HYPERGRAPHS WITH A UNIQUE PERFECT MATCHING

AARON SPINDEL

Under the direction of
Dr. John S. Caughman

A Math 501 Project
Submitted in partial fulfillment of the
requirements for the degree of

Master of Science in Mathematics
at
Portland State University

Abstract. Following the article "On the maximum number of edges in a k-uniform hypergraph with a unique perfect matching" by Deepak Bal, Andrzej Dudek, and Zelealem B. Yilma, this paper states and proves a tight upper bound for the number of edges in a hypergraph that has a unique perfect matching. The two main focuses of this paper are constructing a hypergraph realizing the given bound and proving that the bound applies to all possible hypergraphs with a unique perfect matching.

Contents

Date: February 26, 2012.

1. Introduction

A reader possessing basic knowledge of graph theory may read this section to understand the general goals of this paper. Terms will be formally defined in Section 2.

Imagine a complete graph on six vertices as depicted below.

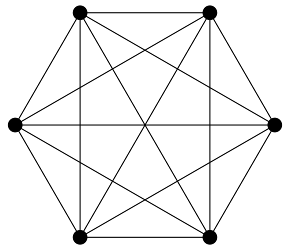

This graph has many perfect matchings, some of which are depicted by red dashed lines below.

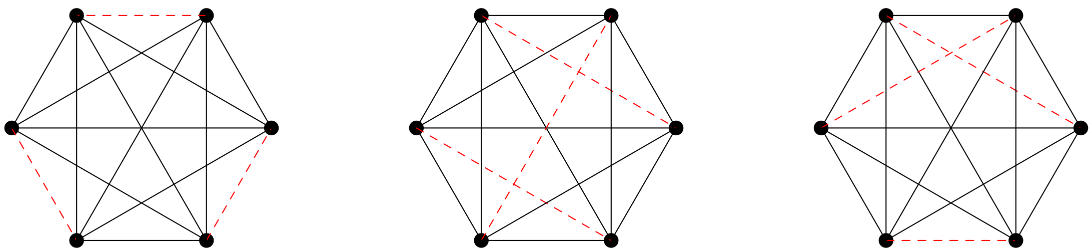

Simply put, having many edges gives a large amount of choices. These choices create enough freedom where perfect matchings can be constructed in many ways. In this paper, we focus on graphs that have exactly one perfect matching. These graphs cannot have too many edges. Otherwise, the abundance of edges would create more than one perfect matching. In fact, we will show that a graph with unique perfect matching possessing $2m$ vertices can have at most m^2 edges.

This bound applies to all possible graphs. We will also study particular examples of graphs that attain this bound. The graph depicted below will be referred to as $\mathcal{H}_3^*$.

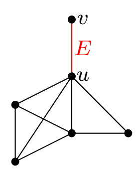

Every perfect matching in $\mathcal{H}_3^*$ must contain edge E since it is the unique edge incident with vertex v. Because E is included in any perfect matching, no other edge incident with vertex u can be included in a perfect matching. Erasing these edges from the picture yields:

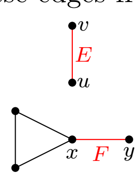

Vertex y is now a pendant vertex and acts similarly to vertex v in the previous discussion. Edge F which is the only edge incident with vertex y must therefore be included in any perfect matching. No other edge in a perfect matching can be incident with vertex x. Erasing such edges yields:

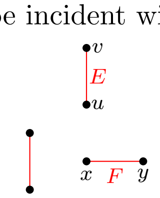

The three remaining edges form a perfect matching in $\mathcal{H}_3^*$. In fact, the above analysis implies that this is the only perfect matching in $\mathcal{H}_3^*$. Additionally, $\mathcal{H}_3^*$ has 6 vertices and 9 edges attaining our bound.

While the primary ideas can be understood by focusing on graphs, this paper proves results in a more general fashion. The main result (Corollary 2.2) applies to all k-uniform hypergraphs. The formula presented in this corollary may seem messy at first. Rather than regarding this formula as many algebraic symbols, it is best to think of it conceptually as a mixture of combinatorial tools. The inclusion-exclusion principle plays a notably important role in this theorem.

2. Background

A **hypergraph** G is a finite set of vertices $\mathcal{V}$ along with a set of edges $\mathcal{E} \subseteq \mathcal{P}\mathcal{V} \setminus \{\emptyset\}$ (where $\mathcal{P}\mathcal{V}$ denotes the power set of $\mathcal{V}$) such that no two edges in $\mathcal{E}$ are equal as sets. Note that multiple edges (pairs of edges containing identical vertices) are prohibited by this definition. Loop edges (edges that contain exactly one vertex) are in general allowed. When convention requires, $\mathcal{V}(G)$ will denote the vertex set $\mathcal{V}$, and $\mathcal{E}(G)$ will denote the edge set $\mathcal{E}$. A k**-uniform** hypergraph (or k**-graph**) is a hypergraph in which every edge $E \in \mathcal{E}$ has cardinality k.

A **pendant** vertex is an element $v \in \mathcal{V}$ that is contained in exactly one edge $E \in \mathcal{E}$. A set of edges $\{E_1, \ldots, E_k\}$ are **incident** with a vertex v if $v \in \bigcup_{i=1}^k E_i$. We abuse notation and assert that edge E (without set braces) is incident with v if $v \in E$. Two vertices v and u are **adjacent** denoted $u \sim v$ if there exists an edge containing both v and u. An **isomorphism** between hypergraphs G and H is a bijection from $\mathcal{V}(G)$ to $\mathcal{V}(H)$ such that for any two vertices $u, v \in \mathcal{V}(G)$, u is adjacent to v in G if and only if $f(u)$ is adjacent to $f(v)$ in H. If there exists an isomorphism from G to H we say G is **isomorphic** to H denoted $G \cong H$ without specific reference to the isomorphism.

While hypergraphs are defined abstractly as sets, we often represent them graphically. Figure 1 below depicts a hypergraph G. Black and red dots represent vertices while finite closed regions of the plane represent edges. Vertices u and v are both pendant vertices because they are only contained in edge E. Moreover, $u \sim v$ since edge E is incident with both u and v. This hypergraph is not k-uniform for any value of k since it contains edges of varying cardinality.

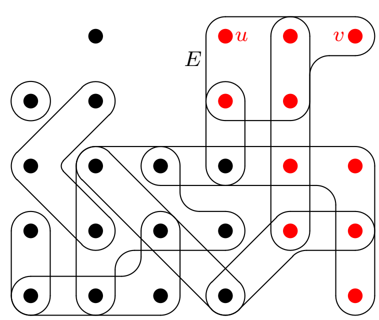

Figure 1

A **subgraph** of a hypergraph G is a hypergraph H with $\mathcal{V}(H) \subseteq \mathcal{V}(G)$ and $\mathcal{E}(H) \subseteq \mathcal{E}(G)$. Given $\mathcal{V}' \subseteq \mathcal{V}$, an **induced subgraph** of a hypergraph $G = (\mathcal{V}, \mathcal{E})$ denoted $G[\mathcal{V}']$ is the hypergraph with vertex set $\mathcal{V}'$ and edge set $\{E \in \mathcal{E} : E \subseteq \mathcal{V}'\}$. The induced subgraph $G[\mathcal{V}']$ contains every edge in G that is not incident with a vertex outside of $\mathcal{V}'$. If $E_1, \ldots, E_k$ are sets of vertices (such as edges in a hypergraph), $G[E_1, \ldots, E_k]$ is defined to mean $G[\bigcup_{i=1}^k E_i]$. If $v_1, \ldots, v_k$ are vertices, $G - \{v_1, \ldots, v_k\}$ denotes $G[\mathcal{V}(G) \setminus \{v_1, \ldots, v_k\}]$. If $E_1, \ldots, E_k$ are edges, $G - \{E_1, \ldots, E_k\}$ is the hypergraph with vertex set $\mathcal{V}(G)$ and edge set $\mathcal{E}(G) \setminus \{E_1, \ldots, E_k\}$. For a vertex v and edge E, we abuse notation and use $G - v$ and $G - E$ to denote $G - \{v\}$ and $G - \{E\}$ respectively.

If G is the hypergraph in Figure 1 above and $\mathcal{V}'$ is the set of vertices depicted in black, $G[\mathcal{V}']$ is drawn below in Figure 2.

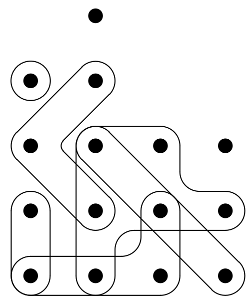

Figure 2

A **matching** in a hypergraph $G = (\mathcal{V}, \mathcal{E})$ is a set of pairwise disjoint edges $\{M_1, \ldots, M_m\}$. Any edge M_i in a matching is referred to as a **matching edge**. A **perfect matching** is a matching $\{M_1, \ldots, M_m\}$ such that $\mathcal{V} = \bigcup_{i=1}^{m} M_i$. In other words, a perfect matching is a collection of edges that partition the vertex set. In a k-uniform hypergraph, $\bigcup_{i=1}^{m} M_i$ has cardinality km. Hence $|\mathcal{V}|$ must be a multiple of k in a k-uniform hypergraph that has a perfect matching. In a 2-uniform hypergraph (which is commonly referred to as a graph), a **nearly perfect matching** is a matching $\{M_1, \ldots, M_m\}$ such that $\bigcup_{i=1}^{m} M_i = \mathcal{V} \backslash \{v_0\}$ for some vertex $v_0 \in \mathcal{V}$. In other words, a nearly perfect matching is a set of pairwise disjoint edges that are incident with every vertex in the graph except v_0. Note that a graph must have an odd number of vertices in order to have a nearly perfect matching.

Figure 3a below depicts a portion of a hypergraph. The entire vertex set is shown. However, some edges have been omitted. The displayed edges are pairwise disjoint and therefore form a matching. The displayed edges are not incident with every vertex. Therefore, this is not a perfect matching. Figure 3b below depicts a portion of a 2-uniform hypergraph. In the case of a 2-uniform hypergraph, dots represent vertices while edges are depicted as line segments. Similar to the previous hypergraph, the entire vertex set is shown, but some edges have been omitted. The edges displayed are pairwise disjoint and therefore form a matching. Since these edges are incident with every vertex except for v_0, this is a nearly perfect matching.

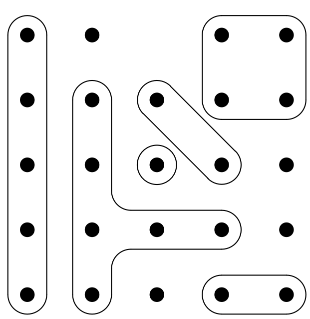

Figure 3a

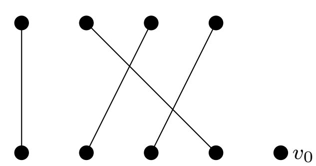

Figure 3b

The following language is used in the special case of a 2-uniform hypergraph. A **path** $\mathcal{P}$ is a graph with distinct vertices $\{v_1, \ldots, v_k\}$ such that $\mathcal{E}(\mathcal{P}) = \{\{v_i, v_{i+1}\} : 1 \leq i \leq k-1\}$. For vertices u and v in a graph G, a u, v-path is a subgraph of G that is a path with pendant vertices u and v. A **cycle** $\mathcal{C}$ is a graph with distinct vertices $\{v_1, \ldots, v_k\}$ such that $\mathcal{E}(\mathcal{C}) = \{\{v_i, v_{i+1}\} : 1 \leq i \leq k-1\} \bigcup \{\{v_1, v_k\}\}$. A graph is **connected** if given any two distinct vertices u and v, there is a u, v-path. A **component** of G is a maximal connected subgraph of G. A **cut-edge** is an edge whose deletion increases the number of components of G. An **endpoint** of an edge E is a vertex contained in E.

The graph in Figure 4 below has two components. Edge E is a cut edge.

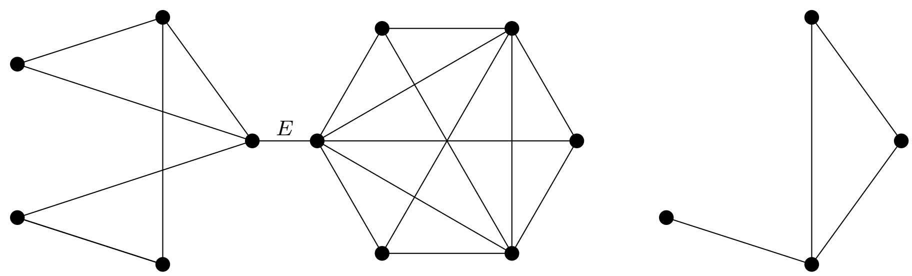

Figure 4

In a k-uniform hypergraph $G = (\mathcal{V}, \mathcal{E})$, any edge $E \in \mathcal{E}$ is defined to be a k-set (set with cardinality k) of elements of $\mathcal{V}$. While the notions may coincide in set-theoretical terms, this paper carefully uses edge and k-set to imply different meanings. An *edge* is an actual element of the edge set of a specific hypergraph. A *k-set* is merely any set with k elements. For example, if $F \subseteq \mathcal{V}$ is a k-set, F might be an element of the edge set $\mathcal{E}$ (in which case it represents an edge present in the hypergraph) or might not be an element of $\mathcal{E}$ (in which case it does not represent an edge present in the hypergraph). It is best to think of k-sets as potential edges. This is contrary to the notation used in [1].

Let $G = (\mathcal{V}, \mathcal{E})$ be a k-uniform hypergraph with km vertices for some $m > 1$. If G is empty (with $\mathcal{E} = \emptyset$), then G lacks perfect matchings. Suppose G is complete ($\mathcal{E} = \{E \in \mathcal{P}\mathcal{V} \setminus \{\emptyset\} : |E| = k\}$) with vertex set $\{1, \ldots, km\}$. One example of a perfect matching is the family of edges $M_i = \{k(i-1)+1, k(i-1)+2, \ldots, ki\}$ for $1 \le i \le m$. Applying any permutation of vertices to the matching $\{M_1, \ldots, M_m\}$ (permuting the sets elementwise) yields another set of disjoint edges (which are in the hypergraph because it is complete) that covers all vertices. Because $m > 1$, switching vertices in between two edges of the matching potentially creates a new matching. This gives many perfect matchings.

Adding edges to a hypergraph G never removes any perfect matchings that initially existed in G. In fact, adding edges to G makes G closer to a complete hypergraph with many possible perfect matchings. This suggests that any hypergraph having exactly one perfect matching can only have a limited number of edges. Above this threshold, a hypergraph is forced to have multiple perfect matchings. The following makes this intuitive notion precise:

Theorem 2.1. *Let $\mathcal{H}_m = (\mathcal{V}_m, \mathcal{E}_m)$ be a k-uniform hypergraph with km vertices and unique perfect matching $\mathcal{M} = \{M_1, \ldots, M_m\}$. Let $\mathcal{B}_\ell$ be the set of edges that intersect exactly ℓ matching edges for $1 \le \ell \le k$. That is, $\mathcal{B}_\ell = \{E \in \mathcal{E}_m : |Q_E| = \ell\}$ where $Q_E = \{M_i \in \mathcal{M} : M_i \bigcap E \neq \emptyset\}$. Then $|\mathcal{B}_1| = m$ and $|\mathcal{B}_\ell| \le b_{k,\ell}\binom{m}{l}$ for $2 \le \ell \le k$ where*

$$b_{k,\ell} = \frac{\ell - 1}{\ell} \sum_{i=0}^{\ell-1} (-1)^i \binom{\ell}{i} \binom{k(\ell - i)}{k}.$$

Corollary 2.2. *Using the same notation as Theorem 2.1, let*

$$f(k, m) = m + b_{k,2}\binom{m}{2} + b_{k,3}\binom{m}{3} + \cdots + b_{k,k}\binom{m}{k}$$

for $k \ge 2$ and $m \ge 1$. Then $|\mathcal{E}_m| \le f(k, m)$. (Remember that $|\mathcal{E}_m|$ is the size of the edge set of the k-uniform hypergraph $\mathcal{H}_m$ that has a unique perfect matching as defined in Theorem 2.1.) Moreover, this bound is tight.

3. Construction

This section presents an infinite family of k-uniform hypergraphs $\mathcal{H}_m^*$ with a unique perfect matching that attain the bound presented in Corollary 2.2. Demonstrating a perfect matching is sufficient to show that $\mathcal{H}_m^*$ has at least one perfect matching. However, showing that no other perfect matchings exist poses a greater challenge. The key observation is that pendant vertices limit possible matchings.

Suppose we are looking for possible perfect matchings in a hypergraph $G = (\mathcal{V}, \mathcal{E})$ that has pendant vertex v. Any perfect matching must contain an edge incident with every vertex of the hypergraph. Since there is only one edge E_v incident with v, E_v must be in every perfect matching. Because edges of a perfect matching are disjoint, no edge nontrivially intersecting E_v can be in the matching. This suggests an iterative approach to the problem of finding a perfect matching. Let G' be the hypergraph with vertex set $\mathcal{V} \setminus E_v$ and edge set $\{E \in \mathcal{E} : E \bigcap E_v = \emptyset\}$. The above reasoning implies that any perfect matching in G corresponds to a perfect matching in G' along with edge E_v. Thus, we have reduced the problem to finding perfect matchings in a hypergraph with fewer vertices. If G' also has a pendant vertex, we can repeat exactly the same argument to reduce the problem further.

With this is mind, we define the hypergraphs $\mathcal{H}_m^*$ recursively in such a way that they predictably have pendant vertices. First, let $\mathcal{H}_1^*$ be the k-graph with vertices $\{1, \dots, k\}$ and exactly one edge $M_1 = \{1, \dots, k\}$. Every vertex in $\mathcal{H}_1^*$ is a pendant vertex. Next, suppose that $\mathcal{H}_{m-1}^*$ is a known hypergraph with vertex set $\{1, \dots, k(m-1)\}$ and edge set $\mathcal{E}_{m-1}$. To create $\mathcal{H}_m^*$ we must add k vertices expanding the vertex set to $\{1, \dots, km\}$. We keep all of the edges in $\mathcal{E}_{m-1}$. In order to attain the bound in Corollary 2.2, we need to add as many edges as we can without destroying the uniqueness of the perfect matching. Adding an edge that only contains vertices in the set $\{1, \dots, k(m-1)\}$ would be like adding an edge to $\mathcal{H}_{m-1}^*$. Since $\mathcal{H}_{m-1}^*$ attains the bound in Corollary 2.2, adding an edge would create a hypergraph that exceeds the bound. Corollary 2.2 (which we have yet to prove) would then imply that the hypergraph would have at least two perfect matchings. Since this is undesirable, we must avoid adding this type of edge. So, every new edge ought to contain at least one of the new vertices. Moreover, because pendant vertices control possible perfect matchings, we should specify one of the new vertices (say km) to be a pendant.

The two main constraints are: all new edges must contain at least one new vertex and exactly one edge should be incident with vertex km. Other than these constraints, we add every edge possible to maximize the number of edges in the hypergraph. We are now ready to precisely define $\mathcal{H}_m^*$. Start with $\mathcal{H}_{m-1}^*$ and add $k-1$ new vertices. Include every possible edge that is incident with at least one of these new vertices. Finally, add another new vertex to act as the pendant vertex. We must add exactly one edge containing this new vertex. This edge must become part of the perfect matching in $\mathcal{H}_m^*$. The idea is to use the unique perfect matching in $\mathcal{H}_{m-1}^*$ along with this new edge to form a unique perfect matching in $\mathcal{H}_m^*$. Because edges in a perfect matching are disjoint, the new edge should not intersect $\mathcal{H}_{m-1}^*$. The only remaining possibility is to add the edge containing the k new vertices that were just added to the hypergraph.

Formally for $m \geq 2$ let $\mathcal{H}_m^* = (\mathcal{V}_m, \mathcal{E}_m)$ be a k-graph with vertices

$$\mathcal{V}_m = \{1, \dots, km\}$$

In order to specify the edge set, it is convenient to define a family of k-sets. Let

$$M_i = \{k(i-1)+1, \dots, ki\} \text{ for } 1 \leq i \leq m.$$

Note that M_m represents the set of k new vertices that have been added to $\mathcal{H}_{m-1}^*$. Then, the edges of $\mathcal{H}_m^* = (\mathcal{V}_m, \mathcal{E}_m)$ are specified by

$$\mathcal{E}_m = \mathcal{E}_{m-1} \bigcup \left\{ E \in \mathcal{P}\mathcal{V}_m \setminus \{\emptyset\} : |E| = k \text{ and } E \bigcap M_m \neq \emptyset \text{ and } km \notin E \right\} \bigcup \{M_m\}$$

Example 3.1. *The following depicts the first three graphs described by this construction in the 2-uniform case ($k = 2$):*

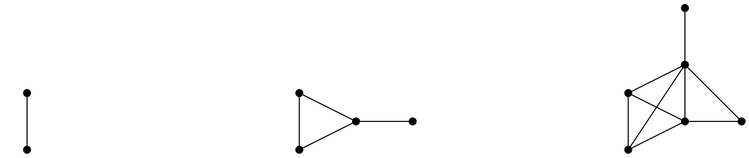

$\mathcal{H}_1^*$ $\mathcal{H}_2^*$ $\mathcal{H}_3^*$

Example 3.2. *The following depicts $\mathcal{H}_1^*$ as described by this construction in the 3-uniform case ($k = 3$). Note that edges are represented as triangles in these hypergraphs.*

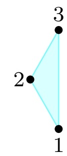

The following sequence of images shows how $\mathcal{H}_2^*$ is constructed starting with $\mathcal{H}_1^*$.

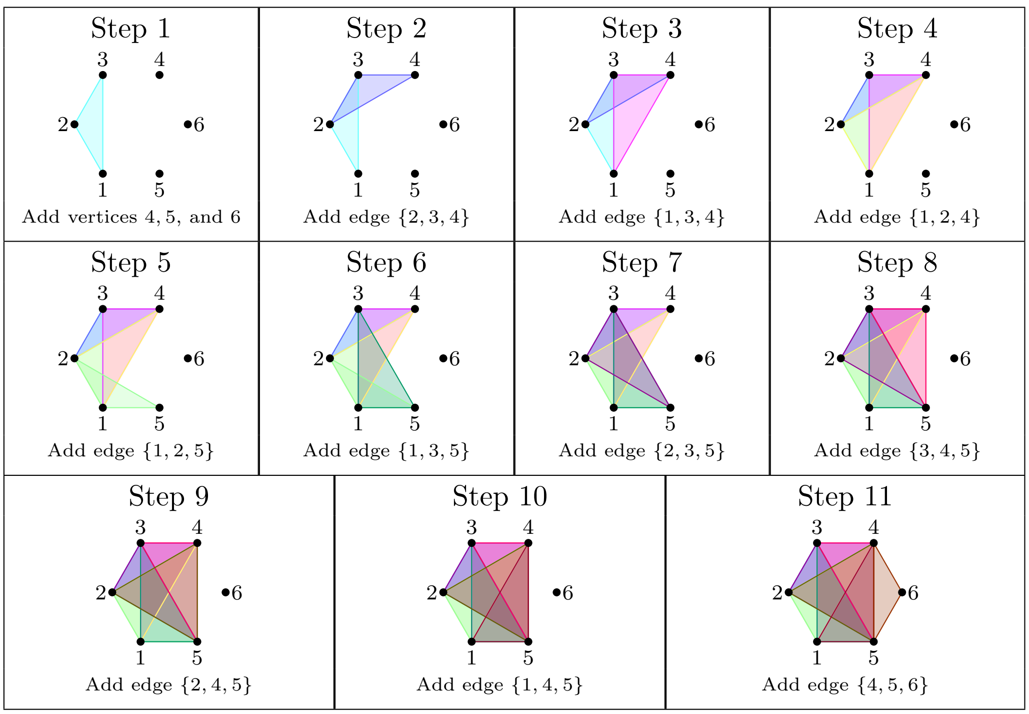

Claim 3.3. *The k-graph $\mathcal{H}_m^*$ has a unique perfect matching $\{M_1, \dots, M_m\}$.*

Proof. We proceed by induction on m.
Base case: $\mathcal{H}_1^*$ has only one edge M_1. This means, it has only one possible perfect matching. Moreover, this edge covers all vertices in $\mathcal{H}_1^*$. Therefore, M_1 is the unique perfect matching in $\mathcal{H}_1^*$.
Inductive step: By construction, M_m is added to the edge set of $\mathcal{H}_m^*$. (This is needed to guarantee M_m is an edge in the hypergraph since it was defined as a k-set.) Any perfect matching in $\mathcal{H}_m^*$ must contain M_m which is the unique edge incident with the pendant vertex km. No other edge in the matching can intersect $M_m = \{k(m-1)+1, \dots, km\}$. Every other matching edge must only contains vertices in the set $\{1, \dots, k(m-1)\}$. All of these edges are therefore present in the hypergraph $\mathcal{H}_{m-1}^*$ and form a perfect matching for this hypergraph. By the induction hypothesis, $\mathcal{H}_{m-1}^*$ has unique perfect matching $\{M_1, \dots, M_{m-1}\}$. Combining this with M_m creates a perfect matching in $\mathcal{H}_m^*$. Since we never had any choice of which edges to include in the perfect matching, it is unique. $\square$

Claim 3.4. *The k-graph $\mathcal{H}_m^* = (\mathcal{V}_m, \mathcal{E}_m)$ attains the bound in Corollary 2.2. That is, $|\mathcal{E}_m| = f(k,m)$.*

Proof. Let $\mathcal{M} = \{M_1, \dots, M_m\}$ be the unique perfect matching as above, and let B_ℓ be the set of edges in the hypergraph that intersect exactly ℓ matching edges for $1 \le \ell \le k$. That is, $\mathcal{B}_\ell = \{E \in \mathcal{E}_m : |Q_E| = \ell\}$ where $Q_E = \{M_i \in \mathcal{M} : M_i \bigcap E \ne \emptyset\}$. Since the edges of a perfect matching partition the vertex set of the hypergraph, every edge in the hypergraph nontrivially intersects at least one matching edge. Additionally, edges which contain k vertices can intersect at most k matching edges. So, every edge in the hypergraph is contained in some $\mathcal{B}_\ell$ for $1 \le \ell \le k$. Since the $\mathcal{B}_\ell$ are also disjoint, they partition the edge set of the hypergraph. We count the edges in each $\mathcal{B}_\ell$ in order to count the total number of edges in the hypergraph.

An edge E is in $\mathcal{B}_1$ if it intersects exactly one matching edge M_i. If E contained a vertex not present in M_i, E would intersect M_i and another distinct matching edge. Since this is disallowed, $E \subseteq M_i$. Also, $|E| = |M_i| = k$. Hence, $E = M_i$ and $\mathcal{B}_1 = \{M_1, \dots, M_m\}$. (Note that each M_i truly is an edge in the hypergraph by Claim 3.3.) Hence $|\mathcal{B}_1| = m$.

We wish to determine $|\mathcal{B}_\ell|$ for $2 \le \ell \le k$. To do this, let $\mathcal{L} = \{M_{i_1}, \dots, M_{i_\ell}\}$ be a set of ℓ distinct edges from the matching with $1 \le i_1 < i_2 < \cdots < i_\ell \le m$. Let $\mathcal{V}_\mathcal{L} = \bigcup_{j=1}^{\ell} M_{i_j}$ be the set of vertices contained in some matching edge in $\mathcal{L}$. Because edges in a matching are disjoint, $|\mathcal{V}_\mathcal{L}| = k\ell$. Let $\mathcal{G}$ be the set of k-sets that only contain elements in $\mathcal{V}_\mathcal{L}$ and nontrivially intersect every edge in $\mathcal{L}$. In symbols, $\mathcal{G} = \{F : |F| = k \text{ and } F \subseteq \mathcal{V}_\mathcal{L} \text{ and } \forall M_i \in \mathcal{L}, F \bigcap M_i \ne \emptyset\}$. Caution: elements of $\mathcal{G}$ need not be edges in the hypergraph.

We use the inclusion-exclusion principle to compute $|\mathcal{G}|$. A k-element subset of $\mathcal{V}_\mathcal{L}$ can be formed by choosing k elements from $k\ell$ possibilities. Naively, there are $\binom{k\ell}{k}$ ways to create such a k-set. However, some of these possibilities do not intersect every edge in $\mathcal{L}$. We must exclude the k-sets that miss at least one edge. A k-set that doesn't intersect M_{i_j} for some j draws k elements from $\mathcal{V}_\mathcal{L} \setminus M_{i_j}$ which has cardinality $k(\ell-1)$. There are $\ell = \binom{\ell}{1}$ choices for j, giving $\binom{\ell}{1}\binom{k(\ell-1)}{k}$ k-sets that miss at least one edge.

The inclusion-exclusion principle asserts that "all possible k-sets" minus "k-sets that miss at least one edge" double counts some k-sets, subtracting "k-sets that miss at least two edges" twice. We must compensate by adding the number of k-sets that miss at least two edges. Suppose a k-set is not incident with M_{i_u} and M_{i_v} for some u, v with $u \ne v$. Then the k-set draws k vertices from $\mathcal{V}_\mathcal{L} \setminus \{M_{i_u} \bigcup M_{i_v}\}$ which has cardinality $k(\ell-2)$. There are $\binom{\ell}{2}$ choices for u and v yielding $\binom{\ell}{2}\binom{k(\ell-2)}{k}$ k-sets that miss at least two edges. Proceeding this way, we find

$$\begin{aligned}
|\mathcal{G}| &= \sum_{i=0}^{\ell-1} (-1)^i \binom{\ell}{i}\binom{k(\ell-i)}{k} \\
&= \underbrace{\binom{k\ell}{k}}_{\substack{\text{Include all} \\ \text{possible} \\ k\text{-sets}}} \underbrace{-\binom{\ell}{1}\binom{k(\ell-1)}{k}}_{\substack{\text{Exclude } k\text{-sets} \\ \text{that miss at} \\ \text{least one edge}}} + \underbrace{\binom{\ell}{2}\binom{k(\ell-2)}{k}}_{\substack{\text{Include } k\text{-sets} \\ \text{that miss at} \\ \text{least two edges}}} - \cdots + \underbrace{(-1)^{\ell-1}\binom{\ell}{\ell-1}\binom{k}{k}}_{\substack{\text{Include/Exclude} \\ k\text{-sets that miss at} \\ \text{least } \ell-1 \text{ edges}}}
\end{aligned}$$

Since no k-set in $\mathcal{G}$ misses all ℓ edges in $\mathcal{M}$, the maximum possible index in the summation is $\ell - 1$.

The edges in $\mathcal{H}_m^*$ that only contain vertices in $\mathcal{V}_\mathcal{L}$ are exactly the same as the edges in $\mathcal{H}_{i_\ell}^*$ that only contain vertices in $\mathcal{V}_\mathcal{L}$. When constructing the hypergraph $\mathcal{H}_{i_\ell}^*$, every edge that intersected M_{i_ℓ} and did not contain ki_ℓ was added to the hypergraph. Every k-set in $\mathcal{G}$ intersects M_{i_ℓ} and is therefore an edge unless it contains vertex ki_ℓ.

Conversely, consider a k-set E in $\mathcal{G}$ that does contain vertex ki_ℓ. Since E intersects ℓ edges in the matching for some $\ell \ge 2$, $E \ne M_{i_\ell}$. However, M_{i_ℓ} is the only edge in $\mathcal{H}_{i_\ell}^*$ that contains vertex

ki_ℓ which was constructed to be a pendant in $\mathcal{H}^*_{i_\ell}$. So, E cannot be an edge in the hypergraph. In total, a k-set in $\mathcal{G}$ is an edge in $\mathcal{H}^*_m$ if and only if it does not contain ki_ℓ.

Every vertex in $\mathcal{V}_\mathcal{L}$ appears symmetrically in the construction of k-sets in $\mathcal{G}$. Therefore, every vertex belongs to the same number of k-sets of $\mathcal{G}$. Suppose every vertex belongs to η k-sets. Using this, we count the number of ordered pairs (x, E) where $x \in E \in \mathcal{G}$ in two ways. Any element of $\mathcal{V}_\mathcal{L}$ may be chosen for x, giving $k\ell$ choices. Once we pick x, there are η k-sets that contain it. Therefore, there are $k\ell\eta$ ordered pairs. Alternately, we can choose a k-set first, giving $|\mathcal{G}|$ choices. Once a k-set is known, there are k vertices it contains. So, there are $|\mathcal{G}|k$ ordered pairs. Equating these yields $\eta = |\mathcal{G}|/\ell$.

The number of edges in $\mathcal{G}$ equals $|\mathcal{G}|$ minus the number of k-sets in $\mathcal{G}$ that contain vertex ki_ℓ. By definition, η k-sets contain ki_ℓ in $\mathcal{G}$, giving a formula for the number of the edges in $\mathcal{B}_\ell$ on vertex set $\mathcal{V}_\mathcal{L}$:

$$|\{E \in \mathcal{E}_M : E \in \mathcal{B}_\ell \text{ and } E \subseteq \mathcal{V}_\mathcal{L}\}| = |\mathcal{G}| - \eta = |\mathcal{G}| - |\mathcal{G}|/\ell = \frac{\ell - 1}{\ell}|\mathcal{G}| = b_{k,\ell}$$

where $b_{k,\ell}$ was defined in Theorem 2.1.

Every edge E in $\mathcal{B}_\ell$ intersects exactly ℓ edges in the matching. First, choose which ℓ edges E will intersect in $\binom{m}{\ell}$ possible ways. Once you have fixed these ℓ possible matching edges, the above argument shows there are $b_{k,\ell}$ possible edges E. So, $|\mathcal{B}_\ell| = b_{k,\ell}\binom{m}{\ell}$ for $2 \le \ell \le k$.

Finally,

$$|\mathcal{E}_m| = \sum_{\ell=1}^{k} |\mathcal{B}_\ell| = f(k, m)$$

where $f(k, m)$ was defined in Corollary 2.2. □

The formula given for $f(k, m)$ in Corollary 2.2 uses a double summation. We now have a concrete interpretation of $f(k, m)$ as the number of edges in $\mathcal{H}^*_m$. Using this, we construct another formula for $f(k, m)$ without the double summation.

Corollary 3.5. *For $k \ge 2$ and $m \ge 1$,*

$$f(k, m) = m + \sum_{i=1}^{m-1} \left[\binom{k(i+1) - 1}{k} - \binom{ki}{k} \right]$$

Proof. The edges in $\mathcal{H}^*_m = (\mathcal{V}_m, \mathcal{E}_m)$ are constructed by adding edges to $\mathcal{E}_{m-1}$, the edge set of $\mathcal{H}^*_{m-1}$. This process can be described as follows: start with the vertices of $\mathcal{H}^*_{m-1}$ and append $k - 1$ new vertices. To create $\mathcal{E}_m$, first include every single possible edge (of cardinality k) on this set of $km - 1$ vertices. (Essentially, we start by considering a complete hypergraph on $km - 1$ vertices.) This gives $\binom{km-1}{k}$ edges. However, we have included too many edges in $\mathcal{E}_m$. We only want edges that contain at least one of the $k - 1$ new vertices. To remedy this, discard edges that only use vertices originally present in $\mathcal{H}^*_{m-1}$. Since $\mathcal{H}^*_{m-1}$ has $k(m - 1)$ vertices, there are $\binom{k(m-1)}{k}$ edges that must be discarded. In symbols,

$$\left|\left\{E \in \mathcal{P}\mathcal{V}_m \setminus \{\emptyset\} : |E| = k \text{ and } E \bigcap M_m \ne 0 \text{ and } km \notin E\right\}\right| = \binom{km - 1}{k} - \binom{k(m-1)}{k}$$

where $M_m = \{k(m - 1) + 1, \dots, km\}$ as previously defined. Then,

$$\begin{aligned} |\mathcal{E}_m| &= \left|\mathcal{E}_{m-1} \bigcup \left\{E \in \mathcal{P}\mathcal{V}_m \setminus \{\emptyset\} : |E| = k \text{ and } E \bigcap M_m \ne \emptyset \text{ and } km \notin E\right\} \bigcup \{M_m\}\right| \\ &= |\mathcal{E}_{m-1}| + \binom{km - 1}{k} - \binom{k(m-1)}{k} + 1 \end{aligned}$$

since the sets in the above union are disjoint. Remembering that $|\mathcal{E}_m| = f(k,m)$, this equation becomes

$$f(k,m) = f(k,m-1) + \binom{km-1}{k} - \binom{k(m-1)}{k} + 1$$

This equation is only valid when $m \geq 2$ since the index of $\mathcal{E}_{m-1}$ must be at least one. Summing both sides of this equation from $i = 2$ to m yields

$$\begin{aligned}\sum_{i=2}^{m} f(k,i) &= \sum_{i=2}^{m} f(k,i-1) + \sum_{i=2}^{m}\left[\binom{ki-1}{k} - \binom{k(i-1)}{k}\right] + \sum_{i=2}^{m} 1 \\ &= \sum_{i=1}^{m-1} f(k,i) + \sum_{i=1}^{m-1}\left[\binom{k(i+1)-1}{k} - \binom{ki}{k}\right] + m - 1\end{aligned}$$

Hence

$$f(k,m) + \sum_{i=2}^{m-1} f(k,i) = \sum_{i=2}^{m-1} [f(k,i)] + f(k,1) + \sum_{i=1}^{m-1}\left[\binom{k(i+1)-1}{k} - \binom{ki}{k}\right] + m - 1$$

Using the fact that $f(k,1) = 1$, this becomes

$$f(k,m) = m + \sum_{i=1}^{m-1}\left[\binom{k(i+1)-1}{k} - \binom{ki}{k}\right]$$

□

4. Nonuniqueness of Hypergraphs Attaining the Edge Bound

While the construction in Section 3 does provide a k-uniform hypergraph with unique perfect matching that attains the bound of Corollary 2.2 for every sensible number of vertices, it does not provide a comprehensive list of all k-uniform hypergraphs with a unique perfect matching that attain this bound. For example, consider the case $k \geq 3$ and $m = 2$. As previously shown, $\mathcal{H}_2^* = (\mathcal{V}_2, \mathcal{E}_2)$ has a unique perfect matching $\{M_1, M_2\}$ and attains the maximum number of possible edges. Create another hypergraph $\mathcal{H}_2^a$ by modifying $\mathcal{H}_2^*$. Fix an edge $E \in \mathcal{E}_2 \setminus \{M_1, M_2\}$. Note that $\overline{E} = \{1, 2, \dots, 2k\} \setminus E$ has cardinality k meaning that it could be allowed as an edge in a k-uniform hypergraph. If $\overline{E} \in \mathcal{E}_2$, $\overline{E}$ and E would form a perfect matching in $\mathcal{H}_2^*$ distinct from $\{M_1, M_2\}$. This would violate the uniqueness of the perfect matching in $\mathcal{H}_2^*$. Therefore $\overline{E} \notin \mathcal{E}_2$. Additionally, since $2k$ is a pendant in $\mathcal{H}_2^*$, $2k \notin E$. Hence, $2k \in \overline{E}$.

Define $\mathcal{E}' = \{\overline{E}\} \bigcup (\mathcal{E}_2 \setminus \{E\})$. We have removed one edge that previously was in $\mathcal{E}_2$ and added one edge that previously was absent in $\mathcal{E}_2$. Thus, $|\mathcal{E}'| = |\mathcal{E}_2|$. Let $\mathcal{H}_2^a$ be the hypergraph with vertex set $\mathcal{V}_2$ and edge set $\mathcal{E}'$. Since $\mathcal{H}_2^a$ and $\mathcal{H}_2^*$ have the same number of edges and vertices, $\mathcal{H}_2^a$ also attains the bound in Corollary 2.2.

For vertices $v \in \{1, \dots, k\}$ in $\mathcal{H}_2^*$, let $v^+ = v + 1$ if $v < k$ and $v^+ = 1$ if $v = k$. (This is similar to arithmetic mod k, but v^+ never equals 0.) Then $M_1, \{k+1\} \bigcup (M_1 \setminus \{v^+\})$, and $\{k+2\} \bigcup (M_1 \setminus \{v^+\})$ are three edges in $\mathcal{H}_2^*$ that contain vertex v. (Notice that $k \geq 3$ implies that $2k = k + k > k + 2$. Hence, $k + 2$ is not the pendant vertex and these k-sets truly are edges in $\mathcal{H}_2^*$.) For $v \in \{k+1, \dots, 2k-1\}$, $\{1, k+1, k+2, \dots, 2k-1\}, \{2, k+1, k+2, \dots, 2k-1\}$, and $\{3, k+1, k+2, \dots, 2k-1\}$ are three edges in $\mathcal{H}_2^*$ containing v. (Since $k \geq 3$, $k+1 > 3$ and each of these sets contains k distinct elements. These k-sets intersect edge M_2 and do not contain vertex $2k$. Thus, they are edges in $\mathcal{H}_2^*$.) Finally, vertex $2k$ is on edge M_2 and no others in $\mathcal{H}_2^*$.

Any vertex in $\{1, \dots, 2k-1\}$ is contained in at least 3 edges in $\mathcal{H}_2^*$. When creating $\mathcal{H}_2^a$, we remove exactly one edge E. Hence any vertex in $\{1, \dots, 2k-1\}$ is contained in at least 2 edges in $\mathcal{H}_2^a$. Vertex $2k$ is contained in edges $\overline{E}$ and M_2 in $\mathcal{H}_2^a$. Since every vertex in $\mathcal{H}_2^a$ is contained in at least two edges, $\mathcal{H}_2^a$ lacks pendant vertices. Therefore, $\mathcal{H}_2^a$ cannot be isomorphic to $\mathcal{H}_2^*$ which has pendant vertex $2k$.

Any perfect matching in $\mathcal{H}_2^a$ would necessarily contain an edge incident with vertex $2k$. The only edges in $\mathcal{H}_2^a$ incident with vertex $2k$ are $\overline{E}$ and M_2. This means that any perfect matching in $\mathcal{H}_2^a$ must use M_2 or $\overline{E}$ as an edge. The edges M_1 and M_2 form a perfect matching in $\mathcal{H}_2^a$. The only edge that could possibly form a perfect matching with $\overline{E}$ is E. Since $E \notin \mathcal{E}'$, no edge in $\mathcal{H}_2^a$ pairs with $\overline{E}$ to form a perfect matching. Thus, $\mathcal{H}_2^a$ has a unique perfect matching.

The case $k = 3$ is depicted below. Figure 1 shows $\mathcal{H}_2^*$ as constructed in Example 3.2. Figure 2 shows the unique perfect matching $\{M_1, M_2\}$ in $\mathcal{H}_2^*$. Figure 3 shows an edge E distinct from M_1 and M_2 and its complement $\overline{E}$. Swapping edge E for $\overline{E}$ creates a new hypergraph $\mathcal{H}_2^a$ depicted in Figure 4. The hypergraph $\mathcal{H}_2^a$ is not isomorphic to $\mathcal{H}_2^*$. However, $\mathcal{H}_2^a$ still contains a unique perfect matching and attains the edge bound in Corollary 2.2.

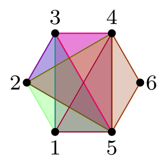

Figure 1

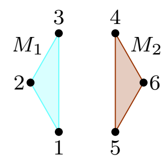

Figure 2

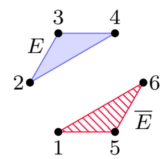

Figure 3

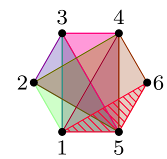

Figure 4

Theorem 4.1. *For $k \geq 3$ and $m \geq 2$, there exists a k-uniform hypergraph $\mathcal{H}_m^a$ with a unique perfect matching $\mathcal{M}^a = \{M_1^a, \ldots, M_m^a\}$ that attains the edge bound of Corollary 2.2 and is not isomorphic to the hypergraph $\mathcal{H}_m^*$ which was constructed in Section 3.*

Proof. Define the family of hypergraphs $\mathcal{H}_m^a$ recursively. $\mathcal{H}_2^a$ was defined above. Once $\mathcal{H}_{m-1}^a$ is known with $m \geq 3$, create $\mathcal{H}_m^a$ by following the same steps used to create $\mathcal{H}_m^*$ from $\mathcal{H}_{m-1}^*$. Formally,

$$\mathcal{V}(\mathcal{H}_m^a) = \{1, \ldots, km\}$$

and

$$\mathcal{E}(\mathcal{H}_m^a) = \mathcal{E}(\mathcal{H}_{m-1}^a) \bigcup \Big\{E \in \mathcal{P}\mathcal{V}_m \setminus \{\emptyset\} : |E| = k \text{ and } E \bigcap M_m \neq \emptyset \text{ and } km \notin E\Big\} \bigcup \{M_m\}$$

where

$$M_m = \{k(m-1)+1, \ldots, km\}$$

An induction argument shows that $\mathcal{H}_m^a$ has a unique perfect matching. The base case was checked above when we verified that $\mathcal{H}_2^a$ has a unique perfect matching. The inductive step is identical to the inductive step in Claim 3.3.

Next we verify that $\mathcal{H}_m^a$ has the correct number of vertices and edges to attain the edge bound in Corollary 2.2. Again, we proceed by induction. As a base case, we have verified $|\mathcal{V}(\mathcal{H}_2^a)| = |\mathcal{V}(\mathcal{H}_2^*)|$ and $|\mathcal{E}(\mathcal{H}_2^a)| = |\mathcal{E}(\mathcal{H}_2^*)|$ when discussing $\mathcal{H}_2^a$ above. Next suppose $|\mathcal{V}(\mathcal{H}_{m-1}^a)| = |\mathcal{V}(\mathcal{H}_{m-1}^*)|$ and $|\mathcal{E}(\mathcal{H}_{m-1}^a)| = |\mathcal{E}(\mathcal{H}_{m-1}^*)|$ and consider $\mathcal{H}_m^a$. Because the recursive process used to construct $\mathcal{H}_m^a$ from $\mathcal{H}_{m-1}^a$ is identical to the process used to construct $\mathcal{H}_m^*$ from $\mathcal{H}_{m-1}^*$,

$$\mathcal{V}(\mathcal{H}_m^a) \setminus \mathcal{V}(\mathcal{H}_{m-1}^a) = \mathcal{V}(\mathcal{H}_m^*) \setminus \mathcal{V}(\mathcal{H}_{m-1}^*) = M_m$$

and

$$\begin{aligned}\mathcal{E}(\mathcal{H}_m^a) \setminus \mathcal{E}(\mathcal{H}_{m-1}^a) &= \mathcal{E}(\mathcal{H}_m^*) \setminus \mathcal{E}(\mathcal{H}_{m-1}^*) \\ &= \Big\{E \in \mathcal{P}\mathcal{V}_m \setminus \{\emptyset\} : |E| = k \text{ and } E \bigcap M_m \neq \emptyset \text{ and } km \notin E\Big\} \bigcup \{M_m\}\end{aligned}$$

Applying the induction hypothesis gives:

$$\begin{aligned}|\mathcal{V}(\mathcal{H}_m^a)| &= |\mathcal{V}(\mathcal{H}_{m-1}^a)| + |\mathcal{V}(\mathcal{H}_m^a) \setminus \mathcal{V}(\mathcal{H}_{m-1}^a)| \\ &= |\mathcal{V}(\mathcal{H}_{m-1}^*)| + |\mathcal{V}(\mathcal{H}_m^*) \setminus \mathcal{V}(\mathcal{H}_{m-1}^*)| \\ &= |\mathcal{V}(\mathcal{H}_m^*)|\end{aligned}$$

and

$$\begin{aligned}|\mathcal{E}(\mathcal{H}_m^a)| &= |\mathcal{E}(\mathcal{H}_{m-1}^a)| + |\mathcal{E}(\mathcal{H}_m^a) \setminus \mathcal{E}(\mathcal{H}_{m-1}^a)| \\ &= |\mathcal{E}(\mathcal{H}_{m-1}^*)| + |\mathcal{E}(\mathcal{H}_m^*) \setminus \mathcal{E}(\mathcal{H}_{m-1}^*)| \\ &= |\mathcal{E}(\mathcal{H}_m^*)|\end{aligned}$$

Finally, it remains to show that $\mathcal{H}_m^a \ncong \mathcal{H}_m^*$ for $m \geq 2$. As usual with a recursive construction, we proceed by induction. When constructing $\mathcal{H}_2^a$ we noted that $\mathcal{H}_2^a \ncong \mathcal{H}_2^*$ establishing the base case. Now suppose that $\mathcal{H}_{m-1}^a \ncong \mathcal{H}_{m-1}^*$. Also, suppose for a contradiction that $\mathcal{H}_m^a \cong \mathcal{H}_m^*$. Let $f : \mathcal{V}(\mathcal{H}_m^a) \to \mathcal{V}(\mathcal{H}_m^*)$ be an isomorphism. We are allowed to assume $m \geq 3$ since the base case already establishes the result when $m = 2$. When $m \geq 3$, $\mathcal{H}_m^a$ has a pendant vertex km by construction. Because isomorphisms preserve pendant vertices and $\mathcal{H}_m^*$ only has one pendant vertex, we must have $f(km) = km$. The vertices in $M_m \setminus \{km\} = \{k(m-1)+1, \ldots, km-1\}$ are adjacent to km in $\mathcal{H}_m^a$. Hence, all of the vertices in $\{f(k(m-1)+1), \ldots, f(km-1)\}$ must be adjacent to km in $\mathcal{H}_m^*$. Since km is a pendant vertex in $\mathcal{H}_m^*$, the vertices adjacent to it are exactly the vertices in $M_m \setminus \{km\}$. In total, $\{f(k(m-1)+1), \ldots, f(km-1)\} \subseteq \{k(m-1)+1, \ldots, km-1\} \Rightarrow \{f(k(m-1)+1), \ldots, f(km-1)\} = \{k(m-1)+1, \ldots, km-1\}$ since f is a bijection.

Let g be the restriction of f to domain $\{1, \ldots, k(m-1)\}$ and codomain $\{1, \ldots, k(m-1)\}$ which is a well-defined bijection by the above argument. Since g is the restriction of an isomorphism, g is an isomorphism between $\mathcal{H}_{m-1}^a$ and $\mathcal{H}_{m-1}^*$. This contradicts our inductive hypothesis. □

Contrasting this, uniqueness does hold for the case $k = 2$ as demonstrated by the following results. First, we start with a lemma proved in [4]:

Lemma 4.2. *In a 2-uniform hypergraph G, an edge is a cut-edge if and only if it is not contained in any cycle.*

Proof. Let $E = \{x, y\}$ be an edge in G. Focus on the component H of G that contains E. Deleting edge E does not effect any component other than H. Hence, we prove that $H - E$ is connected if and only if E is contained in a cycle.

If $H - E$ is connected, it contains an x, y-path. Adding edge E to this path completes a cycle.

Conversely, suppose that E is contained in a cycle C. Let u and v be vertices in H. Since H is connected, it contains a u, v-path $\mathcal{P}$. If $\mathcal{P}$ avoids edge E, it is still a u, v-path in $H - E$. Otherwise, E is contained in path $\mathcal{P}$. In this case, suppose without loss of generality that vertex x is in between u and y in $\mathcal{P}$. Then create a u, v-walk in $H - E$ as follows: first, start at u and follow path $\mathcal{P}$ until you reach vertex x. Next, follow cycle C from x to y while avoiding edge E. Finally, follow path $\mathcal{P}$ from y to v. Since $H - E$ contains a u, v-walk, it also contains a u, v-path and is connected.

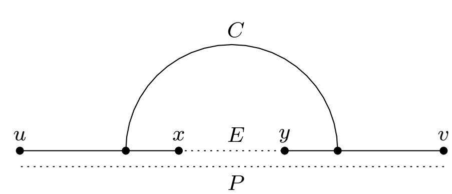

□

The remaining results in this section are due to Lovász in [2].

Definition 4.3. *Let G be a connected 2-uniform hypergraph that lacks cut-edges. Two edges E_1 and E_2 in G are* ***equivalent*** *if $E_1 = E_2$ or if removing edges E_1 and E_2 from G disconnects the graph.*

Lemma 4.4. *Edge equivalence is an equivalence relation.*

Proof. By definition, equivalence of edges is reflexive and symmetric. To see it is also transitive, suppose E_1 is equivalent to E_2 and E_2 is equivalent to E_3. If $E_1 = E_2$ or $E_2 = E_3$, we trivially have E_1 is equivalent to E_3. Otherwise, E_1, E_2, and E_3 are distinct.

Since G lacks cut-edges, both $G - E_1$ and $G - E_2$ are connected. However, $G - \{E_1, E_2\}$ is disconnected. Hence $G - \{E_1, E_2\}$ consists of two components C_1 and C_2. Both E_1 and E_2 connect C_1 to C_2 as depicted below.

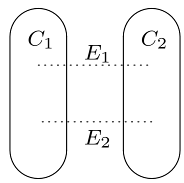

Edge E_i has an endpoint in C_1 and an endpoint in C_2 for $i = 1, 2$. Removing edge E_3 from $G - \{E_1, E_2\}$ cannot join distinct components. Hence, endpoints of E_i are in different components of $G - \{E_1, E_2, E_3\}$ for $i = 1, 2$. Similarly, the endpoints of E_3 also appear in different components of $G - \{E_1, E_2, E_3\}$.

If $G - \{E_1, E_2, E_3\}$ only had two components C_1 and C_2, E_1, E_2, and E_3 would all connect C_1 to C_2. In this case, $G - \{E_1, E_2\}$ would consist of C_1 connected to C_2 by E_3. Then, $G - \{E_1, E_2\}$ would still be connected contrary to the equivalence of E_1 and E_2. Hence, $G - \{E_1, E_2, E_3\}$ must have at least three components.

Since G lacks cut-edges, every component of $G - \{E_1, E_2, E_3\}$ must be incident with at least two of E_1, E_2, and E_3. (A component is incident with an edge if there is some vertex in that component incident with the edge.) The only way this can happen is if G has exactly three components as depicted below:

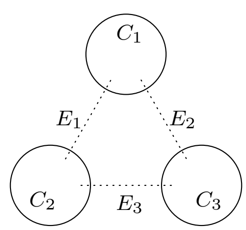

Therefore, removing edges E_1 and E_3 disconnects the graph. That is, E_1 is equivalent to E_3. $\square$

We partition the edges of a graph G into equivalence classes under the above equivalence relation. Denote the equivalence classes by $K_1, \ldots, K_r, L_1, \ldots, L_p$ where $|K_i| = 1$ for $1 \le i \le r$ and $|L_j| > 1$ for $1 \le j \le p$. Note that any two distinct edges E_1 and E_2 in some common L_i are equivalent. Hence $G - \{E_1, E_2\}$ is disconnected.

Lemma 4.5. *Let G be a connected 2-uniform hypergraph that lacks cut-edges and let $K_1, \ldots, K_r$, $L_1, \ldots, L_p$ be equivalence classes of edges as above. Create a new graph by removing all the edges of L_i from G for some fixed i. Then, the components of $G - L_i$ have no cut-edges and are incident with exactly two edges in L_i.*

Proof. As noted in Lemma 4.4, any edge in L_i lies in between two components of $G - L_i$. Create a new graph G' which has one vertex for every component of $G - L_i$. Two vertices in G' are adjacent if there is an edge $E \in L_i$ connecting the corresponding components in $G - L_i$. Since removing a single edge from L_i cannot disconnect G, G' lacks cut-edges. However, any two edges from L_i are equivalent. Hence, removing any two edges in G' disconnects the graph.

Since G' lacks cut-edges, Lemma 4.2 implies every edge in G' is part of a cycle. Fix some edge E_1 in G', and let C_1 be a cycle containing E_1. Suppose there exists an edge E_2 outside of C_1. Because E_2 is contained in some cycle, removing E_2 does not disconnect G'. Moreover, C_1 remains in the graph even after removing edge E_2. So, E_1 is on a cycle in $G' - E_2$. This means that removing E_1 from $G - E_2$ leaves a connected graph. In total, $G' - \{E_1, E_2\}$ is connected graph, contrary to our assumptions. Hence, every edge in the graph must be contained in cycle C_1. Because G' is connected, G' itself is a cycle. Hence, every component of $G - L_i$ (represented by a vertex in G') is incident with exactly two edges in L_i.

It remains to show that a component of $G - L_i$ lacks cut-edges. Suppose for a contradiction that component G_0 of $G - L_i$ has a cut-edge E. The previous part of this proof implies that G_0 is incident with exactly two edges in L_i. Call them E_1 and E_2. If E_1 and E_2 are incident with the same component of $G_0 - E$, E would be a cut edge in G as depicted below:

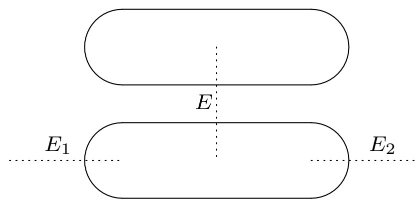

This contradicts the assumption that G lacks cut edges. Alternately, E_1 and E_2 may be incident with separate components of $G_0 - E$:

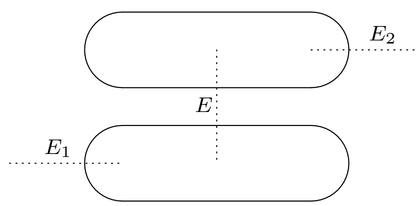

In this case, $G - \{E_1, E\}$ is disconnected, so E_1 is equivalent to E. Hence, E should be in L_i. However, E is in $G - L_i$ by assumption. In all cases we reach a contradiction and no such cut-edge E can exist. □

Definition 4.6. *An edge E in a 2-uniform hypergraph G is an **allowed** edge if it is contained in some maximum matching. Otherwise, E is a **forbidden** edge.*

Theorem 4.7. *Let G be a 2-uniform hypergraph. If G has a unique perfect matching $\mathcal{M}$, then there is an edge $E \in \mathcal{M}$ that is a cut-edge of G.*

Proof. Suppose for a contradiction that there exists a graph with a unique perfect matching which lacks allowed cut-edges. Let G be such a counterexample with the smallest possible number of edges. If G is not connected, there would be a component G_0 of G with fewer edges than G. Then $\{E \in \mathcal{M} : E \in \mathcal{E}(G_0)\}$ forms a perfect matching in G_0. Moreover, any perfect matching in G_0 can be extended to a perfect matching in G by including matchings in the other components. Hence, G_0 must also have a unique perfect matching. Since G_0 has fewer edges than the minimal counterexample, G_0 would necessarily have a cut-edge as part of its unique perfect matching which is a subset of $\mathcal{M}$. However, this would also be a cut-edge in G which is disallowed. Hence, G must be connected.

Suppose G contained a forbidden cut-edge E. Then, a component G_0 of $G - E$ would have a unique perfect matching and fewer edges than G. Hence G_0 would contain a cut-edge that was also a matching edge. This edge would still be both a cut-edge and a matching edge in G. Hence, G lacks forbidden cut-edges. Since G also lacks allowed cut-edges, G does not contain any cut-edge at all.

Let K_i and L_i be equivalence classes of edges of G as previously defined prior to Lemma 4.5. We claim the K_i's do not contain any forbidden edges. To see this, suppose $E \in K_i$ is a forbidden edge for some i. Since $|K_i| = 1$, edge E does not pair with any other edge to disconnect the graph. Removing edge E from G does not create any new cut-edges. Moreover, removing edge E from G does not affect the unique perfect matching in G since E is a forbidden edge. However, by minimality, $G - E$ has an allowed cut-edge F. Hence F must be a cut-edge in the original graph G which lacks cut-edges. Therefore every forbidden edge has to be contained in some L_i.

Next I show that L_i for a fixed i contains at least as many allowed edges as forbidden edges. Let $G_0, \ldots, G_s$ be the components of $G - L_i$. Fix some j with $0 \leq j \leq s$ and consider component G_j. Lemma 4.5 implies that G_j must be incident with exactly two edges E and E' in L_i. Suppose both E and E' are forbidden edges so that the perfect matching in G avoids using them. Then, $\{E \in \mathcal{M} : E \in \mathcal{E}(G_j)\}$ forms a unique perfect matching in the component G_j. By the minimality of G, G_j must contain a cut-edge belonging to its unique perfect matching contradicting Lemma 4.5. Thus any component G_j is incident with at least one allowed edge in L_i. Accordingly, any component G_j is incident with at most one forbidden edge in L_i. Since any edge in L_i is incident with exactly two components of $G - L_i$,

$$\begin{aligned}
2|\{E \in L_i : E \text{ is forbidden}\}| &= \sum_{j=0}^{s} |\{\text{Edges } E \in L_i \text{ incident with } G_j : E \text{ is forbidden}\}| \\
&\leq \sum_{j=0}^{s} |\{\text{Edges } E \in L_i \text{ incident with } G_j : E \text{ is allowed}\}| \\
&= 2|\{E \in L_i : E \text{ is allowed}\}|
\end{aligned}$$

We can extend this result to all of G. The K_i's and L_i's partition $E(G)$, and the K_i's lack forbidden edges. So,

$$\begin{aligned}
\text{number of forbidden edges in } G &= \sum_{i=1}^{p} \text{number of forbidden edges in } L_i \\
&\leq \sum_{i=1}^{p} \text{number of allowed edges in } L_i \\
&\leq \text{number of allowed edges in } G
\end{aligned}$$

This means that G itself has at least as many allowed edges as forbidden edges.

Because G has a unique perfect matching, there are $|V(G)|/2$ allowed edges. There are at most $|V(G)|/2$ forbidden edges in G. Hence G has at most $|V(G)|$ edges. Additionally, G lacks cut-edges so that every edge is contained in some cycle by Lemma 4.2. Since G is connected, the only remaining possibility is that G itself is a cycle. Because G has a perfect matching, it must be an even cycle. However, even cycles have two perfect matchings, contradicting our assumptions. Hence no such counterexample exists. □

Theorem 4.8. *Let G be a 2-uniform hypergraph with $2m$ vertices and m^2 edges. If G has a unique perfect matching, $G \cong \mathcal{H}_m^*$.*

Proof. I proceed by induction on m.
Base case: When $m = 1$, G contains 2 vertices and 1 edge connecting the two vertices. This is exactly the same as $\mathcal{H}_1^*$.
Inductive step: Suppose that the statement holds for all graphs with fewer than $2m$ vertices, and let G be a graph with $2m$ vertices and unique perfect matching $\mathcal{M}$. By Theorem 4.7, there is some edge $E = \{x, y\} \in \mathcal{M}$ that is a cut-edge in G. Form a new graph $G' = G - \{x, y\}$ by removing vertices x and y from G. Note that $\mathcal{M} \setminus \{E\}$ is a perfect matching in G'. Moreover, any perfect matching in G' can be extended to a perfect matching in G by appending edge E. By the uniqueness of the perfect matching in G, this means G' also has a unique perfect matching. Then, by the induction hypothesis, $G' \cong \mathcal{H}_{m-1}^*$.

Reform G from G' by adding in vertices x and y and edges incident with these vertices. We must add $|\mathcal{E}(G)| - |\mathcal{E}(\mathcal{H}_{m-1}^*)| = m^2 - (m-1)^2 = 2m - 1$ edges incident with x or y. One of these edges is E connecting x to y. The $2m - 2$ other edges connect x and y to vertices in $\mathcal{H}_{m-1}^*$. We cannot simultaneously create edges $\{x, v\}$ and $\{y, u\}$ for some $u, v \in \mathcal{V}(\mathcal{H}_{m-1}^*)$. Otherwise, E would no longer be a cut-edge since $\mathcal{H}_{m-1}^*$ is connected as depicted below:

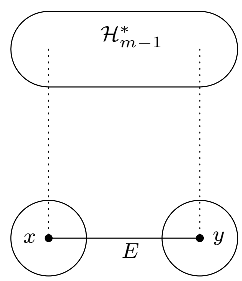

Hence, at most one of x and y is connected by an edge to a vertex in $\mathcal{H}_{m-1}^*$. After renaming, suppose x is not on a common edge with any vertex of $\mathcal{H}_{m-1}^*$. Then, all of the remaining $2m - 2$ edges must be incident with y. Because $\mathcal{H}_{m-1}^*$ only contains $2m - 2$ vertices, an edge must join y to every previous vertex that exists in $\mathcal{H}_{m-1}^*$. Since this exactly mirrors the construction of $\mathcal{H}_m^*$, $G \cong \mathcal{H}_m^*$. □

5. Coverings and Multiple Matchings

In this section we create some tools that will be useful when proving Theorem 2.1. We start by considering situations which lead to multiple perfect matchings.

Example 5.1. *Consider a 2-uniform hypergraph G with unique perfect matching that has at least two matching edges. The image below depicts a portion of G. The solid lines depict two matching edges. The dotted lines represent 2-sets that might be edges in G.*

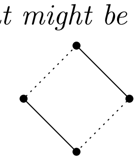

One perfect matching in G exists using the solid edges. If the dashed 2-sets were included in G as well, we could construct a second perfect matching. Start with the original perfect matching and discard the solid edges. Instead, trade them for the dashed edges in the picture. This is a perfect matching in G distinct from the original one, violating the uniqueness of the matching. Since we are not allowed to have both of the dashed edges in G, the total number of edges in G becomes constrained.

Example 5.2. *Example 5.1 directly generalizes. Let G be a 6-uniform hypergraph with a unique perfect matching $\{M_1, \ldots, M_m\}$ where $m \geq 5$. Let $\mathcal{V} = \bigcup_{i=1}^5 M_i$ be the vertices that appear in the* *first five matching edges. These vertices along with $M_1, \ldots, M_5$ are depicted on the left below. The edges $M_1, \ldots, M_5$ partition the vertices in $\mathcal{V}$.*

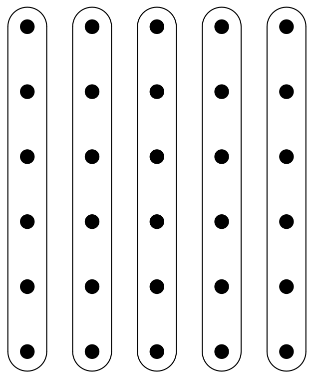

Edges from a perfect matching

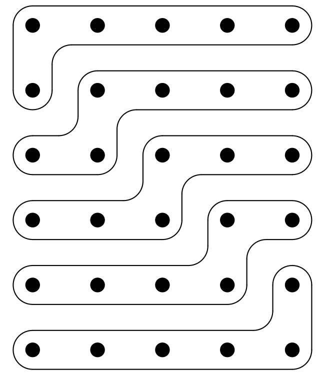

Edges form a covering of the perfect matching

The same vertex set $\mathcal{V}$ is depicted in the right above along with 6-sets $E_1, \ldots, E_5$. Suppose $E_1, \ldots, E_5$ are edges in the hypergraph G. These edges are distinct from the original matching edges and partition $\mathcal{V}$. Informally, $\{E_1, \ldots, E_5\}$ is said to be a covering of $\{M_1, \ldots, M_5\}$ because they both partition the same vertex set. Using the same strategy as before, start with the perfect matching $\{M_1, \ldots, M_m\}$ in G. Discard edges $M_1, \ldots, M_5$ and include edges $E_1, \ldots, E_5$. Both $\{E_1, \ldots, E_5, M_6, \ldots, M_m\}$ and $\{M_1, \ldots, M_m\}$ are perfect matchings violating uniqueness. Thus, we cannot have all of $E_1, \ldots, E_5$ as edges in G. Essentially, coverings are disallowed as they lead to multiple perfect matchings.

In order to employ this strategy in the proof of our main results, it is helpful to establish better language. Throughout the remainder of this section let $\mathcal{H}_m = (\mathcal{V}_m, \mathcal{E}_m)$ be an arbitrary k-uniform hypergraph on km vertices with unique perfect matching $\mathcal{M} = \{M_1, \ldots, M_m\}$. We must formalize the intuitive notion of covering.

Definition 5.3. *Suppose $\mathcal{L} = \{E_1, \ldots, E_\ell\}$ with $1 \leq \ell \leq k$ is a collection of disjoint edges in $\mathcal{H}_m$. A collection of k-sets $\mathcal{C} = \{C_1, \ldots, C_\ell\}$ is a **covering** of $\mathcal{L}$ if*

- $C_i \bigcap E_j \neq \emptyset$ *for all* $1 \leq i, j \leq \ell$ *and*
- $\bigcup_{i=1}^{\ell} C_i = \bigcup_{i=1}^{\ell} E_i$.

Since $|C_1| = k$ and the E_i are disjoint, C_1 can intersect at most k possible E_i. Since C_1 must intersect every E_i, $\ell = |\{E_1, \ldots, E_\ell\}| \leq k$. Additionally, note that $\left|\bigcup_{i=1}^{\ell} C_i\right| = \left|\bigcup_{i=1}^{\ell} E_i\right| = \ell k$. If the k-sets in $\mathcal{C}$ intersected, their union would contain fewer that ℓk elements. Therefore, the elements of $\mathcal{C}$ must be pairwise disjoint.

Definition 5.4. *Define $\mathcal{L}$ as in Definition 5.3 and let $F \subseteq \cup\mathcal{L}$ be a k-set that intersects every edge of $\mathcal{L}$. The **ordered type** of F is $\vec{b} = (b_1, \ldots, b_\ell)$ where $b_i = |F \bigcap E_i|$ for $1 \leq i \leq \ell$.*

Definition 5.5. *Define $\mathcal{L}$ as in Definition 5.3 and let $F \subseteq \cup\mathcal{L}$ be a k-set that intersects every edge of $\mathcal{L}$. The **unordered type** (abbreviated **type**) of F is $\vec{a} = (a_1, \ldots, a_\ell)$ where $(a_1, \ldots, a_\ell)$ is the unique rearrangement of the ordered type $(b_1, \ldots, b_\ell)$ of F such that the entries appear in nonincreasing order. Formally, let $a_i = b_{\sigma(i)}$ where σ is a permutation of $\{1, \ldots, \ell\}$ such that $b_{\sigma(i)} \geq b_{\sigma(i+1)}$ for $1 \leq i \leq \ell - 1$.*

Example 5.6. *Consider an 8-uniform hypergraph and let $\mathcal{L} = \{E_1, E_2, E_3, E_4\}$ be the edges depicted below. Additionally consider the 8-sets F_1 and F_2 also depicted below.*

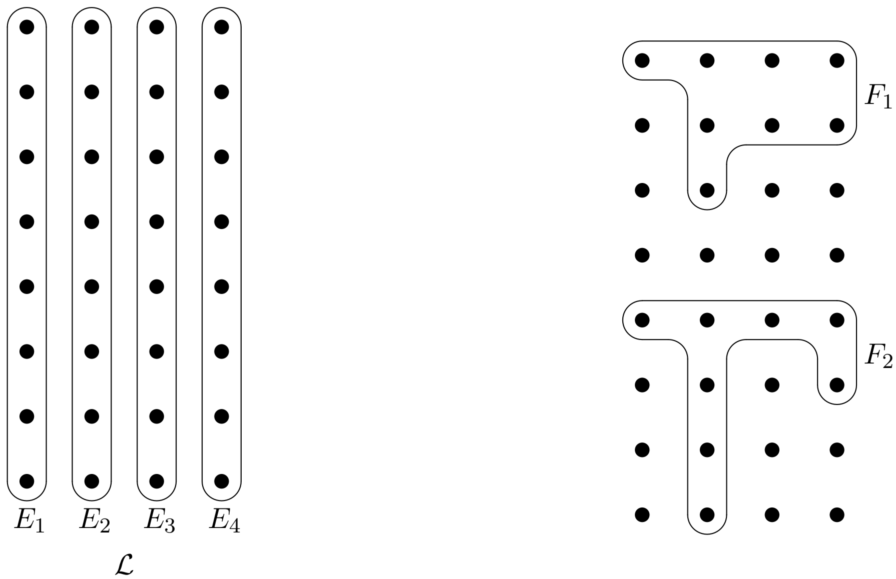

The ordered type of F_1 is $(1, 3, 2, 2)$, and the (unordered) type of F_1 is $(3, 2, 2, 1)$. The ordered type of F_2 is $(1, 4, 1, 2)$, and the (unordered) type of F_2 is $(4, 2, 1, 1)$. Note that the ordered type depends upon how we labeled the edges of $\mathcal{L}$. For example, switching the labels E_1 and E_2 would make F_1 have ordered type $(3, 1, 2, 2)$.

Define $\mathcal{L}$ as in Definition 5.3 and let $F \subseteq \cup\mathcal{L}$ be a k-set that intersects every edge of $\mathcal{L}$. Suppose F is of type $(a_1, \ldots, a_\ell)$. Note that $a_i > 0$ for $1 \leq i \leq \ell$ since F intersects every edge in $\mathcal{L}$. Additionally, since $F \subseteq \cup\mathcal{L}$ and the elements of $\mathcal{L} = \{E_1, \ldots, E_\ell\}$ are disjoint, $F \bigcap E_1, F \bigcap E_2, \ldots, F \bigcap E_\ell$ is a partition of F. Hence $\sum_{i=1}^{\ell} a_i = |F| = k$.

Let $A_{k,\ell} = \{(a_1, \ldots, a_\ell) : a_1 \geq a_2 \geq \cdots \geq a_\ell \geq 1 \text{ and } a_1 + a_2 + \cdots + a_\ell = k\}$. That is, let $A_{k,\ell}$ be the set of vectors that could potentially be the type of a k-set. Let $\mathcal{C}_{\vec{a}}$ be the set of all coverings of $\mathcal{L}$ that only contain k-sets of type $\vec{a}$. In symbols, $\mathcal{C}_{\vec{a}} = \{\mathcal{C} : \mathcal{C} \text{ is a covering of } \mathcal{L} \text{ and every } C \in \mathcal{C} \text{ is of type } \vec{a}\}$.

Example 5.7. *In Example 5.2, $\mathcal{C} = \{E_1, \ldots, E_5\}$ forms a covering of $\mathcal{L} = \{M_1, \ldots, M_5\}$. Since every 6-set in $\mathcal{C}$ is of type $(2, 1, 1, 1, 1)$, $\mathcal{C} \in \mathcal{C}_{(2,1,1,1,1)}$.*

Claim 5.8. *For every $\vec{a} \in A_{k,\ell}$, $\mathcal{C}_{\vec{a}} \neq \emptyset$.*

Proof. Given a vector $\vec{a} = (a_1, \ldots, a_\ell)$, we construct a cover of $\mathcal{L} = \{E_1, \ldots, E_\ell\}$ that is of type $\vec{a}$. Create a table that contains all ℓ cyclic permutations of $(a_1, \ldots, a_\ell)$:

a_1	a_2	a_3	$\ldots$	$a_{\ell-1}$	a_ℓ
a_ℓ	a_1	a_2	$\ldots$	$a_{\ell-2}$	$a_{\ell-1}$
$a_{\ell-1}$	a_ℓ	a_1	$\ldots$	$a_{\ell-3}$	$a_{\ell-2}$
$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$
a_2	a_3	a_4	$\ldots$	a_ℓ	a_1

Notice that every row and column sums to k. Let $w_{i,j}$ be the entry in the i^{th} row and j^{th} column of this table. Form C_1 by picking $w_{1,j}$ elements from E_j for $1 \leq j \leq \ell$. The cardinality of C_1 is the sum of the entries in the first row of the table which is k. Then, form C_2 by picking $w_{2,j}$ elements from $E_j \setminus C_1$ for $1 \leq j \leq \ell$. Proceeding in this fashion, form C_i by picking $w_{i,j}$ elements from $E_j \setminus (C_1 \bigcup C_2 \bigcup \cdots \bigcup C_{i-1})$ for $1 \leq j \leq \ell$. Because the columns sum to k, $|E_j \setminus (C_1 \bigcup C_2 \bigcup \cdots \bigcup C_{i-1})| > 0$ and there are always a sufficient number of elements remaining to form C_i for $1 \leq i \leq \ell$.

The C_i's are chosen to be disjoint subsets of $\bigcup_{i=1}^{\ell} E_i$ of type $\vec{a}$. Moreover, $|\bigcup_{i=1}^{\ell} C_i| = k\ell = |\bigcup_{i=1}^{\ell} E_i|$. Thus the C_i's must contain every vertex in $\bigcup_{i=1}^{\ell} E_i$. Hence $\mathcal{C} = \{C_1, \ldots, C_\ell\} \in \mathcal{C}_{\vec{a}}$. $\square$

6. Example of the Upper Bound

We illustrate the tools of Section 5 by proving Corollary 2.2 in the special case of a 3-uniform hypergraph on 9 vertices. Throughout the paper, we often use the notion of k-set to represent every possible edge. When discussing a concrete example, we change our language slightly. Instead of discussing k-sets, we consider complete hypergraphs which contain every possible edge. We do not lose any generality since every hypergraph is a subgraph of a complete hypergraph. Let G be the complete 3-uniform hypergraph ($\mathcal{E}(G) = \{E \in \mathcal{PV}(G) \setminus \{\emptyset\} : |E| = 3\}$) on 9 vertices and let $\mathcal{M} = \{M_1, M_2, M_3\}$ be a perfect matching in G as depicted below.

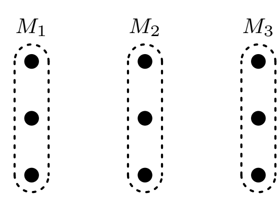

There are currently many perfect matchings in G. We wish to create a hypergraph G' that has unique perfect matching $\mathcal{M}$ by removing edges from G. In particular, we wish to remove all coverings of sets of matching edges from $\mathcal{M}$. To do this in an organized fashion, we study edges based upon their type. We first focus on removing coverings of $\mathcal{M} = \{M_1, M_2, M_3\}$. The possible types of edges (relative to $\mathcal{M}$) are $(1,1,1), (2,1,0)$, and $(3,0,0)$ as depicted below.

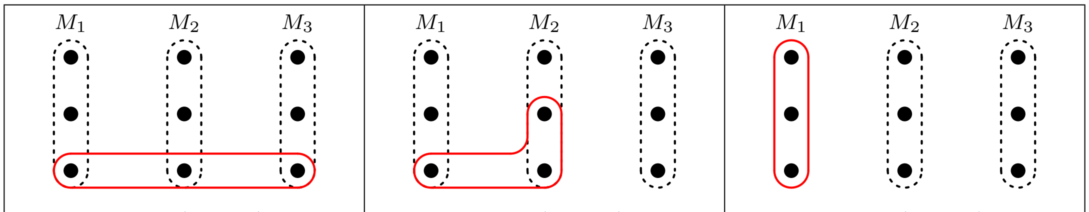

Type $(1,1,1)$ | Type $(2,1,0)$ | Type $(3,0,0)$

Edges in a covering of $\{M_1, M_2, M_3\}$ must nontrivially intersect every edge in $\{M_1, M_2, M_3\}$. Hence, the only possible type of an edge in such a covering is $(1,1,1)$.

We count the number of coverings of $\{M_1, M_2, M_3\}$ that only use edges of type $(1,1,1)$. Let $\{A, B, C\}$ be a covering of $\{M_1, M_2, M_3\}$. Suppose we have three edges that form a covering of $\{M_1, M_2, M_3\}$. There are multiple ways to assign the labels A, B, and C to these edges. However, we only wish to count this covering once. We must be careful not to double count a particular covering. To avoid this, we fix the labels A, B, and C in such a way that they cannot be permuted. In the following figure, we label each vertex by the covering edge that contains it. After possibly renaming the covering edges, we assume that the vertices in M_1 are labeled as depicted. This prevents swapping labels. For example, A is no longer interchangeable with B.

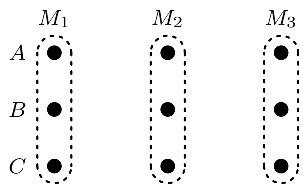

Every edge in $\{A, B, C\}$ is of type $(1,1,1)$. This means that one vertex in edge M_2 must be labeled A, one vertex in edge M_2 must be labeled B, and one vertex in M_2 must be labeled C. (We cannot duplicate or omit any label.) There are $3! = 6$ ways to assign labels to edge M_2. Similarly, there are 6 ways to assign labels to edge M_3. Every labeling corresponds to a distinct covering, giving $6 \cdot 6 = 36$ coverings of $\{M_1, M_2, M_3\}$.

Let C be a fixed edge of type $(1,1,1)$ as depicted below.

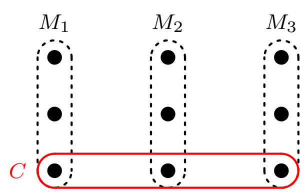

We count the number of coverings $\{A, B, C\}$ that contain edge C and only use edges of type $(1,1,1)$. Again, we do this by assigning labels to vertices. After possibly renaming the edges in the covering (to avoid double counting), we assume the vertices in M_1 are labeled as depicted below.

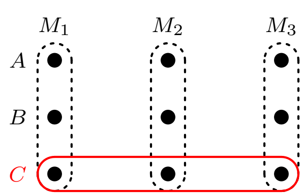

There are 2 ways to assign labels to M_2 and two ways to assign labels to M_3 yielding $2 \cdot 2 = 4$ coverings that contain edge C.

By symmetry, every edge of type $(1,1,1)$ is contained in exactly 4 coverings. Removing 1 edge removes 4 coverings from the hypergraph. Removing 2 edges removes at most 8 coverings from the hypergraph. Note that removing 2 edges can possible remove fewer than 8 coverings if the coverings removed by the first edge overlap with the coverings removed by the second edge. Removing 8 edges from the hypergraph removes at most 32 coverings. This is not enough to vanquish all 36 coverings. Hence, we must remove at least 9 edges of type $(1,1,1)$ to remove all coverings from the hypergraph.

We must also consider coverings of subsets of matching edges. For example, a covering of $\{M_1, M_2\}$ is also forbidden as it can be used to create multiple perfect matchings. We follow the same procedure as before: we organize our search by type, count the total number of coverings, and count how many coverings contain a fixed edge.

As depicted below, every edge in a covering of $\{M_1, M_2\}$ is of type $(2,1)$. Note that type $(3,0)$ is disallowed since edges in a covering must intersect every edge in $\{M_1, M_2\}$.

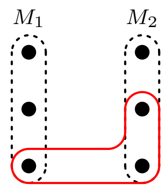

Before we proceed counting coverings, it's worth observing how useful the notions of covering and type are in this discussion. Previously when discussing covers of $\{M_1, M_2, M_3\}$, we ruled out the type $(2,1,0)$. Now we are rectifying this by considering the type $(2,1)$. Every type does get considered at the proper time. However, no type gets counted twice. Essentially, we have organized edges into disjoint sets based upon type. This helps explain the condition in Definition 5.3 requiring k-sets in a covering of $\mathcal{L}$ to nontrivially intersect every edge in $\mathcal{L}$ where $\mathcal{L}$ was described in the definition.

We now count the number of coverings of $\{M_1, M_2\}$ that only use edges of type $(2,1)$. Let E be an edge with $|E \bigcap M_1| = 1$ and $|E \bigcap M_2| = 2$. In order to specify such an edge E, we pick one vertex from M_1 in $\binom{3}{1} = 3$ ways and two vertices from M_2 in $\binom{3}{2} = 3$ ways giving 9 possible edges. Define a function f with domain $\{E \in \mathcal{E}(G) : |E \bigcap M_1| = 1 \text{ and } |E \bigcap M_2| = 2\}$ and codomain coverings of $\{M_1, M_2\}$ in which every edge is of type $(2,1)$ by $E \mapsto \{E, \overline{E}\}$.

Conversely, consider a covering of $\{M_1, M_2\}$ in which every edge is of type $(2,1)$. Exactly one edge E in a given covering has the property $|E \bigcap M_1| = 1$ and $|E \bigcap M_2| = 2$. Hence, we can define a function g with domain coverings of $\{M_1, M_2\}$ in which every edge is of type $(2,1)$ and codomain $\{E \in \mathcal{E}(G) : |E \bigcap M_1| = 1 \text{ and } |E \bigcap M_2| = 2\}$ by selecting the edge in the covering that has the proper intersection sizes. Since f and g are inverses of each other, we see that f is a bijection. Hence, the number of coverings equals the number of edges E with $|E \bigcap M_1| = 1$ and $|E \bigcap M_2| = 2$. As previously counted, this means there are 9 total coverings.

Any edge E of type $(2,1)$ is contained in exactly one covering $\{E, \overline{E}\}$. Hence, removing n edges of type $(2,1)$ removes at most n coverings from the hypergraph. In order to remove all 9 coverings, we must remove at least 9 edges. While this discussion focused on coverings of $\{M_1, M_2\}$, a similar statement holds for any set of two edges from the matching $\mathcal{M}$. There are $\binom{3}{2}$ ways to pick a set of two edges from $\mathcal{M}$. Hence we must remove $\binom{3}{2} \cdot 9 = 27$ edges of type $(2,1)$ from G.

The complete 3-uniform hypergraph on 9 vertices has $\binom{9}{3} = 84$ edges. We must remove 9 edges of type $(1,1,1)$ and 27 edges of type $(2,1)$. There are at most $84 - 9 - 27 = 48$ edges remaining in the hypergraph. This exactly agrees with the bound in Corollary 2.2 since

$$f(3,3) = 3 + 9\binom{3}{2} + 18\binom{3}{3} = 48$$

where $f(k,m)$ was defined in the corollary.

The above discussion also justifies the edge bound in the case of a 3-uniform hypergraph on 6 vertices with unique perfect matching $\{M_1, M_2\}$. As before, every covering of this perfect matching is of type $(2,1)$. There are 9 such coverings in a complete hypergraph. Because every edge of type $(2,1)$ is present in exactly 1 covering, removing n edges of type $(2,1)$ removes at most n coverings. Hence, we must remove at least 9 edges of type $(2,1)$ from the complete hypergraph in order to remove all coverings. There are $\binom{6}{3} = 20$ edges in the complete 3-uniform hypergraph on 6 vertices. After removing at least 9 edges, at most 11 edges remain. This agrees with the bound in Corollary 2.2 since

$$f(3,2) = 2 + 9\binom{2}{2} + 18\binom{2}{3} = 11$$

where $f(k,m)$ was defined in the corollary. Note that $\binom{2}{3} = 0$.

7. Proof of Upper Bound

The construction in Section 3 establishes part of Corollary 2.2 by showing the bound is attainable by a hypergraph that has a unique perfect matching. In this section, we prove Theorem 2.1 and complete the proof of Corollary 2.2.

Throughout the entirety of this section, let $\mathcal{H}_m = (\mathcal{V}_m, \mathcal{E}_m)$ be an arbitrary k-uniform hypergraph on km vertices with unique perfect matching $\mathcal{M} = \{M_1, \ldots, M_m\}$.

Proof of Theorem 2.1. As in the statement of the theorem, let $\mathcal{B}_\ell$ be the set of edges that intersect exactly ℓ matching edges for $1 \le \ell \le k$. That is, $\mathcal{B}_\ell = \{E \in \mathcal{E}_m : |Q_E| = \ell\}$ where $Q_E = \{M_i \in \mathcal{M} : M_i \bigcap E \neq \emptyset\}$. As in the proof of Claim 3.4, $\mathcal{B}_1 = \{M_1, \ldots, M_m\}$ and $|\mathcal{B}_1| = m$.

Suppose for the purpose of a contradiction that $|\mathcal{B}_\ell| > b_{k,\ell}\binom{m}{\ell}$ for some $2 \le \ell \le k$ where $b_{k,\ell}$ was defined in the statement of Theorem 2.1. There are m edges in the matching $\mathcal{M}$ giving $\binom{m}{\ell}$ ways to pick ℓ matching edges. Let $\mathcal{L}_1, \ldots, \mathcal{L}_{\binom{m}{\ell}}$ be all the possible sets of ℓ matching edges. For every $E \in B_\ell$, E intersects exactly ℓ matching edges. In symbols, $|Q_E| = \ell$. So, $Q_E = \mathcal{L}_i$ for some $1 \le i \le \binom{m}{\ell}$. We partition the edges $E \in \mathcal{B}_\ell$ based upon which matching edges E intersects:

$$\mathcal{B}_\ell = \underbrace{\left(\mathcal{B}_\ell \bigcap \mathcal{E}\left(\mathcal{H}_m[\mathcal{L}_1]\right)\right)}_{\substack{\text{Edges } E \in \mathcal{B}_\ell \\ \text{where } Q_E = \mathcal{L}_1}} \bigcup \underbrace{\left(\mathcal{B}_\ell \bigcap \mathcal{E}\left(\mathcal{H}_m[\mathcal{L}_2]\right)\right)}_{\substack{\text{Edges } E \in \mathcal{B}_\ell \\ \text{where } Q_E = \mathcal{L}_2}} \bigcup \cdots \bigcup \underbrace{\left(\mathcal{B}_\ell \bigcap \mathcal{E}\left(\mathcal{H}_m[\mathcal{L}_{\binom{m}{\ell}}]\right)\right)}_{\substack{\text{Edges } E \in \mathcal{B}_\ell \text{ where} \\ Q_E = \mathcal{L}_{\binom{m}{\ell}}}}$$

If $|\mathcal{B}_\ell \bigcap \mathcal{E}\left(\mathcal{H}_m[\mathcal{L}_i]\right)| \le b_{k,\ell}$ for every $1 \le i \le \binom{m}{\ell}$, we would have

$$\begin{aligned} |\mathcal{B}_\ell| &= \left|\mathcal{B}_\ell \bigcap \mathcal{E}\left(\mathcal{H}_m[\mathcal{L}_1]\right)\right| + \left|\mathcal{B}_\ell \bigcap \mathcal{E}\left(\mathcal{H}_m[\mathcal{L}_2]\right)\right| + \cdots + \left|\mathcal{B}_\ell \bigcap \mathcal{E}\left(\mathcal{H}_m[\mathcal{L}_{\binom{m}{\ell}}]\right)\right| \\ &\le b_{k,\ell} + b_{k,\ell} + \cdots + b_{k,\ell} \\ &= \binom{m}{\ell} b_{k,\ell} \end{aligned}$$

contrary to our assumptions. Thus (by the pigeonhole principle) there exists at least one i with $1 \le i \le \binom{m}{\ell}$ and

$$\left|\mathcal{B}_\ell \bigcap \mathcal{E}\left(\mathcal{H}_m[\mathcal{L}_i]\right)\right| \ge b_{k,\ell} + 1 \tag{1}$$

Using similar counting techniques as Claim 3.4, let $\mathcal{G}$ be the collection of k-sets of vertices in $\mathcal{H}_m[\mathcal{L}_i]$ that intersect all matching edges in $\mathcal{L}_i$. That is

$$\mathcal{G} = \{A : |A| = k, A \subseteq \mathcal{V}(\mathcal{H}_m[\mathcal{L}_i]), \text{ and } \forall M \in \mathcal{L}_i, A \bigcap M \neq \emptyset\}$$

Let $\mathcal{G}_{\vec{a}}$ be the set of k-sets in $\mathcal{G}$ of type $\vec{a}$ and note that $\{\mathcal{G}_{\vec{a}} : \vec{a} \in A_{k,\ell}\}$ partitions $\mathcal{G}$. Using the formula in Claim 3.4

$$b_{k,\ell} = \frac{\ell-1}{\ell}|\mathcal{G}| = \frac{\ell-1}{\ell} \sum_{\vec{a} \in A_{k,\ell}} |\mathcal{G}_{\vec{a}}|$$

where $b_{k,\ell}$ was defined in Theorem 2.1. Substituting into Equation 1 yields

$$\left|\mathcal{B}_\ell \bigcap \mathcal{E}\left(\mathcal{H}_m[\mathcal{L}_i]\right)\right| \geq \frac{\ell-1}{\ell} \sum_{\vec{a} \in A_{k,\ell}} |\mathcal{G}_{\vec{a}}| + 1 \tag{2}$$

Suppose $|\mathcal{B}_\ell \bigcap \mathcal{G}_{\vec{a}}| \leq \frac{\ell-1}{\ell}|\mathcal{G}_{\vec{a}}|$ for all $a \in A_{k,\ell}$. Note that $\mathcal{B}_\ell \bigcap \mathcal{E}\left(\mathcal{H}_m[\mathcal{L}_i]\right) \bigcap \mathcal{G}_{\vec{a}} \subseteq \mathcal{B}_\ell \bigcap \mathcal{G}_{\vec{a}}$. Then

$$\begin{aligned}
\left|\mathcal{B}_\ell \bigcap \mathcal{E}\left(\mathcal{H}_m[\mathcal{L}_i]\right)\right| &= \sum_{a \in A_{k,\ell}} \left|\mathcal{B}_\ell \bigcap \mathcal{E}\left(\mathcal{H}_m[\mathcal{L}_i]\right) \bigcap \mathcal{G}_{\vec{a}}\right| \\
&\leq \sum_{a \in A_{k,\ell}} \left|\mathcal{B}_\ell \bigcap \mathcal{G}_{\vec{a}}\right| \\
&\leq \sum_{a \in A_{k,\ell}} \frac{\ell-1}{\ell}|\mathcal{G}_{\vec{a}}| \\
&= \frac{\ell-1}{\ell}|\mathcal{G}|
\end{aligned}$$

contradicting Equation 2. So (by the pigeonhole principle) there is at least one vector $\vec{a} \in A_{k,\ell}$ with

$$|\mathcal{B}_\ell \bigcap \mathcal{G}_{\vec{a}}| \geq \frac{\ell-1}{\ell}|\mathcal{G}_{\vec{a}}| + 1 \tag{3}$$

As before, let $\mathcal{C}_{\vec{a}}$ be the nonempty set of coverings of $\mathcal{L}_i$ such that every $A \in \mathcal{C} \in \mathcal{C}_{\vec{a}}$ is of type $\vec{a}$. Recall that $|\mathcal{C}| = \ell$ for every $\mathcal{C} \in \mathcal{C}_{\vec{a}}$.

By symmetry, every k-set $A \in \mathcal{G}_{\vec{a}}$ belongs to λ coverings $\mathcal{C} \in \mathcal{C}_{\vec{a}}$ where λ is some constant. We count the number of ordered pairs $(A, \mathcal{C})$ where $A \in \mathcal{C} \in \mathcal{C}_{\vec{a}}$ in two ways. There are $|\mathcal{G}_{\vec{a}}|$ ways to pick a k-set A. Once A is fixed, there are λ coverings that could pair with A. Thus, there are $|\mathcal{G}_{\vec{a}}|\lambda$ ordered pairs. Alternately, there are $|C_{\vec{a}}|$ ways to pick a covering $\mathcal{C} \in \mathcal{C}_{\vec{a}}$. Once a covering is known, there are ℓ ways to pick a k-set contained in the covering. This yields $|\mathcal{C}_{\vec{a}}|\ell$ possible ordered pairs. Together, this implies every k-set in $\mathcal{G}_{\vec{a}}$ belongs to exactly

$$\lambda = \frac{|\mathcal{C}_{\vec{a}}|\ell}{|\mathcal{G}_{\vec{a}}|}$$

coverings $\mathcal{C} \in \mathcal{C}_{\vec{a}}$.

No covering $\mathcal{C} \in \mathcal{C}_{\vec{a}}$ can be contained in $\mathcal{E}_m$. Otherwise $(\mathcal{M} \setminus \mathcal{L}_i) \bigcup \mathcal{C}$ would form a perfect matching in $\mathcal{H}_m$ contradicting the uniqueness of $\mathcal{M}$. So there must be at least one k-set in every covering $\mathcal{C} \in \mathcal{C}_{\vec{a}}$ that is not an edge in $\mathcal{H}_m$.

One k-set is contained in λ coverings $\mathcal{C} \in \mathcal{C}_{\vec{a}}$. Two k-sets are contained in at most 2λ coverings $\mathcal{C} \in \mathcal{C}_{\vec{a}}$. (Each k-set is individually contained in λ coverings, but there may be a covering containing both of the k-sets that gets double counted.) In general, n k-sets are contained in at most $n\lambda$ coverings. In order to have a k-set from every covering, we must have $n\lambda \geq |\mathcal{C}_{\vec{a}}|$. Hence $n \geq \dfrac{|\mathcal{C}_{\vec{a}}|}{\lambda} = \dfrac{|\mathcal{G}_{\vec{a}}|}{\ell}$.

Since $\mathcal{H}_m$ lacks coverings, there must be at least $\dfrac{|\mathcal{G}_{\vec{a}}|}{\ell}$ k-sets of type $\vec{a}$ that are not edges in the hypergraph. Hence

$$\underbrace{|\mathcal{B}_\ell \bigcap \mathcal{G}_{\vec{a}}|}_{\substack{\text{Edges } E \in \mathcal{E}(\mathcal{H}_m) \\ \text{with } Q_E = \mathcal{L}_i \\ \text{and type } \vec{a}}} \leq \underbrace{|\mathcal{G}_{\vec{a}}|}_{\substack{\text{All } k\text{-sets } F \text{ with} \\ F \subseteq \cup \mathcal{L}_i \text{ and} \\ \forall M \in \mathcal{L}_i, F \bigcap M \neq \emptyset \\ \text{that have type } \vec{a}}} - \underbrace{\frac{|\mathcal{G}_{\vec{a}}|}{\ell}}_{\substack{\text{Remove at least} \\ \text{one } k\text{-set from} \\ \text{each covering of } \mathcal{L}_i}}$$

$$= \frac{\ell - 1}{\ell} |\mathcal{G}_{\vec{a}}|$$

contradicting Equation 3. □

Proof of Corollary 2.2. Define $\mathcal{B}_\ell$ for $1 \leq \ell \leq k$ as in Theorem 2.1. As noted in the proof of Claim 3.4, $\{\mathcal{B}_1, \ldots, \mathcal{B}_k\}$ forms a partition of the edges of $\mathcal{H}_m$. Then

$$|\mathcal{E}_m| = \sum_{i=1}^{k} |\mathcal{B}_i|$$

By Theorem 2.1

$$|\mathcal{E}_m| \leq m + \sum_{i=2}^{k} b_{k,\ell} \binom{m}{\ell}$$

□

8. Miscellaneous Related Results

This section presents some related results that extend Corollary 2.2 to other situations.

Theorem 8.1. *Let $G = (\mathcal{V}, \mathcal{E})$ be a 2-uniform hypergraph with $2m+1$ vertices and a nearly perfect matching $\mathcal{M} = \{M_1, \ldots, M_m\}$. If G lacks isolated vertices, then G has at least two nearly perfect matchings.*

Proof. Suppose $\bigcup_{i=1}^{m} M_i = \mathcal{V} \setminus \{v\}$ for some $v \in \mathcal{V}$. Then there is a matching edge incident with every vertex in the graph other than v (as depicted below with solid lines). Since v is not isolated, there is an edge E connecting v to some other vertex u (depicted by a dotted line below). Let $F = \{u, w\}$ be the matching edge incident with u.

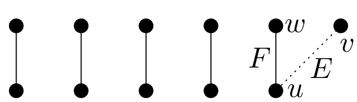

Create a new nearly perfect matching by trading edge F for E. That is, $\mathcal{M}' = \{E\} \bigcup (\mathcal{M} \setminus \{F\})$ is a distinct nearly perfect matching which contains an edge incident with every vertex except w. □

Corollary 8.2. *Let G be a k-uniform hypergraph with km vertices and exactly two perfect matchings $\mathcal{M}_1$ and $\mathcal{M}_2$. Then $|\mathcal{E}(G)| \leq f(k, m) + 1$ where $f(k, m)$ was defined in Corollary 2.2.*

Proof. Since $\mathcal{M}_1$ and $\mathcal{M}_2$ both contain m edges and $\mathcal{M}_1 \neq \mathcal{M}_2$, there exists some edge $E \in \mathcal{M}_2 \setminus \mathcal{M}_1$. Create a new hypergraph G' be removing edge E from G. The matching $\mathcal{M}_2$ is no longer present in G'. Hence, G' has a unique perfect matching $\mathcal{M}_1$. Then, by Corollary 2.2, $|\mathcal{E}(G')| \leq f(k, m)$. Since G has one more edge than G', $|\mathcal{E}(G)| \leq f(k, m) + 1$. □

Theorem 8.3. *Let G be a k-uniform hypergraph with $2k$ vertices that lacks a perfect matching. Then $|\mathcal{E}(G)| \leq \binom{2k-1}{k}$. Furthermore, this bound is tight.*

Proof. Let G be a k-uniform hypergraph with $2k$ vertices that lacks a perfect matching. Fix a vertex $v \in \mathcal{V}(G)$ and enumerate the edges incident with v as $\{E_1, E_2, \ldots, E_m\}$. For each E_i with $1 \leq i \leq m$, $\overline{E_i}$ cannot be an edge in the hypergraph. Otherwise, $\{E_i, \overline{E_i}\}$ would form a perfect matching in G. Let G' be the hypergraph with vertex set $\mathcal{V}(G)$ and edge set $(\mathcal{E}(G) \setminus \{E_1, E_2, \ldots, E_m\}) \bigcup \{\overline{E_1}, \overline{E_2}, \ldots, \overline{E_m}\}$. That is, trade edge E_i for edge $\overline{E_i}$ for $1 \leq i \leq m$. Removing one edge E_i and adding another edge $\overline{E_i}$ to the edge set does not change the total number of edges in a hypergraph. Hence, $|\mathcal{E}(G)| = |\mathcal{E}(G')|$. Note that v is an isolated vertex in G'. Therefore, every edge in G' only contains elements from $\mathcal{V}(G') \setminus \{v\}$ which has cardinality $2k - 1$. There are at most $\binom{2k-1}{k}$ such edges.

We construct a hypergraph to demonstrate that the bound is attainable. Let H be the complete k-uniform hypergraph on $2k - 1$ vertices. That is $|\mathcal{V}(H)| = 2k - 1$ and $\mathcal{E}(H) = \{E \in \mathcal{P}\mathcal{V}(H) \setminus \{\emptyset\} : |E| = k\}$. Appending an isolated vertex to H creates a hypergraph with $2k$ vertices and $\binom{2k-1}{k}$ edges. This hypergraph cannot have a perfect matching because it contains an isolated vertex. $\square$

9. Application

As discussed in [3], enumerating perfect matchings arises naturally when studying properties of molecules. To illustrate this, we consider the molecule benzene (C_6H_6). Carbon molecules tend to bond to four other molecules while hydrogen molecules only bond to one other molecule. If we think of the physical structure of this molecule as a graph, this means we must have 6 vertices of degree 4 and 6 vertices of degree 1. This alone is not sufficient to determine the structure of the molecule. Experimental data has suggested that the molecule is roughly a planar ring. It also seems that all of the carbons are in equivalent positions in the molecule. Similarly the hydrogens seem to be in equivalent positions as well. One potential model is to place the carbons in a cycle of length 6 and attach one hydrogen to each carbon as depicted below.

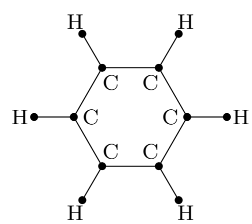

This symmetric planar ring configuration is further supported by experimental evidence. Hydrogen atoms in this molecule can be replaced by OH groups in order to form new molecules. As depicted below, there are three distinct ways to exchange two hydrogens for two OH groups.

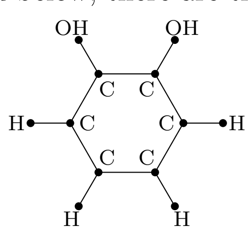

Type I

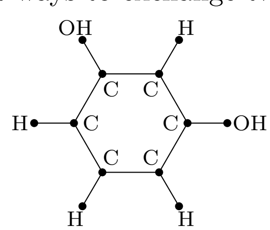

Type II

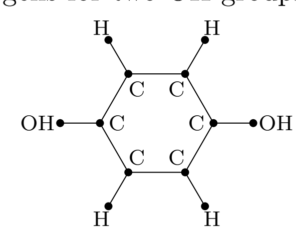

Type III

In Type I, the two OH groups are bonded to adjancent carbons. In Type II, the two carbon atoms that are bonded with the OH groups are distance 2 apart. In Type III, the two carbon atoms bonded with the OH groups are distance 3 apart. These three configurations lead to three distinct isomers of this molecule, each of which exhibits slightly different properties.

Unfortunately, this model as currently stated has an obvious flaw. The carbon molecules which bond 4 times must be represented by degree 4 vertices in the graph. As previously depicted, the carbon vertices only have degree 3. We are missing some bonds. Specifically, there must be additional bonds between the carbon atoms. To rectify this, we introduce double bonds into the cycle of carbons. There are two ways to do this:

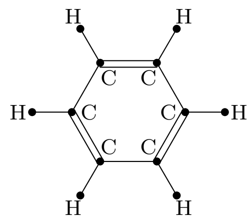

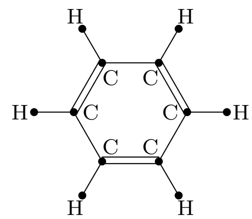

Note that the double bonds between the carbon atoms correspond to a perfect matching in the 6-cycle. The different possibilities for the structure arise because the 6-cycle does not have a unique perfect matching.

While these new bonds do rectify the degree constraints of the graph, they unfortunately remove some of the symmetry previously present in the graph. In particular, there are now four distinct ways to remove two hydrogen atoms and replace them with two OH groups:

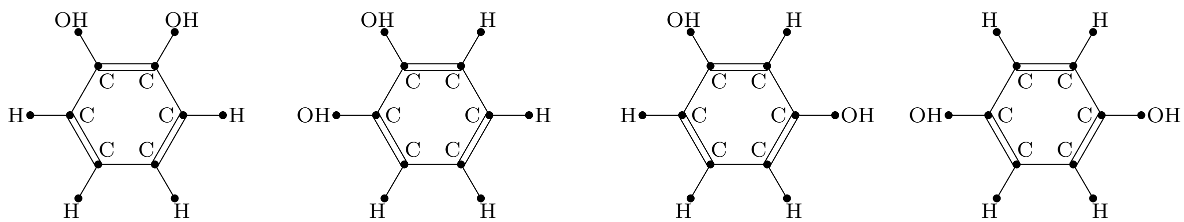

Type I Type II Type III Type IV

Let C_1 and C_2 be the two carbons adjacent with the OH group. In Type I, C_1 and C_2 are adjacent and double bonded. In Type II, C_1 and C_2 are adjacent and single bonded. In Type III, the distance between C_1 and C_2 is 2. In Type IV, the distance between C_1 and C_2 is 3. Because of the double bonding, there are now two types of adjacent carbon atoms.

Unfortunately, fixing the degrees of the carbon vertices has created a nonphysical result. In nature, replacing two hydrogens by two OH groups only creates three types of molecules, not four. The above model must also be flawed. To rectify this, the carbons in the benzene molecule are believed to resonate between the two possible states. Suppose carbon atom C_1 is adjacent to carbon atoms C_2 and C_3. Sometimes, C_1 is double bonded with C_2. Other times, C_1 is double bonded with C_3. This means that any pair of adjacent carbon atoms are sometimes double bonded and sometimes single bonded. This merges Types I and II above, reducing the number of possibilities to three. This once again agrees with physical observations.

In graph theoretical terms, the bonds in the carbon cycle in benzene resonate in such a way that corresponds to oscillating between the two perfect matchings in the 6-cycle. This configuration seems to correspond to physical properties of the molecule. Benzene is incredibly stable. It seems that alternating between perfect matchings increases the stability of a molecule. In fact, the resonance conjecture states that the stability of molecules resembling benzene is directly proportional to the number of perfect matchings present in the molecular structure. Such molecules with a unique perfect matching tend to be less stable.

10. Further Study

Many interesting related questions remain. One natural extension is to ask for a formula bounding the maximum number of edges allowed in a hypergraph with exactly k perfect matchings for some natural number k. This paper addresses the cases $k = 1$ in Corollary 2.2 and $k = 2$ in Corollary 8.2.

In the 2-uniform case, the idea of a k-factor provides an alternate way to generalize the problem. The **degree** of a vertex v in a 2-uniform hypergraph G is the number of vertices adjacent to v. A **k-regular** graph is a graph where every vertex has degree k. A **k-factor** of a graph G is a k-regular subgraph of G that has vertex set $\mathcal{V}(G)$. A 1-factor is a decomposition of a graph into paths of length 1. That is, a 1-factor is essentially the same as a perfect matching. A 2-factor decomposes a graph into disjoint cycles. We may also consider an upper bound on the number of edges present in a graph with a unique k-factor.

References

[1] Bal, Deepak, Andrzej Dudek, Zelealem B. Yilma, On the maximum number of edges in a hypergraph with a unique perfect matching, Discrete Mathematics 311 (2011) 2577-2580. 5

[2] Lovász, L., On the structure of factorizable graphs. I, Acta Mathematica Academiae Scientiarum Hungaricae 23 (1972) 179-195. 12

[3] Merris, Russell, Graph Theory, Wiley-Interscience, New York, 2001. 25

[4] West, Douglas B., Introduction to Graph Theory (Second Edition), Prentice Hall, Upper Saddle River, NJ, 2001. 12