TOPICS IN GRAPH THEORY

LEONID GLADKOV

Abstract. This paper is an exposition of some classic results in graph theory and their applications. A proof of Tutte's theorem is given, which is then used to derive Hall's marriage theorem for bipartite graphs. Some compelling applications of Hall's theorem are provided as well. In the final section we present a detailed proof of Menger's theorem and demonstrate its power by deriving König's theorem as an immediate corollary.

Contents

1. Definitions

Before presenting the proofs, we provide some prerequisite definitions from graph theory. In the definitions given here and later on, we borrow a lot of the language and notation introduced in the first chapter of Diestel's classic text [1].

Definition 1.1. Given a set A, we define $[A]^k := \{ S \subseteq A \mid |S| = k \}$.

Definition 1.2. A graph $G := (V, E)$ is a pair of sets V and $E \subseteq [V]^2$. Elements of V are called *vertices* and elements of E *edges*. We define $V(G) := V$ and $E(G) := E$ to refer to these sets.

We write uv for edges $\{ u, v \} \in E$ and say that uv *joins* the vertices u and v. Instead of explicitly writing $v \in V$ and $e \in E$ for vertices and edges, we may simply write $v \in G$ and $e \in G$.

We call the elements of an edge its *ends*, and vertices that are joined by an edge are said to be *adjacent*. Adjacent vertices are also called *neighbors*. Similarly, distinct edges that share an end are adjacent. Sets of pairwise non-adjacent vertices or edges are *independent*. Vertices that are joined by an edge are *incident* with that edge, and an edge is incident with the vertices it joins. If $A, B \subseteq V$ and $e = ab \in E$ such that $a \in A$ and $b \in B$, then we call e an A–B edge.

A graph is represented visually by points in $\mathbb{R}^2$ for elements of V and curves between pairs of points for elements of E. That is, $u, v \in V$ are joined by a curve

Date: November 15, 2015.

iff $uv \in E$. This geometric interpretation is remarkably intuitive and serves as a powerful aid in reasoning about graphs.

Definition 1.3. Given a graph $G = (V, E)$, we define its *order* by $|G| := |V|$ and denote its number of edges by $\|G\| := |E|$.

A graph is called even or odd depending on whether $|G|$ is even or odd.

Definition 1.4. Given $G = (V, E)$ and $V' \subseteq V$, we define

$$N_G(V') := \left\{ u \in V \setminus V' \;\middle|\; uv \in E \text{ for some } v \in V' \right\}.$$

That is, $N_G(V')$ denotes the neighbors of vertices in V' that lie outside of V'. For $v \in V$, we define $d_G(v) := \left|N_G(\{ v \})\right|$ to be the *degree* of v in G. (If it is clear which graph is being referred to, we omit the subscript G.)

Now, having established the rudimentary features of graphs, it is also important to formalize set-like relations and operations on them:

Definition 1.5. Given graphs H and G, we call H a *subgraph* of G and write $H \subseteq G$ whenever $V(H) \subseteq V(G)$ and $E(H) \subseteq E(G)$. The graph G is then called a *supergraph* of H. We also express this relation by saying that H *lies in* G or that H is *in* G.

Definition 1.6. Let G and H be graphs. We define $G \cup H$ to be the graph with vertex set $V(G) \cup V(H)$ and edge set $E(G) \cup E(H)$. Intersections of graphs are defined in an analagous manner. Given $V' \subseteq V(G)$, the graph $G - V'$ is obtained from G by deleting the vertices in V' and any edges incident with these vertices. As a notational convenience, we also define $G - H := G - V(H)$. Likewise, given $E' \subseteq [V(G)]^2$, we obtain $G \pm E'$ from G by deleting or adding the edges in E'.

In general, we omit brackets for sets of a single vertex or edge. For example, we shorten $G + \{ e \}$ to $G + e$ and $N_G(\{ v \})$ to $N_G(v)$.

Definition 1.7. The graphs $G_1, \dots, G_n$ are *disjoint* if $V(G_i) \cap V(G_j) = \emptyset$ for all $i \neq j$.

It is also useful to be able to talk about the portion of a graph over some particular subset of its vertices:

Definition 1.8. Given $G = (V, E)$ and $V' \subseteq V$, we define the *induced* subgraph of G over V' to be

$$G[V'] := \Big(V', \left\{ uv \in E \;\middle|\; u, v \in V' \right\}\Big).$$

The subgraph relation provides a natural means to discuss maximality and minimality with respect to graph properties:

Definition 1.9. A graph G is said to be *maximal* (*minimal*) for a given property (for example, connectedness; see below) if adding (deleting) a vertex or an edge creates a supergraph (subgraph) that violates the property. As a special case of this, graphs are called *edge-maximal* (*edge-minimal*) for a property if adding (deleting) an edge between existent vertices creates a graph that violates the property.

We can also talk about (edge-)maximal and (edge-)minimal subgraphs of G if we restrict ourselves to adding and deleting vertices and edges that lie in G.

We also wish to formalize the notions of paths, cycles, and walks:

Definition 1.10. A graph P is called a *path* if $V(P) = \{ v_1, \dots, v_n \}$ and $E(P) = \{ v_i v_{i+1} \mid 1 \leq i < n \}$. A graph C is called a *cycle* if $C = P + v_n v_1$ for some path P. A *walk* is a path whose vertices may repeat.

We often denote a path by a natural sequence of its vertices: $P = v_1 \dots v_n$. We also write $v_i P v_j$ for $P\Big[\{ v_i, \dots, v_j \}\Big]$. For example, we can combine these shorthand notations to express $P + v_n u$ as $v_1 P v_n u$. We may also write $aPbQc$ for $aPb \cup bQc$. The *interior* of a path P refers to the subgraph $v_2 P v_{n-1}$.

Definition 1.11. Given a graph G and $A, B \subseteq V(G)$, a path $P = v_1 \dots v_n \subseteq G$ is called an A–B path in G if $v_1 \in A$ and $v_n \in B$ are its only vertices in $A \cup B$.

An a–b path is said to *link* the vertices a and b. This language brings us to the notion of connectedness, which is the last topic that must be formalized before we are able to state and prove the theorems of interest.

Definition 1.12. A graph G is *connected* if every pair of vertices in $V(G)$ is linked by a path in G.

Definition 1.13. A *component* of a graph G is a maximal connected subgraph of G. The set of components of G is denoted by $\mathscr{C}_G$.

Definition 1.14. Given a graph G and $A, B \subseteq V(G)$, we say that $S \subseteq V$ is an A–B *separator*, or, equivalently, that S *separates* A from B in G, if every A–B path in G has a vertex in S.

2. Tutte's theorem

In both practical applications and other branches of mathematics, we often wish to create pairings in the context of some kind of relationship, whether it be matching prospective clients with tutors or, as we shall discover later, determining if distinct partitions of a set can have the same representatives. In graph theory, this problem reduces to asking whether or not we can assign every vertex to a pair with one of its neighbors. That is, we attempt to select edges in such a way that every vertex is incident with exactly one edge in our selection. Then every edge uv that is selected pairs the vertices u and v. Clearly, any independent set of edges succeeds in matching at least some vertices: the edges do not have any ends in common, so each edge determines a pair. For this reason, independent sets of edges are called *matchings*:

Definition 2.1. Given a graph G, a matching M in G is an independent subset of $E(G)$. If every $v \in V' \subseteq V(G)$ is incident with an edge in M then M is called a *matching of* V' and V' is said to be *matched by* M.

Tutte's theorem characterizes graphs that contain matchings of their entire vertex set. Before we are able to state the theorem, however, we require an alternative means of considering matchings.

Definition 2.2. A subgraph $H \subseteq G$ is *spanning* if $V(H) = V(G)$.

Definition 2.3. A *k-regular* graph G satisfies $d(v) = k$ for every $v \in V(G)$.

Definition 2.4. A k-regular spanning subgraph of G is called a *k-factor* of G.

In fact, 1-factors are intimately related to matchings:

Proposition 2.5. *A graph G contains a matching of $V(G)$ iff it contains a 1-factor.*

Proof. Suppose $H \subseteq G$ is a 1-factor. Then, since every vertex in H has degree 1, it is clear that every $v \in V(G) = V(H)$ is incident with exactly one edge in $E(H)$. Thus, $E(H)$ forms a matching of $V(G)$.

On the other hand, if $V(G)$ is matched by $M \subseteq E(G)$, it is easy to see that the subgraph $H \coloneqq \big(V(G), M\big)$ must be a 1-factor. The subgraph H is spanning by definition, and since M is a matching of $V(G)$, every $v \in V(H) = V(G)$ is incident with exactly one edge in $E(H) = M$. Hence every vertex in H has degree 1. Thus H is a 1-factor of G. □

We are now ready to state Tutte's theorem.

Definition 2.6. Given a graph G, we define

$$q(G) \coloneqq \big|\{\, C \in \mathscr{C}_G \mid C \text{ is odd} \,\}\big| \,.$$

Theorem 2.7 (Tutte's theorem)**.** *A graph $G = (V, E)$ contains a 1-factor iff $q(G - S) \leq |S|$ for every $S \subseteq V$.*

The requirement that $q(G - S) \leq |S|$ for every S is henceforth referred to as *Tutte's condition.* We provide a proof of the theorem based on Diestel's adaptation [1] of a proof due to Lovász [2]. We begin by demonstrating the easier necessity of Tutte's condition as a lemma:

Lemma 2.8. *A graph G contains a 1-factor only if $q(G - S) \leq |S|$ for every $S \subseteq V(G)$.*

Proof. Suppose that G contains a a 1-factor. Then by Proposition 2.5, it contains a matching M of $V(G)$. For any $S \subseteq V$, let $C \in \mathscr{C}_{G-S}$ be an odd component. Since C has an odd number of vertices, these cannot all be matched by edges contained in C. Hence there must be at least one C–S edge in M, as there can be no edges between components of $G - S$. Thus, there are at least $q(G - S)$ edges in M between S and odd components of $G - S$. All of these edges must be incident with distinct vertices in S since M is independent. Therefore, $|S| \geq q(G - S)$. □

Next, we prove a lemma which provides alternative criteria for sets that violate Tutte's condition in even edge-maximal graphs with no 1-factor. As we shall see, this lemma is crucial to the proof of Tutte's theorem that follows.

Lemma 2.9. *Let $G = (V, E)$ be an even edge-maximal graph without a 1-factor. Let $S \subseteq V$ be such that every component of $G - S$ is complete and every $s \in S$ is adjacent to every other vertex. Then S satisfies $q(G - S) > |S|$.*

Proof. Suppose, by way of contradiction, that $q(G - S) \leq |S|$. We wish to show that we can then find a matching of V in E. Finding such a matching would violate the fact that G has no 1-factor (Proposition 2.5). Since $|S|$ is large enough, and every $s \in S$ is adjacent to every other vertex, we can first select $q(G - S)$ independent edges joining a vertex from every odd component of $G - S$ with distinct vertices in S. This leaves an even number of vertices unpaired since $|G|$ is even. These can then be matched because $G[S]$ and all components of $G - S$ are complete. Hence we reach the desired contradiction, so $q(G - S) > |S|$. □

We are now ready to complete our proof of Tutte's theorem.

Proof of Theorem 2.7. By Lemma 2.8, it remains to show the sufficiency of Tutte's condition. We do so by proving the contrapositive: if G does not contain a 1-factor, then there exists $S \subseteq V$ with $q(G-S) > |S|$.

We first show that G can be assumed edge-maximal with no 1-factor. If G is not edge-maximal, let us continue to add edges until the resultant supergraph is. Call this graph G'. Suppose we find a set $S \subseteq V$ which violates Tutte's condition in G'. Every component in $\mathscr{C}_{G'-S}$ is the union of components in $\mathscr{C}_{G-S}$. The components of $G-S$ whose union is an odd component of $G'-S$ cannot all be even, so $q(G-S) \geq q(G'-S) > |S|$. Thus, S also violates Tutte's condition in G. Therefore, we can add edges to any graph G which is not edge-maximal until we obtain a graph G' which is, and finding a set S which violates Tutte's condition in G reduces to finding such a set in G'.

Therefore, let us suppose that G is edge-maximal without a 1-factor to begin with. If $|G|$ is odd, it must contain an odd number of odd components, so $S = \emptyset$ satisfies

$$q(G-S) = q(G) \geq 1 > 0 = |S| \,.$$

If $|G|$ is even, then by Lemma 2.9 our task reduces to finding a set S such that every component of $G-S$ is complete and every $s \in S$ is adjacent to every other vertex in G. Let

$$S = \left\{ s \in V \;\middle|\; N(s) = V \setminus \{s\} \right\}. \tag{2.10}$$

If every component of $G-S$ is complete, then we are done. If not, then there exists $C \in \mathscr{C}_{G-S}$ such that $xx' \notin C$ for some vertices $x, x' \in C$. Since C is connected, we can find a shortest path $P \subseteq C$ linking x and x'. Let x, y, z be the first three vertices on this path. Since P is a shortest path, $xz \notin E$. Because $y \in G-S$, there exists $w \in V \setminus \{\, y \,\}$ such that $yw \notin E$. And since G is edge-maximal without a 1-factor, the supergraphs $G+xz$ and $G+yw$ contain matchings M_1 and M_2, respectively, of G's entire vertex set. Note that $xz \in M_1$ and $yw \in M_2$. Now, we construct a maximal path $Q = w_1 \ldots w_n \subseteq G$ such that $w_1 = w$, the edge $w_1w_2 \in M_1$, and the edges of Q alternate between M_1 and M_2. Suppose that $w_{n-1}w_n \in M_1$. Then the maximality of Q implies that the M_2 edge at w_n is not in G or that it creates a walk if added to Q. In the first case, this edge must be yw as it is the only M_2 edge not in G. In the second, this edge must join w_n with one of $w_1, \ldots, w_{n-1}$, but the M_2 edges incident with $w_2, \ldots, w_{n-1}$ are already in $Q - w_n$, so again the edge must be $w_nw_1 = yw$. In particular, $w_n = y$.

If $w_{n-1}w_n \in M_2$, then adding to Q the M_1 edge incident with w_n would not create a walk: such an edge would have to join w_n with one of $w_1, \ldots, w_{n-1}$, but the M_1 edges at these vertices are already in $Q - w_n$. Hence the M_1 edge at w_n must not be in G, and the only such edge is xz. Hence w_n is x or z.

Depending on whether the last edge of Q is in M_1 or M_2, we set $C = w_1Qw_{n-1}w_1$ or $C = w_1Qw_nyw_1$, respectively. We claim that in each case, C is an even cycle in $G+yw$ such that every second edge of C lies in M_2. Recall that the first edge of Q lies in M_1, and its edges alternate between M_1 and M_2. In the first case, the last edge of Q lies in M_1, so $|Q|$ is odd and $|C| = |Q|+1$ is even. In addition, the edge $w_nw_1 = yw \in M_2$ added to Q to construct C ensures that the alternating pattern of edges is preserved. Since yw is the only edge added to $Q \subseteq G$, we have that $C \subseteq G+yw$. In the second case, the last edge of Q lies in M_2, so $|Q|$ is even and so is $|C| = |Q|+2$. Furthermore, $w_{n-1}w_n \in M_2$ implies that $w_ny \notin M_2$. And since $yw_1 = yw \in M_2$, we find once again that every second edge of C lies in M_2. Finally, since $w_n = x$ or z, the edge $w_n y \in G$, so the only edge of C not in G is yw. Hence $C \subseteq G + yw$.

Therefore, we may construct a matching $M \subseteq E$ of V from M_2: this can be done by replacing the M_2 edges in C with the other edges in C, which must all lie in G. But then G would contain a 1-factor (Proposition 2.5), which contradicts the initial assumption. Hence the set S given by (2.10) must violate Tutte's condition. □

3. Hall's marriage theorem

The task of finding matchings is most natural in the context of *bipartite* graphs, whose vertex set admits a bipartition such that every edge joins vertices from different subsets:

Definition 3.1. A graph $G = (V, E)$ is called bipartite if $V = A \cup B$ for disjoint sets A and B and every edge in E is an A–B edge. We express a bipartite graph with bipartition $V = A \cup B$ as $G(A, B)$.

Hall's marriage theorem completely characterizes bipartite graphs $G(A, B)$ in which we can find a matching of A or B. In fact, Hall's theorem is a bipartite analogue of Tutte's theorem: if $|A| = |B|$, then finding a matching of A or B is equivalent to finding a matching of $V(G)$.

Theorem 3.2 (Hall's marriage theorem)**.** *Let $G(A, B)$ be a bipartite graph. Then G contains a matching of A iff $\big|N(S)\big| \geq |S|$ for all $S \subseteq A$.*

The requirement that $\big|N(S)\big| \geq |S|$ for every S is referred to as the *marriage condition*.

We proceed to derive the marriage theorem from Tutte's theorem, an approach which may be indirect but one which provides great insight into the relationship between the two results.

Proof of Theorem 3.2. Clearly, every $S \subseteq A$ must have enough neighbors in B in order for the vertices of S to be matched, so $\big|N(S)\big| \geq |S|$. Hence the necessity of the marriage condition is trivial.

We proceed to derive the sufficiency from the stronger Tutte's theorem. To do so, we must ensure that finding a matching of A in G can be reduced to finding a 1-factor, perhaps in some graph G' constructed from G. To construct G', we first replace B with

$$B' = \begin{cases} B & \text{if } G \text{ is even} \\ B \cup \{b^*\} & \text{if } G \text{ is odd.} \end{cases}$$

In other words, if need be, we add a vertex b^* to B to ensure that the resultant graph is even. The desired graph G' is

$$G' = G + \left\{ bb' \;\middle|\; b, b' \in B' \right\}.$$

That is, we add edges to make $G'[B']$ complete. Suppose we find a matching $M \subseteq E(G')$ of $V(G')$. Let

$$M' = \{ uv \in M \mid u \in A \text{ or } v \in A \}.$$

Since every $e \in M'$ must lie in $E(G)$, and every $a \in A \subseteq V(G')$ is incident with an edge in M, the subset M' forms the desired matching of A in G. Thus, our task is to find a 1-factor in G' (Proposition 2.5). Therefore, we wish to show that the marriage condition in G implies Tutte's condition in G'.

Let $S \subseteq V(G')$. Set $S_A = S \cap A$ and $S_{B'} = S \cap B'$. We can consider the subtraction $G - S$ as occurring in two steps: we first remove $S_{B'}$ and then S_A. Thus, we verify Tutte's condition in two steps. First, we will show that $q(G' - S_{B'}) \leq |S_{B'}|$. Then we will demonstrate that $q(G' - S) \leq q(G' - S_{B'}) + |S_A|$. Combining these inequalities will provide the desired result:

$$q(G' - S) \leq q(G' - S_{B'}) + |S_A| \leq |S_{B'}| + |S_A| = |S| . \tag{3.3}$$

Let us proceed with the first step. Because every $a \in A$ has at least one neighbor in $B \subseteq B'$ by the marriage condition, and $G'[B']$ is complete, the graph G' is connected. Also G' is even, so $q(G') = 0$. Since $G[B' - S_{B'}]$ is also complete, $S_{B'}$ only separates in G' the vertices in A all of whose neighbors lie in $S_{B'}$. Let A^* be the set of such vertices:

$$A^* \coloneqq \{\, a \in A \mid N_{G'}(a) \subseteq S_{B'} \,\}.$$

There are no edges between vertices of A^*, so the components of $G' - S_{B'}$ are the singleton vertices in A^* and the subgraph $G' - S_{B'} - A^*$. Since $N_{G'}(A^*) \subseteq S_{B'}$ by definition, and $\left|N_{G'}(A^*)\right| = \left|N_G(A^*)\right| \geq |A^*|$ by the marriage condition, it follows that

$$|A^*| \leq \left|N_{G'}(A^*)\right| \leq |S_{B'}| . \tag{3.4}$$

Hence we have a bound on the number of singleton components in $G' - S_{B'}$. The component $G' - S_{B'} - A^*$ may also be odd. If the inequality in (3.4) is strict, then Tutte's condition holds regardless:

$$q(G' - S_{B'}) \leq |A^*| + 1 \leq |S_{B'}| - 1 + 1 = |S_{B'}| .$$

Otherwise, we have $|A^*| = |S_{B'}|$. But then

$$\left|G' - S_{B'} - A^*\right| = \left|G'\right| - 2|A^*| \equiv 0 \bmod 2,$$

so the only odd components are vertices in A^*. Thus

$$q(G - S_{B'}) = |A^*| \leq |S_{B'}|$$

by (3.4), so again Tutte's condition holds.

It remains to show that $q(G' - S) \leq q(G - S_{B'}) + |S_A|$. If S_A is empty then we are done. If not, consider removing S_A from $G' - S_{B'}$. Recall that the components of $G' - S_{B'}$ are $G' - S_{B'} - A^*$ and singleton vertices in A^*. Set

$$S_{A_1} \coloneqq S_A \cap A^*$$

and

$$S_{A_2} \coloneqq S_A \setminus S_{A_1}.$$

Removing S_{A_1} from $G' - S_{B'}$ decreases the number of odd (singleton) components:

$$q(G' - S_{B'} - S_{A_1}) = q(G' - S_{B'}) - |S_{A_1}| .$$

Having removed S_{A_1}, we know that the vertices of S_{A_2} still lie in the unaltered component $G' - S_{B'} - A^*$. However, S_{A_2} does not separate any vertices in this component since every $v \in G' - S_{B'} - A^*$ has a neighbor in $B' \setminus S_{B'}$, and $G'[B' \setminus S_{B'}]$ is complete. Therefore, removing S_{A_2} from $G' - S_{B'} - S_{A_1}$ does not introduce any new components. It is possible, however, that $\left|G' - S_{B'} - A^* - S_{A_2}\right|$ is odd, so removing S_{A_2} can increase the number of odd components by at most 1. Therefore, we arrive at the desired result:

$$\begin{aligned} q(G' - S) &= q(G' - S_{B'} - S_A) \\ &= q(G' - S_{B'} - S_{A_1} - S_{A_2}) \\ &\leq q(G' - S_{B'} - S_{A_1}) + 1 \\ &= q(G' - S_{B'}) - |S_{A_1}| + 1 \\ &\leq q(G' - S_{B'}) + 1 \\ &\leq q(G' - S_{B'}) + |S_A| . \end{aligned}$$

The last inequality holds because $S_A \neq \emptyset$. By (3.3), this completes the proof. □

Now, with Hall's theorem in our toolkit, we are ready to demonstrate some surprising applications. The theorems that follow are both prompted by exercises in Diestel's text [1], and both reduce to showing that the marriage condition holds in a bipartite graph. In each case, the difficulty lies in constructing the appropriate graph.

Theorem 3.5. *Let S be a finite set. If $P = \{ P_1, \ldots, P_n \}$ and $Q = \{ Q_1, \ldots, Q_n \}$ are two distinct partitions of S into k-sets, then P and Q admit a common choice of representatives.*

Proof. Set $A = \{ 1, \ldots, n \}$ and $B = A$. Construct the bipartite graph $G(A, B)$ such that the A–B edge $ij \in E(G)$ iff $P_i \cap Q_j \neq \emptyset$. If we can find a matching M of A, then this matching will determine a common choice of representatives. For every $ij \in M$, we would choose any $x \in P_i \cap Q_j$ to represent $P_i \in P$ and $Q_j \in Q$.

To show that such a matching exists, we need to verify that the marriage condition holds. Let $S \subseteq A$, and let $x \in P_i$ for $i \in S$. There exists j such that $x \in Q_j$. Then $ij \in E(G)$, so $j \in N(S)$. Hence,

$$\bigcup_{i \in S} P_i \subseteq \bigcup_{j \in N(S)} Q_j .$$

Thus,

$$k|S| = \Big| \bigcup_{i \in S} P_i \Big| \leq \Big| \bigcup_{j \in N(S)} Q_j \Big| = k \big| N(S) \big| .$$

Dividing by k gives the desired result. □

Theorem 3.6. *Let X be a finite set, with $X_1, \ldots, X_n \subseteq X$. Set $d_1, \ldots, d_n \in \mathbb{N}$. We can choose $D_i \subseteq X_i$ such that $|D_i| = d_i$ for every i and $D_i \cap D_j = \emptyset$ for $i \neq j$ iff*

$$\Big| \bigcup_{i \in I} X_i \Big| \geq \sum_{i \in I} d_i \tag{3.7}$$

for every $I \subseteq \{ 1, \ldots, n \}$.

The necessity of (3.7) is trivial: if there exist disjoint $D_i \subseteq X_i$ with $|D_i| = d_i$, then clearly

$$\Big| \bigcup_{i \in I} X_i \Big| \geq \Big| \bigcup_{i \in I} D_i \Big| = \sum_{i \in I} d_i$$

for any I.

The necessity of (3.7) is analogous to the necessity of the marriage condition in Hall's theorem. The marriage condition is vital because enough neighbors must exist in order for a matching to be possible. Likewise, the inequality here is necessary because the sets X_i must each contain enough elements to make it possible for the disjoint subsets D_i to exist. This analogy suggests that we may be able to apply Hall's theorem if we construct the proper bipartite graph. The structure of (3.7) suggests that B should equal X, and that for every $S \subseteq A$, the set $N(S)$ must equal $\bigcup_I X_i$ for some index I. Having made these observations, we are ready to complete the proof.

Proof of Theorem 3.6. As discussed above, the necessity of (3.7) is clear. To prove the sufficiency, we construct the appropriate bipartite graph. First, let

$$A = \bigcup_{i=1}^{n} A_i,$$

where $|A_i| = d_i$ for every i and the A_i are pairwise disjoint. Set $B = X$. Now, construct the bipartite graph $G(A, B)$ such that

$$E(G) = \{\, ax \mid a \in A_i \text{ and } x \in X_i,\ 1 \le i \le n \,\}. \tag{3.8}$$

Thus, for all i, every $a \in A_i$ is joined to every element of X_i. We claim that choosing the desired subsets $D_i \subseteq X_i$ reduces to finding a matching of A in G. Suppose A is matched by $M \subseteq E(G)$, and set

$$D_i = \{\, x \in X_i \mid ax \in M \text{ for } a \in A \,\},\ 1 \le i \le n.$$

Since $N(X_i) = A_i$ for every i,

$$D_i = \{\, x \in X_i \mid ax \in M \text{ for } a \in A_i \,\},\ 1 \le i \le n.$$

Thus, every subset D_i is determined by the A_i–X_i edges of M. Because $A_i \subseteq A$ is matched by M, there are d_i such edges, and each must be incident with distinct vertices in X_i since M is independent. Therefore, $|D_i| = d_i$ for all i. Furthermore, suppose that $x \in D_i \cap D_j$ for some $i \neq j$. But then there exist $a_i \in A_i$ and $a_j \in A_j$ such that $a_i x$ and $a_j x$ lie in M. But A_i and A_j are disjoint, so $a_i x \neq a_j x$ are adjacent edges in M, which contradicts the fact that M is a matching. Hence, $D_i \cap D_j = \emptyset$ for $i \neq j$, as desired.

To show that a matching of A exists in G, we need to verify that the marriage condition holds. Let $S \subseteq A$. For every $a \in A$ exists i such that $a \in A_i$, so $N(a) = X_i$ by (3.8). Thus,

$$N(S) = \bigcup_{a \in S} N(a) = \bigcup_{i \in I} X_i,$$

where $I \subseteq \{\, 1, \dots, n \,\}$. Also, note that $S \subseteq \bigcup_{i \in I} A_i$. By (3.7), the desired result quickly follows:

$$\left|N(S)\right| = \left|\bigcup_{i \in I} X_i\right| \ge \sum_{i \in I} d_i = \left|\bigcup_{i \in I} A_i\right| \ge |S|\,.$$

□

4. Menger's theorem

We devote this section to Menger's theorem, a classic result which relates the task of separating vertex sets (Definition 1.14) to the task of finding disjoint paths (Definition 1.7).

Definition 4.1. Given a graph G and $A, B \subseteq V(G)$, let $\mathbb{S}$ be the set of A–B separators in G. We define

$$k(G, A, B) \coloneqq \min_{X \in \mathbb{S}} |X| \,.$$

In other words, $k(G, A, B)$ is the minimum number of vertices required to separate A from B in G. We are now ready to state Menger's theorem.

Theorem 4.2 (Menger's theorem)**.** *Let $G = (V, E)$ and $A, B \subseteq V$. The maximum number of disjoint A–B paths in G equals $k(G, A, B)$.*

It is simple to show that there can be no more than $k(G, A, B)$ disjoint A–B paths:

Lemma 4.3. *Let $G = (V, E)$ and $A, B \subseteq V$. Any set of disjoint A–B paths has no more than $k(G, A, B)$ elements.*

Proof. Set $k = k(G, A, B)$, and let $X \subseteq V$ be an A–B separator of minimum size. Suppose we find more than k disjoint A–B paths. Every such path has a vertex in X. But $|X| = k$, so it follows by the pigeonhole principle that at least two of the paths share a vertex in X. Hence a contradiction. □

Thus, to prove Menger's theorem it suffices to demonstrate that $k(G, A, B)$ disjoint A–B paths exist. In this paper we proceed to do so by induction on $\|G\|$, following an outline provided by Diestel [1, p. 83, Exercise 16]. We proceed with a lemma that completes the proof of Menger's theorem in one of its major cases.

Definition 4.4. Given a graph G and $X, Y \subseteq V(G)$, define $G_{X,Y} \subseteq G$ to be induced by the union of X and the components of $G - X$ that meet Y.

Lemma 4.5. *Let $G = (V, E)$ be a graph. Let $X \subseteq V$ be a minimum A–B separator in G. Suppose Menger's theorem holds for graphs H with $\|H\| < \|G\|$. Then if $\|G_{X,A}\| < \|G\|$ and $\|G_{X,B}\| < \|G\|$, there exist $|X|$ disjoint A–B paths in G.*

Proof. Let U be an A–X separator in $G_{X,A}$, and let P be an A–X path in G. The interior of P lies in a component of G–X that meets A, so P is also an A–X path in $G_{X,A}$. Thus, P has a vertex in U, so U is an A–X separator in G. Therefore, U is also an A–B separator in G, so $|U| \geq |X|$. Since $\|G_{X,A}\| < \|G\|$, we can apply Menger's theorem to find a set $\mathscr{A}$ of $|X|$ disjoint A–X paths in $G_{X,A} \subseteq G$. Similarly, we can find a set $\mathscr{B}$ of $|X|$ disjoint X–B paths in G. The paths in $\mathscr{A}$ cannot intersect with paths in $\mathscr{B}$ outside of X, since otherwise X would not separate A from B. Hence, each $x \in X$ is the single vertex shared by a unique path in $\mathscr{A}$ with a unique path in $\mathscr{B}$. Therefore, we can join these pairs of paths in X to create the desired $|X|$ disjoint A–B paths in G. □

We are now ready to prove Menger's theorem.

Proof of Theorem 4.2. Set $k = k(G, A, B)$. Recall that by Lemma 4.3, our task is to show that k disjoint A–B paths exist in G. We proceed by induction on $\|G\|$. If

G has no edges, then $A \cap B$ is an A–B separator, so $|A \cap B| \geq k$. Thus, the vertices in $A \cap B$ provide k disjoint (trivial) A–B paths.

Now let $xy \in E$. Suppose that S is a minimum A–B separator in $G - xy$. By the inductive hypothesis, $G - xy$ contains $|S|$ disjoint A–B paths. Call the set of these paths $\mathscr{P}$. Every A–B separator in G must have at least $|S|$ vertices, since otherwise $|S|$ would not be a smallest separator in $G - xy$. Thus, if S' with $|S'| = |S|$ separates A from B in G, it follows that $|S'| = k$ and the $|S|$ paths in $\mathscr{P}$ suffice. If no such separator exists, then in particular S fails to separate A from B in G. Thus, there exists a path $P \subseteq G$ that avoids S. Because S is an A–B separator in $G - xy$, it follows that $xy \in P$. Hence $x, y \notin S$. However, P cannot avoid $S_x \coloneqq S \cup \{ x \}$ or $S_y \coloneqq S \cup \{ y \}$. Thus, S_x and S_y are both smallest A–B separators in G, with $|S_x| = |S_y| = |S| + 1 = k$.

If for one of $X \in \{ S_x, S_y \}$, it is true that $\|G_{X,A}\| < \|G\|$ and $\|G_{X,B}\| < \|G\|$, then an application of Lemma 4.5 provides $|X| = k$ disjoint A–B paths in G. Otherwise, for every $X \in \{ S_x, S_y \}$, we have $\|G_{X,A}\| = \|G\|$ or $\|G_{X,B}\| = \|G\|$.

We claim that if $\|G_{X,A}\| = \|G\|$, then $X \subseteq B$. Suppose not. Then there exists $v \in X$ such that $v \notin B$. We wish to reach a contradiction by showing that X is not minimum. Thus, it suffices to demonstrate that every A–B path that passes through v also passes through another $w \in X \setminus \{ v \}$. Then $X \setminus \{ v \}$ would be a smaller A–B separator in G. Let $Q = v_1 \ldots v_n$ be an A–B path in G with $v_1 \in A$ and $v \in Q$. Thus $v = v_j$ for some $1 \leq j < n$. Since $E = E(G_{X,A})$, we have that $N_G(v) \subseteq G_{X,A}$. It follows that $v_{j+1} \in G_{X,A}$. If $v_{j+1} \in X$, we are done. If not, then v_{j+1} lies in some component $C \in \mathscr{C}_{G-X}$ that meets A. Set $a \in V(C) \cap A$. Let R be an a–v_{j+1} path in C. Then the A–B path $aRv_{j+1}Qv_n \subseteq G$ must have a vertex $w \in X$. Since $V(R) \cap X = \emptyset$, it follows that $w = v_k$ for some $k > j + 1$. Thus, $w \in V(Q) \cap \big(X \setminus \{ v \}\big)$ is the desired vertex.

Analogously, if $\|G_{X,B}\| = \|G\|$, then $X \subseteq A$. Thus, S_x and S_y are each contained in A or B. In particular, $x \in A$ or $x \in B$, and $y \in A$ or $y \in B$. Suppose $x, y \in A \setminus B$ or $x, y \in B \setminus A$. But then no A–B path can contain xy, so S suffices to separate A from B in G, contradicting our assumption. Thus, $x \in A$ and $y \in B$ or vice versa, so the edge xy forms an A–B path of length 1. Recall that no path in $\mathscr{P}$ contains x or y. Therefore, $\mathscr{P} \cup \{ xy \}$ is a set of $|S| + 1 = k$ disjoint A–B paths in G. □

To demonstrate the power of Menger's theorem, we proceed to derive König's theorem about matchings. Although König's theorem is nontrivial to prove directly, it is nothing but a special case of the much stronger Menger's theorem. Hence, we state it here as a corollary.

Definition 4.6. Let $G = (V, E)$. Given $V' \subseteq V$ and $E' \subseteq E$, we say that V' *covers* E' if every $e \in E'$ is incident with some $v \in V'$. The set V' is then called a *cover* of E'.

Corollary 4.7 (König's theorem)**.** *Given a bipartite graph $G(A, B)$, the maximum size of a matching in G equals the minimum number of vertices that cover $E(G)$.*

Proof. Since G is bipartite, a matching in G is equivalent to a set of disjoint A–B paths. Furthermore, a cover of $E(G)$ is equivalent to an A–B separator. Therefore, a direct application of Menger's theorem suffices. □

Acknowledgments. This paper would not be possible without the patient guidance of my mentor, John Wilmes, who has provided me with countless hints, suggestions, and advice. I would also like to thank all REU instructors for sharing their knowledge in the many fascinating lectures held this summer. Finally, I would like to express gratitude to Professor May for organizing the REU.

References

[1] Reinhard Diestel. Graph Theory. Springer. 2012.
[2] L. Lovász. Three Short Proofs in Graph Theory. Journal of Combinatorial Theory. 1975.