2 Graphic, linear, and transversal matroids

Matroids underlie several mathematical structures, such as graphs, and a set of vectors in a vector space. In this section, we see how matroids arise in these settings.

First, we need to recall several notions that you will have seen before, relating to linear algebra and graphs. In particular, graph terminology can vary, so watch out for any differences from what you are used to!

Field preliminaries

First, we make the notion of a *multiset* precise. This is a set with repetition of elements allowed. More formally:

Definition 2.1. A multiset is a pair (S, ρ), where S is a set and ρ is a function from S to the positive integers. If s is in S, then $\rho(s)$ tells us how many times s appears in the multiset. If (S, ρ) and (S', ρ') are multisets, then we say that (S', ρ') is a *subset* of (S, ρ) if $S' \subseteq S$ and $\rho'(s) \leq \rho(s)$ for all $s \in S'$.

When (S, ρ) is a multiset, we abuse notation and write S to represent the multiset.

Recall that if F is a set, then a *binary operation* on F is a function,

$$*\colon F \times F \to F,$$

from the Cartesian product $F \times F$, to F. If a and b are in F, then $a * b$ denotes the image of (a, b) under $*$. Familiar examples of binary operations include addition and multiplication. A *field* $\mathbb{F} = (F, +, \cdot, 0, 1)$ consists of a set F, with binary operations $+$ (*addition*) and $\cdot$ (*multiplication*), such that F has distinct members 0 (the *additive identity*) and 1 (the *multiplicative identity*) such that the following statements are true for any $a, b, c \in F$.

F1. $a + (b + c) = (a + b) + c$.

F2. $a + 0 = a = 0 + a$

F3. There is an element $-a \in F$ (the *additive inverse* of a) such that $a + (-a) = 0 = (-a) + a$.

F4. $a + b = b + a$.

F5. $a \cdot (b \cdot c) = (a \cdot b) \cdot c$.

F6. $a \cdot 1 = a = 1 \cdot a$.

F7. If $a \neq 0$, then there is an element $a^{-1} \in F$ (the *multiplicative inverse* of a) such that $a \cdot a^{-1} = 1 = a^{-1} \cdot a$.

F8. $a \cdot b = b \cdot a$.

F9. $a \cdot (b + c) = (a \cdot b) + (a \cdot c)$.

We refer to the members of F as the *elements* of the field.

To simplify notation, we will often use $\mathbb{F}$ to stand both for the field, and for the set of elements in that field (that is, we write $\mathbb{F}$ in place of F). We also refer to the additive and multiplicative identities of $\mathbb{F}$ as 'zero' and 'one' respectively. Moreover, we omit the $\cdot$ sign, so $a \cdot b$ is written ab. Note that $+$ and $\cdot$ may not be the same as addition and multiplication in, say, the real numbers. For example, if we let $F = \{0, 1\}$, and define $+$ and $\cdot$ to be addition and multiplication modulo 2 (so that $1 + 1 = 0$), then $(F, +, \cdot, 0, 1)$ is a field. In general, the integers modulo a prime p form a field, which we denote $\mathrm{GF}(p)$. The rational numbers, real numbers, and complex numbers are other examples of fields. Moreover, there are other fields, with only finitely many elements, that do not arise from considering the integers modulo a prime. We will not discuss their structure here, however, we will note the following fundamental theorem from abstract algebra.

Theorem 2.2. *Suppose that q is a positive integer. There is a field with q elements if and only if q is a power of a prime. Moreover, if q is a power of a prime, then all fields with q elements are isomorphic.*

Two fields $(F, +, \cdot, 0, 1)$ and $(F', \oplus, \odot, 0', 1')$ are *isomorphic* if there is a bijection $\psi \colon F \to F'$ such that $\psi(a + b) = \psi(a) \oplus \psi(b)$ and $\psi(a \cdot b) = \psi(a) \odot \psi(b)$ for all $a, b \in F$.

Theorem 2.2 tells us that there are (for example) no fields of size six, ten, or twelve. On the other hand, there are unique (up to isomorphism) fields of size two, three, and four. If q is a power of a prime, then $\mathrm{GF}(q)$ denotes the field with q elements.

If $\mathbb{F} = (F, +, \cdot, 0, 1)$ is a field, and $F' \subseteq F$ has the property that $\mathbb{F}' = (F', +, \cdot, 0, 1)$ is a field, then we say $\mathbb{F}'$ is a *subfield* of $\mathbb{F}$.

Exercise 2.3. Prove that the intersection of any family of subfields is another subfield.

Let $\mathbb{F}$ be a field. Because of the previous exercise, the intersection of all subfields of $\mathbb{F}$ is itself a subfield. We call this the *prime subfield* of $\mathbb{F}$. The prime subfield consists of the numbers produced by repeatedly adding 1 to itself and repeatedly subtracting 1 from 0, and then taking the multiplicative inverses of all the numbers we construct in this way. One of the key ingredients in the proof of Theorem 2.2 is showing that the prime subfield is either isomorphic to the field of rational numbers, or to the integers modulo a prime number. If the prime subfield of $\mathbb{F}$ is isomorphic to the integers modulo p, for some prime p, then we say that the *characteristic* of $\mathbb{F}$ is p. In this case, adding together p copies of any number in $\mathbb{F}$ produces 0. If the prime subfield is isomorphic to the field of rationals, we say that $\mathbb{F}$ has characteristic zero.

Now suppose that $\mathbb{F}$ is a field, and that n is some positive integer. Then $\mathbb{F}^n$ is the set of ordered n-tuples of the form $(v_1, v_2, \ldots, v_n)$, where each v_i is an element of $\mathbb{F}$. If $\mathbf{u} = (u_1, \ldots, u_n)$ and $\mathbf{v} = (v_1, \ldots, v_n)$ are members of $\mathbb{F}^n$, then $\mathbf{u} + \mathbf{v}$ is defined to be

$$(u_1 + v_1, u_2 + v_2, \ldots, u_n + v_n).$$

Similarly, if α is some element of $\mathbb{F}$, then $\alpha\mathbf{u}$ is

$$(\alpha u_1, \alpha u_2, \ldots, \alpha u_n).$$

The members of $\mathbb{F}^n$ are called *vectors*. The set of vectors, along with these operations of *addition* and *scalar multiplication*, forms a *vector space*.

We use $\mathbf{0}$ to denote the member $(0, 0, \ldots, 0)$ of $\mathbb{F}^n$. Suppose that $V = \{\mathbf{v}_1, \ldots, \mathbf{v}_t\}$ is a multiset of vectors in $\mathbb{F}^n$. We say that a *linear combination* of V is a sum $\alpha_1\mathbf{v}_1 + \alpha_2\mathbf{v}_2 + \cdots + \alpha_t\mathbf{v}_t$ where each $\alpha_1, \ldots, \alpha_t$ is in $\mathbb{F}$ and at least one of the elements α_i is non-zero. We say that V is *linearly dependent* if there exists a linear combination of V that is equal to $\mathbf{0}$. If V is not linearly dependent then it is *linearly independent*. Obviously a set that contains $\mathbf{0}$ is linearly dependent. A multiset that contains more than one copy of any vector is also linearly dependent. If $\mathbf{v}$ is a member of $\mathbb{F}^n$, and $\mathbf{v}$ can be expressed as a linear combination

$$\mathbf{v} = \alpha_1\mathbf{v}_1 + \alpha_2\mathbf{v}_2 + \cdots + \alpha_t\mathbf{v}_t$$

for some sequence $\alpha_1, \ldots, \alpha_t$ of elements of $\mathbb{F}$, then we say that V *spans* $\mathbf{v}$.

Representable matroids.

Theorem 2.4. *Let A be a matrix over a field $\mathbf{F}$, where the columns are labelled by elements of a set E. Let $\mathcal{I}$ be the set of subsets X of E for which the columns labelled by X form a linearly independent set. Then there is a matroid on ground set E whose family of independent sets is $\mathcal{I}$.*

Proof. The empty set of vectors is trivially linearly independent. It is also easy to see that any subset of a linearly independent set is itself linearly independent (since if a set of vectors has no linear combination that is $\mathbf{0}$, then neither will a subset of those vectors). Therefore **I1** and **I2** are satisfied. It remains to prove that **I3** holds. Suppose that I_1 and I_2 are independent subsets of E, and that $|I_2| < |I_1|$. We abuse notation and also use I_1 and I_2 to refer to the set of vectors labelled by I_1 and I_2, respectively. Assume for a contradiction that there is no element $\mathbf{v} \in I_1 - I_2$ such that $I_2 \cup \mathbf{v}$ is linearly independent.

Let $I_2 = \{\mathbf{v}_1, \ldots, \mathbf{v}_t\}$, and let $\mathbf{v}$ be an arbitrary vector in $I_1 - I_2$. Then $I_2 \cup \mathbf{v}$ is linearly dependent, by our assumption. Therefore

$$\alpha_1\mathbf{v}_1 + \alpha_2\mathbf{v}_2 + \cdots + \alpha_t\mathbf{v}_t + \alpha\mathbf{v} = \mathbf{0}$$

for some elements (not all zero) $\alpha_1, \ldots, \alpha_t, \alpha$ in $\mathbb{F}$. If $\alpha = 0$, then I_2 is linearly dependent, which is a contradiction. Therefore $\alpha \neq 0$, so α has a multiplicative inverse, α^{-1}, and

$$\mathbf{v} = (-\alpha^{-1}\alpha_1)\mathbf{v}_1 + \cdots + (-\alpha^{-1}\alpha_t)\mathbf{v}_t.$$

Thus I_2 spans every vector in $I_1 - I_2$. Obviously I_2 spans every vector in I_2, so I_2 spans every vector in $I_1 \cup I_2$.

We have shown that there is at least one subset of $I_1 \cup I_2$ with size t that spans $I_1 \cup I_2$. Suppose that U is such a set, and that U has been chosen so that $|U \cap I_1|$ is as large as possible. Let $U = \{\mathbf{u}_1, \ldots, \mathbf{u}_t\}$. Since $|I_1| > |I_2| = t = |U|$, there is some vector $\mathbf{v} \in I_1 - U$. Then $\mathbf{v}$ is spanned by U, so

$$\mathbf{v} = \alpha_1\mathbf{u}_1 + \cdots + \alpha_t\mathbf{u}_t$$

for some elements $\alpha_1, \ldots, \alpha_t \in \mathbb{F}$. Therefore

$$\alpha_1\mathbf{u}_1 + \cdots + \alpha_t\mathbf{u}_t + (-1)\mathbf{v} = \mathbf{0}. \tag{2.1}$$

If $\alpha_1 = \alpha_2 = \cdots = \alpha_t = 0$, then $\mathbf{v} = \mathbf{0}$. But this is contradictory, as no linearly independent set can contain the zero vector, and $\mathbf{v}$ is contained in the linearly independent set I_1. Hence, at least one of $\alpha_1, \ldots, \alpha_t$ is non-zero. If $\alpha_i = 0$ whenever $\mathbf{u}_i \notin I_1$, then we could remove the terms in (2.1) that correspond to vectors not in I_1, and deduce that I_1 contains a linearly dependent set. This is a contradiction, so there must be some i such that $\mathbf{u}_i \notin I_1$ and $\alpha_i \neq 0$. This means that, by relabelling, we can assume that α_1 is non-zero, and that $\mathbf{u}_1$ is not in I_1.

Now

$$\mathbf{u_1} = (-\alpha_1^{-1}\alpha_2)\mathbf{u}_2 + \cdots + (-\alpha_1^{-1}\alpha_t)\mathbf{u}_t + \alpha_1^{-1}\mathbf{v}$$

so $\mathbf{u}_1$ is spanned by $(U - \mathbf{u}_1) \cup \mathbf{v}$. Since every vector in $I_1 \cup I_2$ is spanned by U, it is easy to see that this means that every vector in $I_1 \cup I_2$ is spanned by $(U - \mathbf{u}_1) \cup \mathbf{v}$. But $|(U - \mathbf{u}_1) \cup \mathbf{v}| = t$, and $(U - \mathbf{u}_1) \cup \mathbf{v}$ intersects I_1 in one more element than U does. This contradicts our choice of U.

Therefore **I3** is satisfied, and $\mathcal{I}$ is the family of independent sets of a matroid, as desired. □

Let $\mathbb{F}$ be a field, and let A be a matrix with entries from $\mathbb{F}$, whose columns are labelled by elements of a set X. Let m be the number of rows in A. Then Theorem 2.4 says that there is a matroid with X as its ground set, whose independent sets are the subsets of X that correspond to linearly independent sets of columns of A. We use the notation $M[A]$ to denote this matroid, and we say that $M[A]$ is *representable* over $\mathbb{F}$, or $\mathbb{F}$-*representable*. A matroid is *linear* (or *representable*) if it is $\mathbb{F}$-representable for some field $\mathbb{F}$. A matroid is *binary* if it is representable over GF(2) and it is *ternary* if it is representable over the field GF(3). Characterizing the $\mathbb{F}$-representable matroids, for various fields $\mathbb{F}$, is one of the oldest and most difficult problems in matroid theory.

Note that any multiset of vectors from a vector space $\mathbb{F}^m$, for some field $\mathbb{F}$, can be put in a matrix A (over $\mathbb{F}$), and given some column labels. So, for any multiset of vectors, there is a corresponding matroid.

Example. Figure 5 shows a matrix over the finite field GF(2) and a geometric representation of the corresponding binary matroid. You should verify that a 3-element set of points in the matroid is a basis if and only if the corresponding set of columns is linearly independent in the matrix.

This rank-3 matroid is called the *Fano matroid*, and it is denoted by F_7. It plays a very important role in the study of binary matroids. ♢

Exercise 2.5. Prove that $U_{2,4}$ is not a binary matroid.

Graphic preliminaries

Definition 2.6. A graph G consists of a set V, a set E, and an incidence function ϕ that maps elements of E to a subset of V of size one or two.

An equivalent viewpoint of graphs is a pair $G = (V, E)$ where V is a set and E is a multiset such that each member of E is a subset of V of size one

$$\begin{array}{ccccccc} a & b & c & d & e & f & g \end{array}$$
$$\begin{bmatrix} 1 & 0 & 0 & 1 & 1 & 0 & 1 \\ 0 & 1 & 0 & 1 & 0 & 1 & 1 \\ 0 & 0 & 1 & 0 & 1 & 1 & 1 \end{bmatrix}$$

A

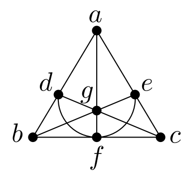

$M[A] \cong F_7$

Figure 5: A matrix over GF(2) and its matroid.

or two. (There is no fundamental difference to these viewpoints, but there are some pros and cons to either approach. We go with the first option, though introduce terminology that often avoids explicitly referring to the incidence function ϕ.)

We say that V and E are the *vertex set* and the *edge set* of G, respectively. If G is a graph, then $V(G)$ and $E(G)$ denote its vertex set and edge set respectively. If $e \in E$ and $\phi(e) = \{v\}$, then e is called a *loop* (at v). If $e \in E$ and $\phi(e) = \{u, v\}$ for distinct $u, v \in V$, then e is a *non-loop* edge, and e *joins* u and v. We say that u and v are the *ends* of e. If e and e' are non-loop edges that join the same pair of vertices, then we say that e and e' are *parallel edges*. Although the study of infinite graphs is very well developed, we will always assume that E and V are finite.

A graph is often represented by a drawing in which the vertices are dots and an edge e with $\phi(e) = \{u, v\}$ is represented by a line joining u and v. Note that if there are parallel edges e and e' joining u and v, then these are represented by two lines that join u and v. A loop at v is represented by a small circle touching the vertex v. There may be more than one loop at a vertex.

Example. The drawing in Figure 6 represents a graph G with $V(G) = \{v_1, \ldots, v_4\}$, $E(G) = \{e_1, \ldots, e_7\}$, and incidence function ϕ such that $\phi(e_1) = \{v_1, v_2\}$, $\phi(e_2) = \{v_2, v_3\}$, $\phi(e_3) = \{v_3, v_4\}$, $\phi(e_4) = \{v_1, v_4\}$, $\phi(e_5) = \{v_1, v_4\}$, $\phi(e_6) = \{v_1, v_3\}$, $\phi(e_7) = \{v_3\}$. $\diamondsuit$

If $v \in \phi(e)$ for some $v \in V$ and $e \in E$, then we say that e and v are *incident*. A vertex that is not incident with any edges is an *isolated* vertex. If u and v are distinct vertices and there is an edge e such that $\phi(e) = \{u, v\}$, then we say that u and v are *adjacent*.

Next we define *subgraphs*. Assume that $G = (V, E)$ is a graph and that $E' \subseteq E$. Let

$$V' = \bigcup_{e \in E'} \phi(e)$$

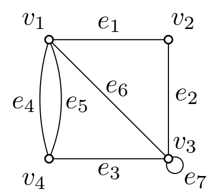

Figure 6: An example of a graph drawing.

be the set of all vertices that are incident with an edge in E'. Then $G[E']$ is the graph (V', E'). Similarly, if V' is a subset of V, then let

$$E' = \{e \in E\colon \phi(e) \subseteq V'\}$$

be the set of edges that are only incident with vertices in V'. Now $G[V']$ denotes the graph (V', E'). This is sometimes called the *subgraph induced by* V'.

Example. If G is the graph shown in Figure 6, then $G[\{e_4, e_6, e_7\}]$ and $G[\{v_1, v_3, v_4\}]$ are the subgraphs shown in Figure 7. $\diamondsuit$

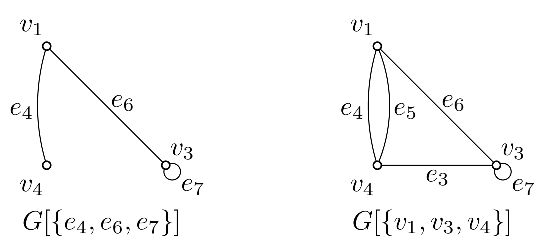

Figure 7: Two subgraphs

A *walk* in a graph G is an alternating sequence of vertices and edges of G,

$$v_0, e_0, v_1, e_1, v_2, \ldots, v_{t-1}, e_{t-1}, v_t,$$

such that $\phi(e_i) = \{v_i, v_{i+1}\}$ for $0 \le i \le t-1$. We say that this is a walk from v_0 to v_t. If the vertices $v_0, \ldots, v_t$ are pairwise distinct, then the walk is said to be a *path*. If $v_1, \ldots, v_t$ are pairwise distinct, and $v_0 = v_t$, then the walk is said to be a *cycle*. Note that if e is a loop at the vertex v, then v, e, v is a cycle.

Exercise 2.7. Let G be a graph with vertices u and v. Prove that if there is a walk from u to v, then there is a path from u to v.

Note that an implication of Exercise 2.7 is that if u and v are vertices in a connected graph G, then there is a path from u to v.

We define a relation $\sim$ on the vertices of a graph. Let u and v be vertices. Then $u \sim v$ if there is a walk from u to v.

Exercise 2.8. Let G be a graph with vertex set V. Prove that $\sim$ is an equivalence relation on V.

Let G be a graph. The equivalence classes of $\sim$ are called *connected components* of G. If G has only one connected component, then it is a *connected* graph. A graph that contains no cycles is called a *forest*. A connected forest is a *tree*.

Let G be a graph with vertex set V and edge set E, and let $v \in V$ be a vertex. The *degree* of $v \in V$ is given by the formula

$$|\{e \in E\colon v \in \phi(e),\ e \text{ is a non-loop edge}\}| + 2|\{e \in E\colon v \in \phi(e),\ e \text{ is a loop edge}\}|.$$

Intuitively, the degree of v is the number of edges incident with v, but where loops count for double.

Proposition 2.9. *Every tree with at least one vertex contains a vertex with degree at most one.*

Proof. Let G be a tree, and let $v_0, e_0, v_1, e_1, v_2, \ldots, v_{t-1}, e_{t-1}, v_t$ be a path in G with as many vertices as possible. If $t = 0$, then this means that G has no edges, or else we would choose a path with at least one edge. In this case, G has only one vertex, and this vertex has degree zero, so we are done. Therefore we assume that $t > 0$.

Now v_0 is incident with at least one edge, e_0, so v_0 has degree at least one. Assume that the degree of v_0 is greater than one. Since G has no loops, there is an edge, e, incident with v_0, such that e is not equal to e_0. Let the other vertex incident with e be v. If v is not in $v_1, \ldots, v_t$, then $v, e, v_0, e_0, v_1, \ldots, v_{t-1}, e_{t-1}, v_t$ is a longer path, so v is equal to v_i for some $i \in \{1, \ldots, t\}$. Now $v, e, v_0, e_0, v_1, \ldots, v_{i-1}, e_{i-1}, v_i$ is a cycle of G, which contradicticts that G is a tree. Therefore the degree of v_0 is exactly one. □

Let $G = (V, E)$ be a graph. A *forest of* G is a subset $F \subseteq E$ such that $G[F]$ is a forest. In other words, F is a forest if $G[F]$ contains no cycles.

Proposition 2.10. *Let F be a non-empty forest of the graph G. Let k be the number of components of $G[F]$. Then the number of vertices in $G[F]$ is $|F| + k$.*

Proof. The proof is by induction on $|F|$. Suppose $|F| = 1$. Then $G[F]$ contains exactly one component. As the edge in F is not a loop, the number of vertices in $G[F]$ is two, so the result holds when $|F| = 1$. We now assume that $|F|$ is greater than one.

Let F' be a maximal non-empty subset of F such that $G[F']$ is connected. Since $G[F']$ is connected and contains no cycle, it is a tree. By Proposition 2.9, we can choose a vertex v that has degree at most one in $G[F']$. As $G[F']$ is connected and contains at least one edge, it follows that v has degree exactly one in $G[F']$. Let f be the edge of $G[F']$ that is incident with v.

We claim that $G[F' - f]$ is connected. Let u and u' be two arbitrary vertices of $G[F' - f]$. Notice that v is not a vertex of $G[F' - f]$, so neither u nor u' is equal to v. As $G[F']$ is connected, there is a path P from u to u' in $G[F']$. The path P cannot contain v, or else it would be forced to use the edge f twice, in which case P is not a path. So the path P from u to u' does not use v, and it follows that P is a path from u to u' in $G[F' - f]$. Since there is a path between any two vertices of $G[F' - f]$, it follows that $G[F' - f]$ is connected, as claimed.

First assume F' is a single edge. Then $G[F - f]$ contains one fewer component than $G[F]$. By induction, the number of vertices in $G[F - f]$ is $|F - f| + (k - 1) = |F| + k - 2$. As $G[F - f]$ has two fewer vertices than $G[F]$, we deduce that the number of vertices in $G[F]$ is $|F| + k$, as required. Now assume that F' has more than one edge. Then the number of connected components of $G[F - f']$ is k. By induction the number of vertices in $G[F - f]$ is $|F - f| + k = |F| + k - 1$. As $G[F - f']$ has one fewer vertex than $G[F]$, since it does not contain v, the number of vertices in $G[F]$ is $|F| + k$, which completes the proof. ☐

Let G be a graph. If F is a forest of G and F is not properly contained in any other forest of G, then it is a *maximal forest*.

Proposition 2.11. *Let F be a maximal forest of the graph G, and let u and v be distinct vertices of G. If there is a walk from u to v in G, then there is a walk from u to v in $G[F]$.*

Proof. Suppose there is a walk from u to v in G. Assume for a contradiction that there is no walk from u to v in $G[F]$. Then u and v are not in the same component of $G[F]$. Let P be a path from u to v in G. Let the vertices of P be $v_0, \ldots, v_t$, where $u = v_0$ and $v = v_t$. Let i be the first index such that v_i is not in the same component of $G[F]$ as v_{i+1}. Let e be the edge of P that joins v_i to v_{i+1}. Then e is not in F. As F is a maximal forest, it follows that $G[F \cup e]$ contains a cycle. This cycle must use e, since $G[F]$ does not contain a cycle. This means that $G[F]$ contains a path from v_i to v_{i+1}, but this is impossible as these vertices are not in the same component of $G[F]$. This contradiction completes the proof. □

Corollary 2.12. *Let F be a maximal forest of the graph G. Let n be the number of vertices in G, and let k be the number of components in G. Then $|F| = n - k$.*

Proof. Note that the maximal forest F does not contain any loop edges, so we lose no generality by assuming that G has no loops. Let s be the number of isolated vertices in G. Each of these vertices is a component of G, so by deleting these vertices we obtain a graph G' with s fewer vertices and s fewer components. Let n' and k' be the number of vertices of G', and the number of components of G', respectively. Since $n' - k' = (n - s) - (k - s) = n - k$, and a maximal forest of G is also a maximal forest of G', it suffices to show the result holds for a graph with no isolated vertices. Henceforth, we assume that G has no isolated vertices. Now each vertex of G is joined to a different vertex by a walk in G. Therefore every vertex is joined to a different vertex by a walk in $G[F]$, by Proposition 2.11. Hence every vertex of G is also a vertex of $G[F]$; in particular, $G[F]$ has n vertices.

Let u and v be two distinct vertices of G. If they are not joined by a walk in G, then they are certainly not joined by a walk in the subgraph $G[F]$. On the other hand, if u and v are joined by a walk in G, then they are joined by a walk in $G[F]$, by Proposition 2.11 again. This shows that the equivalence classes defined by the relation $\sim$ are exactly the same for the vertices of G as they are for the vertices of $G[F]$; in particular, the number of components in $G[F]$ is k. Now, by Proposition 2.10, $|F| = n - k$, as required. □

Graphic matroids.

Now we can show that every graph gives rise to a matroid.

Theorem 2.13. *Let G be a graph with edge set E, and let*

$$\mathcal{B} = \{B \subseteq E \colon B \textit{ is a maximal forest of } G\}.$$

Then there is a matroid with ground set E whose family of bases is $\mathcal{B}$.

Proof. It suffices to show that $\mathcal{B}$ satisfies the axioms **B1** and **B2**. So first we check that $\mathcal{B}$ is non-empty. For any graph, the empty set is a forest, since it contains no cycles. Hence every graph contains at least one forest, and thus at least one maximal forest. This shows $\mathcal{B}$ is not empty, so **B1** holds.

Next we show that $\mathcal{B}$ obeys the axiom **B2**. Let B_1 and B_2 be maximal forests of G, and let x be an edge in $B_1 - B_2$. Assume that x joins the vertices u and v. If there were a path from u to v in the subgraph $G[B_1 - x]$, then this path, together with x, would form a cycle in $G[B_1]$, which contradicts that B_1 is a forest. Hence there is no such path, so there is no walk from u to v in $G[B_1 - x]$. Thus u and v are in different components of $G[B_1 - x]$.

Because B_2 is a maximal forest, and x is not in B_2, it follows that $B_2 \cup x$ is not a forest, and hence $G[B_2 \cup x]$ contains a cycle. This cycle must have x as an edge, or else $G[B_2]$ would contain a cycle. Therefore $G[B_2]$ contains a path from u to v. Let $v_0, v_1, \ldots, v_t$ be the vertices of such a path, where $u = v_0$ and $v = v_t$. Let i be the smallest integer such that v_i and v_{i+1} are not in the same component of $G[B_1 - x]$. Let y be an edge in B_2 that joins v_i to v_{i+1}. Note that $y \neq x$, since x is not in B_2. Furthermore, y is not in B_1, as v_i and v_{i+1} are not in the same component of $G[B_1 - x]$.

We will show that $(B_1 - x) \cup y$ is a maximal forest. For a contradiction, assume that $(B_1 - x) \cup y$ is not a forest. Then $G[(B_1 - x) \cup y]$ contains a cycle, and this cycle must contain y, since $G[B_1 - x]$ does not contain a cycle. Therefore $G[B_1 - x]$ contains a path from v_i to v_{i+1}. But v_i and v_{i+1} are not in the same component of $G[B_1 - x]$, so this is contradictory. Therefore $(B_1 - x) \cup y$ is a forest. We note that $|(B_1 - x) \cup y| = |B_1|$. By Corollary 2.12, each maximal forest of G has the same size, so $(B_1 - x) \cup y$ is a maximal forest, and **B2** is satisfied. □

Theorem 2.13 tells us that if $G = (V, E)$ is a graph, there is a corresponding matroid with E as its ground set and the maximal forests of G as its bases. We denote this matroid by $M(G)$, and we call it the *cycle matroid* of G. We say that a matroid M is *graphic* if there exists some graph G such that $M = M(G)$.

Exercise 2.14. Let $G = (V, E)$ be a graph, and let $X \subseteq E$. Prove that

(a) X is independent in $M(G)$ if and only if $G[X]$ is a forest, and

(b) X is a circuit of $M(G)$ if and only if X is a cycle in G.

Example. Figure 8 shows K_4, the complete graph on four vertices, and two geometric representations of its cycle matroid, $M(K_4)$. You should verify that a basis in the matroid corresponds to the edges of a maximal forest of K_4. In the second representation we have used a curved line (so it is no longer a representation in Euclidean space). Curved lines and planes are common features of matroid representations. ♢

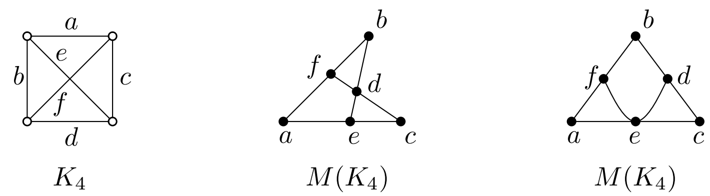

Figure 8: A graph and two geometric representations of its cycle matroid.

Exercise 2.15. Prove that the uniform matroid $U_{2,4}$ is not graphic.

Transversal matroids. In this section we introduce one more fundamental class of matroids (although we will spend less time considering this family). Let G be a *bipartite* graph with *bipartition* (E, A). This means that the vertex set of G is $E \cup A$, where $E \cap A = \emptyset$, and every edge of G joins a vertex in E to a vertex to A. If $X \subseteq E$, then we say that X is *matchable* if there is a set T of edges of G such that every vertex in X is incident to an edge of T, and no vertex of G is incident to more than one edge of T. We say that T *certifies* that X is matchable.

Exercise 2.16. Let G be a graph with no loops where no vertex is incident to more than two edges. Prove that every component of G is either a path or a cycle.

Theorem 2.17. *Let G be a bipartite graph with bipartition (E, A). Let $\mathcal{I}$ be the collection of matchable subsets of E. Then there is a matroid with ground set E whose family of independent sets is $\mathcal{I}$.*

Proof. The empty set of edges certifies that the empty subset is matchable. Therefore **I1** is satisfied. If T is a set of edges certifying that $I_1 \subseteq E$ is matchable, then T also certifies that any subset of I_1 is matchable, so **I2** is satisfied.

To prove that **I3** holds, we let I_1 and I_2 be two matchable subsets of E, and assume that $|I_2| < |I_1|$. Let T_1 and T_2 be sets of edges that certify I_1 and I_2, respectively, are matchable. By removing redundant edges from T_1 and T_2, we can assume that $|I_1| = |T_1|$ and $|I_2| = |T_2|$. Consider the subgraph $G[T_1 \cup T_2]$. Note that G has no loops. Any vertex in $G[T_1 \cup T_2]$ is incident with at most one edge in T_1 and at most one edge in T_2. Therefore every component of $G[T_1 \cup T_2]$ is either a path or a cycle. This means that if S is the set of edges of a component of $G[T_1 \cup T_2]$, then $-1 \leq |S \cap T_1| - |S \cap T_2| \leq 1$. For at least one component, S must satisfy $|S \cap T_1| - |S \cap T_2| = 1$, since $|T_1| = |I_1| > |I_2| = |T_2|$. This component must be a path, $v_0, e_1, v_1, \ldots, v_{t-1}, e_t, v_t$, where e_1 and e_t are in T_1 and v_0 is in $I_1 - I_2$. Note that $e_2, e_4, e_6, \ldots, e_{t-1}$ are in T_2. Let T' be

$$(T_2 - \{e_2, e_4, e_6, \ldots, e_{t-1}\}) \cup \{e_1, e_3, e_5, \ldots, e_t\}.$$

No vertex is in more than one edge in T', and every vertex in $I_2 \cup v_0$ is in one edge in T'. Now the set of edges T' certifies that $I_2 \cup v_0$ is matchable, so **I3** holds, and thus $\mathcal{I}$ is the collection of independent sets of a matroid, as desired. □

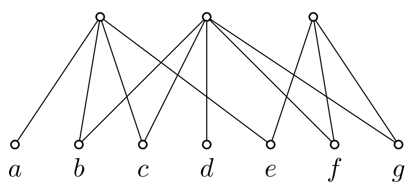

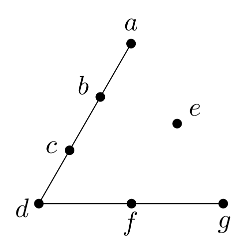

Figure 9: A bipartite graph and its corresponding transversal matroid.

Any matroid that arises from a bipartite graph as in Theorem 2.17 is said to be a *transversal* matroid.

In Figure 9 we see a bipartite graph, and a geometric representation of the corresponding transversal matroid.

Exercise 2.18. Find examples of two non-isomorphic bipartite graphs that correspond to isomorphic transversal matroids.