1)\cdots(\lambda-[n-1]),$$

the *falling factorial* of degree n. (K_0 is included since $(\lambda)_0 = 1$). Expanding the falling factorial in powers of λ, the coefficients are the well-known Stirling numbers of the first kind, $s(n,k)$, defined by the equation

$$(\lambda)_n := \sum_{k=0}^{n} s(n,k)\lambda^k.$$

Thus, the Stirling numbers of the first kind are the Whitney numbers of the first kind of complete graphs: $w_i(K_n) = s(n, n-i)$. (That is how these Whitney numbers came to be called "of the first kind".) The "unsigned" Stirling numbers, $(-1)^k s(n,k)$, count the permutations of $[n]$ that have k cycles.

The Stirling numbers of the second kind are the coefficients in the reverse expansion:

$$\lambda^n = \sum_{k=0}^{n} S(n,k)(\lambda)_k.$$

They, unlike those of the first kind, are nonnegative. There are Whitney numbers of the second kind as well: $W_i(\Gamma)$ is the number of closed edge sets with $n-i$ components. For complete graphs, $W_i(K_n) = S(n, n-i) =$ the number of partitions of $[n]$ into $n-i$ parts.

The opposite indexing of Stirling and Whitney numbers is due to the fact that Whitney numbers were introduced in terms of matroids and geometric lattices, where the natural indexing is in terms of rank. We'll see more about that in Chapter V [**matroid chapter**].

Next we come to three other exciting properties of $\chi_\Gamma(\lambda)$. For connection we have the number of components, $c(\Gamma)$; for 2-connection we define $\mathrm{bl}^*(\Gamma)$ to be the number of blocks that are not loops or isolated vertices.

Proposition H.5. *[*[LABEL T:0912 cp factors]] *For a link graph Γ, the highest power of λ that divides the chromatic polynomial is $\lambda^{c(\Gamma)}$. The highest power of $\lambda - 1$ that divides it is $(\lambda-1)^{\mathrm{bl}^*(\Gamma)}$.*

Exercise H.1. [[LABEL Ex:0912 cp block]] Prove Proposition H.5. Prove the second half by induction. (Later we'll have a better proof of the second half that explains why it is true. [**INSERT FORWARD REFERENCE.**])

Proposition H.6. *[*[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. □

Digression: The Möbius function on closed edge sets.

The Möbius function on a poset P is $\mu_P : P \times P \to \mathbb{Z}$ defined by

$$\mu_P(a,b) = \begin{cases} 0 & \text{if } a \not\leq b, \\ \delta_{ab} - \sum_{c:\ a\leq c<b} \mu_P(a,c) & \text{if } a \leq b, \end{cases}$$

where δ_{ab} is the Kronecker delta. Thus, $\mu_P(a,a) = 1$, and if b covers a, then $\mu_P(a,b) = -1$.

For a closure operator on a set E the poset is the family of closed sets, ordered by set containment. We supplement the preceding definition by declaring that $\mu(\varnothing, A) = 0$ for any closed set A, if $\varnothing$ is not closed. The case we are interested in, of course, is that in which $E = E(\Gamma)$ and the poset is $\operatorname{Lat}\Gamma$. We call this Möbius function μ_Γ.

The basic property of the Möbius function is the following pair of well-known facts (see a combinatorics book, e.g., [EC1]).

Theorem H.7 (Möbius Inversion)**.** *[*[LABEL T:0912muinversion]] *Consider two functions f and g on a poset P. Then*

$$g(a) = \sum_{b\leq a} f(b) \iff f(b) = \sum_{a\leq b} \mu_P(a,b)g(a)$$

and

$$g(a) = \sum_{b\geq a} f(b) \iff f(b) = \sum_{a\geq b} \mu_P(b,a)g(a).$$

Möbius expansion and improper edge sets.

The third exciting property is the connection between the chromatic polynomial and the lattice of closed edge sets. Write μ_Γ for the Möbius function on $\operatorname{Lat}\Gamma$.

Proposition H.8. *[[LABEL P:0912muchromatic]] The chromatic polynomial has the formula*

$$\chi_\Gamma(\lambda) = \sum_{A \in \operatorname{Lat}\Gamma} \mu_\Gamma(\varnothing, A)\lambda^{c(A)}.$$

[INSERT $I(\gamma)$ exposition from later.]

Lemma H.9. *[[LABEL L:0912improperclosed]] The improper edge set $I(\gamma)$ of a coloration is a closed edge set.*

Proof. By transitivity along improper edges of equality of colors. □

Proof of Proposition H.8. Here is the standard proof (due to Rota [FCT]). Let $f(S) :=$ the number of λ-colorations whose improper edge set $I(\gamma)$ is exactly S (thus $f(S) = 0$ if S is not closed) and let $g(S) :=$ the number for which $I(\gamma) \supseteq S$. Clearly, for any closed set A, $g(A) = \sum_{B \geq A} f(B)$. Therefore, $f(B) = \sum_{A \geq B} \mu_\Gamma(B, A)g(A)$ by Möbius inversion. Now we set $B = \varnothing$ and deduce that

$$f(\varnothing) = \sum_{A \in \operatorname{Lat}\Gamma} \mu_\Gamma(\varnothing, A)g(A).$$

(For the time being we assume $\varnothing$ is closed.)

Now we interpret the terms. Plainly, $f(\varnothing) = \chi_\Gamma(\lambda)$. As for $g(A)$, it counts all colorations that are constant on each component of A; that means $g(A) = \lambda^{c(A)}$. This establishes the proposition. □

We infer from Proposition H.8 that the Whitney numbers of the first kind are sums of Möbius function values:

$$w_i(\Gamma) = \sum_{A \in \operatorname{Lat}\Gamma:\ c(A)=n-i} \mu(\varnothing, A).$$

Later we'll see that the Whitney numbers' signs alternate (in Theorem H.29).

H.4. **Maximal forests.** [[LABEL 1.treecount]]

Sept 12b: Simon Joyce

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

$$f(\Gamma) := \text{ the number of maximal forests in a graph } \Gamma.$$

Proposition H.10. *[[LABEL L:0912 tree dc]] The number of maximal forests in a graph has the deletion-contraction property*

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

for any edge e that is not a loop or an isthmus, the multiplicative property

$$f(\Gamma_1 \cupdot \Gamma_2) = f(\Gamma_1)f(\Gamma_2),$$

and the invariance property

$$\Gamma \cong \Gamma' \implies f(\Gamma) = f(\Gamma').$$

Furthermore,

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

Exercise H.2. [[LABEL Ex:0912 tree dc]] Prove Proposition H.10.

The maximal forest number is truly different from the chromatic polynomial since no evaluation of the latter can give the former. We show that with two small examples.

Example H.2. [[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, $f(K_1^\bullet) = 1$ but $\chi_{K_1}(\lambda) = 0$, so evaluating the chromatic polynomial cannot give the maximal forest 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 $f(K_2) = 1$ while $f(2K_2) = 2$.

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 maximal forests, but I can't give an example.

We'll have more to say about maximal forests—and spanning trees—after we meet two two-variable polynomials.

H.5. **Polynomials with two variables.** [[LABEL 1.dichromatic]]

You will have noticed that the three valuable properties of the chromatic polynomial are shared by the maximal forest number (modulo a couple of slight differences). There is a function that encompasses both the chromatic polynomial and the spanning-tree number and has the algebraic properties of both. That is the dichromatic polynomial.

The dichromatic polynomial.

The dichromatic polynomial generalizes (approximately) 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.6 is what tells us the dichromatic polynomial does specialize to the chromatic polynomial; specifically, $\chi_\Gamma(\lambda) = (-1)^n Q_\Gamma(-\lambda, -1)$. (How do we know to make this substitution? The minus sign in the chromatic deletion-contraction formula tells us that we need a sign correction, which can only be $(-1)^{|V|}$ if it is to give the correct signs of both the deletion, where $|V|$ does not change, and the contraction, where $|V|$ decrements by 1. The $(-1)^{|S|}$ in the chromatic formula obliges us to replace v by -1. The rest is then obvious.) Many of the algebraic properties of the chromatic polynomial also generalize; especially, the fundamental deletion-contraction identity.

Theorem H.11. *[*[LABEL P:0912 dichrom dc]] *The dichromatic polynomial of a graph is additive under deletion and contraction:*

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

for any edge e that is not a loop. It is multiplicative:

$$Q_{\Gamma_1 \mathbin{\dot\cup} \Gamma_2}(u,v) = Q_{\Gamma_1}(u,v) Q_{\Gamma_2}(u,v).$$

It is a graph invariant:

$$\Gamma \cong \Gamma' \implies Q_\Gamma(u,v) = Q_{\Gamma'}(u,v).$$

And it is trivial on the empty graph:

$$Q_\varnothing(u,v) = 1.$$

Proof. Regarding the deletion-contraction formula, 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. (We did exactly this to prove Proposition H.6.)

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 $\subseteq E\setminus e$; the first summation is now obviously what we want:

$$= Q_{\Gamma\setminus e}(u,v) + \sum_{T\subseteq E(\Gamma/e)} u^{c_{\Gamma/e}(T)} v^{|T|-(n-1)+c_{\Gamma/e}(T)}$$

but the second is not as obvious:

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

if e is a link, because $|V(\Gamma/e)| = n-1$ and contracting an edge does not change the number of components; but this step fails if e is a loop,

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

The rest of the proof can safely be left to the reader. □

Exercise H.3. [[LABEL Ex:0912Qproperties]]

This proof of the deletion-contraction formula is nice in part because it tells us exactly when the formula is and is not true. It holds for all links and fails for all loops.

Sept 17a: Peter Cohen et al.

The *vertex amalgamation* of two graphs, specifically the *amalgamation at* v, is defined as

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

where Γ_1 and Γ_2 share a vertex v and have no other vertex or edge in common. This occurs, e.g., when a graph has an isthmus or when an isthmus is contracted. See Figure H.1.

ADD GRAPH DIAGRAM.

FIGURE H.1. A graph that is a vertex amalgamation, and the amalgamation of two graphs at a vertex v.
[[LABEL F:0917vamalg]]

The dichromatic polynomial is almost multiplicative over vertex amalgamations; it needs only a slight correction:

Proposition H.12. *[[LABEL P:0917Q-vamalg]]* $Q_{\Gamma_1\cup_v\Gamma_2} = v^{-1} Q_{\Gamma_1} Q_{\Gamma_2}$.

I leave the proof as an exercise for the reader.

Sept 12: Zaslavsky

Example H.3. [[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—additivity in Theorem H.11. 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 Theorem H.11.

Example H.4. [[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.5. [[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.3. 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.6. [[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.

Example H.7. [[LABEL X:0912maxforest no.]] We haven't seen how to obtain the maximal forest number $f(\Gamma)$ from the dichromatic polynomial. Let's see how it should be done. If by clever substitution in Q_Γ we can make the term of S take the value 1 for every maximal forest and 0 for every other edge set, then we get $f(\Gamma)$. Now, let's recall that $|S| = n - c(S)$ for every forest, but for any other set, $|S| > n - c(S)$. Rewrite the definition slightly:

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

Setting $v = 0$ eliminates every set that is not a forest! Therefore, we should be looking at

$$Q_\Gamma(u,0) = \sum_{S\text{: forest}} u^{c(S)}.$$

Here we run into a problem if we try the same trick: every term has a positive power of u (we're assuming Γ is not the empty graph) so by setting $u = 0$ we get sum 0. "Obviously", we should first divide by u. Then

$$u^{-1} Q_\Gamma(u,0)\big|_{u=0} = \sum_{S\text{: forest}} 0^{c(S)-1} = \sum_{S\text{: spanning tree}} 1,$$

the number of spanning trees. That is all right if Γ is connected but it is the wrong number to look at in general (because it fails multiplicativity). In general,the largest power of u that divides every term is $u^{c(\Gamma)}$; thus, what we should do is divide Q_Γ through by $u^{c(\Gamma)}$ and only then set $u = 0$. This works:

$$u^{-c(\Gamma)} Q_\Gamma(u,0)\big|_{u=0} = \sum_{S\text{: forest}} 0^{c(S)-c(\Gamma)} = \sum_{S\text{: maximal forest}} 1 = f(\Gamma).$$

One could suspect that the polynomial $u^{-c(\Gamma)} Q_\Gamma(u,v)$ is important. And it is …

The corank-nullity polynomial. [[LABEL 1.corank-nullity]]

Sept 15a: Yash Lodha

In the definition of the dichromatic polynomial there are implicit 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 cyclomatic number.[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

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

Theorem H.13. *[[LABEL P:0915BR]] The corank-nullity polynomial satisfies the additive relation*

$$R_\Gamma = R_{\Gamma\setminus e} + R_{\Gamma/e}$$

for an edge e that is neither a loop nor an isthmus. It satisfies multiplicativity,

$$R_{\Gamma_1 \mathbin{\dot\cup} \Gamma_2} = R_{\Gamma_1 \cup_v \Gamma_2} = R_{\Gamma_1} R_{\Gamma_2},$$

[1]The names come from the oriented incidence matrix and the lattice of closed edge sets. See Section E.1 for nullity. *Corank* is rank computed in $\operatorname{Lat}\Gamma$ from the top down, the rank of an arbitrary edge set S being the rank in that lattice of $\operatorname{clos} S$. This rank is the rank of the submatrix of $\mathrm{H}(\Gamma)$ composed of the columns of S; that is the dual theorem to Theorem G.11.

for both disjoint union and vertex amalgamation. It is invariant:

$$\Gamma \cong \Gamma' \implies R_\Gamma(u,v) = R_{\Gamma'}(u,v).$$

And it is trivial on an edgeless graph:

$$R_\varnothing(u,v) = R_{K_1}(u,v) = 1.$$

[HOW MUCH PROOF IS NEEDED? Can this be derived from Q? Should it be, though?]

Proof of additivity. 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. □

Proof of multiplicativity. Consider the case of a vertex amalgamation, $\Gamma = \Gamma_1 \cup_v \Gamma_2$. 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 $c(S) = c(S_1) + c(S_2)$. ☐

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

Sept 15b: Yash Lodha

[REWRITE. Integrate parts of this into earlier section H.4 on max forests?]

Let $f(\Gamma)$ be the number of maximal forests of Γ (introduced in Section H.4 and $t(\Gamma)$ the number of spanning trees. To better understand the two two-variable polynomials 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

Theorem H.14. [[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 1: Γ has a loop e. Then $\Gamma = (\Gamma \setminus e) \cup_v K_1^\circ$. By multiplicativity in Theorem H.13, $R_\Gamma = (1+v)R_{\Gamma\setminus e}$. 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 2: 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 3: Γ has a circle C of length greater than one. Let $e \in C$. Then e is not a loop or isthmus, so by additivity in Theorem H.13 and Proposition 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 $\Gamma = \Gamma_1 \,\dot\cup\, \Gamma_2$, where our theorem holds for Γ_1 and Γ_2 is a connected graph. Then by multiplicativity in Theorem H.13 and Proposition H.10 we get our inductive step:

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

Here are a few examples that illustrate the theorem.

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

Spanning trees are necessarily more complex because, unlike the number of maximal forests, that of spanning trees does not satisfy additivity under deletion and contraction. Still, there is a formula based on the dichromatic polynomial. Sept 22a: Jackie Kaminski

Theorem H.15. *[*[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.16. *[*[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, $S_0 \cup S_1 = S$. 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

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

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

$$= \sum_{S_0 \subseteq E_0} \sum_{S_1 \subseteq E_1} u^{n-|S_1|-c(\Gamma)} \cdot v^{|S_0|}$$
$$= u^{n-c(\Gamma)} \left(\sum_{S_0 \subseteq E_0} v^{|S_0|} \right) \left(\sum_{S_1 \subseteq E_1} \left(\frac{1}{u}\right)^{|S_1|} \right)$$
$$= u^{n-c(\Gamma|E_1)} \left(\sum_{S_0 \subseteq E_0} v^{|S_0|} \right) \left(\sum_{S_1 \subseteq E_1} \left(\frac{1}{u}\right)^{|S_1|} \right)$$

$$
\begin{aligned}
&= u^{|E_1|}(v+1)^{|E_0|}\Big(\frac{1}{u}+1\Big)^{|E_1|}\\
&= (v+1)^{|E_0|}(1+u)^{|E_1|}. \qquad \square
\end{aligned}
$$

Proof of Theorem H.15. In preparation, recall from Lemma H.10 that $t(\Gamma) = t(\Gamma\setminus e)+t(\Gamma/e)$ for all links e, and from Proposition 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 $t(\Gamma_1 \cup \Gamma_2) = 0$. 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 additivity in Theorem H.11.

The proof of the theorem is by induction on $|E|$. First we look at the base cases, where $|E| \le 1$; 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.10, 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 additivity in Theorem H.11, 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.15. $\square$

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

Sept 22b: Jackie Kaminski

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 $\mathrm{X}_\Gamma(k, \cdot)$ (capital 'Chi') to be the *generating function of k-colorations by the number of improper edges*; that is,

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

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 $\mathrm{X}_\Gamma(k, 0) = \chi_\Gamma(k)$, since χ_Γ counts the number of k-colorations that are proper. We now prove a theorem of Tutte's which implies that X_Γ is also a polynomial in k.

Theorem H.17 (Tutte)**.** *[[LABEL T:0922 Tutte]] For a graph Γ, the generating function of colorations by the number of improper edges satisfies*

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

Proof. First we reformulate the dichromatic polynomial:

(H.1) [[LABEL E:0922Qreduced]] $$Q_\Gamma(u, v) = \sum_{S \subseteq E} u^{c(S)} v^{|S|-n+c(S)} = v^{-n} \sum_{S \subseteq E} (uv)^{c(S)} v^{|S|}.$$

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

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

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

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

which by Proposition H.6

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

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 $T = R\setminus S$:

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

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

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

and by Equation (**??**) this is

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

We end with an example.

Example H.8. [[LABEL P:0922 Cycle]] For $n\geq 1$,

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

This follows from Example H.5 and Theorem H.17.

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.

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

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

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. In the process we get a formula for the number of acyclic orientations.

Acyclic vs. cyclic orientations.

Define $\mathrm{AO}(\Gamma)$ to be the set of acyclic orientations of Γ and let $a(\Gamma)$ denote the number of them, i.e., $a(\Gamma) := |\,\mathrm{AO}(\Gamma)|$. There is a deletion-contraction formula for this number.

Consider an orientation $\vec{\Gamma}$ of Γ. If we delete an edge $e{:}vw$, $\Gamma \setminus e$ inherits an orientation from Γ in the obvious way. If we contract e, identifying v and w in the contraction, we get an orientation of Γ/e; the only possible ambiguity is in how to orient an edge f with an endpoint v or w. We orient f by treating the contracted vertex v_e as representing the endpoint v or w of f and letting f inherit its orientation from Γ. This is straightforward unless f is parallel to e; then it becomes a loop in Γ/e whose two ends, which are distinguishable according to our conventions, remain its head and tail.

Lemma H.18. *[*[LABEL L:0919 AO]] *Given a graph Γ and a link $e \in E(\Gamma)$, there is a bijection*

$$\mathrm{AO}(\Gamma \setminus e) \mathbin{\dot\cup} \mathrm{AO}(\Gamma/e) \longleftrightarrow \mathrm{AO}(\Gamma).$$

Thus, for any link e, $a(\Gamma) = a(\Gamma \setminus e) + a(\Gamma/e)$.

Sublemma H.19. *[*[LABEL L:0919 ao opposite paths]] *In an acyclic orientation $\vec{\Gamma}$, there can be no two vertices such that there is a coherent path from the first to the second and another from the second to the first.*

Proof. In an oriented graph, the notation $\vec{P}{:}\vec{vw}$ means a coherent path from v to w; that is, every edge in P is oriented from v to w along P.

If coherent paths $\vec{P}{:}\vec{vw}$ and $\vec{Q}{:}\vec{wv}$ both exist in $\vec{\Gamma}$, we can find a cycle in $\vec{\Gamma}$. Follow $\vec{P}$ from v to the first vertex where it intersects $\vec{Q}$, then follow $\vec{Q}$ back to v. The concatenation of these paths is a coherent walk from v to v that repeats no vertex except when it ends at v. □

Proof. We establish a natural 1:1/2:2 correspondence, or *sesquijection* (I made that up), between $\mathrm{AO}(\Gamma \setminus e) \mathbin{\dot\cup} \mathrm{AO}(\Gamma/e)$ and $\mathrm{AO}(\Gamma)$.

Consider $\alpha_0 \in \mathrm{AO}(\Gamma \setminus e)$. If α_0 is also an acyclic orientation of Γ/e, then there is no coherent path either from v to w or from w to v in $\Gamma \setminus e$, for such a path would be a cycle in the contraction. Thus, α_0 can be extended to an acyclic orientation of Γ by adding either $e{:}\vec{vw}$ or $e{:}\vec{wv}$. On the other hand if α_0 is not acyclic in Γ/e, then there is a coherent path in $\Gamma \setminus e$ from v to w or w to v. (Lemma H.19 tells us that both cannot exist.) If a coherent path exists only from v to w, then $e{:}\vec{vw}$ makes Γ acyclic, but $e{:}\vec{wv}$ does not. If such a path exists only from w to v, then similarly we get one acyclic orientation of Γ from α_0.

Suppose $\alpha_1 \in \mathrm{AO}(\Gamma/e)$. Then α_1 applied to $\Gamma \setminus e$ cannot have a cycle. If it did have a cycle $\vec{C}$, then C would have to contain both v and w, or else it would become a cycle in Γ/e upon contraction. Therefore, C consists of two coherent paths, one from v to w and one in the opposite direction. After contracting e, each path becomes a cycle in Γ/e. That proves α_1 is an acyclic orientation of $\Gamma \setminus e$ and therefore in the previous paragraph we counted all acyclic orientations of the deletion and the contraction. Since an acyclic orientation of Γ necessarily comes from one of $\Gamma \setminus e$, we have counted all acyclic orientations of all three graphs. That establishes the sesquijection we want. □

Proper and compatible pairs.

Sept 19:
Simon Joyce

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 $\gamma(v) \leq \gamma(w)$.

Given Γ and a number of colors k, we define $\mathrm{CP}(\Gamma)$ to be the set of all compatible pairs and we let

$$p_k(\Gamma) := |\,\mathrm{CP}(\Gamma)| = \#\text{ of compatible pairs.}$$

Lemma H.20. *[[LABEL L:0919 c pairs DC]] Given a graph Γ and a link $e \in E(\Gamma)$, then*

$$p_k(\Gamma) = p_k(\Gamma \setminus e) + p_k(\Gamma/e).$$

Proof. For fixed $k \geq 0$, we will prove there exists a natural sesquijection (a 1:1/2:2 correspondence) between $\mathrm{CP}(\Gamma \setminus e) \cup \mathrm{CP}(\Gamma/e)$ and $\mathrm{CP}(\Gamma)$. Fix the link $e{:}vw$ and an orientation $\alpha_0 \in \mathrm{AO}(\Gamma \setminus e)$.

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

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 $e{:}\vec{vw}$ and $e{:}\vec{wv}$. 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.21 (Stanley)**.** *[[LABEL T:0919 Stanley's]] For a graph Γ and $k \in \mathbb{Z}_{\geq 0}$,*

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

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

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

So in this case we have equality.

If Γ contains a link then we use Lemma H.20, deletion-contraction of χ_Γ and induction. □

The geometry of proper and compatible pairs.

Sept 19: T.Z.

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

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

The number of these points is given by a function of k,

$$E^{\circ}_{[0,1]^n,\mathcal{H}[\Gamma]}(k+1) := \Big|\Big((0,1)^n \setminus \bigcup \mathcal{H}[\Gamma]\Big) \cap \frac{1}{k+1}\mathbb{Z}^n\Big|,$$

known as the *open Ehrhart polynomial* of $([0,1]^n, \mathcal{H}[\Gamma])$ because it is a polynomial (a theorem of Ehrhart, extended to pairs $(P, \mathcal{H})$ in [IOP, Section 5]. The *closed Ehrhart polynomial* is

$$E_{[0,1]^n,\mathcal{H}[\Gamma]}(k+1) = \sum_{\gamma \in \cap [0,k]^n} m(\gamma),$$

where $m(x) :=$ number of closed regions of $\mathcal{H}[\Gamma]$ that contain x. A theorem proved by Ehrhart and I.G. Macdonald shows that the open and closed polynomials are very close to each oterh.

Theorem H.22 (Ehrhart Reciprocity)**.** *[[LABEL T:0919Erecip]] $E^{\circ}_{[0,1]^n,\mathcal{H}[\Gamma]}(t) = (-1)^n E_{[0,1]^n,\mathcal{H}[\Gamma]}(-t)$.*

Now, suppose γ is a k-coloration of Γ, i.e., $\gamma \in \mathbb{Z}^n \cap [k]^n$. Each $\alpha \in \mathrm{AO}(\Gamma)$ corresponds (by Theorem G.14) to an open region $R(\alpha)$ of $\mathcal{H}[\Gamma]$, defined by $x_i < x_j$ when $\exists\ v_i\vec{v}_j$ in α, and also a closed region $\bar{R}(\alpha)$ defined by $\gamma_i \leq \gamma_j$ when $\exists\ v_i\vec{v}_j$ in α. The number of proper pairs $(\alpha, \gamma) =$ the number of open regions that contain γ, which is 1 if γ is proper and 0 if it is not, so

$$E^{\circ}_{[0,1]^n,\mathcal{H}[\Gamma]}(k+1) = \text{the number of proper pairs } (\alpha,\gamma) = \chi_\Gamma(k).$$

Tracking the definitions shows that (α, γ) is compatible iff $\gamma \in \bar{R}(\alpha)$ (because both compatibility and closed regions are determined by weak, not strong, inequalities). So, for each k-coloration γ, the number of compatible pairs $(\alpha, \gamma) =$ the number of closed regions of $\mathcal{H}[\Gamma]$ that contain γ, which is $m(\gamma)$; in other words,

$$E_{[0,1]^n,\mathcal{H}[\Gamma]}(k-1) = \text{the number of compatible pairs } (\alpha,\gamma).$$

Now we do a short computation.

$$\begin{aligned}(-1)^n\chi_\Gamma(-k) &= (-1)^n E^{\circ}_{[0,1]^n,\mathcal{H}[\Gamma]}(-k+1) \\ &= E_{[0,1]^n,\mathcal{H}[\Gamma]}(k-1) = \text{the number of compatible pairs } (\alpha,\gamma).\end{aligned}$$

This is the geometrical proof of Stanley's theorem. (It is essentially the same as one of the three proofs Stanley gave, but the presentation differs in that we treat all the regions together instead of separately.)

The curious fact that there is really only one geometric polynomial that counts both proper and compatible pairs is no surprise if we already know Stanley's theorem, but it hints at vast generalizations, which I will mention when we get to the signed-graphic Stanley theorem in Section **??**.

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. Universality means that every function satisfying those relations is an evaluation of the Tutte polynomial. We've already seen a form of the Tutte polynomial, as it is equivalent to the corank-nullity polynomial. However, its definition is quite different. Sept 24: Nate Reff

Let's begin with a review of the relations.

Tutte–Grothendieck invariants.

We found that:

- Deletion-Contraction Property:
  $Q_\Gamma = Q_{\Gamma\setminus e} + Q_{\Gamma/e}$ if e is not a loop.
  $R_\Gamma = R_{\Gamma\setminus e} + R_{\Gamma/e}$ if e is not a loop or isthmus.
- Disjoint Graph Multiplicativity:
  $Q_{\Gamma_1 \,\dot\cup\, \Gamma_2} = Q_{\Gamma_1} Q_{\Gamma_2}$ and $R_{\Gamma_1 \,\dot\cup\, \Gamma_2} = R_{\Gamma_1} R_{\Gamma_2}$.
- Cutpoint Multiplicativity:
  $R_{\Gamma_1 \cup_v \Gamma_2} = R_{\Gamma_1} R_{\Gamma_2}$.
- 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:**
$F(\Gamma) = F(\Gamma \setminus e) + F(\Gamma/e)$ if e is not a loop or isthmus.

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

(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 all such functions are.

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 multiplicativity property (M), $F(\Gamma) =$ the product of F(blocks). Due to the property of invariance (I), $F(\text{loop}) =$ a value y that is the same for all loops, and also $F(\text{isthmus}) =$ a value x that is the same for all isthmi. Lastly, there is a simple form for a basic special case.

Lemma H.23. *[*[LABEL L:0924lemma1 Loop Isthmus Lemma]] *Suppose* Γ *has* l *loops and* i *isthmi and no other edges. Then* $F(\Gamma) = x^i y^l$.

Another side comment: If the codomain of F is an integral domain, then (U) is almost superfluous; that is, it can be deduced from the other properties with the exception of a handful of functions F.

Exercise H.4. [[LABEL Ex:0924 Uexceptions]] Find all the functions F into an integral domain that satisfy (DC, M, I) but not (U). (Hint: Attempt to derive (U) from (M) and (I); find the exceptional cases.)

Returning to Lemma H.23, let's look as a simple example. Suppose we define a graph G as in Figure H.2. The digon C_2 is a circle of length 2; it consists of two vertices and two parallel edges between those vertices. Calculating $F(C_2)$ using the deletion-contraction method (as seen in Figure H.2) we get the following:

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

Theorem H.24 (Universality of R_Γ)**.** *[[LABEL T:0924Theorem1 Main Theorem]] Suppose F is a Tutte–Grothendieck invariant of graphs. Let $x = F(isthmus)$, and let $y = F(loop)$. Then*
(a) *$F(\Gamma) = R_\Gamma(x - 1, y - 1)$, a polynomial function of x and y,*
(b) *the polynomial has nonnegative integral coefficients, and*
(c) *any evaluation of $R_\Gamma(u, v)$ 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.25. *[[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.24. □

The Tutte polynomial.

The Tutte polynomial is the expression we get for a graph when it is fully reduced by the properties that define a Tutte–Grothendieck invariant of graphs. Let's write those properties in a function-independent way, as reduction rules for graphs themselves: Sept 24: T.Z.

(H.2)

$$\begin{aligned} \Gamma &\mapsto (\Gamma \setminus e) + (\Gamma/e) \quad \text{if } e \text{ is neither a loop nor an isthmus,} \\ \Gamma_1 \mathbin{\dot\cup} \Gamma_2 &\mapsto \Gamma_1 \cdot \Gamma_2, \\ \Gamma_1 \cup_v \Gamma_2 &\mapsto \Gamma_1 \cdot \Gamma_2, \\ \varnothing, K_1 &\mapsto 1, \\ \Gamma_1 \cong \Gamma_2 &\implies \Gamma_1 \longleftrightarrow \Gamma_2, \end{aligned}$$ [[LABEL E:0924TGreductions]]

where the two-headed arrow means either graph can be converted to the other. The only graphs that cannot be reduced to smaller graphs by these rules are blocks in which every

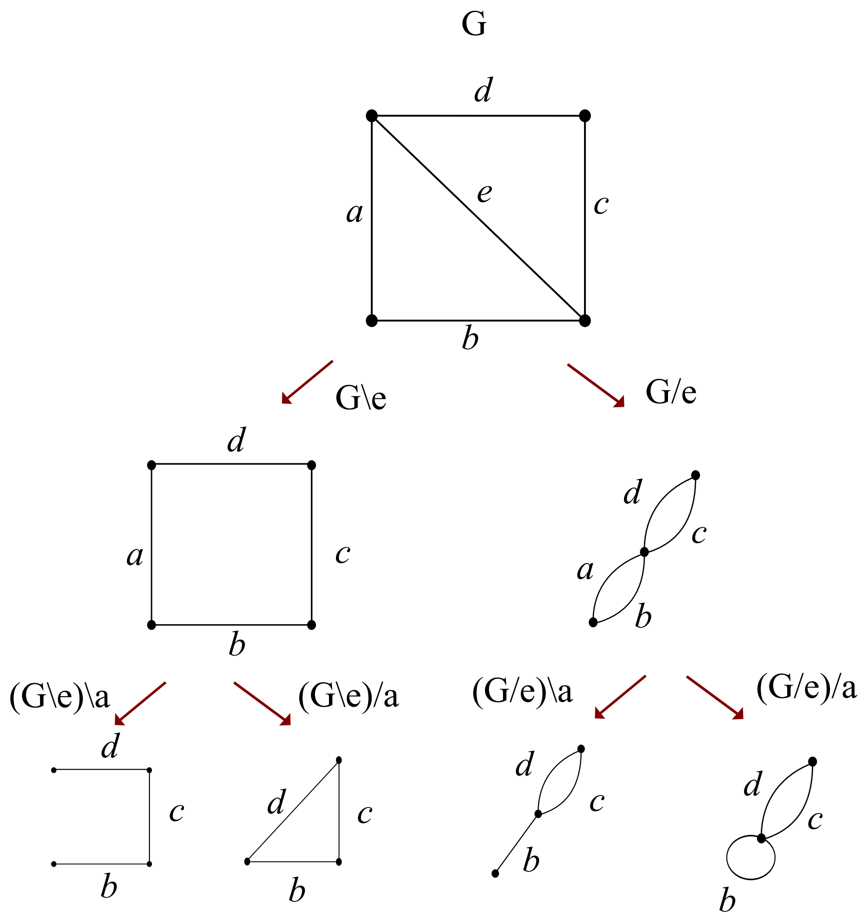

FIGURE H.2. Calculation of $F(G)$ by the method of deletion and contraction. [[LABEL F:0924Figure1]]

edge is a loop or isthmus; that is, $\varnothing$, K_1, K_1° (the loop graph: a vertex with a loop), and K_2 (the isthmus graph). The first two reduce to 1;[2] the latter do not reduce, but we replace them by y and x, respectively (to make things look more algebraic, I suppose).

We apply these rules to a graph Γ until no more reduction is possible. The order in which we apply them, however, is extremely flexible and involves choice after choice after choice. We may have to choose between applying deletion-contraction of an edge e and multiplicative reduction, or between deletion-contraction of e and of a different edge e'. There is no intrinsic guarantee that the results of different reduction processes will be the same. However, they are, and that is the justification for the Tutte polynomial and a great deal that follows from it.

[2]We are virtually forced by multiplicativity to make that reduction. In disjoint multiplicativity Γ_2 could be $\varnothing$; in cutpoint multiplicativity it could be K_1. In either case the union is Γ_1 again, so $\Gamma_1 \cdot \varnothing = \Gamma_1$ and similarly for K_1. The best—in view of the coming analysis the only—way to satisfy these equations is to reduce $\varnothing$ and K_1 to 1 in the Tutte ring.

But wait! How can I add and multiply graphs? The answer is that we are operating in the polynomial ring $\mathbb{Z}[\{\text{graphs}\}]$ formally generated by all graphs,[3] as if the graphs were indeterminates. In this ring we introduce the relations $\Gamma = (\Gamma \setminus e) + (\Gamma/e)$ and so on; technically, we take the *Tutte ideal*

$$\begin{aligned} I := \Big\langle\ & \Gamma - (\Gamma \setminus e) - (\Gamma/e) \text{ (when } e \text{ is not a loop or isthmus)}, \\ & (\Gamma_1 \uplus \Gamma_2) - \Gamma_1 \cdot \Gamma_2,\ (\Gamma_1 \cup_v \Gamma_2) - \Gamma_1 \cdot \Gamma_2, \\ & \Gamma_1 - \Gamma_2 \text{ (when } \Gamma_1 \cong \Gamma_2), \\ & \varnothing - 1,\ K_1 - 1\ \Big\rangle \end{aligned}$$

and we work in $\mathbb{Z}[\{\text{graphs}\}]/I$; that makes sense of the algebraic operations. This quotient is the *Tutte ring*. It is generated by the only two graphs that don't reduce: $x = K_2$ and $y = K_1^\circ$. It follows that the Tutte ring is a homomorphic image of $\mathbb{Z}[x, y]$. The image $\bar{\Gamma}$ of a graph Γ in the Tutte ring is an element of this homomorphic image.

On the other hand, the Tutte ring maps onto $\mathbb{Z}[u, v]$ by the mapping $\bar{\Gamma} \mapsto R_\Gamma(u, v)$, which carries $x = K_2 \mapsto R_{K_2}(u, v) = u + 1$ and $y = K_1^\circ \mapsto R_{K_1^\circ}(u, v) = v + 1$; that mapping is well defined because R_Γ is a Tutte–Grothendieck invariant. Since the generators x, y of the Tutte ring map to generators of the polynomial ring $\mathbb{Z}[u, v]$, the Tutte ring must be the whole polynomial ring $\mathbb{Z}[x, y]$ and not a proper homomorphic image of it. Therefore, $\bar{\Gamma}$ is a polynomial belonging to $\mathbb{Z}[x, y]$; it is called the *Tutte polynomial* of Γ and is denoted by $T_\Gamma(x, y)$.

We now know the following:

Theorem H.26. *[[LABEL T:0924tuttepoly]]*

1. *(Well definition) The result of reducing a graph Γ by the reduction rules* (H.2) *is a polynomial $T_\Gamma(x, y)$ in the indeterminates $x = K_2$ and $y = K_1^\circ$, that is the same for every choice of reduction process.*
2. *(Algebraic expression) $T_\Gamma(x, y) = R_\Gamma(x - 1, y - 1)$. [[LABEL T:0924tuttepoly R]]*
3. *The Tutte polynomial is a Tutte–Grothendieck invariant of graphs, and so is every evaluation of it.*
4. *(Universality of the Tutte polynomial) Every Tutte–Grothendieck invariant F of graphs is the evaluation of the Tutte polynomial with $x = F(K_2)$ and $y = F(K_1^\circ)$.*

Algebra explained.

How does this differ from Theorem H.24 and Corollary H.25? Notice that we have passed from (Tutte–Grothendieck invariant) functions on graphs to algebraic expressions (in the Tutte ring) in terms of graphs. The functions are then defined on the Tutte ring, within which all graphs live (as their homomorphic images, $\bar{\Gamma} = T_\Gamma$)—that is possible because every Tutte–Grothendieck invariant contains the Tutte ideal in its kernel. Previously, with R_Γ, we were working with the dual of the Tutte ring, that is, the algebra of functions from it into various codomains; now we are working with the core object itself, which is independent of any choice of codomain. [4]

[3] We overlook the technicality that "all graphs" is not a set.

[4] This paragraph is entirely owed to Gian-Carlo Rota, who introduced me to Tutte–Grothendieck invariants of matroids and graphs and, indeed, coined the TG name. Later we liked to say Grothendieck had nothing to do with it, as Tutte's development of the Tutte ring and polynomial greatly preceded Grothendieck's similar development of the Grothendieck group.

A little bit more explanation: We can now treat Tutte–Grothendieck invariant of graphs with values in a commutative, unital ring A as a homomorphism $F : \mathbb{Z}[x, y] \to A$. Any homomorphism from a polynomial ring is completely determined by its values on the generators, in our case x (that is, K_2) and y (that is, K_1°), and those values may be assigned arbitrarily. Thus, the explanation of Corollary H.25 is that the Tutte ring is a polynomial ring. (This is important. There are generalizations where the Tutte ring is no longer a polynomial ring; then the generalized Tutte–Grothendieck invariants are much harder to describe.)

Properties of the Tutte polynomial.

Sept 24: Nate Reff

Using previous results we can now say that

$$T_\Gamma(1, 1) = R_\Gamma(0, 0) = f(\Gamma),$$

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. But we can also talk about the properties of T_γ itself. Write

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

Theorem H.27. *[[LABEL T:0924Theorem2]]*

1. *The degree of x equals the rank* $\mathrm{rk}(\Gamma) = n - c(\Gamma)$ *and the degree of y equals the nullity of Γ, that is, $|E| - n + c(\Gamma)$. [[LABEL T:0924Theorem2 deg]]*
2. *All coefficients in the Tutte polynomial are nonnegative integers. [[LABEL P:0924tuttecoefficients]]*
3. *$b_{00} = 0$ if $|E| > 0$. [[LABEL T:0924Theorem2 b00]]*

Proof. Part (1) is an immediate consequence of Theorem H.26(2).

Parts (2) and (3) can be proved by induction. For the former, since we always add, no negatives are ever introduced into T_Γ. □

In view of Theorem H.27(2) it is natural to ask whether the coefficients b_{ij} count anything. Indeed they do, but the description of what they count is complicated. We postpone it to the matroid chapter. [**MAKE SURE THIS IS DONE.**] It's worth mentioning that the counting description was Tutte's first way to define the Tutte polynomial—and it was tedious to prove the polynomial was independent of the order of reduction operations. Only later did he discover the easier approach through the closed-form definitions of the dichromatic and corank-nullity polynomials.

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

$$R_\Gamma(u, v) = \sum_S u^{c(S)-c(F)} v^{|S|-n+c(F)} = \sum_{k,l} a_{kl} u^k v^l, \text{ [[LABEL E:0924Tutte1]]} \tag{H.3}$$

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 $l = |S| - n + c(S)$. We deduce from the correspondence between the Tutte polynomial and the corank-nullity polynomial that

$$R_\Gamma(u, v) = T_\Gamma(u + 1, v + 1) = \sum_{i,j \geq 0} b_{ij} (u + 1)^i (v + 1)^j$$

$$= \sum_{i,j \geq 0} b_{ij} \sum_k \binom{i}{k} u^k \sum_l \binom{j}{l} v^l$$

$$= \sum_{i,l\geq 0} u^k v^k \underbrace{\sum_{i,j\geq 0} b_{ij} \binom{i}{k}\binom{j}{l}}_{\text{coefficient of } u^k v^k \text{ in (H.3)}}$$

$$= \sum_{k,l} a_{kl} u^k v^l.$$

This string of equalities shows that:

Proposition H.28. *[*[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 and in the next subsubsection: For instance, $|w_i|$ bounds (Dowling–Wilson for matroids), same for inseparable graphs (Brylawski for matroids), ...]**

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.

Let's take a second look at the chromatic polynomial. We can now say that:

$$\sum_{k=0}^{n} (-1)^{n-k} w_{n-k} \lambda^k = (-1)^n \chi_\Gamma(-\lambda) = Q_\Gamma(\lambda, -1) = T_\Gamma(1+\lambda, 0)$$

$$= \sum_{i,j\geq 0} (1+\lambda)^i 0^j b_{ij} = \sum_i (1+\lambda)^i b_{i0}$$

$$= \sum_k \lambda^k \sum_i \binom{i}{k} b_{i0}.$$

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

There is a direct proof by induction and deletion-contraction, but this one is simpler—once we know about the Tutte polynomial.

More can be said about the Whitney numbers with further study involving the Tutte polynomial, but we stop here.

I. LINE GRAPHS

[[LABEL 1.lg]]

The *line graph* of Γ, denoted by $L(\Gamma)$, is defined as follows: Sept 26: Yash Lodha

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

(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 nonisomorphic graphs can produce the same line graph? We now provide two 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!]**

I.1. **Eigenvalues.** [[LABEL 1.lg-evalues]]

In Section E we computed the product of an incidence matrix times its transpose. Now we reverse the product. First, the unoriented incidence matrix: $B^{\mathrm{T}}B$, an $E \times E$ matrix. In this matrix the entry x_{ij} is the number of edges between the vertices v_i and v_j, and x_{ii} is the degree of vertex v_i. It is clear that $B^{\mathrm{T}}B = A(L) + 2I$, where $L = L(\Gamma)$. Since $B^{\mathrm{T}}B$ is positive semidefinite, the eigenvalues are greater than or equal to zero.

Theorem I.1. *[*[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 $\lambda \geq -2$. □

People like to use eigenvalues (or eigenvectors) of graph matrices, especially the adjacency matrix, to deduce graph properties. The reason is that eigenvalues are easy to compute, but many graph properties are not; if eigenvalues give information about those properties, it's useful. So when we notice that line graphs have the striking property of a universal constant lower bound on eigenvalues, we wonder whether this is a special property of line graphs. In other words, are there any graphs besides line graphs whose least eigenvalue is ≥ -2?

The answer: A few, in fact a small finite number. This was proved by Cameron, Goethals, Seidel, and Schult (1976a???) by applying the theory of positive semidefinite matrices.

[MORE; ALSO, ADD POS. SEMIDEF. TO MATRIX SECTION.]

Assume G is a graph such that $A(G)$ has least egienvalue ≥ -2. Then $A(G) + 2I$ has nonnegative eigenvalues. That implies it is a Gram matrix, i.e., a matrix of inner products of vectors in some vector space. Let's suppose $A(G) + 2I = M^{\mathrm{T}}M.$, where M has columns $m_1, \ldots, m_n$. Then

[ADD MORE TO THIS. Root syst, CGSS, s.g. needed to get the full picture.]

J. Cycles, Cuts, and their Spaces

[[LABEL 1.cyclecut]]

There is linear algebra about graphs. One point of view is based on looking at the row and null spaces of the incidence matrix as aspects of a homology theory. Another viewpoint treats the power set of E as a vector space over the field $\mathbb{F}_2$, within which live special subsets like circles and cuts and the subspaces they generate. In both we follow Rota's dictum that combinatorics is vector spaces with bases. We'll look at both of these as vector spaces over a field, in preparation for signed graphic spaces (where we may have to generalize to modules over a commutative ring). In both we see the sense of Rota's dictum that "Combinatorics is vector spaces with bases."[5] 2014 Oct 19: T.Z.

[MAKE SURE TO EDIT THIS SUITABLY!]

J.1. **Graphic vector spaces.** [[LABEL 1.fieldcyclecut]]

In what follows, $\mathbf{F}$ will be a field and E will be a set. Graphs are assumed to be ordinary; the edges are links and loops. October 13, 2014: Richard Behr

Definition J.1. [[LABEL D:20141013 linearcombinationsofelements]] We write $\mathbf{F}E$ for the set of all linear combinations of elements of E with coefficients in $\mathbf{F}$. We also write $\mathbf{F}^E := \{g \mid g : E \to \mathbf{F}\}$.

Now, $\mathbf{F}E$ is a vector space with basis E, and $\mathbf{F}^E$ is a vector space with basis $\{1_e : e \in E\}$ where 1_e is defined by the rule $1_e(f) = \delta_{ef}$ for $f \in E$. Each element $\varphi \in \mathbf{F}^E$ extends (uniquely) to a linear function on all elements $\alpha = \sum_{i=1}^m \alpha_i e_i$ of $\mathbf{F}E$ via the definition $\varphi(\alpha) = \sum_{i=1}^m \alpha_i \varphi(e_i)$. In other words, $\mathbf{F}^E$ is the set of linear functionals for the vector space $\mathbf{F}E$, and thus $(\mathbf{F}E)^* = \mathbf{F}^E$.

If E is a set of oriented edges, we say that $\mathbf{F}E$ is the *edge space* of E over the field $\mathbf{F}$, and $\mathbf{F}^E$ is the *co-edge space*. Notice that $\mathbf{F}E$ and $\mathbf{F}^E$ are canonically isomorphic. The isomorphism is given by the map from $\mathbf{F}E$ to $\mathbf{F}^E$ that sends a basis element $e \in \mathbf{F}E$ to the basis element $1_e \in \mathbf{F}^E$. Thus we may think of the edge space and the co-edge space as being essentially the same thing, and it will be convenient to change our perspective at times. Thus it may occur that we write 1_e where we mean e, and vice versa. We can say the same things about $\mathbf{F}V$ and $\mathbf{F}^V$; these two objects are also canonically isomorphic (but there is no orientation of vertices). We refer to them as the *vertex space* and *co-vertex space*, respectively. We shall write $\mathbf{F}E = C_1(\Gamma)$, $\mathbf{F}V = C_0(\Gamma)$, $\mathbf{F}^E = C^1(\Gamma)$, and $\mathbf{F}^V = C^0(\Gamma)$ to agree with the convention in algebraic topology. (When it is necessary to specify $\mathbf{F}$, we write $C_1(\Gamma; \mathbf{F})$ et al.)

Definition J.2. [[LABEL D:20141013 boundarymap]] For an oriented edge $\vec{e}$, the head $h(\vec{e})$ is the endpoint in the direction the arrow is pointing and the tail $t(\vec{e})$ is the other endpoint. The mapping $\partial : \mathbf{F}E \to \mathbf{F}V$ given by $\partial(\vec{e}) = h(\vec{e}) - t(\vec{e})$ is called the 1-*boundary mapping.* (Since there are other boundary mappings, we may write this one as ∂_1.)

This is the boundary mapping from algebraic topology, if we consider the graph as a topological space. Keeping in that spirit, we say that an $\mathbf{F}$-*cycle* is an element f of $\mathbf{F}E$ whose boundary is 0, or in other words $f \in \operatorname{Ker} \partial$. We write $Z_1(\Gamma) := \operatorname{Ker}(\partial)$. Also, we write $B_0(\Gamma) = \operatorname{Im}(\partial)$, the set of $\mathbf{F}$-0-*boundaries.*

[5] I don't have a citation, but I heard it with my own ears.

Along with our boundary operator ∂, we have a coboundary operator $\delta = \delta_0 : \mathbf{F}^V \to \mathbf{F}^E$ which is determined by

$$\delta(1_v) := \sum_{e \in E(v)} \eta(v,e) 1_e = \sum_{e \in E} \eta(v,e) 1_e,$$

where $E(v) := \{\text{all edges incident with } v\}$ and

$$\eta(v,e) = \begin{cases} 1 & \text{if the head of } e \text{ is } v, \\ -1 & \text{if the tail of } e \text{ is } v, \\ 0 & v \text{ is not an endpoint of } e. \end{cases}$$

We write $Z^0(\Gamma) := \mathrm{Ker}(\delta)$ and call it the 0-*cocycle space*. Note that the vertices in $Z^0(\Gamma)$ are the ones that have the same number of edges going in as they do going out. We write $B^1(\Gamma) := \mathrm{Im}(\delta)$.

The *indicator function* of $W \subseteq V$ is $1_W : V \to \mathbf{F}$ defined by $1_W(v) = 1$ if $v \in W$ and 0 if $v \notin W$.

Theorem J.1. *[[LABEL T:20141013 bdycomps]]* $B_0 = \mathrm{span}\{1_W : W = V(\textit{component of } \Gamma)\}^\perp$.

Proof. For the statement in the theorem to make sense, we have to think of the 1_W as members of $\mathbf{F}V$, which we do by using the canonical isomorphism between $\mathbf{F}^V$ and $\mathbf{F}V$. Let $Y = \mathrm{span}\{1_W : W = V(\text{component of } \Gamma)\}$. We will have $B_0 = Y^\perp \iff B_0^\perp = Y \iff \forall w \in W, \forall(\partial(\alpha)), 1_w \cdot \partial\alpha = 0$, where $\alpha = \sum_{e \in E} \alpha_e e \in C_1(\Gamma)$. Thus,

$$1_w \cdot \partial\alpha = \sum_{v \in V} 1_w(v) v \cdot \sum_{u, \vec{e}} \eta(u, \vec{e}) \alpha_e u = \sum_u \sum_v \sum_{\vec{e}} 1_w(v) \alpha_e \eta(u, \vec{e})(v \cdot u).$$

Notice that $v \cdot u$ is 1 if and only if $v = u$, since this is an inner product of basis vectors. Thus the above triple sum is equal to

$$\sum_v \sum_{\vec{e}} 1_w(v) \eta(v, \vec{e}) \alpha_e = \sum_{\vec{e}} \alpha_e \sum_v \eta(v, \vec{e}) = \sum_e \alpha_e \sum_{v \in W} \eta(v, \vec{e})$$

Note that if $e \notin E : W$, $\sum_{v \in W} \eta(v, \vec{e}) = 0$. Additionally, if $e \in E{:}W$, we see $\sum_{v \in W} \eta(v, \vec{e}) = 0$, since each edge has both endpoints in W (it is a component). Thus, the statement is proved. □

Now let us consider $z = \sum_{e \in E} \alpha_e e \in Z_1(\Gamma)$. Then $\partial(z) = 0$. Thus,

$$0 = \sum_{e \in E} \alpha_e \partial(e) = \sum_{e \in E} \alpha_e (h(\vec{e}) - t(\vec{e})) = \sum_{v \in V} v(\sum_{\vec{e}: h(e) = v} \alpha_e + \sum_{\vec{e}: t(e) = v} (-\alpha_e)).$$

Thus,

$$\sum_{\vec{e}: h(e) = v} \alpha_e + \sum_{\vec{e}: t(e) = v} (-\alpha_e) = 0 = \sum_e \eta(v, \vec{e}) \alpha_e.$$

(This final quantity is the net inflow of α at vertex v, if we think of α as an amount of stuff flowing through the edges.) We have uncovered the following theorem:

Theorem J.2. *[[LABEL T:20141013 netinflow]]* $Z_1(\Gamma) = \mathrm{Ker}(\partial) = \{\alpha \in \mathbf{F}E : \text{the net inflow at any vertex}$ $\{\alpha \in \mathbf{F}E : \sum_{\vec{e}} \eta(v, \vec{e}) \alpha_e = 0 \ \forall v\}$.

We can also find another characterization of Z_1. To do this we need the concept of an *oriented circle* $\vec{C}$, which is a circle with a direction (not the same as an orientation; the edges are not oriented, only the circle is oriented). In an orientation of Γ, the (oriented) indicator function of $\vec{C}$ is

$$1_{\vec{C}}(\vec{e}) = \begin{cases} 1 & \text{if the orientation of } \vec{e} \text{ agrees with that of } \vec{C}, \\ -1 & \text{if the orientation of } \vec{e} \text{ disagrees with that of } \vec{C}, \\ 0 & e \text{ is not in } \vec{C}. \end{cases}$$

Although $1_{\vec{C}} \in \mathbf{F}^E$, we can think of it as lying in $\mathbf{F}E$ by writing it as $\sum_{\vec{e}} 1_{\vec{C}}(\vec{e})\vec{e}$. We are now ready to state the next theorem:

Theorem J.3. *[[LABEL T:20141013 spancircle]]* $Z_1 = \operatorname{span}\{1_{\vec{C}} : C \text{is a circle of } \Gamma\}$.

To prove this theorem, we will first prove the following lemma:

Lemma J.4. *Suppose that T is a maximal forest in Γ, $C_T(e)$ is the fundamental circle of e with respect to T, and $\vec{C}_T(\vec{e})$ is $C_T(e)$ oriented to agree with the orientation of e. Then*

$$\alpha \in Z_1 \iff \alpha = \sum_{e \notin T} \alpha_e 1_{\vec{C_T}(\vec{e})}.$$

Proof. Notice that if $f \notin T$,

$$[\alpha - \sum_{e \notin T} \alpha_e 1_{\vec{C_T}(e)}]f = \alpha f - \alpha_f f = 0$$

Now suppose that u is a vertex in T with degree 1, and let e_0 be its adjacent edge. Then, $(\partial\alpha)u = 0$, and thus $\alpha e_0 = 0$. Thus, we see that

$$[\alpha - \sum_{e \notin T} \alpha_e 1_{\vec{C_T}(e)}]e_0 = \alpha e_0 - 0 = 0.$$

Now we can forget about e_0 to get a new edge e_1 that has degree 1 (ignoring e_0), and replace e_0 with e_1 in the above formula. Thus $[\alpha - \sum_{e \notin T} \alpha_e 1_{\vec{C_T}(e)}]$ is 0 on all edges in Γ. $\square$

Corollary J.5. *The fundamental circles of Γ with respect to a maximal forest T are a basis for Z_1.*

Proof. The above lemma shows that any element of Z_1 can be written as a linear combination of indicator vectors for fundamental circles, and we know that the set of fundamental circles generate all cycles. $\square$

Similarly to what we did before, we can characterize $Z^0(\Gamma)$ and $B^1(\Gamma)$. October 15, 2014: Ting Su

Theorem J.6. *[[LABEL T:20141015 coboundarykernel]]* $Z^0 = \operatorname{span}\{1_W : W \in \{\text{components}\}\}$. *(So $Z^0 = B_0^\perp$.)*

Proof. By definition,

$$\delta(1_v) := \sum_{e \in E} \eta(v, \vec{e}) 1_e \longleftrightarrow \delta(1_v)(\vec{e}) = \eta(v, \vec{e})$$

Suppose $\beta \in C^0(\Gamma)$, $\beta = \sum_v \beta(v) 1_v$, thus

$$\delta\beta = \sum_v \beta(v)\delta(1_v) = \sum_v \beta(v) \sum_e \eta(v, \vec{e}) 1_e = \sum_e 1_e (\sum_v \eta(v, \vec{e})\beta(v))$$

We can see that the coefficients

$$\sum_v \eta(v,\vec{e})\beta(v) = \eta(h(\vec{e}),\vec{e})\beta(h(\vec{e})) + \eta(t(\vec{e}),\vec{e})\beta(t(\vec{e})) = \beta(h(\vec{e})) - \beta(t(\vec{e})).$$

Thus, $\beta \in Z^0 \iff \delta\beta = 0 \iff$ all coefficients $= 0 \iff \beta(v) - \beta(w) = 0, \forall \vec{e}{:}v\vec{w} \iff \beta$ is constant on components. $\iff \beta =$ linear combination of 1_W's $\iff \beta \in \mathrm{span}\{1_W : W \in \{\text{components}\}\}$. So, $Z^0 = \mathrm{span}\{1_W : W \in \{\text{components}\}\}$. ☐

Before giving the characterization of B^1, we need to first introduce cutsets.

Definition J.3. [[LABEL D:20141015 cut]] A *cut* or *cutset* is the set of edges between a vertex set $X \in V$ and its complement X^c (if this set is nonempty).

For $X, Y \in V$ 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.4. [[LABEL D:20141015 vertex cut]] A *vertex cut* is the set of all edges incident to a vertex, i.e., $E(\{v\}, V \setminus v)$.

Also, we have the directed cut.

Definition J.5. [[LABEL D:20141015 directedcut]] A *directed cut* is a cut $E(X,X^c)$ with a direction (either from X or from X^c). The notation is $\vec{E}(X,X^c)$. The indicator function is $1_{\vec{E}(X,X^c)}$ defined by

$$1_{\vec{E}(X,X^c)}(\vec{e}) := \begin{cases} 1 & \text{if } \vec{e} \in \text{cut and oriented from } X \text{ to } X^c, \\ -1 & \text{if } \vec{e} \in \text{cut and oriented from } X^c \text{ to } X, \\ 0 & \text{if } \vec{e} \notin \text{cut.} \end{cases}$$

[{cases} does not take the colon : .]

It's easy to see that $1_{\vec{E}(X,X^c)} = -1_{\vec{E}(X^c,X)}$.
Now let's find the image of the coboundary function.

Theorem J.7. *[*[LABEL T:20141015 cobdyim]] $B^1 = \mathrm{span}\{\text{oriented vertex cut functions } 1_v\} = \mathrm{span}\{\text{directed cut functions}\}$.

Proof. $\gamma \in B^1 \iff \gamma = \delta\beta$ for some $\beta \in C^0$ (By definition, $\beta = \sum_v \beta(v)1_v$.) $\iff \gamma = \sum_v \beta(v)\delta 1_v \iff \gamma \in \mathrm{span}\{\delta 1_v : v \in V\}$ ☐

Also we can find a basis of this cut space. To do this we will first choose any vertex $v_W \in W$ to be the root of the component $\Gamma{:}W$. A basis of B^1 is $\{\delta 1_v : v \neq \text{root of a component}\} = \{1_{E(v,V\setminus v)} : v \neq \text{root of a component}\}$. To prove this, we need the following lemma.

Lemma J.8. *[*[LABEL L:20141015 cutfuncomp]] $\delta 1_W = \sum_{v\in W} \delta 1_v = 0$ *if* $W = V$(component).

Proof. **[THIS PROOF NEEDS WORDS TO EXPLAIN WHAT THE FORMULAS ARE FOR AND HOW THEY ARE RELATED.]**

$$\delta 1_W = \sum_{v\in W} \delta 1_v = \sum_{v\in W}\sum_e \eta(v,\vec{e})1_e = \sum_e 1_e \sum_{v\in W}\eta(v,\vec{e})$$

$$\sum_{v\in W}\eta(v,\vec{e}) = \begin{cases} 0 & \text{if e} \notin \text{E: W} \\ \eta(h(\vec{e}),\vec{e}) + \eta(t(\vec{e}),\vec{e}) = 1 - 1 = 0 & \text{if e} \in \text{E: W} \end{cases}$$

$$\begin{array}{cc}
C_2(\Gamma;F) = \langle \text{circles} \rangle = F\mathcal{C} & C^2(\Gamma;F) = F^{\mathcal{C}} = \langle 1_{\vec{C}} \rangle \\
\downarrow \partial_2 & \uparrow \delta_1 \\
C_1(\Gamma;F) = \langle \text{edges} \rangle = FE & C^1(\Gamma;F) = F^E = \langle 1_{\vec{e}} \rangle \\
\downarrow \partial_1 & \uparrow \delta_0 \\
C_0(\Gamma;F) = \langle \text{vertices} \rangle = FV & C^0(\Gamma;F) = F^V = \langle 1_V \rangle \\
\downarrow \partial_0 & \uparrow \delta_{-1} \\
C_{-1}(\Gamma;F) = \langle \text{components} \rangle & C^{-1}(\Gamma;F) = \langle 1_{\text{component}} \rangle
\end{array}$$

FIGURE J.1. The chain and cochain complexes of a graph over a field F. [[LABEL F:20141017chainmaps]]

So, $\delta 1_W = 0$. □

Thus $\{\delta 1_v : v \in W,\ W = V(\text{component})\}$ is dependent. But if we omit the root r, the edges of $E(r, W \setminus r)$ don't zero out. In fact, for any X such that $\varnothing \subset X \subset W$,

$$\delta 1_X = \sum_{v \in X} \delta 1_v = 1_{\vec{E}(X, W\setminus X)} \neq 0$$

because $\vec{E}(X, W \setminus X) \neq \varnothing$. So $\{\delta 1_v : v \neq \text{root of a component}\}$ is a basis of B^1.
[This will be (or is) part of a theorem. TZ will take care of it.]

***In class we mentioned the following four theorems: October 17, 2014: Melissa Fuentes

Theorem J.9. $B_0 = \left\langle \sum_{v\in W} v \in C_0 : W = V(\text{component}) \right\rangle^{\perp} = \{\sum_{v \in \text{component}} \alpha_v v\} = Z_0$.

Theorem J.10. $Z_1 = \left\langle \sum_{e \in C} 1_{\vec{C}}(\vec{e})e : C \in \mathcal{C} \right\rangle = B_1$.

Theorem J.11. $Z^0 = \langle 1_W \in C^0 : W = V(\text{component}) \rangle$.

Theorem J.12. $B^1 = \langle 1_{E(v,v^c)} \rangle$.

Let C^k be the space of linear functionals on C_k. That means for any $(\alpha, \varphi) \in C_k \times C^k$, $(\alpha, \varphi) \mapsto \varphi(\alpha) \in F$. For example, when $k = 0$, the pair $(w, 1_v) \mapsto 1_v(w) = \delta_{vw}$. Thus, we have the following theorem.

Theorem J.13. *[[LABEL T:20141017 zerocyclespace]]*
$Z_0 = \{\alpha \in C_0 : 1_W(\alpha) = 0\ \forall\ W = V(\text{component})\} = \{\alpha \in C_0 : \varphi(\alpha) = 0\ \forall\ \varphi \in B^0\}$.

The last part of Theorem J.13 [**USE LABEL, not number A.5**] says that B^0 is $\text{Ann}(Z_0)$, the annihilator of Z_0 in $C_0* = C^0$. We will denote by $\text{Ann}(Z_0)$ [**denote** $\text{Ann}(Z_0)$ **to be IS BAD GRAMMAR.**] the set of all functions φ of C_0 such that $\varphi(Z_0) = 0$, that is, $Z_0 \subseteq \text{Ker}(\varphi)$. If we identify $C_0 \leftrightarrow C^0$, then notice that we can identify $\text{Ann}(Z_0) \leftrightarrow Z_0^{\perp}$. This setup, but with $C^k = C_k^*$ generalizes better to rings (e.g. $\mathbb{Z}$, $\mathbb{Z}_n$).

We go on to prove some useful properties of the function $1_{\vec{E}(X,X^c)} : E \to F$.

Lemma J.14. *[[LABEL L:20141017 indcutspace]] Suppose $v, w \in V$. Then* $1_{\vec{E}(v,v^c)} + 1_{\vec{E}(w,x^c)} = 1^c_{\vec{E}(\{v,w\},\{v,w}$$)\}$.

Proof. Recall that for $X \subseteq V$ we have $1_{\vec{E}(X,X^c)} = \sum_{v\in X} 1_{\vec{E}(v,v^c)}$. Thus, simply let $X = \{v, w\}$ to obtain the result. □

Lemma J.15. *[[LABEL L:20141017 indcutspace2]] If V(component) $= W = \{v_1, \ldots, v_k\}$ then $1_{\vec{E}(W,W^c)} = 0$, that is, $\sum_{v\in W} 1_{\vec{E}(v,v^c)} = 0$.*

Proof. This is true since $1_{\vec{E}(v,v^c)} = 0$ for every $v \in W$. □

Lemma J.16. *[[LABEL L:20141017 kernelindicator]] $1_W(v) = 1$ if and only if $v \notin \mathrm{Ker}(1_W)$ if and only if $v \in W$.*

Proof. Recall that

$$1_W(v) := \begin{cases} 1 & : \text{if } v \in W, \\ 0 & : \text{if } v \notin W. \end{cases}$$

Since $\mathrm{Ker}(1_W) = \{v \in V : 1_W(v) = 0\}$, it is clear that $1_W(v) = 1$ if and only if $v \notin \mathrm{Ker}(1_W)$. Note that $\mathrm{Ker}(1_W) = \{v \in V : v \notin W\}$. Hence, $v \notin \mathrm{Ker}(1_W)$ if and only if $v \in W$. □

J.1.1. *The cycle and cut spaces over a field – Old draft.* [[LABEL 1.fieldcyclecut]]

[THIS WILL BE INTEGRATED INTO THE PRECEDING PART OF THE SUBSECTION]

Oct 3 (draft): Peter Cohen & T.Z.

Now we let F be any field.

Definition J.6. A *directed circle* is a circle given a direction. This is separate from any orientation of the edges of the circle. Generally, any walk may be directed.

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

Definition J.7. 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 $I_C(e) = 0$. This definition applies as well to any directed trail. For a directed walk W, if an edge is repeated its values are summed to get $I_W(e)$.

Reversing the direction of a 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.8. The *cycle space* of Γ over F, $Z_1(\Gamma; F)$, is the span of the indicator vectors of the circles in Γ.

Definition J.9. A *directed cut*, denoted by $\vec{E}(X, Y)$, is a cut with a direction from X to Y.

Lemma J.17 (2)**.** *[[LABEL L:1003 2]] $I_C \cdot I_D = 0$ for any directed walk W and directed cut D.*

Proof. A closed directed walk W (not necessarily a circle) that crosses a cut (i.e., traverses it from one side to the other) must then return to the first side by another edge. In fact, each time the walk crosses from the first side to the second, the next of its edges that crosses the cut must be in the opposite direction. In each such pair $\{e, f\}$ of crossing edges, one is directed by W with the cut D and one is directed against it. Compare the signs: If e and f are both oriented with W, then $I_W(e) = I_W(f) = 1$ and $I_D(e) = -I_D(f)$, so $I_W(e)I_D(e) + I_W(f)I_D(f) = 0$. If we change the orientation of e, for instance, we negate both $I_W(e)$ and $I_D(e)$; that does not change the sum. Since $I_W \cdot I_D$ is a sum of such pair sums of products, it totals to zero. □

Theorem J.18. *[*[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**] □

J.2. **The binary cycle and cut spaces.** [[LABEL 1.binarycyclecut]]

A binary vector space—i.e., a vector space over the two-element field $\mathbb{F}_2$ (or $\mathbb{Z}_2$, if you prefer)—with a basis indexed by a set S can be treated as the power set $\mathcal{P}(S)$ with symmetric difference as addition. (That is why I write $\oplus$ for symmetric difference.) A vector $a \in \mathbb{F}_2^S$ is determined by its support; addition of vectors corresponds to symmetric difference of their supports; and there is no real scalar multiplication when you can multiply only by 0, which simply eliminates what you multiply, and 1, which has no effect. In particular, a linear combination $\sum_i \alpha_i a_i$ of vectors $a_i \in \mathbb{F}_2^S$ with coefficients $\alpha_i \in \mathbb{F}_2$ is algebraically the same as the sum $\sum_{i:\alpha_i=1} A_i$ where $A_i = \operatorname{supp} a_i$. All this implies that binary graphic spaces are truly combinatorial. That may be why most graph theorists who use graphic spaces use the binary ones only. There's justice to that, but I prefer to also put them in the setting of spaces over arbitrary fields; hence the first treatment of graphic spaces was a general one, over an arbitrary field. Now we look into the combinatorial development of binary spaces. 2014 Oct 19: T.Z.

The binary edge space.

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. Oct 3: Peter Cohen, T.Z.

Definition J.10. 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 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.

The first essential subspace of the binary edge space is the cycle space. Sept 15: Yash Lodha

Definition J.11. [[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.19. [[LABEL P:1003z1even]] *The binary cycles are the even-degree subsets of E.*

In fact, the real proposition is stronger; it has Proposition J.19 as an immediate corollary.

Proposition J.20. [[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.12. [[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.21. [[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, though.)

Proof. Let C be any circle, with cotree edges $e_1, e_2, \ldots, e_k$. Then $A := C \oplus \bigoplus_{i=1}^k C_T(e_i) \subseteq T$; that is, A is (the edge set of) a forest. Being a set sum of circles, A has even degree everywhere; but a nontrivial forest has at least one monovalent vertex (Theorem B.6(c)), which is a contradiction unless $A = \varnothing$. But that implies $C = \bigoplus_{i=1}^k C_T(e_i)$.

Suppose C is a sum $\bigoplus_{i=1}^j C_T(f_i)$. The cotree edges in that sum are $f_1, \ldots, f_j$, so they must be the cotree edges of C; that proves uniqueness. □

The binary cut space.

Sept 26: Yash Lodha

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.

Definition J.13. [[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.14. [[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.15. [[LABEL D:1003binarycutspace]] The *binary cut space* of Γ, written $B^1(\Gamma; \mathbb{F}_2)$, is the set $\{\text{cuts}\} \cup \varnothing$ in $\mathcal{P}(E) \cong \mathbb{F}_2^E$.
[MUST BE UPDATED.]

[THE FOLLOWING THEOREM MUST BE SOMEWHERE IN HERE.]

Theorem J.22. *[[LABEL T:1003cuts-space]] The set of cuts, with the empty set, is a subspace of $\mathcal{P}(E)$; it corresponds to the subspace $B^1(\Gamma; \mathbb{F}_2)$ in $\mathbb{F}_2 E$.*

Proof. □

Proposition J.23. *[[LABEL P:1003cutbonds]] Every cut is a disjoint union of bonds in a unique way.*

This is remarkably similar to its dual, Proposition J.20, but the uniqueness is a difference; it fails to hold true in Proposition J.20.

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.

Here is a list of the principal properties of the binary cycle and cut spaces, other than those already mentioned. Oct 3 (draft): Peter Cohen

Theorem J.24. *[[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 cuts, hence by the vertex bonds.*
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.

Readings and Bibliography

References in the form Name (yearx) will be found in:

[BSG] Thomas Zaslavsky, A mathematical bibliography of signed and gain graphs and allied areas. *Electronic J. Combin.*, Dynamic Surveys in Combinatorics (1998), No. DS8 (electronic). MR 2000m:05001a. Zbl 898.05001.

The most up-to-date version is on my Web site at

`http://www.math.binghamton.edu/zaslav/Bsg/`

A. Background

[Pearls] Nora Hartsfield and Gerhard Ringel, *Pearls in Graph Theory.* Academic Press, Boston, 1990. MR 91f:05001. Zbl 703.05001. Reprinted by Dover, Mineola, N.Y., 2003.

Ch. 1, §§ 2.1–2, § 3.1, Ch. 8 for elementary background in graph theory. Very readable. Not always precisely correct, so make sure you understand the proofs.

[BM] J.A. Bondy and U.S.R. Murty, *Graph Theory with Applications.* American Elsevier (North-Holland), New York, 1976. MR 54 #117.

An advanced introduction, conceptually oriented. The best graph theory textbook to find background for this course. In connection with our course, see especially Chapters 1–3, 6, 8, 9, 12, and some of 10. The table of contents and links to the entire book are on the Web at

`http://www.ecp6.jussieu.fr/pageperso/bondy/bondy.html`

[Tbook] W.T. Tutte, *Graph Theory.* Encyc. Math. Appl., Vol. 21. Addison–Wesley, Reading, Mass., 1984. MR 87c:05001. Zbl 554.05001. Reprinted by Cambridge Univ. Press, Cambridge, Eng., 2001. MR 2001j:05002. Zbl 964.05001.

A rigorous advanced introduction from the viewpoint of one of the great graph theorists. A large part of Tutte's research, presented in readable form—if you pay close attention. Includes Tutte's definition of connectivity, his decomposition theory for graphs that are not 3-connected, the basics of his treatment of graph invariants, and much more.

[PG] Martin Charles Golumbic, *Algorithmic Graph Theory and Perfect Graphs.* Academic Press, New York, 1980. MR 81e:68081. Zbl 541.05054.

Readable treatment of special topics, e.g., chordal graphs and comparability graphs.

[Cox] H.S.M. Coxeter, *Introduction to Geometry*, second ed. Wiley, New York, 1969. MR 49 #11369. Zbl 181, 481a (e: 181.48101).

Ch. 12–14, §§ 15.1–3. Basic geometrical background for use in the course.

[Oxley] J.G. Oxley, *Matroid Theory.* Oxford Univ. Press, Oxford, 1992. MR 94d:05033. Zbl 784.05002.

The standard place to learn matroid theory, to which I refer occasionally.

[FUTA] Thomas Zaslavsky, *Facing Up to Arrangements: Face-Count Formulas for Partitions of Space by Hyperplanes.* Mem. Amer. Math. Soc., No. 154 (= Vol. 1, Issue 1). American Mathematical Society, Providence, R.I., 1975. MR 50 #9603. Zbl 296.50010.

Real hyperplane arrangements.

[AH] Peter Orlik and Hiroaki Terao, *Arrangements of Hyperplanes*. Grundl. math. Wiss., Vol. 300. Springer-Verlag, Berlin, 1992. MR 94e:52014. Zbl 757.55001.

Complex hyperplane arrangements.

B. Signed Graphs

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

Basic (but not necessarily elementary) concepts and properties, including the frame matroid (that is optional), examples, and incidence and adjacency matrices (Section 8).

[SGC] T.Z., Signed graph coloring. *Discrete Math.* **39** (1982), 215–228. MR 84h:05050a. Zbl 487.05027.

Basic (but not necessarily elementary) concepts and properties of coloring and the chromatic polynomial.

[OSG] T.Z., Orientation of signed graphs. *European J. Combin.* **12** (1991), 361–375. MR 93a:05065. Zbl 761.05095.

Oriented signed graphs (bidirected graphs) and their hyperplane geometry.

[MTS] T.Z., Matrices in the theory of signed simple graphs. *Proc. International Conference on Discrete Mathematics* (Mysore, 2009), to appear.

Introductory survey.

[CW] Beifang Chen and Jue Wang, The flow and tension spaces and lattices of signed graphs. *European J. Combin.* **30** (2009), 263–279.

C. Geometry of Signed Graphs

[GRS] T.Z., The geometry of root systems and signed graphs. *Amer. Math. Monthly* **88** (1981), 88–105. MR 82g:05012. Zbl 466.05058.

A readable introduction to some of the connections between graph theory and geometry.

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

A classic research paper in the background of part of our topic. Not introductory.

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

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

[E8] G.R. Vijayakumar, Representation of signed graphs by root system E_8. *Graphs Combin.* 10 (1994), 383–388. MR 96a:05128. Zbl 821.05040.

These are the three main papers on vertex representation of signed graphs.

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

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

Textbook. Ch. 10 on strongly regular graphs (§§ 10.1–2), Ch. 11 on two-graphs and equiangular lines, and Ch. 12 (and § 1.7) on line graphs (§§ 12.1–3). Presents much

of [CGSS], with its background, in a somewhat more systematic but more advanced way.

D. Gain Graphs

[BG1] T.Z., Biased graphs. I. Bias, balance, and gains. *J. Combin. Theory Ser. B* **47** (1989), 32–52. MR 90k:05138. Zbl 714.05057.
§ 5: Fundamentals of gain graphs.

[BG3] T.Z., Biased graphs. III. Chromatic and dichromatic invariants. *J. Combin. Theory Ser. B* **64** (1995), 17–88. MR 96g:05139. Zbl 857.05088.
§ 5: General theory of coloring gain graphs.

E. Biased Graphs and Gain Graphs

[BG1] T.Z., Biased graphs. I. Bias, balance, and gains. *J. Combin. Theory Ser. B* **47** (1989), 32–52. MR 90k:05138. Zbl 714.05057.
Fundamentals of gain graphs and biased graphs.

[BG2] T.Z., Biased graphs. II. The three matroids. *J. Combin. Theory Ser. B* **51** (1991), 46–72. MR 91m:05056. Zbl 763.05096.
The closure operations that are basic to the theory, in terms of frame matroids (called bias matroids), in §§ 2, 3. Important open problems.

[BG3] T.Z., Biased graphs. III. Chromatic and dichromatic invariants. *J. Combin. Theory Ser. B* **64** (1995), 17–88. MR 96g:05139. Zbl 857.05088.
Chromatic polynomial et al., with and without colorings.

[BG4] T.Z., Biased graphs IV: Geometrical realizations. *J. Combin. Theory Ser. B* **89** (2003), no. 2, 231–297. MR 2005b:05057. Zbl 1031.05034.
Extensive treatment of the geometry of gain graphs. Fundamentals in §§ 2, 4. More advanced topics throughout, especially in § 7.

F. General References

[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).