Freie Universität Berlin

Extremal Hypergraphs for Ryser's Conjecture

Dissertation vorgelegt von

Lothar Narins

Betreut durch

Prof. Tibor Szabó, PhD

Eingereicht am

Fachbereich der Mathematik und Informatik

der Freien Universität Berlin

Im September 2014

Erstgutachter und Betreuer: Prof. Tibor Szabó, PhD

Zweitgutachter: Prof. Ron Aharoni, PhD

Tag der Disputation: 15.12.2014

Contents

Preface

Ryser's Conjecture from the year 1971 is that the inequality $\tau(\mathcal{H}) \leq (r-1)\nu(\mathcal{H})$ holds for every r-partite r-uniform hypergraph $\mathcal{H}$, where $\tau(\mathcal{H})$ and $\nu(\mathcal{H})$ represent the vertex cover number and the matching number, respectively. The conjecture is still wide open, though advances in various directions have been made by Aharoni, Berger, Füredi, Haxell, Lovász, Mansour, Scott, Song, Tuza, Yuster, and Ziv, among others. In 1999, Aharoni gave a proof of the conjecture for the case $r = 3$.

The main result of this thesis is the characterization of all 3-partite 3-uniform hypergraphs $\mathcal{H}$ for which $\tau(\mathcal{H}) = 2\nu(\mathcal{H})$, in other words, the extremal hypergraphs for Ryser's Conjecture for $r = 3$. These all have a special form, which we call "home-base" hypergraphs. They consist of $\nu(\mathcal{H})$ subhypergraphs, each with $\tau = 2$ and $\nu = 1$, together with possibly some extra hyperedges that intersect these parts in a very particular fashion. Along the way towards the proof of this characterization, we also find a characterization of all bipartite graphs that are extremal for a certain topological problem.

For both characterizations, we utilize knowledge about the topology of the independence complex $\mathcal{I}$ of line graphs. For this reason, we next investigate a lower bound on the connectedness of $\mathcal{I}(L(\mathcal{H}))$ with respect to $\tau(\mathcal{H})$. We conjecture that this bound can be improved in the case of r-partite r-uniform hypergraphs, and we verify the conjecture for the special cases $r = 3$ and $\tau(\mathcal{H}) \leq 12$.

A theorem of Meshulam that concerns the connectedness of the independence complex of a graph plays an important role in our proofs. The proof of this theorem that one finds in the literature is rather algebraic. We give a more geometric proof using certain triangulation techniques. The correctness of these methods, which were used for instance by Szabó and Tardos, has recently come into question. In the last part of this thesis, we provide a thorough proof of their correctness.

Acknowledgements

First and foremost, I would like to thank my advisor, Tibor Szabó, for all of his help guiding me throughout the past three years of research, for guiding me towards interesting and fruitful problems, for organizing incredibly stimulating workshops, and for motivating me during hard times.

I am grateful to the graduate school Methods for Discrete Structures for generously providing my funding throughout my PhD research period, and also for supporting me on several research visits. I also greatly appreciate the Berlin Mathematical School for bringing me to Berlin and funding my first two years here.

I would like to thank my co-authors Penny Haxell, Alexey Pokrovskiy, and Tuan Tran for enjoyable times working together on various projects. I extend my gratitute as well to all of the people in Tibor's working group for providing a nice work environment and making me feel so included. I would also like to thank Anita Liebenau for invaluable input on the summary of this very thesis.

I also wish to thank Jiří Matoušek and the group of Emo Welzl for a wonderful research stay in Zurich.

I am grateful also to Tillmann Miltzow and Peter Patzt for being good friends and for sharing their math with me.

And of course I could not have done this without my parents' loving support. You guys are amazing.

Chapter 1

Introduction

1.1 Ryser's Conjecture

A hypergraph $\mathcal{H}$ is a pair (V, E), where $V = V(\mathcal{H})$ is the set of *vertices*, and $E = E(\mathcal{H})$ is a **multiset** of subsets of vertices called the *edges* of $\mathcal{H}$. The number of times a subset $e \subseteq V$ appears in E is called the *multiplicity* of e. If the cardinality of every edge is r, we call $\mathcal{H}$ an *r-uniform* hypergraph, or *r-graph* for short. A 2-graph is called a *graph*. We mostly have no restriction on the multiplicity of edges; whenever we want to assume that each multiplicity is at most 1, we will explicitly say *simple hypergraph*, *simple r-graph*, or *simple graph*. An edge $e \in E$ is called *parallel* to an edge $f \in E$ if their underlying vertex subsets are the same. In particular, every edge is parallel to itself.

Let $\mathcal{H}$ be a hypergraph. A *matching* in $\mathcal{H}$ is a set of disjoint edges of $\mathcal{H}$, and the *matching number*, $\nu(\mathcal{H})$, is the size of the largest matching in $\mathcal{H}$. If $\nu(\mathcal{H}) = 1$, then $\mathcal{H}$ is called *intersecting*. A *vertex cover* of $\mathcal{H}$ is a set of vertices which intersects every edge of $\mathcal{H}$. The size of the smallest vertex cover is called the *vertex cover number* of $\mathcal{H}$ and is denoted by $\tau(\mathcal{H})$. It is immediate to see that if $\mathcal{H}$ is r-uniform, then the following bounds always hold:

$$\nu(\mathcal{H}) \leq \tau(\mathcal{H}) \leq r\nu(\mathcal{H}).$$

Both inequalities are easily seen to be tight for general hypergraphs. Ryser's Conjecture (see e.g. [29]), which appeared first in the early 1970's, states that the upper bound can be lowered by considering only r-partite hypergraphs. (An even stronger conjecture was made around the same time by Lovász [20].) An r-graph is called *r-partite* if its vertices can be partitioned into r parts called *vertex classes* such that every edge intersects each vertex class in exactly one vertex.

Conjecture 1 (Ryser's Conjecture)**.** *If $\mathcal{H}$ is an r-partite r-graph, then*

$$\tau(\mathcal{H}) \leq (r-1)\nu(\mathcal{H}).$$

This conjecture turned out to be extremely difficult to attack. It is solved completely only for $r = 2$ and 3, and a few partial results exist for $r = 4$ and 5. The conjecture is wide open for $r \geq 6$. In particular, when $r = 2$, the conjecture is just the well known König's Theorem. It has been proven for intersecting hypergraphs when $r \leq 5$ by Tuza ([28, 29]), with $r \geq 6$ still open. The general case of the conjecture for $r = 3$ was solved by Aharoni via topological methods [2]. Fractional versions of the conjecture have also been studied, and it was shown by Füredi [12] that $\tau^* \leq (r-1)\nu$, and shown by Lovász [20] that $\tau \leq \frac{r}{2}\nu^*$, where τ^* and ν^* are the fractional vertex cover and matching numbers, respectively. Aharoni and Berger [3] also formulated a generalization of the conjecture to matroids, which has been partially solved in a special case by Berger and Ziv [8]. Mansour, Song, and Yuster [21] have found bounds on the minimum number of edges for an intersecting r-partite r-graph to be tight for Ryser's conjecture, with exact numbers known only for the cases $r \leq 5$. Haxell and Scott [17] have proven that for $r = 4, 5$ there is an $\epsilon > 0$ such that $\tau(\mathcal{H}) \leq (r-\epsilon)\nu(\mathcal{H})$ for any r-partite r-graph $\mathcal{H}$.

One plausible approach to Ryser's Conjecture for 4-graphs is via studying the 3-uniform link hypergraphs. Given three of the four vertex classes V_1, V_2, V_3 of a 4-partite 4-graph $\mathcal{H}$, the link hypergraph of V_4 in $\mathcal{H}$ is the multiset of those 3-element sets which are the intersection of an edge of $\mathcal{H}$ with $V_1 \cup V_2 \cup V_3$. Having structural information on the links would be helpful in understanding the situation for 4-graphs. Aharoni's proof however does not provide information on the 3-graphs which are extremal for his theorem. Our eventual aim is to give a complete characterization of them.

We say that a 3-partite 3-graph $\mathcal{H}$ is *Ryser-extremal* if $\tau(\mathcal{H}) = 2\nu(\mathcal{H})$. There are two types of intersecting Ryser-extremal 3-graphs. One is the truncated Fano plane F, shown below:

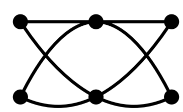

Figure 1.1: The truncated Fano plane F.

One may remove any edge from F, and the resulting hypergraph is still Ryser-extremal. Call this hypergraph H, and note that H has three vertices of degree 2. Note that every edge of H contains two of these three vertices. Adding edges that intersect at least two of these three vertices yields another Ryser-extremal 3-graph, as pictured below:

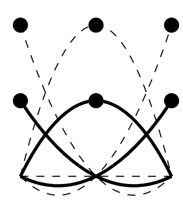

Figure 1.2: The 3-graph H, with possible additional edges drawn in dashed lines.

The main result of the thesis is that all Ryser-extremal 3-partite 3-graphs are built out of these two types of hypergraphs. This motivates the definition of a *home-base hypergraph*:

Definition 1.1.1. A *home-base hypergraph* is a 3-partite 3-graph $\mathcal{H}$ consisting of $\nu(\mathcal{H})$ disjoint copies of F and H, possibly together with some additional edges, each of which contain two degree 2 vertices of some copy of H.

We call hypergraphs of this form home-base hypergraphs because every edge has a unique copy of F or H that it can consider its "home." In Chapter 3, we will give a slightly different definition of home-base hypergraphs, which can easily be seen to be equivalent to the one given here. The definition in Chapter 3 is designed to highlight the parts of the structure that are uniquely determined. The main part of this thesis is devoted to proving the following theorem:

Theorem 1.1.2. *Let $\mathcal{H}$ be a 3-partite 3-graph. Then $\tau(\mathcal{H}) = 2\nu(\mathcal{H})$ if and only if $\mathcal{H}$ is a home-base hypergraph.*

In Chapter 2 we develop the necessary knowledge about the link graphs of Ryser-extremal 3-graphs. First we show that these link graphs are extremal with respect to a natural extremal graph theoretic problem of topological nature, namely the topological connectedness of the independence complexes of their line graphs is lowest possible for their matching number (the topological terms will be explained in Chapter 2). In Chapter 2, we characterize all those bipartite graphs that are extremal for this problem (Theorem 2.1.4). The structure we derive from this characterization theorem will be an integral part of our proof of Theorem 1.1.2 in Chapter 3. Nevertheless, we find the extremal graph theory problem interesting in its own right.

In Chapter 4 we discuss a related extremal problem of a topological nature, which may possibly yield some insight into the case of 4-partite 4-graphs. In it, we find a lower bound on the connectedness of the independence complex of line graphs in terms of the vertex cover number of the hypergraph. The bound is tight for general r-graphs, but there is hope to improve it for r-partite r-graphs.

Then, in Chapter 5 we give a solid foundation to certain proof techniques involving triangulations of spheres, and give a triangulation proof of Meshulam's Theorem (Theorem 2.1.5), which relates the connectedness of the independence complex of a graph to that of certain modifications of the graph.

1.1.1 Extremal Problems with Many Extremal Structures

Before beginning the main work of this thesis, we take a moment here to reflect on why it is that our task of characterizing Ryser-extremal hypergraphs seems to require rather complex arguments. For many of the questions of extremal combinatorics that are solved, there is a unique example that provides the extremal value. In such cases, a proof of optimality can be guided by the properties and features of this extremal structure. The situation becomes more complex for problems in which there there are two or more very different extrema. Then a purely combinatorial argument is less and less likely to succeed, because the proof must eventually consider all the extremal structures. For our characterization problem, the number of extremal structures is infinite for every fixed value of the benchmark parameter. This is one of the few cases in which the full characterization of the extremal structures of an extremal combinatorial problem with infinitely many extrema is known.

On rare occasions, the difficulties posed by multiple extremal examples can be mitigated by realizing that the combinatorial problem, or rather its extremal structures, hide the features and concepts of another mathematical discipline in the background. In such cases, the extremal structures can be described more naturally in "another mathematical language," making a translation back to the language of combinatorics at least a possibility.

A simple example of a problem of this sort is the one described by the famous Oddtown Theorem of Berlekamp [9]. The problem asks for the maximum size of a family of subsets of odd cardinality in an n-element base set, such that the intersection of any two members of the family has even cardinality. It turns out that this problem can easily be solved by a simple application of linear algebra, even though there are superexponentially many extremal structures [7, Exercise 1.1.14]. A combinatorial characterization of the extremal families however is still outstanding, and it is questionable whether it is feasible at all.

Another prominent example is the extremal problem known as Sidorenko's Conjecture [25, 24]. Roughly speaking, Sidorenko's Conjecture asks for the minimum number of copies of a fixed bipartite graph H in a "large" graph on n vertices with $m = \Theta(n^2)$ edges. The conjecture states that for every bipartite graph H the minimum is essentially taken by quasirandom graphs. Sidorenko's conjecture is known to hold for many bipartite graphs, for example trees, even cycles, the hypercube, complete bipartite graphs, but wide open in general; see [19] and its references. Since the random graph $G(n, m)$ is conjectured to be essentially extremal for the problem, it is then also plausible to expect that there are many combinatorially different extremal or close to extremal constructions and hence a combinatorial characterization of the extremal examples seems out of reach. However, in the analytic language of graph limits, where graphs are interpreted as symmetric measurable functions on the unit square (called *graphons*), the asymptotically equivalent formulation of Sidorenko's Conjecture is conjectured to have a unique extremal graphon (for every bipartite graph H with a cycle): the constant function $2m/n^2$. This stronger uniqueness statement, called the forcing conjecture, is also known to hold for all cases when Sidorenko's Conjecture is known to be true [19]. So it seems that the concept of graph limits provide the proper, now analytic, language for Sidorenko's Conjecture and it would probably be futile to try to give a combinatorial description of the various (almost) extremal structures, because they are unique only in the language of analysis.

Aharoni [2] invoked topological considerations to prove Ryser's Conjecture for 3-graphs and hence overcame the combinatorial difficulty of having infinitely many extremal structures. Our main tasks, the characterization of the extremal 3-graphs for Ryser's Conjecture (in Chapter 3) and their link-graphs (in Chapter 2), go a step further in this direction: they show that the extremal structures naturally live in the field of topology and hence it is not unexpected that their combinatorial characterization is complicated.

Chapter 2

Connectedness of Line Graphs of Bipartite Graphs

Joint work with Penny Haxell and Tibor Szabó.

2.1 Introduction

2.1.1 Connectedness of Line Graphs of Bipartite Graphs

The connectedness of the independence complex will be our main parameter to describe the line graphs of the link graphs of Ryser-extremal 3-graphs.

Let $k \geq -1$ be an integer. A topological space X is said to be *k-connected* if for any integer j with $-1 \leq j \leq k$, any continuous map from the j-dimensional sphere S^j into the space X can be extended to a continuous map from the $(j+1)$-dimensional ball B^{j+1} to X. The *connectedness* of X, denoted $\mathrm{conn}(X)$ is the largest k for which X is k-connected.

A simplicial complex $\mathcal{K}$ is a family of simplices in $\mathbb{R}^N$ such that (1) if τ is a face of a simplex $\sigma \in \mathcal{K}$ then $\tau \in \mathcal{K}$ and (2) if $\sigma, \sigma' \in \mathcal{K}$ then $\sigma \cap \sigma'$ is a face of both σ and σ'. The connectedness of a simplicial complex $\mathcal{K}$ is just the connectedness of its body $\|\mathcal{K}\|$ (the union of its simplices).

An *abstract simplicial complex* $\mathcal{C}$ is a simple hypergraph that is closed under taking subsets. The simple hypergraph consisting of the vertex sets of simplices of a simplicial complex $\mathcal{K}$ (called the vertex scheme of $\mathcal{K}$) is an abstract simplicial complex. Every abstract simplicial complex $\mathcal{C}$ has a geometric realization, that is a simplicial complex whose vertex scheme is $\mathcal{C}$. The geometric realization is unique up to homeomorphism. The connectedness of an abstract simplicial complex is just the connectedness of its geometric realization.

For a graph G, we define the *independence complex* $\mathcal{I}(G)$ to be the abstract simplicial complex on the vertices of G whose simplices are the independent sets of G. We will simply write $\mathrm{conn}(G)$ for $\mathrm{conn}(\mathcal{I}(G))$, and refer to this as the *connectedness* of G.

One of the basic parameters of a simplicial complex is its *dimension*, that is, the largest dimension that occurs among its simplices. The connectedness of an arbitrary simplicial complex, or even of an arbitrary graph's independence complex can be arbitrarily small while its dimension is large: just consider the complete bipartite graph $K_{d+1,d+1}$, having an independence complex with dimension d and connectedness -1.

Comparing dimension and connectedness becomes more interesting if we introduce restrictions on the graphs we consider. For *line graphs* for example, a lower bound on the connectedness in terms of the dimension is implicit in the work of Aharoni and Haxell [6]. The *line graph* $L(\mathcal{H})$ of a hypergraph $\mathcal{H}$ is the simple graph $L(\mathcal{H})$ on the vertex set $E(\mathcal{H})$ with $e, f \in V(L(\mathcal{H}))$ adjacent if $e \cap f \neq \emptyset$. With foresight, we state the lower bound of [6] in a more general format, which will be necessary for our investigations. Note that the dimension of the independence complex of a line graph of a hypergraph is just its matching number minus 1.

Theorem 2.1.1. *Let $\mathcal{G}$ be an r-graph, and let $J \subseteq L(\mathcal{G})$ be a subgraph of the line graph of $\mathcal{G}$. Let $M \subseteq V(J)$ be a matching in $\mathcal{G}$. Then*

$$\mathrm{conn}(J) \geq \frac{|M|}{r} - 2.$$

In particular, for any graph G we have $\mathrm{conn}(L(G)) \geq \frac{\nu(G)}{2} - 2$.

Aharoni and Haxell [6] essentially proved that the connectedness of the line graph is at least the so called *independent set domination number* $i\gamma$ of the line graph minus 2 (where $i\gamma(G)$ is the smallest number x, such that every independent set of G can be dominated with x vertices.) Theorem 2.1.1 then follows from $i\gamma(L(\mathcal{H})) \geq \frac{\nu(\mathcal{H})}{r}$, which is immediate from the definitions.

We begin our study of Ryser-extremal 3-graphs with their link graphs.

Definition 2.1.2. Let $\mathcal{H}$ be a 3-partite 3-graph with parts V_1, V_2, and V_3. Let $S \subseteq V_i$ for some $i = 1, 2, 3$. Then the *link graph* $\mathrm{lk}_{\mathcal{H}}(S)$ is the bipartite graph with vertex classes V_j and V_k (where $\{i, j, k\} = \{1, 2, 3\}$) whose edge multiset is $\{e \setminus V_i : e \in E(\mathcal{H}), e \cap V_i \subseteq S\}$.

Note that a pair of vertices appears as an edge in $\mathrm{lk}_{\mathcal{H}}(S)$ with the same multiplicity as the number of edges in $\mathcal{H}$ that contain it together with a vertex from S.

First we will show that the link graphs of Ryser-extremal 3-graphs attest that Theorem 2.1.1 is optimal for $r = 2$, that is, among bipartite graphs they minimize the connectedness of the independence complex of the line graph.

Theorem 2.1.3. *If $\mathcal{H}$ is a 3-partite 3-graph with vertex classes V_1, V_2, and V_3, such that $\tau(\mathcal{H}) = 2\nu(\mathcal{H})$, then for each i we have*

(i) $\mathrm{conn}(L(\mathrm{lk}_{\mathcal{H}}(V_i))) = \nu(\mathcal{H}) - 2$.

(ii) $\nu(\mathrm{lk}_{\mathcal{H}}(V_i)) = \tau(\mathcal{H})$.

In particular

$$\operatorname{conn}(L(\operatorname{lk}_{\mathcal{H}}(V_i))) = \frac{\nu(\operatorname{lk}_{\mathcal{H}}(V_i))}{2} - 2. \tag{2.1.1}$$

We prove Theorem 2.1.3 in Section 2.3. On the way, we also give a proof of Aharoni's Theorem [2], that is somewhat different from the original argument. We also mention here that in Chapter 3 we derive, as a consequence of Theorem 1.1.2, a sort of converse of Theorem 2.1.3: *every* bipartite graph which is optimal for Theorem 2.1.1 is the link of some Ryser-extremal 3-graph.

In the main theorem of this chapter, proven in Section 2.4, we characterize those bipartite graphs which are extremal for Theorem 2.1.1 and hence we also obtain valuable structural information about the link graphs of Ryser-extremal 3-graphs.

Theorem 2.1.4. *Let G be a bipartite graph. Then* $\operatorname{conn}(L(G)) = \frac{\nu(G)}{2} - 2$ *if and only if G has a collection of $\nu(G)/2$ pairwise vertex-disjoint subgraphs, each of them a C_4 or a P_4, such that every edge of G is parallel to an edge of one of the C_4's or is incident to an interior vertex of one of the P_4's.*

To be precise, in this chapter, we will in fact only prove the "only if" direction of this theorem. While it is possible to prove the "if" direction directly by finding some generalized octahedra in the independence complex that cannot be filled, we will make use of the easy direction of Theorem 1.1.2. Thus, the other direction will be proven in Chapter 3.

2.1.2 Topological Tools

The proofs of Theorems 2.1.3 and 2.1.4, as well as the proof of Theorem 1.1.2 (given in Chapter 3) use two tools to bound the topological connectedness of graphs.

The first one is a non-homological version of a theorem of Meshulam [22], which is particularly well-suited for inductive arguments. Let G be a graph, and let e be an edge of G. We denote by $G - e$ the graph G with the edge e deleted. We denote by $G \divideontimes e$ the graph G with both endpoints of e and their neighbors deleted. $G \divideontimes e$ is called G with e *exploded*. We will often write edges with endpoints x and y as xy.

Theorem 2.1.5. *Let G be a graph and let $e \in E(G)$. Then we have*

$$\operatorname{conn}(G) \geq \min\left\{\operatorname{conn}(G - e), \operatorname{conn}(G \divideontimes e) + 1\right\}. \tag{2.1.2}$$

Meshulam proved a homological version of this theorem, where everywhere in the statement conn is replaced by the homological connectedness conn_H. As $\operatorname{conn}_H(G)$ could be strictly larger than $\operatorname{conn}(G)$, these two statements do not immediately imply each other. In Section 2.2 we indicate how to extend Meshulam's argument using the approach of Adamaszek and Barmak [1] and obtain (2.1.2). It is also possible to give a homology-free proof of Theorem 2.1.5 via triangulations along the lines of [27] (see Chapter 5). Theorem 2.1.5 in this formulation but with a modified (non-topological) definition of conn was also stated in [14] and proved without direct reference to topology.

Our second tool makes a direct connection between the size of the largest hypergraph matching and the connectedness of the link.

Theorem 2.1.6. *Let $d \geq 0$ be an integer and let $\mathcal{H}$ be a 3-uniform 3-graph with vertex classes V_1, V_2, and V_3. If we have that $\operatorname{conn}(L(\mathrm{lk}_{\mathcal{H}}(S))) \geq |S| - d - 2$ for every $S \subseteq V_i$, then $\nu(\mathcal{H}) \geq |V_i| - d$.*

For $d = 0$ this theorem is implicit in [6] and was stated explicitly in [3]. For our application we will need the deficiency version with $d \geq 0$. We prove it by constructing a special colored triangulation of the simplex and using Sperner's Lemma. The argument works naturally in the following more general setup about colored simplicial complexes.

A *coloring* of the vertices of a simplicial complex $\mathcal{C}$ by colors from a set X is a function $\chi : V(\mathcal{C}) \to X$. For a subset $S \subseteq X$ of colors, denote by $\mathcal{C}|_S$ the subcomplex of $\mathcal{C}$ induced by the vertices which have colors from S: that is, let $V(\mathcal{C}|_S) = \chi^{-1}(S)$ and $\mathcal{C}|_S = \{\sigma \in \mathcal{C} : \chi(\sigma) \subseteq S\}$.

Theorem 2.1.7. *Let $\mathcal{C}$ be a simplicial complex whose vertices are colored with colors from a set X, and let $d \geq 0$ be an integer. If for every $S \subseteq X$ we have that $\operatorname{conn}(\mathcal{C}|_S) \geq |S| - d - 2$, then $\mathcal{C}$ has a rainbow simplex with $|X| - d$ vertices.*

For the proof of Theorem 2.1.6 the crucial thing to note is that if for each hyperedge $xyz \in E(\mathcal{H})$ we color the corresponding edge xy of the link graph $\mathrm{lk}_{\mathcal{H}}(V_i)$ with the third vertex $z \in V_i$, then a matching in the hypergraph $\mathcal{H}$ corresponds to a *rainbow matching* (a matching with edges having pairwise distinct colors) in the link graph $\mathrm{lk}_{\mathcal{H}}(V_i)$. Then Theorem 2.1.6 is an immediate consequence of Theorem 2.1.7 applied with the independence complex $\mathcal{I}(L(\mathrm{lk}_{\mathcal{H}}(V_i)))$ of the link graph. Indeed, $\mathcal{I}(L(\mathrm{lk}_{\mathcal{H}}(V_i)))|_S = \mathcal{I}(L(\mathrm{lk}_{\mathcal{H}}(S)))$ and the vertices of a rainbow simplex in the independence complex of $L(\mathrm{lk}_{\mathcal{H}}(V_i))$ correspond to pairwise disjoint edges in the link $\mathrm{lk}_{\mathcal{H}}(V_i))$, which extend to pairwise distinct vertices in V_i, and hence form a hypergraph matching.

2.1.3 The Structure of the Chapter

In Section 2.2, we prove Theorem 2.1.7 using triangulations. As we have seen above, Theorem 2.1.6 is a corollary. We also discuss here the proof of Theorem 2.1.5 and include an argument to derive Theorem 2.1.1 from it.

In Section 2.3 we go on to prove Theorem 2.1.3, and on the way we reprove Aharoni's Theorem for the 3-partite case of Ryser's Conjecture.

In Section 2.4 we prove the main theorem of the chapter, Theorem 2.1.4. We show that those bipartite graphs whose line graphs are optimal for Theorem 2.1.1 must have a certain form, which we call a CP-decomposition. We show a slightly more general statement involving any *subgraph* of the line graph of a bipartite graph. The precise definition of CP-decomposition in this general setup is given in Section 2.4.

In Section 2.5 we prove a theorem that will be crucial for our proof of Theorem 1.1.2 in Chapter 3. We define the notion of good sets. Good sets will turn out to be very useful to have in one of the link graphs of a Ryser-extremal 3-graph. In the main theorem of Section 2.5 we show that the lack of good sets in a bipartite graph imposes very strong restrictions on its structure. The proof of this theorem is included in this chapter because it uses several of the technical definitions and lemmas introduced for the proof of our main theorem in Section 2.4.

In the final section we collect several remarks and open problems.

2.2 Topological Preliminaries

2.2.1 Rainbow Simplices

We now briefly introduce a couple of topological notions which we need for the proof of Theorem 2.1.7.

The *join* of two abstract simplicial complexes $\mathcal{C}$ and $\mathcal{D}$ is the abstract simplicial complex $\mathcal{C} * \mathcal{D} = \{(\sigma \times \{0\}) \cup (\tau \times \{1\}) : \sigma \in \mathcal{C}, \tau \in \mathcal{D}\}$. A useful fact relating connectedness to joins is the following:

Proposition 2.2.1 (Lemma 2.3 in [23])**.** *If $\mathcal{C}$ and $\mathcal{D}$ are abstract simplicial complexes, then*

$$\operatorname{conn}(\mathcal{C} * \mathcal{D}) \geq \operatorname{conn}(\mathcal{C}) + \operatorname{conn}(\mathcal{D}) + 2.$$

A map $f : V(\mathcal{C}) \to V(\mathcal{D})$ is a *simplicial map* if the image of each simplex of $\mathcal{C}$ is a simplex of $\mathcal{D}$.

If $\mathcal{K}$ is a simplicial complex, then a *subdivision* of $\mathcal{K}$ is a simplicial complex $\mathcal{K}'$ with $\|\mathcal{K}'\| = \|\mathcal{K}\|$ such that every simplex in $\mathcal{K}'$ is contained in a simplex in $\mathcal{K}$.

To determine the connectedness of a simplicial complex, it is sufficient to consider simplicial maps into subdivisions of the simplex.

Proposition 2.2.2 ([27, Proposition 2.8])**.** *A given simplicial complex $\mathcal{C}$ is k-connected if and only if for every j with $-1 \leq j \leq k$ and for every simplicial map $f : V(\mathcal{S}) \to V(\mathcal{C})$, where $\mathcal{S}$ is a subdivision of the boundary of a $(j+1)$-simplex, there is a subdivision $\mathcal{B}$ of a $(j+1)$-simplex with $\mathcal{S}$ as its boundary, and a simplicial map $\hat{f} : V(\mathcal{B}) \to V(\mathcal{C})$ extending f.*

We prove Theorem 2.1.7 by constructing an appropriate colored triangulation of the simplex and then using Sperner's Lemma. This type of approach was introduced in [6].

Lemma 2.2.3 (Sperner's Lemma [26])**.** *Let $\mathcal{T}$ be a subdivision of a simplex Δ of dimension n. Let $c : V(\mathcal{T}) \to A$ be a coloring of the vertices of the subdivision such that*

(1) Each vertex of Δ receives a different color,

(2) The vertices of $\mathcal{T}$ on a face σ of Δ are colored by the colors of the vertices of σ.

Then there is an n-dimensional rainbow simplex in $\mathcal{T}$.

Proof of Theorem 2.1.7. We will prove the statement by induction on d. Let first $d = 0$.

Let $\mathcal{C}$ be a simplicial complex with a coloring $c : V(\mathcal{C}) \to X$ of its vertices satisfying the conditions of the theorem and let Δ be an $(|X| - 1)$-dimensional simplex (so with $|X|$ vertices). The *k-skeleton* of Δ is the subcomplex containing all faces of dimension up to k. By induction on k, we construct a subdivision $\mathcal{T}_k$ of the k-skeleton of Δ for every $k = 0, 1, \ldots, |X| - 1$, together with a simplicial map $f_k : V(\mathcal{T}_k) \to V(\mathcal{C})$ so that coloring each vertex $v \in V(\mathcal{T}_k)$ of the subdivision by the color $c(f_k(v))$ produces a coloring which has property (1) of Sperner's Lemma, as well as property (2) for each face σ of Δ up to dimension k. (Such a coloring of will be called a *Sperner coloring.*)

We start with the 0-skeleton $\mathcal{T}_0 = \Delta^{(0)}$, which consists of just the vertices of Δ. We choose a simplicial map $f_0 : V(\mathcal{T}_0) \to V(\mathcal{C})$ so that every vertex is sent to a vertex with a different color. This is possible because we have as many vertices as there are colors and, most importantly, because the assumption on the connectedness requires that there is a vertex of every color in $\mathcal{C}$. Indeed, for any $x \in X$, we have $\operatorname{conn}(\mathcal{C}|_{\{x\}}) \geq |\{x\}| - 2 = -1$, hence the subcomplex $\mathcal{C}|_{\{x\}}$ is nonempty.

Now suppose that we have already defined a subdivision $\mathcal{T}_k$ of the k-skeleton of Δ and a simplicial map $f_k : V(\mathcal{T}_k) \to V(\mathcal{C})$ such that if one colors the vertices of the subdivision by the colors of their images under f_k, we get a Sperner coloring. We will extend $\mathcal{T}_k$ and f_k to the $(k + 1)$-skeleton of Δ by defining the extensions independently for each $(k + 1)$-face σ of Δ. The boundary $\partial\sigma$ of σ is contained in the k-skeleton, so $\mathcal{T}_k$ contains a subdivision $\mathcal{D}$ of $\partial\sigma$. Let $S = c(f_k(V(\sigma))) \subseteq X$ be the set of colors of the images of the vertices of σ under f_k. Because f_k induces a Sperner coloring, we must have that $|S| = k+2$ and $f_k(V(\mathcal{D})) \subseteq \mathcal{C}|_S$. By assumption, $\operatorname{conn}(\mathcal{C}|_S) \geq |S| - 2 = k$, and since $\mathcal{D}$ is a subdivision of the boundary of a $(k + 1)$-simplex, by Proposition 2.2.2 there is a subdivision $\mathcal{E}$ of σ with $\mathcal{D}$ as its boundary, and a simplicial map $f_\sigma : V(\mathcal{E}) \to V(\mathcal{C}|_S)$ extending f_k. Doing this for each $(k + 1)$-simplex one after another, we obtain a subdivision $\mathcal{T}_{k+1}$ of the $(k + 1)$-skeleton and a map $f_{k+1} : V(\mathcal{T}_{k+1}) \to V(\mathcal{C})$ defined as the union of all the maps f_σ with σ ranging over the $(k + 1)$-faces of Δ. Since each f_σ agrees with f_k on the boundary, the union agrees with f_k on the k-skeleton and it is well-defined. Also, f_{k+1} induces a Sperner coloring by construction.

Continuing in this manner, we end up with a subdivision $\mathcal{T}_{|X|-1} = \mathcal{T}$ of the entire simplex Δ and a simplicial map $f : V(\mathcal{T}) \to V(\mathcal{C})$ inducing a Sperner coloring. Hence, by Sperner's Lemma, there is a rainbow simplex τ in $\mathcal{T}$ with $|X|$ vertices. The colors of $V(\tau)$ were defined as the colors of its image via f, hence the simplex of $\mathcal{C}$ with vertices $f(V(\tau))$ must also have $|X|$ vertices with all different colors. So we found our rainbow simplex, which concludes the proof for $d = 0$.

Let now $d \geq 1$ and let $\mathcal{C}$ be a simplicial complex with a coloring $c : V(\mathcal{C}) \to X$ of its vertices such that for every $S \subseteq X$ we have that $\mathrm{conn}(\mathcal{C}|_S) \geq |S| - d - 2$. Our strategy is to add some new vertices and new simplices to $\mathcal{C}$ to get a complex $\hat{\mathcal{C}}$ and extend the coloring c to $\hat{\mathcal{C}}$ such that $\hat{\mathcal{C}}$ satisfies the conditions of the theorem with $d_{\hat{\mathcal{C}}} = d - 1$. We will then apply the induction hypothesis to find a rainbow simplex in $\hat{\mathcal{C}}$, and since it will turn out that it may contain at most one new vertex, removing it will yield a rainbow simplex in $\mathcal{C}$ with at least $|X| - d$ vertices.

For each $x \in X$, let $v^{(x)}$ be a new vertex which we color by x. Let $\mathcal{M}$ be the simplicial complex consisting of the isolated vertices $\{v^{(x)} : x \in X\}$, and let $\hat{\mathcal{C}} = \mathcal{C} * \mathcal{M}$. We claim that $\hat{\mathcal{C}}$ fulfills the conditions of the theorem with $d_{\hat{\mathcal{C}}} = d-1$. Indeed, applying Proposition 2.2.1 we get that $\mathrm{conn}(\hat{\mathcal{C}}|_S) \geq (|S|-d-2)-1+2 = |S| - (d-1) - 2$ for every $S \subseteq X$. Here we used that $\hat{\mathcal{C}}|_S = \mathcal{C}|_S * \mathcal{M}|_S$ and that $\mathrm{conn}(\mathcal{M}|_S) = -1$, as each color is represented among the new vertices, so $\mathcal{M}|_S$ is non-empty. Thus, by induction, $\hat{\mathcal{C}}$ contains a rainbow simplex τ with $|X| - d + 1$ vertices. To complete the proof of the theorem we just need to recall that no two vertices of $\mathcal{M}$ form a simplex, hence τ can contain at most one of the new vertices. Thus there is a face of τ spanned by at least $|X| - d$ vertices from $\mathcal{C}$, providing the rainbow simplex we were looking for. $\square$

2.2.2 The Independence Complex

Meshulam [22] proved a homological version of Theorem 2.1.5, where everywhere in the statement conn is replaced by the homological connectedness conn_H. He used the Mayer-Vietoris sequence and the observation that, provided G is simple, $\mathcal{I}(G - e) = \mathcal{I}(G) \cup (e * \mathcal{I}(G \divideontimes e))$ and $\mathcal{I}(G) \cap (e * \mathcal{I}(G \divideontimes e))$ is the suspension of $\mathcal{I}(G \divideontimes e)$. (Once proved for simple graphs, Theorem 2.1.5 follows easily for arbitrary G.) Adamaszek and Barmak [1], mostly concerned with a question of Aharoni, Berger, and Ziv [4], proved that the conn on the right hand side of inequality (2.1.2) can be replaced with the following function ψ:

$$\psi(G) = \begin{cases} -2 & G = \emptyset \\ +\infty & V(G) \neq \emptyset, E(G) = \emptyset \\ \max_{e \in E(G)} \min\{\psi(G - e), \psi(G \divideontimes e) + 1\} & \text{otherwise.} \end{cases}$$

It can be easily seen by induction on $|E(G)|$ that Theorem 2.1.5 implies the theorem of Adamaszek and Barmak [1], but there seems to be no direct way to derive the implication in the other direction. However, the *proof* in [1] can easily be modified to give Theorem 2.1.5. One simply takes e to be an arbitrary edge, defines $k = \min(\mathrm{conn}(G - e), \mathrm{conn}(G \divideontimes e) + 1)$, and proceeds as in [1] to show that the homological connectedness of G is at least k. To conclude that $\mathrm{conn}(G) \geq k$, one only needs to show that $k \geq 1$ implies that $\mathcal{I}(G)$ is simply connected and then appeal to the Hurewicz Theorem. This can be done in an argument identical to the one in [1].

One can apply Theorem 2.1.5 to prove Theorem 2.1.1.

Proof of Theorem 2.1.1. We proceed by induction on $|E(J)|$. If J contains an isolated vertex, the lemma is trivially true, since then $\operatorname{conn}(J) = \infty$. Thus we may assume that every vertex of J has a neighbor. If $M = \emptyset$, the lemma is trivially true, since the connectedness of anything is at least -2, so assume $|M| \geq 1$. Now consider an edge $m \in M \subseteq V(J)$. This edge (vertex of J) has a neighbor e in J. Since $M \subseteq V(J - me)$, by induction we have $\operatorname{conn}(J - me) \geq |M|/r - 2$. Now consider what happens when we explode me. We remove from $V(J)$ all neighbors of m and e. Since $m \in M$, none of the neighbors of m are in M, and since e has size at most r, it intersects at most r edges of M (one of them being m). Therefore, $V(J \divideontimes me)$ still contains a matching of size at least $|M| - r$. By induction, we then have $\operatorname{conn}(J \divideontimes me) \geq (|M| - r)/r - 2 = |M|/r - 3$. Applying Theorem 2.1.5, we obtain

$$\operatorname{conn}(J) \geq \min\{\operatorname{conn}(J - me), \operatorname{conn}(J \divideontimes me) + 1\} \geq \frac{|M|}{r} - 2,$$

which is what was wanted. □

2.3 Connectedness of the Link Graph

In this section we prove Theorem 2.1.3, which states that the link graph of any Ryser-extremal 3-graph minimizes the connectedness of the independence complex of its line graph. On the way we also give a proof of Aharoni's Theorem, which is somewhat different from the original argument.

Let $\mathcal{H}$ be a 3-partite 3-graph with vertex classes V_1, V_2, and V_3. We aim to show that $\tau(\mathcal{H}) \leq 2\nu(\mathcal{H})$. To do this, we will consider the link graph (recall Definition 2.1.2). An important thing to note is that if each edge of a matching in the link graph $\operatorname{lk}_{\mathcal{H}}(V_i)$ extends to a different vertex of V_i, then the extended edges form a matching in $\mathcal{H}$. Thus, we can color each edge of the link graph by the vertex to which it extends (since we are considering the link graph as a multigraph, that vertex is uniquely determined for each edge) so that a rainbow matching (a matching with each edge of a different color) in the link graph corresponds to a matching in the hypergraph $\mathcal{H}$. Now we will use the vertex cover number of $\mathcal{H}$ to find a lower bound on the connectedness of the line graphs of the link graphs, and we will use the matching number of $\mathcal{H}$ to find an upper bound for at least one link. Combining these bounds will yield the desired inequality $\tau(\mathcal{H}) \leq 2\nu(\mathcal{H})$. So let's calculate.

Proposition 2.3.1. *Let $\mathcal{H}$ be a 3-partite 3-graph with vertex classes V_1, V_2, and V_3. Then for each $i \in \{1, 2, 3\}$ we have the following:*

(i) For all $S \subseteq V_i$ we have

$$\operatorname{conn}(L(\operatorname{lk}_{\mathcal{H}}(S))) \geq \frac{\tau(\mathcal{H}) - (|V_i| - |S|)}{2} - 2.$$

(ii) There is some $S \subseteq V_i$ such that

$$\operatorname{conn}(L(\operatorname{lk}_{\mathcal{H}}(S))) \leq \nu(\mathcal{H}) - (|V_i| - |S|) - 2.$$

(iii) For every $S \subseteq V_i$ for which the inequality in (ii) holds we have

$$|S| \geq |V_i| - (2\nu(\mathcal{H}) - \tau(\mathcal{H})).$$

Proof. Let $S \subseteq V_i$. We want to show that the line graph $L(\mathrm{lk}_{\mathcal{H}}(S))$ has sufficiently high connectedness. We construct a vertex cover T_S of $\mathcal{H}$ by taking the vertices in $V_i \setminus S$ and a minimum vertex cover of $\mathrm{lk}_{\mathcal{H}}(S)$. This is clearly a vertex cover of $\mathcal{H}$ because any edge not incident to S intersects $V_i \setminus S$ and any edge incident to S induces an edge in the link of S, and hence intersects the vertex cover of the link. We have $|T_S| = |V_i| - |S| + \tau(\mathrm{lk}_{\mathcal{H}}(S))$, and since this is a vertex cover, we thus have

$$|V_i| - |S| + \tau(\mathrm{lk}_{\mathcal{H}}(S)) \geq \tau(\mathcal{H}) \tag{2.3.1}$$

for all subsets $S \subseteq V_i$. By König's Theorem, we have $\tau(\mathrm{lk}_{\mathcal{H}}(S)) = \nu(\mathrm{lk}_{\mathcal{H}}(S))$. We therefore have a lower bound on the matching number of the link graph, and so by Theorem 2.1.1, we have

$$\mathrm{conn}(L(\mathrm{lk}_{\mathcal{H}}(S)) \geq \frac{\nu(\mathrm{lk}_{\mathcal{H}}(S))}{2} - 2 \geq \frac{\tau(\mathcal{H}) - (|V_i| - |S|)}{2} - 2,$$

which is the inequality in statement (i).

Now we want to show that the inequality in statement (ii) holds for some S. Suppose to the contrary that for every $S \subseteq V_i$ we had $\mathrm{conn}(L(\mathrm{lk}_{\mathcal{H}}(S))) \geq \nu(\mathcal{H}) - (|V_i| - |S|) - 1$. We will aim to apply Theorem 2.1.7 with $X = V_i$ and $\mathcal{C} = \mathcal{I}(L(\mathrm{lk}_{\mathcal{H}}(V_i)))$ to find a large rainbow matching in $\mathrm{lk}_{\mathcal{H}}(V_i)$ and hence a large matching in $\mathcal{H}$. By our supposition, for each $S \subseteq V_i$ we have $\mathrm{conn}(L(\mathrm{lk}_{\mathcal{H}}(S))) \geq |S| - (|V_i| - \nu(\mathcal{H}) - 1) - 2$, and hence we can apply Theorem 2.1.7 with $d = |V_i| - \nu(\mathcal{H}) - 1$ to get a rainbow matching of size $|V_i| - (|V_i| - \nu(\mathcal{H}) - 1) = \nu(\mathcal{H}) + 1$, which is a contradiction. Thus some $S \subseteq V_i$ must indeed satisfy the inequality in (ii).

Now consider such an S. Combining the inequalities in (i) and (ii), we get

$$\frac{\tau(\mathcal{H}) - (|V_i| - |S|)}{2} - 2 \leq \nu(\mathcal{H}) - (|V_i| - |S|) - 2,$$

from which the inequality in (iii) follows after some rearranging. □

Now Aharoni's Theorem follows in one line from the above proposition: there is an $S \subseteq V_i$ such that $|S| \geq |V_i| - (2\nu(\mathcal{H}) - \tau(\mathcal{H}))$, and hence

$$\tau(\mathcal{H}) + |V_i| - |S| \leq 2\nu(\mathcal{H}).$$

Since $|V_i| \geq |S|$, we thus have $\tau(\mathcal{H}) \leq 2\nu(\mathcal{H})$ as desired.

We also use Proposition 2.3.1 to derive the main theorem of this section.

Proof of Theorem 2.1.3. Applying Proposition 2.3.1 to $\mathcal{H}$, we see by (iii) that in (ii) equality holds if and only if $S = V_i$ for some i. Combining the inequalities in (i) and (ii) for $S = V_i$ with the fact that $\tau(\mathcal{H}) = 2\nu(\mathcal{H})$ immediately gives that $\operatorname{conn}(L(\mathrm{lk}_{\mathcal{H}}(V_i))) = \nu(\mathcal{H}) - 2$, showing part (i) of Theorem 2.1.3. This gives the following chain of inequalities:

$$\begin{aligned}\frac{\tau(\mathcal{H})}{2} - 2 = \nu(\mathcal{H}) - 2 &= \operatorname{conn}(L(\mathrm{lk}_{\mathcal{H}}(V_i))) \\ &\geq \frac{\nu(\mathrm{lk}_{\mathcal{H}}(V_i))}{2} - 2 = \frac{\tau(\mathrm{lk}_{\mathcal{H}}(V_i))}{2} - 2 \\ &\geq \frac{\tau(\mathcal{H})}{2} - 2,\end{aligned}$$

where the first inequality is valid because of Theorem 2.1.1, the equality following it is König's Theorem, and the last inequality is just equation (2.3.1) for $S = V_i$. It follows that every inequality is actually an equality, from which part (ii) of Theorem 2.1.3 follows.

From parts (i), (ii), and the fact that $\nu(\mathcal{H}) = \frac{\tau(\mathcal{H})}{2}$, it follows that the link graphs $\mathrm{lk}_{\mathcal{H}}(V_i)$ of a Ryser-extremal 3-graph $\mathcal{H}$ must be extremal for Theorem 2.1.1:

$$\operatorname{conn}(L(\mathrm{lk}_{\mathcal{H}}(V_i))) = \frac{\nu(\mathrm{lk}_{\mathcal{H}}(V_i))}{2} - 2.$$

□

2.4 The Link Characterization Theorem

In the main theorem of this section we fully characterize those bipartite graphs for which the connectedness of the line graph is as small as possible, that is, it is equal to two less than half its matching number.

For the proof we need to choose our definitions very subtly and in order to make the induction work, we need to consider a carefully formulated more general statement involving arbitrary *subgraphs* of the line graphs.

Definition 2.4.1. Let G be a bipartite graph, and let $J \subseteq L(G)$ be a subgraph of its line graph. Two edges of G are called *J-adjacent* if they are connected by an edge in J, and otherwise *J-nonadjacent*. An edge $e \in V(J)$ is *at home* in a subgraph $T \subseteq G$ if T is a path on 4 vertices, e intersects T in an interior vertex, and e is J-adjacent to some edge of T.

Definition 2.4.2. Let $k \in \mathbb{N}$, let G be a bipartite graph, let $J \subseteq L(G)$ be a subgraph of its line graph, and let $M \subseteq V(J)$ be a matching in G of size $2k$. A *CP-decomposition of J with respect to M* is a set of k vertex-disjoint subgraphs $S_1, \ldots, S_s, T_1, \ldots, T_t$ of G such that

(1) Each S_i is isomorphic to C_4 (a cycle on 4 vertices), contains two edges of M, and every two intersecting edges are J-adjacent.

(2) Each T_j is isomorphic to P_4 (a path on 4 vertices), contains two edges of M, and every two intersecting edges are J-adjacent.

(3) Every edge in $V(J)$ is equal to or parallel to an edge of some S_i, or is at home in some T_j.

We call $k = |M|/2$ the *order* of the CP-decomposition. Observe for property (3) that the edges of any of the subgraphs T_j are themselves at home in T_j by definition.

Theorem 2.4.3 (CP-Decomposition Theorem)**.** *Let G be a bipartite graph, let $J \subseteq L(G)$ be a subgraph of its line graph, and let $M \subseteq V(J)$ be a matching in G. If $\operatorname{conn}(J) \leq \frac{|M|}{2} - 2$, then J has a CP-decomposition with respect to M.*

Note that by Theorem 2.1.1 we must have that $\operatorname{conn}(J) = \frac{|M|}{2} - 2$, so $|M|$ is even and $V(J)$ does not contain a larger matching than M.

First we spell out the special case when J is the entire line graph and prove Theorem 2.1.4. This will provide a characterization of those bipartite graphs G whose line graphs have connectedness as small as possible in terms of the matching number of G.

Proof of Theorem 2.1.4. Suppose that $\operatorname{conn}(L(G)) = \frac{\nu(G)}{2} - 2$. Then by Theorem 2.4.3, $L(G)$ has a CP-decomposition, which is a collection of $\nu(G)/2$ pairwise vertex-disjoint subgraphs, each of them a C_4 or a P_4, such that every edge of G is either an edge of one of the C_4's or is incident to an interior vertex of one of the P_4's.

As was mentioned in the introduction, the converse of this statement is not used at all in our argument. We include it only to provide a full characterization of the extremal graphs. It will be proven in Chapter 3, since the proof uses the concept of home-base hypergraph which is the central concept of that chapter. □

The proof of Theorem 2.4.3 is quite involved and will take up the next two subsections. We start with some auxiliary lemmas.

2.4.1 Lemmas on M-reduced Subgraphs

For the proof of Theorem 2.4.3 and later we will often use Theorem 2.1.5 in its contrapositive form, which we state here as a corollary.

Corollary 2.4.4. *Let H be a graph, let $e \in E(H)$, and let $k \in \mathbb{N}$. If $\operatorname{conn}(H) \leq k$, then either $\operatorname{conn}(H - e) \leq k$ or $\operatorname{conn}(H \circledast e) \leq k - 1$.*

In light of Corollary 2.4.4, the following definitions will be useful.

Definition 2.4.5. An edge $e \in E(H)$ is called *decouplable* if $\operatorname{conn}(H - e) \leq \operatorname{conn}(H)$, and *explodable* if $\operatorname{conn}(H \circledast e) \leq \operatorname{conn}(H) - 1$.

By Corollary 2.4.4 every edge is either decouplable or explodable (or both). In the grand plan of our proof of the CP-decomposition theorem we intend to delete edges of $J \subseteq L(G)$ iteratively until there are no decouplable edges left and hence all edges are explodable (and then we explode one, hence decreasing the connectedness). Crucially, deleting decouplable edges does not increase the connectedness. This explains the following key definition of this subsection.

Definition 2.4.6. Let G be a bipartite graph and let $M \subseteq E(G)$ be a matching of it. A subgraph $J \subseteq L(G)$ of the line graph is called *M-reduced* if

(1) $M \subseteq V(J)$,

(2) $\operatorname{conn}(J) \leq \frac{|M|}{2} - 2$, and

(3) no edge $ef \in E(J)$ is decouplable.

Again, note that by Theorem 2.1.1, if J is M-reduced, then $\operatorname{conn}(J) = \frac{|M|}{2} - 2$ and hence M must have an even number of edges.

It will be important to note that if J is M-reduced, then J is also M'-reduced for any matching $M' \subseteq V(J)$ with $|M'| = |M|$. In particular, if we replace edges of M by parallel edges in $V(J)$, these must share any properties we can deduce for the original edges.

Assumptions. For the remainder of the section let G be a bipartite graph, let $M \subseteq E(G)$ be a matching of size $2k$ in G, and let $J \subseteq L(G)$ be an M-reduced subgraph of the line graph.

Lemma 2.4.7 (Degree Lemma)**.** *For every edge $e \in V(J) \setminus M$ either no edge of M is J-adjacent to e or two edges of M are J-adjacent to e. In particular, if e is parallel to an edge of M, then it is not J-adjacent to that edge.*

Proof. Since J is M-reduced, we have $\operatorname{conn}(J) \leq k - 2$. Clearly an edge can be J-adjacent to at most two edges of M because M is a matching in G and $J \subseteq L(G)$. Suppose for the sake of contradiction that some edge $e \in V(J)$ is J-adjacent to $m \in M$, but not J-adjacent to any other edge of M. Since $me \in E(J)$ and J is M-reduced, by Corollary 2.4.4, upon exploding me we have $\operatorname{conn}(J') \leq k - 3$ for $J' = J \divideontimes me$. Since e is J-adjacent to only one edge from M, the explosion keeps $M \setminus \{m\}$ in J', so J' still contains a matching of size $2k - 1$. Then by Theorem 2.1.1 we have $\operatorname{conn}(J') \geq \frac{2k-1}{2} - 2 > k - 3$, which is a contradiction. Thus every edge in $V(J)$ is J-adjacent to either two edges of M or no edge of M. □

Corollary 2.4.8. *Let x, y, x', and $y' \in V(G)$ be the vertices of a C_4 such that $xy, x'y' \in M$, and $xy', x'y \in V(J)$. Then for every $zy \in V(J)$ with $z \in V(G) \setminus \{x, x'\}$ we have that zy is J-adjacent to xy if and only if it is J-adjacent to $x'y$.*

Proof. Suppose $z \in V(G) \setminus \{x, x'\}$ with $zy \in V(J)$, and zy is J-adjacent to xy. Then by the Degree Lemma there is an edge $zw \in M$ which is J-adjacent to zy. Now consider the matching $M^{\times} = M \cup \{xy', x'y\} \setminus \{xy, x'y'\}$. Note that $|M^{\times}| = |M|$ and $M^{\times} \subseteq V(J)$. Applying the Degree Lemma to $M^{\times}$, we have that since $zw \in M^{\times}$ is J-adjacent to zy, also $x'y \in M^{\times}$ must be J-adjacent to zy. The reverse inclusion can be shown by exchanging the roles of M and $M^{\times}$. □

If M is a matching in a graph, then an M-*exposed* vertex is one not in any edge of M. A path or cycle is M-*alternating* if for every pair of consecutive edges, exactly one of them is in M.

Lemma 2.4.9 (Alternating Lemma)**.** *Let J be M-reduced, and let $e_1, \dots, e_q \in V(J)$ be the edges of an M-alternating path in G starting at an M-exposed vertex or the edges of an M-alternating cycle in G with $e_q e_1 \notin E(J)$. Then in both cases $e_i e_{i+1} \notin E(J)$ for all $i = 1, \dots, q-1$.*

Proof. **Case 1**. $e_1, e_2, \dots, e_q \in V(J)$ are the edges of an M-alternating path starting at an M-exposed vertex.

Suppose the lemma did not hold and let $j = \min\{i : e_i e_{i+1} \in E(J)\}$. If j is odd, then $e_j \notin M$. Since $e_j e_{j+1} \in E(J)$, by the Degree Lemma there must be another edge of M which is J-adjacent to e_j. However, e_1 has an M-exposed vertex, so $j \neq 1$, from which it follows that $e_{j-1} e_j \in E(J)$, which contradicts the minimality of j.

Therefore j is even and $e_j \in M$. Since by assumption J is M-reduced, $e_j e_{j+1}$ is explodable, hence $J' = J \divideontimes e_j e_{j+1}$ satisfies $\operatorname{conn}(J') \leq k-3$. Note that since $e_{j-1} e_j \notin E(J)$, the explosion does not delete e_{j-1}. Thus $M' = M \cup \{e_1, e_3, \dots, e_{j-1}\} \setminus \{e_2, e_4, \dots, e_j, e_{j+2}\} \subseteq V(J')$ is a matching of size $2k-1$ (if $j+2 > q$, let e_{j+2} be the second edge of M that is J-adjacent to e_{j+1}, which exists by the Degree Lemma). This means that by Theorem 2.1.1, $\operatorname{conn}(J') \geq \frac{2k-1}{2} - 2 > k-3$, which is a contradiction. Thus the lemma holds for paths.
Case 2. $e_1, e_2, \dots, e_q \in V(J)$ are the edges of an M-alternating cycle with $e_q e_1 \notin E(J)$.

Since we can reverse the direction of the cycle if necessary, we can assume without loss of generality that $e_q \in M$ and $e_1 \notin M$. If the lemma does not hold, then let $j = \min\{i : e_i e_{i+1} \in E(J)\}$. If j is odd, then a reasoning identical to the one in Case 1 yields a contradiction.

Therefore j is even and $e_j \in M$. By assumption, $e_j e_{j+1}$ is explodable, hence $J' = J \divideontimes e_j e_{j+1}$ satisfies $\operatorname{conn}(J') \leq k-3$. We have a matching $M' = M \cup \{e_1, e_3, \dots, e_{j-1}, e_{j+3}, \dots, e_{q-1}\} \setminus \{e_2, e_4, \dots, e_q\} \subseteq V(J')$ of size $2k-1$, so by Theorem 2.1.1, $\operatorname{conn}(J') \geq \frac{2k-1}{2} - 2 > k-3$, which is a contradiction. Thus the lemma also holds for cycles. □

Given two incident non-parallel edges $m \in M$ and $e \in V(J) \setminus M$, we define $\mathcal{P}_M(m, e)$ to be the set of edges of M which participate in some M-alternating path in G starting with m, continuing with e, and using only edges from $V(J)$. Note that we do not require the edges of the path to be J-adjacent. Also note that $m \in \mathcal{P}_M(m, e)$, and if $me \in E(J)$, then $\mathcal{P}_M(m, e)$ contains at least one more edge of M, namely the other one J-adjacent to e, which exists by the Degree Lemma.

Lemma 2.4.10. *Let $m \in M$, $e \in V(J) \setminus M$ with $me \in E(J)$, and let $m' \in M$ be the other M-edge J-adjacent to e. Let W_1 and W_2 be the vertex classes of the bipartite graph G, and let $m \cap e \subseteq W_i$. Then for every $m^* \in \mathcal{P}_M(m, e) \setminus \{m, m'\}$, there is an edge $g \in V(J)$ for which the following hold:*

(i) g *is* J*-adjacent to* m^*,

(ii) $g \cap m^* \subseteq W_{3-i}$,

(iii) If $\hat{m} \in M$ *is the other* M*-edge (besides* m^**)* J*-adjacent to* g*, then* $\hat{m} \notin \mathcal{P}_M(m, e)$.

Proof. Suppose not. Then fix $m^* \in \mathcal{P}_M(m, e) \setminus \{m, m'\}$ for which the lemma fails. Let $Q = \{g \in V(J) : gm^* \in E(J), g \cap m^* \subseteq W_{3-i}\}$. If Q is not empty, then by assumption every edge $g \in Q$ fails property (iii).

Since J is M-reduced, we have $\text{conn}(J) \leq k - 2$ and when we explode me, we get $\text{conn}(J \divideontimes me) \leq k - 3$. We then iteratively delete decouplable edges of $J \divideontimes me$ in an arbitrary order until no edge is decouplable. This results in an M'-reduced $J' \subseteq J \divideontimes me$, where $M' = M \setminus \{m, m'\}$ and $\text{conn}(J') \leq k-3$ (recall that deleting a decouplable edge does not increase the connectedness). Let a be the vertex in $m' \setminus e$.

Note that if a is not the endpoint of any edge contained in $V(J')$, we are done, since then $\mathcal{P}_M(m, e) = \{m, m'\}$, so there is no m^* to choose, and the statement is vacuously true. Thus, assume this is not the case.

We will arrive at a contradiction by showing that m^* is isolated in J', which implies $\text{conn}(J') = \infty$.

First, we show that m^* has no J'-neighbors incident to it in W_{3-i}. Take an arbitrary edge g which intersects m^* in W_{3-i}. If $m^*g \notin E(J)$, then $m^*g \notin E(J')$, so we are done. Thus assume $m^*g \in E(J)$, which implies $g \in Q$, and this means that $\hat{m}$, the other M-edge J-adjacent to g (which exists by the Degree Lemma for M and J), is in $\mathcal{P}_M(m, e)$ by our assumption on m^*. If $\hat{m} \in \{m, m'\}$, then $m^*g \notin E(J')$ by the Degree Lemma applied to J' (because $m, m' \notin V(J')$). Otherwise, there is an M-alternating path $e_1, \ldots, e_q = \hat{m}$ starting at the vertex $a \in e_1$. This is clearly also an M'-alternating path, and since a is an M'-exposed vertex in J', the Alternating Lemma (Lemma 2.4.9) applied to M' and J' gives that none of the pairs e_i, e_{i+1} are J'-adjacent; in particular, $e_{q-1}\hat{m} \notin E(J')$. Now there are two cases.

Case 1. m^* is on this path.

Then the segment of the path starting at m^* and ending with $\hat{m}$, together with g, forms an M'-alternating cycle. Since $e_{q-1}\hat{m} \notin E(J')$, the Alternating Lemma tells us that m^* and g are not J'-adjacent.

Case 2. m^* is not on the path.

Then $e_1, \ldots, e_q, g, m^*$ is an M'-alternating path, and the Alternating Lemma again tells us that m^* and g are not J'-adjacent.

This proves that m^* has no J'-neighbor which intersects it in W_{3-i}.

We now show that it also has no J'-neighbor intersecting it in W_i. Take an arbitrary edge g which intersects m^* in W_i. We may again assume g is J-adjacent to m^*, and hence there is an $\hat{m} \in M$, which is the other M-edge J-adjacent to g. Again, if $\hat{m} \in \{m, m'\}$, then $m^*g \notin E(J')$ because then $\hat{m} \notin V(J')$ and the Degree Lemma for J' gives that g is not J'-adjacent to any edge of $M' = M \cap V(J')$. There is an M'-alternating path $e_1, \ldots, e_q = m^*$ starting at the vertex $a \in e_1$. Because the path starts at an M'-exposed vertex, no two consecutive edges are J'-adjacent by the Alternating Lemma. Again there are two cases.

Case 1. $\hat{m}$ is on this path.

Then the segment of the path starting at $\hat{m}$ and ending with m^*, together with g, forms an M'-alternating cycle. Since $e_{q-1}m^* \notin E(J')$, the Alternating Lemma tells us that m^* and g are not J'-adjacent.

Case 2. $\hat{m}$ is not on the path.

Then $e_1, \ldots, e_q, g$ is an M'-alternating path, and the Alternating Lemma will again tell us that $m^* = e_q$ and g are not J'-adjacent.

In conclusion, we have shown that m^* does not have any neighbor in J', which was our desired contradiction. Hence no such m^* exists and the proof is complete. □

Lemma 2.4.11. *Let $m, e, m', f \in V(J)$ be the edges of an M-alternating C_4 with $m, m' \in M$. Let $M^\times = M \cup \{e, f\} \setminus \{m, m'\}$. Then $\mathcal{P}_{M^\times}(e, m) = \mathcal{P}_M(m, e) \cup \{e, f\} \setminus \{m, m'\}$; in particular, $|\mathcal{P}_{M^\times}(e, m)| = |\mathcal{P}_M(m, e)|$.*

Proof. Let a be the vertex in $m' \cap f$. Any M-alternating path starting with m, e must continue with m' and then a path starting at a and never again intersect the vertices of the C_4. Similarly, any $M^\times$-alternating path starting with e, m must continue with f and a path starting at a and never again intersect the vertices of the C_4. Thus the edges outside of the C_4 which are reached will be the same, because the matchings are the same outside the C_4. □

Lemma 2.4.12. *Let $m, e, m', f \in V(J)$ be the edges of an M-alternating C_4 with $m, m' \in M$, and let $M^\times = M \cup \{e, f\} \setminus \{m, m'\}$. Then J has a CP-decomposition with respect to $M^\times$ if and only if J has a CP-decomposition with respect to M.*

Proof. Suppose J has a CP-decomposition with respect to M. We will show that it has a CP-decomposition with respect to $M^\times$. Since the roles of M and $M^\times$ are symmetric, the reverse implication is analogous. Note also that J is M-reduced if and only if it is $M^\times$-reduced. There are two cases.

Case 1. $mem'f$ is a C_4 in the CP-decomposition with respect to M.

Then $mem'f$ is still an $M^\times$-alternating 4-cycle and incident edges are J-adjacent, so the same CP-decomposition is also a CP-decomposition with respect to $M^\times$.

Case 2. $mem'f$ is not a C_4 in the CP-decomposition with respect to M.

Then m and m' must be in either a C_4 or a P_4 in this decomposition. Suppose first that $m \in S_1$, where S_1 is a C_4 in the CP-decomposition. Then $m' \notin S_1$, and hence $e, f \notin S_1$. It follows that e and f are neither equal to nor parallel to edges of any C_4 in the CP-decomposition, and thus by property (3) of the CP-decomposition, they are each at home in some P_4 of the CP-decomposition. This means that both endpoints of m' must be interior vertices of some P_4. However this is impossible, since M-edges are the ending edges of the P_4's of the decomposition, so only one endpoint could be interior. So m is not in a C_4 of the decomposition and by symmetry, neither is m'.

From now on we assume that m and m' are in two distinct P_4's (they are not in the same P_4 because then either e or f would not be at home in any P_4). Call these P_4's T_1 and T_2 with edges mgh and $m'g'h'$, respectively. Note that $m \cap g$ contains an interior vertex, as does $m' \cap g'$, and since e and f must be at home somewhere, one of them, say e, is at home in $m \cap g$ and the other (f) in $m' \cap g'$ (since e and f are disjoint).

We claim now that replacing T_1 and T_2 by egh and $fg'h'$ gives us a CP-decomposition with respect to $M^\times$. To check this, we must show that egh and $fg'h'$ are both P_4's whose incident edges are J-adjacent, and that every edge which was at home in T_1 or T_2 is still at home in either egh or $fg'h'$. Both are straightforward consequences of Corollary 2.4.8, which states that the J-neighbors of e and of m at $e \cap m$, which are outside of the 4-cycle are the same (and likewise for f and m'). Thus, e is J-adjacent to g and f is J-adjacent to g'. And any edge which was at home in T_1 because it was J-adjacent to m is J-adjacent also to e, and so still at home in egh (and likewise for T_2 and $fg'h'$). The only edges left to check are m and m' and edges parallel to m, m', e, or f. Here m and m' are at home in egh and $fg'h'$, respectively, because they are J-adjacent to g and g', respectively. Edges parallel to m or m' are J-adjacent to g or g', respectively, since they needed to be at home in some P_4 of the original CP-decomposition and those are the only possibilities (they are not J-adjacent to m or m' by the Degree Lemma, because J is M-reduced). Thus they are at home in the new P_4's. All e-parallel edges are also J-adjacent to g, and f-parallel ones to g' because of Corollary 2.4.8. This means that this is indeed a CP-decomposition with respect to $M^\times$, and it is clearly of the same order. This completes the proof. □

2.4.2 Proof of the CP-Decomposition Theorem

We are now ready to start the proof of Theorem 2.4.3.

Proof of Theorem 2.4.3. We prove this by induction on $|M|$. Recall that $|M|$ must be even, so write $|M| = 2k$ and proceed by induction on k.

For $k = 0$, we have $\operatorname{conn}(J) = -2$, which means $V(J)$ is empty. Thus, it has a CP-decomposition of order 0, which is an empty collection of cycles and paths.

For $k = 1$, we get $\operatorname{conn}(J) = -1$, so $I(J)$ must have at least two components. Thus there exist two disjoint non-empty subsets $E_1, E_2 \subseteq V(J)$ with $V(J) = E_1 \cup E_2$ such that for all $e_1 \in E_1$ and $e_2 \in E_2$, we have $e_1e_2 \in E(J)$. By assumption there is a matching $M = \{m_1, m_2\} \subseteq V(J)$. Since m_1 and m_2 are not J-adjacent (as they are disjoint), they must be in the same component of $I(J)$, and so assume without loss of generality that $m_1, m_2 \in E_1$. Then every edge in E_2 is J-adjacent to both m_1 and m_2, and since G is bipartite, every such edge must intersect m_1 in one vertex class of G and m_2 in the other. Thus the graph formed by the edges in E_2 together with m_1 and m_2 is either a C_4 or a P_4 together with possibly some parallel edges. If it forms a C_4, then the rest of E_1 must consist of edges parallel to m_1 and m_2 because they must be

J-adjacent to both of the non-M edges of the C_4. If the graph is a P_4, then the rest of the edges in E_1 must be J-adjacent to all of the middle edges, and hence are at home in that P_4. Therefore, J has a CP-decomposition consisting of a single C_4 or P_4. This completes the proof for $k = 1$.

Now assume $k \geq 2$. If $|E(J)| = 0$, then $\text{conn}(J) = \infty$, so the statement is vacuously true. So assume $|E(J)| \geq 1$. We may assume J is M-reduced so that all edges of J are explodable. (If J is not M-reduced, iteratively delete decouplable edges of J until the subgraph is M-reduced. A CP-decomposition for the subgraph of J will also be a CP-decomposition of J.)

Case 1. There is an edge $m = ab \in M$ with no J-neighbor incident to a.

Then there must be a J-neighbor e of m incident to b, otherwise m is isolated, and $\text{conn}(J) = \infty$, which is a contradiction. Since J is M-reduced, when we explode $me \in E(J)$, we have that $J' = J \circledast me$ satisfies $\text{conn}(J') \leq k - 3$. By the Degree Lemma for J, there is another edge $m' \in M$ which is J-adjacent to e. Since $M' = M \setminus \{m, m'\} \subseteq V(J')$ is a matching of size $2k - 2$, we have that J' together with M' satisfy the conditions of the theorem for $k' = k - 1$, so by induction, there is a CP-decomposition of J' with respect to M', say $S_1, \ldots, S_s, T_1, \ldots T_t$ with $s + t = k - 1$, where each $S_i \cong C_4$ and each $T_j \cong P_4$.

Define T_{t+1} to be a P_4 consisting of the edges m, e, and m'. We claim that $S_1, \ldots, S_s, T_1, \ldots, T_{t+1}$ is a CP-decomposition of J with respect to M. Since $J' \subseteq J$ and $M' \subseteq M$, the subgraphs S_i form C_4's with two M-edges, with intersecting edges J-adjacent to each other, and the subgraphs T_j with $j < t+1$ form P_4's also with this property. The new path T_{t+1} of course satisfies this as well, so the only thing we still need to check is that the remaining edges are parallel to edges of some S_i or at home in some T_j. Clearly, this is already true of the edges in $V(J')$, so consider an edge $f \in V(J) \setminus V(J')$. Then $f \in N_J(m)$ or $f \in N_J(e)$. If $f \in N_J(e)$, then f is at home in T_{t+1}, because both endpoints of e are interior in T_{t+1}. If $f \in N_J(m)$, then f is also at home in T_{t+1} because m did not have a J-neighbor incident to a, so f must be adjacent to m at b, which is an interior vertex of T_{t+1}. This completes the proof of Case 1.

Case 2. Every edge in M has a J-neighbor on both sides.

Recall that given two incident non-parallel edges $m \in M$ and $e \in V(J) \setminus M$, we define $\mathcal{P}_M(m, e)$ to be the set of edges of M which participate in some M-alternating path in G starting with m, e using edges in $V(J)$. Note that $m \in \mathcal{P}_M(m, e)$, and if $me \in E(J)$, then $\mathcal{P}_M(m, e)$ contains at least one more edge of M, namely the other one J-adjacent to e (which exists by the Degree Lemma).

Let $\mathcal{M} = \mathcal{M}(M, J)$ be the smallest family of all matchings $\hat{M} \subseteq V(J)$ with the properties that

(1) $M \in \mathcal{M}$

(2) For every $\hat{M} \in \mathcal{M}$ and for every C_4 with edges $\hat{m}, \hat{e}, \hat{m}', \hat{f} \in V(J)$, where $\hat{m}, \hat{m}' \in \hat{M}$, we have $\hat{M} \cup \left\{\hat{e}, \hat{f}\right\} \setminus \{\hat{m}, \hat{m}'\} \in \mathcal{M}$.

Obviously, each member of $\mathcal{M}$ can be obtained from M by a finite sequence of the above "C_4-switch" operation. Observe also that J is $\hat{M}$-reduced for every matching $\hat{M} \in \mathcal{M}$.

Let (M_1, m, e) be chosen such that $|\mathcal{P}_{M_1}(m, e)|$ is maximum among

$$\left\{(\hat{M}, \hat{m}, \hat{e}) : \hat{M} \in \mathcal{M}, \hat{m} \in \hat{M}, \hat{e} \in N_J(\hat{m})\right\}.$$

Note that the set we are maximizing over is non-empty because we are in Case 2, so $M \in \mathcal{M}$ has an edge J-adjacent to another edge. Our plan is to find a CP-decomposition with respect to M_1. This will be enough to prove our theorem because we can then "undo" the switches to arrive at our original matching M by repeatedly applying Lemma 2.4.12. For convenience we denote the vertex classes of G by A and B, with $m \cap e \subseteq A$.

Let $m' \in M_1$ be the other M_1-edge J-adjacent to e. If m has no J-neighbor intersecting it in B, we may proceed as in Case 1, and thereby have a CP-decomposition with respect to M_1. Otherwise, m has a J-neighbor on both sides, so let f be a J-neighbor of m with $m \cap f \subseteq B$. By the Degree Lemma, f is J-adjacent to another edge $m^* \in M_1$. We claim that in fact $m^* = m'$, and hence the edges m, e, m', f form a C_4.

Suppose $m^* \neq m'$. If $m^* \notin \mathcal{P}_{M_1}(m, e)$, we immediately arrive at a contradiction, because $\mathcal{P}_{M_1}(m^*, f)$ would then properly contain $\mathcal{P}_{M_1}(m, e)$ (just prepend m^*, f onto any M_1-alternating path starting with m, e), which contradicts the maximality of $|\mathcal{P}_{M_1}(m, e)|$. Thus we must have $m^* \in \mathcal{P}_{M_1}(m, e) \setminus \{m, m'\}$. By Lemma 2.4.10, there is an edge $g \in V(J) \setminus M_1$ which is J-adjacent to m^* with $m^* \cap g \subseteq B$ so that its other J-adjacent matching edge, $\hat{m} \in M_1$, is not in $\mathcal{P}_{M_1}(m, e)$. Then we claim $\mathcal{P}_{M_1}(\hat{m}, g)$ properly contains $\mathcal{P}_{M_1}(m, e)$, which would again be a contradiction.

To see that this is the case, take any matching edge $\tilde{m} \in \mathcal{P}_{M_1}(m, e)$, and we will show that $\tilde{m} \in \mathcal{P}_{M_1}(\hat{m}, g)$. If an M_1-alternating path starting with m, e reaching $\tilde{m}$ contains m^*, then we can start with $\hat{m}, g$ and continue along the segment of this path starting at m^*, since neither $\hat{m}$ nor g could be used in this path (otherwise $\hat{m} \in \mathcal{P}_{M_1}(m, e)$). If, on the other hand, $\tilde{m}$ is reachable from m, e without touching m^*, then we may reach $\tilde{m}$ by a path starting with $\hat{m}, g, m^*, f, m, e$. Thus, $\mathcal{P}_{M_1}(m, e) \subseteq \mathcal{P}_{M_1}(\hat{m}, g)$, and since the latter contains $\hat{m}$, while the former does not, we have the contradictory proper containment we were hoping for. Therefore $m^* = m'$.

Thus m has only f and edges parallel to f as J-neighbors at B. We will show now that similarly, m' has only e-parallel edges as J-neighbors at B. By Lemma 2.4.11 applied to $mem'f$, we have $\left|\mathcal{P}_{M_1^\times}(e, m)\right| = |\mathcal{P}_{M_1}(m, e)|$, so $(M_1^\times, e, m)$ is also a maximizing triple, where $M_1^\times = M_1 \cup \{e, f\} \setminus \{m, m'\}$. Thus, the argument of the previous two paragraphs can be applied to show that e only has m'-parallel edges as J-neighbors at B. By Corollary 2.4.8, this implies that m' also has only e-parallel edges as J-neighbors on that side.

We claim that among m, e, m', f, and all parallel edges we have that every parallel pair is non-J-adjacent and every pair of intersecting non-parallel edges is J-adjacent. To see that two parallel edges are not J-adjacent to each other, one must simply apply the Degree Lemma to M_1, $M_1^\times$, or one of these with a matching edge switched out for a parallel edge. Now suppose on the contrary that edges g parallel to m and h parallel to e are not J-adjacent. Then the Alternating Lemma for $M_1 \cup \{g\} \setminus \{m\}$ would imply that m' and f are not J-adjacent, which would be a contradiction.

Now we distinguish two further cases.

Case 2(a). $mem'f$ and parallel edges form a connected component of J.

Then we explode me to yield $J' = J \circledast me$ with $\mathrm{conn}(J') \leq k-3$. Since J' contains the matching $M' = M_1 \setminus \{m, m'\}$ of size $2k-2$, J' and M' satisfy the conditions of the theorem with $k' = k-1$, so by induction, there is a CP-decomposition with respect to M', say $S_1, \ldots, S_s, T_1, \ldots T_t$ with $s+t = k-1$.

Define S_{s+1} to be the C_4 given by $mem'f$. It is clear, that adding S_{s+1} to this CP-decomposition yields a CP-decomposition of J. This completes the proof of Case 2(a).

Case 2(b). $mem'f$ and parallel edges do not form a component of J.

Suppose without loss of generality that there is an edge $g \in V(J)$ not parallel to any of $mem'f$ which is J-adjacent to m. Note that we must have $m \cap g = m \cap e$ because all the J-neighbors of m intersecting it in $m \cap f$ are parallel to f. Then we explode mg and iteratively delete all decouplable edges to yield an M_1'-reduced $J' \subseteq J \circledast mg$ with $\mathrm{conn}(J') \leq k-3$, where $M_1' = M_1 \setminus \{m, m_1\}$ with m_1 the other M_1-edge J-adjacent to g. Since all J-neighbors of m' are parallel to e, and they are all J-adjacent to m, no J-neighbors of m' are present in J'. So m' has no J'-neighbor at $m' \cap e$. Thus it must have some J'-neighbor g' at $m' \cap f$, otherwise m' would be isolated and $\mathrm{conn}(J') = \infty$, a contradiction. Thus we explode $m'g'$ and get $J'' = J' \circledast m'g'$ with $\mathrm{conn}(J'') \leq k-4$. Let $m_2 \in M_1$ be the other matching edge J-adjacent to g' by the Degree Lemma. Then the matching $M'' = M_1 \setminus \{m, m', m_1, m_2\}$ of size $2k-4$ is contained in J''. Therefore, J'' and M'' satisfy the conditions of the theorem for $k'' = k-2$, so J'' has a CP-decomposition with respect to M'', say $S_1, \ldots, S_s, T_1, \ldots, T_t$ with $s+t = k-2$.

We define T_{t+1} to be the P_4 with edges $\{m, g, m_1\}$, and T_{t+2} to be the P_4 with edges $\{m', g', m_2\}$. Then we claim $S_1, \ldots, S_s, T_1, \ldots, T_{t+2}$ is a CP-decomposition of J with respect to M_1. To see this, we must verify that every edge not in an S_i and not parallel to an edge of an S_i is at home in some T_j. This is already true for all edges in $V(J'')$ (since $J'' \subseteq J$), so we only need to consider the edges we have removed by exploding mg and $m'g'$. However, all of these edges were by definition J-adjacent (or even J'-adjacent) to m, g, m', or g'. The edges J-adjacent to g and g' are automatically at home in T_{t+1} or T_{t+2} because the vertices of g and g' are the interior vertices of the respective P_4's. However, the only edges J-adjacent to m or m' but not at $m \cap g$ or $m' \cap g'$ are parallel to e and f. However, e-parallel edges are J-adjacent to g and f-parallel edges are J-adjacent to g' by Corollary 2.4.8, so they are also at home in T_{t+1} or T_{t+2}. Thus we have a CP-decomposition with respect to M_1.

All we need now is to use this CP-decomposition to get a CP-decomposition with respect to our original M. This is possible by several applications of Lemma 2.4.12 because M_1 is obtainable from M by a sequence of C_4-switches.

□

2.5 Good Sets

This section introduces the concept of good sets, which (as we will later see in Chapter 3) will help us find the substructure we need in our Ryser-extremal hypergraph in order to prove our characterization theorem by induction. The main result of this section implies that we can find good sets inside our link graphs in several cases, and hence if there are no good sets, we will know that the link graphs must have a certain form.

We start with a graph-theoretic definition, which will form the backbone of the definition of a good set.

Definition 2.5.1. Let G be a bipartite graph with vertex classes A and B. A subset $X \subseteq B$ is called *decent* if it satisfies the following conditions:

(1) $\nu(G) = |N(X)| + |B \setminus X|$,

(2) For every $x \in X$ and $y \in N(x)$ the edge xy participates in a maximum matching of G.

Note that if X is decent, then (1) implies that $|N(X)| \leq |X|$.

Lemma 2.5.2. *Let G be a bipartite graph with vertex classes A and B, and let M be a maximum matching in G. Let $X_0 \subseteq B$ be the set of M-unsaturated vertices in B, and let X be the set of vertices in B reachable on an M-alternating path from X_0 (including X_0). Then X is decent, and $|N(X)| = |X| - |X_0|$.*

Proof. Let $Y = N(X)$. Then Y is the set of vertices in A reachable on an M-alternating path from X_0. To see this, consider a vertex $x \in X$ and a neighbor $y \in N(x)$. Either x is unsaturated, in which case $x \in X_0$, so xy is an M-alternating path from X_0 to y, or there is an M-alternating path from X_0 to x, which must end with a matching edge. If y is on this path, we are done. Otherwise, xy is not a matching edge, and hence we can extend our path by the edge xy.

We claim that M saturates Y with (X, Y)-edges. This is because M is maximum, and thus every M-alternating path starting from an unsaturated vertex must end in a saturated vertex, and therefore every vertex of Y is incident to an edge of M. Extending the path by such a matching edge must land us in X by definition. Thus this matching edge is an (X, Y)-edge. Since X contains all M-unsaturated vertices, M saturates Y and $B \setminus X$ with distinct edges, and these are clearly all the edges of M. Thus $\nu(G) = |Y| + |B \setminus X|$, so X satisfies property (1).

We now show that X satisfies (2). Take an edge $e \in E(G)$ between X and Y. If $e \in M$, then we are done. If it has an M-unsaturated vertex, then it is only adjacent to one matching edge $m \in M$, and so $M \cup \{e\} \setminus \{m\}$ is a maximum matching containing e.

Otherwise, e is adjacent to two matching edges $m, m' \in M$. Since e goes between X and Y, the vertices of m and m' are reachable by an M-alternating path starting from X_0. Without loss of generality, the vertex in $m \cap e$ is in X. So consider an M-alternating path from X_0 which ends at that vertex. Note that its last edge is m. If m' is not in this path, then we can extend the path by e and m'. Switching along this extended path will create a maximum matching containing e (since the path starts at an M-unsaturated vertex). If, however, m' was in the original path, then adding e to the path forms an M-alternating cycle. Switching the matching along the cycle produces the desired matching. Therefore X is decent, as desired. □

Definition 2.5.3. Let G be a bipartite graph. A subset X of a vertex class of G is called *equineighbored* if X is nonempty and $|N(X)| = |X|$.

Note that if G has a perfect matching, then each vertex class is an equineighbored set (unless G is the empty graph).

Lemma 2.5.4. *Let G be a bipartite graph with vertex classes A and B and let M be a perfect matching in G. Let $X_0 \subseteq B$, and let X be the set of vertices in B reachable on an M-alternating path from X_0 (including X_0) starting with a non-matching edge. Then X is equineighbored.*

Proof. Let $Y = N(X)$. Since M is a perfect matching, every $y \in Y$ has a partner $x \in B$ matched to it by M. If there is an M-alternating path from X_0 to y starting with an edge not in M, then $x \in X$ because either $x \in X_0 \subseteq X$ or the path can be extended by the matching edge xy. If this holds for every $y \in Y$, then there is a matching from Y to X, so that $|Y| \leq |X|$, from which $|Y| = |X|$ follows by Hall's Theorem.

Therefore, we need to show that every $y \in Y$ can be reached from X_0 by an M-alternating path starting with a non-matching edge. Since $y \in N(X)$, it has a neighbor $x \in X$. By the definition of X, there is such an M-alternating path ending in x. If y is on that path, we are done. Otherwise, xy is not an edge of M (because the path to x ends with the matching edge incident to x), and so the path could be extended by xy, and thus y is on such a path. This concludes the proof. □

Lemma 2.5.5. *Let G be a bipartite graph with vertex classes A and B, and let M be a perfect matching in G. Let $X \subseteq B$ be a minimal equineighbored set in B. Then X is decent.*

Proof. Since G has a perfect matching, there is a matching saturating B, and since $|X| = |N(X)|$, we have $\nu(G) = |B| = |N(X)| + |B \setminus X|$, so X satisfies (1).

We now show that X satisfies (2). Let $Y = N(X)$. Let $x \in X$, $y \in Y$, and let $xy \in E(G)$. Fix a perfect matching M. Because $N(X) = Y$, it must match X to Y. If $xy \in M$, we are done. Otherwise there exist edges $xy', x'y \in M$ adjacent to xy. We claim that these edges participate in an M-alternating cycle with xy, and thus by switching along the cycle we get a new perfect matching which does include xy. To show that this happens, consider all M-alternating paths starting at x' with a non-matching edge. If there is such a path which hits y', then we can extend the path by $y'x$ and xy to give an M-alternating cycle in which xy participates. So assume that no such path hits y'. Let X' be the set of X-vertices which we can hit on such a path. Then X' is a proper ($x \notin X'$) non-empty ($x' \in X'$) equineighbored subset of X by Lemma 2.5.4 applied with $X_0 = \{x'\}$. This is a contradiction because X was chosen to be minimal. □

Definition 2.5.6. Let G be a bipartite graph with vertex classes A and B. A subset $X \subseteq B$ is called *good* if it is decent, and if for all $y \in N(X)$ we have $\mathrm{conn}\left(L\left(G - \{yz \in E(G) : z \in B \setminus X\}\right)\right) > \mathrm{conn}(L(G))$.

Note in particular that if X is good, then $\{yz \in E(G) : z \in B \setminus X\} \neq \emptyset$ for all $y \in N(X)$.

Lemma 2.5.7. *Let G be a bipartite graph with vertex classes A and B. Suppose $\nu(G) = 2k$ for some integer k and $\mathrm{conn}(L(G)) = k - 2$. If G has no good set in A nor in B, then the following hold:*

(i) G has a perfect matching

(ii) For every minimal equineighbored subset $X \subseteq A$ or $X \subseteq B$ we have $|X| = 2$. In particular, $G[X \cup N(X)]$ is a C_4 (possibly with parallel edges).

Note that the minimality requirement in (ii) is well-defined because by (i) both A and B are equineighbored.

Proof. Assume that G has no good sets. First, we show that (i) holds. Suppose G does not have a perfect matching. Let M be a maximum matching in G. By assumption, there are some M-unsaturated vertices in $A \cup B$. Without loss of generality assume that at least one of them is in B. Let X_0 be the set of M-unsaturated vertices in B. Consider all the M-alternating paths in G starting from X_0. Let X be the set of vertices in B reachable on an M-alternating path from X_0 (including X_0), and let $Y = N(X)$. We claim that X is a good subset. By Lemma 2.5.2 X is decent, so we must simply check that for all $y \in Y$ we have $\mathrm{conn}\left(L\left(G - \{yz \in E(G) : z \in B \setminus X\}\right)\right) > \mathrm{conn}(L(G))$.

Let $y \in Y$. Let $G_y = G - \{yz \in E(G) : z \in B \setminus X\}$. Clearly M is still a maximum matching in G_y and X_0 remains the set of M-unsaturated vertices. All of the (X, Y)-edges have been preserved in G_y, so X and Y are still the sets of vertices reachable by an M-alternating path from X_0. Suppose for the sake of contradiction that we had $\mathrm{conn}(L(G_y)) = k-2$. Then we pass to an M-reduced subgraph $J \subseteq L(G_y)$ of the line graph by iteratively deleting all decouplable edges (see Definition 2.4.6). This means $\mathrm{conn}(J) = k-2$, but $\mathrm{conn}(J-e) \geq k-1$ for all $e \in E(J)$).

Claim. *The edges between X and Y form an independent set in J.*

Proof of claim. First, by the Degree Lemma (Lemma 2.4.7), any edge e parallel to an edge of M is not J-adjacent to any edge of M. Next, by the Alternating Lemma (Lemma 2.4.9) any two edges which are together in an M-alternating path from X_0 are not J-adjacent. Now consider a matching edge $m \in M$ and an (X, Y)-edge e which intersects it in a vertex v. Because m hits X, there is an M-alternating path starting at X_0 which has m as its last edge. If this path ends in v, then we can add e to that path to obtain either a longer M-alternating path or to obtain an M-alternating cycle. Either way, the Alternating Lemma gives that e and m are not J-adjacent.

If v is not at the end of this path, then consider the other M-edge m' which intersects e (if this does not exist, then e is not J-adjacent to m by the Degree Lemma (Lemma 2.4.7)). There is an M-alternating path starting at X_0 which has m' as its last edge. This path ends in the intersection of m' and e, so by the previous argument, e and m' cannot be J-adjacent, and so by the Degree Lemma, e and m are not J-adjacent either. Thus we have shown that none of the (X, Y)-edges are J-adjacent to the edges of M.

Now consider two intersecting non-matching edges e and f between X and Y. If they were J-adjacent, then they would be explodable, but because e and f are not J-adjacent to any M-edges, $M \subseteq V(J \circledast ef)$, so by Lemma 2.1.1, $\operatorname{conn}(J \circledast ef) \geq |M|/2 - 2 = k - 2$. This contradicts explodability, so they must not be J-adjacent. $\square$

Now consider the matching edge $m \in M$ containing y. It is isolated in J, because all of the edges intersecting m at all are (X, Y)-edges. This is a contradiction, because m is then an isolated vertex of J, which means $\operatorname{conn}(J) = \infty$, a contradiction. Thus we must have $\operatorname{conn}(L(G_y)) \geq k - 1$ as desired. Thus X is good. This contradicts the assumption that there were no good sets, so G must in fact have a perfect matching, proving (i).

Now we will show (ii) holds. Let $X \subseteq B$ be a minimal equineighbored set. We want to show that $|X| = 2$, from which easily follows that the edges incident to X form a C_4 (possibly with parallel edges). Indeed, if X is a minimal equineighbored set of size 2, then its vertices must both have two neighbors (a vertex with only one neighbor would be a proper equineighbored subset, a vertex with more than two neighbors is ruled out by the fact that $|N(X)| = 2$, and an isolated vertex is ruled out by the fact that we have a perfect matching), which means they both connect to both neighbors of X, which forms a C_4.

So suppose that $|X| \neq 2$. We will show that X is good. By Lemma 2.5.5, X is decent, so we must simply check that for all $y \in N(X)$, the graph G_y formed by erasing from G all edges incident to y and not incident to X has the property that $\operatorname{conn}(L(G_y)) \geq k - 1$.

Indeed suppose it did not. We could then apply Theorem 2.4.3 to get a CP-decomposition of $L(G_y)$. Note that X is still a minimal equineighbored subset of B in G_y.

Claim. *X does not contain any interior vertex of a P_4 in any CP-decomposition of $L(G_y)$ with respect to any perfect matching.*

Proof of claim. Fix a perfect matching M of G_y, and fix a CP-decomposition $S_1, \ldots, S_s, T_1, \ldots, T_t$ of $L(G_y)$ with respect to M. Let X_0 be the set of interior vertices of the paths T_j in X. Then $X \setminus X_0$ is also equineighbored because the endpoints of the paths T_j which are partnered with the vertices of X_0 in the matching M are not in the neighborhood of $X \setminus X_0$ since all edges incident to them must connect to interior vertices of the paths. Since there are $|X_0|$ endpoints in X, we have removed at least as many vertices from the neighborhood as we have removed from X. Note that $X \setminus X_0$ cannot be empty as X could not have consisted entirely of interior vertices of the paths, since those have at least two distinct neighbors each. It follows that X_0 must have been empty and the claim follows. □

Claim. *X does not contain any vertices of a C_4 in any CP-decomposition of $L(G_y)$ with respect to any perfect matching.*

Proof of claim. Fix a perfect matching M of G_y, and fix a CP-decomposition $S_1, \dots, S_s$, $T_1, \dots, T_t$ of $L(G_y)$ with respect to M. Let X_0 be the vertices of some 4-cycle S_i which are contained in X. Then $X \setminus X_0$ is also equineighbored because the two vertices of that S_i which are adjacent to X_0 are not in the neighborhood of $X \setminus X_0$ as X does not contain any interior vertices of any T_j by the previous claim, and the only neighbors of the vertices of S_i are other vertices of S_i and interior vertices of paths T_j by the definition of a CP-decomposition. Therefore we would remove at least as many vertices from the neighborhood of X as we would remove from X. It follows that if X_0 is nonempty, then $|X_0| = 2$, because if $|X_0| = 1$, then we would have $|N(X \setminus X_0)| < |X \setminus X_0|$, which contradicts the fact that G_y has a perfect matching. Since $|X| \neq 2$, we cannot have $X \setminus X_0 = \emptyset$, so $X \setminus X_0$ is a proper equineighbored subset of X, which is a contradiction to the minimality of X. □

Thus we have shown that X consists entirely of endpoints of P_4's (there are no other types of vertices, since we have a perfect matching). Then y is an interior vertex of some P_4. However, y only has neighbors in X, so this cannot be the case (since every interior vertex of a path is adjacent to another interior vertex). Since we have reached a contradiction, it follows that we must have $\text{conn}(L(G_y)) \geq k - 1$. Thus X is a good set, which is a contradiction to the conditions of the lemma. Therefore, we must have $|X| = 2$ and $G[X \cup N(X)]$ is a C_4, which is (ii). This proves the lemma. □

2.6 Remarks and Open Problems

Concerning the tightness of Theorem 2.1.1 several interesting questions remain open. In the main result of this chapter we characterized those *bipartite* graphs for which the theorem is tight when $r = 2$.

What happens with this characterization if one leaves out the restriction of bipartiteness? The graph G consisting of a triangle and a hanging edge is an example of a non-bipartite graph which is tight for Theorem 2.1.1. Indeed, $\nu(G) = 2$ while the line graph is K_4 minus an edge, having a disconnected independence complex. It would be very interesting to obtain a full characterizations of those graphs G which are tight for Theorem 2.1.1.

Another natural direction is to consider hypergraphs with uniformity higher than 2. It is not difficult to see that Theorem 2.1.1 is also best possible for every $r > 2$. Just take a matching of size mr and add m edges that intersect r different matching edges each. However, a characterization of those r-graphs for which $\mathrm{conn}(\mathcal{H}) = \frac{\nu(\mathcal{H})}{r} - 2$ is still outstanding; the case of r-partite r-graphs already being very interesting.

A related question concerns the relationship of Theorem 2.1.1 to Ryser's Conjecture for $r > 2$. We mentioned already that in [16] we complete the proof that a graph is tight for Theorem 2.1.1 if and only if it is the link graph of a Ryser-extremal 3-graph. Is this equivalence or at least one of its directions true for $r > 2$?

Finally, Theorem 2.1.1 has a chance to be best possible only for graphs whose matching number is even. It would be interesting to prove a characterization of 2-graphs with an *odd* matching number and having a line graph with connectedness as small as possible (in terms of the matching number). Is there is a CP-decomposition-type characterization of all (bipartite) graphs with matching number $2k + 1$ and connectedness $k - 1$?

Chapter 3

Home-Base Hypergraphs

Joint work with Penny Haxell and Tibor Szabó.

3.1 Introduction

Our aim in this chapter is to prove Theorem 1.1.2, which we repeat here for convenience:

Theorem 1.1.2. *Let $\mathcal{H}$ be a 3-partite 3-graph. Then $\tau(\mathcal{H}) = 2\nu(\mathcal{H})$ if and only if $\mathcal{H}$ is a home-base hypergraph.*

Home-base hypergraphs have a restricted structure, but are far from being unique: for any given positive integer $k \in \mathbb{N}$ there are infinitely many home-base hypergraphs with matching number k. The precise description is given in the following subsection.

3.1.1 Home-Base Hypergraphs

To motivate our definition of home-base hypergraphs, let us start with some examples of 3-graphs $\mathcal{H}$ with $\tau(\mathcal{H}) = 2 = 2\nu(\mathcal{H})$. A general example of an r-graph, which is tight for Ryser's Conjecture is the *truncated projective plane* $F^{(r)}$. Its vertex set is constructed by taking the projective plane over the $(r-1)$-element field and removing one point from it. The lines of the plane which were incident to this point become the vertex classes of the r-graph, and the rest of the lines become the edges. Since any two lines of the projective plane intersect, we have $\nu(F^{(r)}) = 1$. It is also not difficult to see that the smallest vertex covers are the vertex classes and hence $\tau(F^{(r)}) = r - 1$. Truncated projective planes exist whenever r is one greater than a prime power. Luckily, 3 is such a number, and thus we have the truncated Fano plane. Concretely, the *truncted Fano-plane* is the 3-graph $F^{(3)} = F$ with vertex set $\{a, b, c, x, y, z\}$ and edges abc, ayz, xbz, and xyc (here the vertex classes are $\{a, x\}$, $\{b, y\}$, and $\{c, z\}$).

Adding parallel edges to any hypergraph does not affect the vertex cover number or the matching number. We call any 3-graph a *truncated multi-Fano plane*, if it is obtained from the truncated Fano-plane by adding an arbitrary number of parallel edges.

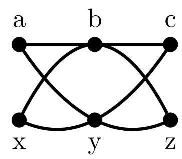

Figure 3.1: The truncated Fano plane.

However, the truncated Fano-plane is not minimal, since removing any edge from it yields another example of an intersecting hypergraph which cannot be covered by a single vertex. To be concrete, let H be the hypergraph on the vertex set $\{a, b, c, x, y, z\}$ and edges ayz, xbz, and xyc. Three of the vertices have degree 2 and three have degree 1. One can extend H by adding edges (perhaps containing new vertices) which contain two of the degree 2 vertices and still obtain an intersecting hypergraph (and obviously the vertex cover number does not decrease). This creates a family of edges which is intersecting simply because they all contain two of the vertices x, y, and z. Thus this family is determined by the set $R = \{x, y, z\}$.

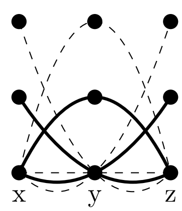

Figure 3.2: The truncated Fano plane minus one edge, with possible additional edges drawn in dashed lines.

We say that a 3-partite 3-graph $\mathcal{H}$ is *Ryser-extremal*, if $\tau(\mathcal{H}) = 2\nu(\mathcal{H})$. Our hope would be that every Ryser-extremal 3-graph is made up of such R-families and truncated multi-Fano-planes. This is indeed the case, but the edges of these substructures can intersect in various intricate ways. How exactly, is made precise in the following series of definitions.

Definition 3.1.1. Let $\mathcal{H}$ be a 3-partite 3-graph. An *FR-partition* of $\mathcal{H}$ is a triple $(\mathcal{F}, \mathcal{R}, W)$ with $\mathcal{F}, \mathcal{R} \subseteq 2^{V(\mathcal{H})}$ and $W \subseteq V(\mathcal{H})$ which satisfies the following conditions:

(1) $\mathcal{F} \cup \mathcal{R} \cup \{W\}$ is a partition of the vertices of $\mathcal{H}$,

(2) For each $F \in \mathcal{F}$, the induced hypergraph $\mathcal{H}|_F$ is isomorphic to a truncated multi-Fano plane,

(3) Each $R \in \mathcal{R}$ is a three-vertex set with one vertex from each vertex class of $\mathcal{H}$,

(4) $|\mathcal{F} \cup \mathcal{R}| = \nu(\mathcal{H})$.

Note that $\mathcal{F}$ is a 6-graph and $\mathcal{R}$ is a 3-graph.

Definition 3.1.2. Let $\mathcal{H}$ be a 3-partite 3-graph with vertex classes V_1, V_2, and V_3, and let $(\mathcal{F}, \mathcal{R}, W)$ be an FR-partition of $\mathcal{H}$. For each vertex class V_i, we define a bipartite graph B_i with vertex classes $\mathcal{R}$ and $W \cap V_i$ and with an edge between $R \in \mathcal{R}$ and $w \in W \cap V_i$ precisely when there is an edge of $\mathcal{H}$ containing w and two vertices of R. The partition $(\mathcal{F}, \mathcal{R}, W)$ is called *matchable* if each B_i has a matching saturating $\mathcal{R}$.

An example of a non-matchable FR-partition is given in the following picture, where the boxes correspond to two R's and the unboxed vertices are in W:

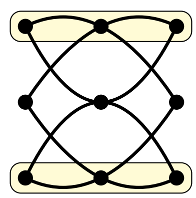

Figure 3.3: An unmatchable FR-partition.

Definition 3.1.3. An FR-partition $(\mathcal{F}, \mathcal{R}, W)$ of $\mathcal{H}$ is said to have the *edge-home* property if every edge of $\mathcal{H}$ is either in $\mathcal{H}|_F$ for some $F \in \mathcal{F}$ or contains two vertices from some $R \in \mathcal{R}$.

Definition 3.1.4. A matchable FR-partition with the edge-home property is called a *home-base partition*. $\mathcal{H}$ is called a *home-base hypergraph* if it has a home-base partition.

Notation. For each $F \in \mathcal{F}$, we call an edge an F-edge if it is in $\mathcal{H}|_F$. For each $R \in \mathcal{R}$, we call an edge an R-edge if it contains two vertices from R. We call an edge an $\mathcal{F}$-edge if it is an F-edge for some $F \in \mathcal{F}$, and call an edge an $\mathcal{R}$-edge if it is an R-edge for some $R \in \mathcal{R}$.

Here follows an example of a home-base hypergraph. The boxes correspond to members of $\mathcal{F}$ or $\mathcal{R}$, and the unboxed vertices are in W. The bolded edges are the edges of $\mathcal{H}|_F$ for some $F \in \mathcal{F}$ or the edges corresponding to the edges of arbitrarily chosen matchings saturating $\mathcal{R}$ in the auxiliary bipartite graphs B_i.

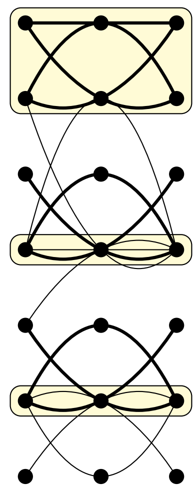

Figure 3.4: A home-base hypergraph with its home-base partitition.

We can easily see one direction of Theorem 1.1.2:

Proposition 3.1.5. *If $\mathcal{H}$ has a home-base partition $(\mathcal{F}, \mathcal{R}, W)$, then $\tau(\mathcal{H}) = 2\nu(\mathcal{H})$.*

Proof. Let $T \subseteq V(\mathcal{H})$ be a vertex cover. We aim to show that it has size at least $2\nu(\mathcal{H}) = 2\,|\mathcal{F} \cup \mathcal{R}|$. Since the partition is matchable, each of the auxiliary bipartite graphs B_1, B_2, and B_3 have matchings saturating $\mathcal{R}$, say M_1, M_2, and M_3, respectively. Then each $R = \{r_1, r_2, r_3\} \in \mathcal{R}$ has three W-vertices, $w_i^R \in V_i$ assigned to it, so that $Rw_i^R \in M_i$, which means that $w_i^R r_j r_k$ are edges for each choice of $\{i, j, k\} = \{1, 2, 3\}$. So consider only the edges of this form together with the edges of $\mathcal{H}|_F$ for each $F \in \mathcal{F}$. Each set of edges for each $R \in \mathcal{R}$ and $F \in \mathcal{F}$ is disjoint from the other sets, so any vertex cover must cover each set with different vertices. Since each such set forms an intersecting 3-partite 3-graph with vertex cover number 2, T must have at least two vertices for each $R \in \mathcal{R}$ and each $F \in \mathcal{F}$, giving a total of at least $2\,|\mathcal{R} \cup \mathcal{F}| = 2\nu(\mathcal{H})$ vertices as required. This shows $\tau(\mathcal{H}) \geq 2\nu(\mathcal{H})$. Since Ryser's Conjecture is true for 3-partite 3-graphs, we have $\tau(\mathcal{H}) = 2\nu(\mathcal{H})$. □

Note that we did not make use of the edge-home property in this proof. This property is necessary however to ensure that if a home-base partition exists, then it is unique. Uniqueness is not necessary for our proof of the main theorem, but we include it here out of interest.

Proposition 3.1.6. *Let $\mathcal{H}$ be a 3-partite 3-graph with home-base partitions $(\mathcal{F}, \mathcal{R}, W)$ and $(\mathcal{F}', \mathcal{R}', W')$. Then $\mathcal{F} = \mathcal{F}'$, $\mathcal{R} = \mathcal{R}'$, and $W = W'$.*

Proof. Consider $F \in \mathcal{F}$. Call its vertices $\{a, b, c, x, y, z\}$ so that abc, ayz, xbz, and xyc are edges of $\mathcal{H}$. Note that no other edge of $\mathcal{H}$ intersects F in more than one vertex by the edge-home property of $(\mathcal{F}, \mathcal{R}, W)$. If $F \notin \mathcal{F}'$, then at least one of these edges is not an $\mathcal{F}'$-edge. By the symmetries of the truncated Fano plane, we may assume without loss of generality that edge is abc. Because $(\mathcal{F}', \mathcal{R}', W')$ has the edge-home property, abc must be an $\mathcal{R}'$-edge. Without loss of generality, let $a, b \in R \in \mathcal{R}'$. Now the edge ayz has one vertex in R, so it cannot be an $\mathcal{F}'$-edge either. There are two possibilities: either it is an R-edge, or it is an R'-edge for some other $R' \in \mathcal{R}'$. If it is an R-edge, then $R = \{a, b, z\}$. Because $(\mathcal{F}', \mathcal{R}', W')$ is matchable, there must be W'-vertices on each side which are in an R-edge. But as we have noted, no edge outside of abc, ayz, xbz, and xyc intersects R in two vertices. Thus it must be the case that $x, y, c \in W'$. But then xyc is an edge which is neither an $\mathcal{F}'$-edge nor an $\mathcal{R}'$-edge — a contradiction. Therefore ayz must have been an R'-edge with $y, z \in R'$. But again by matchability, there must be a vertex $w \in W'$ such that $wyz \in E(\mathcal{H})$. Since $a \notin W'$, we must have $w \neq a$, which cannot happen for the same reason as before. Thus $F \in \mathcal{F}'$, and by symmetry, we thus have $\mathcal{F} = \mathcal{F}'$.

Consider now $R \in \mathcal{R}$. Call its vertices $\{x, y, z\}$, and let $a, b, c \in W$ such that $ayz, xbz, xyc \in E(\mathcal{H})$ (these edges exist because $(\mathcal{F}, \mathcal{R}, W)$ is matchable. In $(\mathcal{F}', \mathcal{R}', W')$ these are all $\mathcal{R}'$-edges, because if there were an $\mathcal{F}'$-edge among them, this would contradict the fact that $\mathcal{F} = \mathcal{F}'$. Thus if $R \notin \mathcal{R}'$, then at least one of the vertices a, b, or c must be in some $R' \in \mathcal{R}'$ such that one of these edges is an R'-edge (otherwise we would quickly conclude that one of the edges is neither an $\mathcal{F}'$-edge nor an $\mathcal{R}'$-edge). By symmetry, we may assume without loss of generality that $a, y \in R'$. Now consider the edge xyc. Again there are two possibilites: either it is an R'-edge, or it is an R''-edge for some other $R'' \in \mathcal{R}'$. If it is an R'-edge, then $R' = \{a, y, c\}$ and by the matchability of $(\mathcal{F}', \mathcal{R}', W')$, there would need to be an edge awc for some $w \in W'$. But this edge cannot exist, because it contains two vertices of W (namely a and c), and hence is neither an $\mathcal{F}$-edge nor an $\mathcal{R}$-edge, which cannot be the case because $(\mathcal{F}, \mathcal{R}, W)$ has the edge-home property. Thus, xyc must be an R''-edge with $x, c \in R''$. But then again by matchability, there must be an edge xwc for some $w \in W'$ (and hence $w \neq y$). This edge cannot exist if $(\mathcal{F}, \mathcal{R}, W)$ is a home-base partition because it contains one R-vertex, one W-vertex and one third vertex which is not in R. This is a contradiction, and thus $R \in \mathcal{R}'$. By symmetry, we then have $\mathcal{R} = \mathcal{R}'$.

Since $W = V(\mathcal{H}) \setminus (\bigcup(\mathcal{F} \cup \mathcal{R})) = V(\mathcal{H}) \setminus (\bigcup(\mathcal{F}' \cup \mathcal{R}')) = W'$, we have shown that these are in fact the same home-base partitions. $\square$

It is clear that given the characterization in Theorem 1.1.2, we can easily enumerate all Ryser-extremal 3-graphs.

3.1.2 Proof Outline

The main topic of this chapter is the proof of Theorem 1.1.2. We have just seen that home-base hypergraphs are Ryser-extremal. The proof of the reverse implication will be done by induction on $\nu(\mathcal{H})$.

The case $\nu(\mathcal{H}) = 0$ is trivial, and even the case $\nu(\mathcal{H}) = 1$ is not difficult to check. Much of the work involved in proving the cases $\nu(\mathcal{H}) \geq 2$ consists of finding an appropriate structure to which we can apply induction. That means a subhypergraph $\mathcal{H}_0 \subseteq \mathcal{H}$ which also satisfies $\tau(\mathcal{H}_0) = 2\nu(\mathcal{H}_0)$ and has $\nu(\mathcal{H}_0) < \nu(\mathcal{H})$. By induction, this will have a home-base partition, but in order to be able to extend this partition to a home-base partition of the whole of $\mathcal{H}$ we will also need the edges outside of $\mathcal{H}_0$ to behave nicely.

A more precise description of the structure of the proof is given by the flow chart in Figure 3.5. Please note that it is intended as a guide to be referred to throughout the proof, and many of the terms will only be introduced in later sections.

In Section 3.2, we collect theorems we have shown in Chapter 2 about the connectedness of the line graphs of the link graphs of Ryser-extremal 3-graphs. Among others, this involves a structural characterization of the link graphs, which we call a CP-decomposition, as well as a theorem about bipartite graphs without so-called good sets. Good sets will turn out to be very useful to have in one of the link graphs of a Ryser-extremal 3-graph, while the lack of good sets in a bipartite graph imposes very strong restrictions on its structure, which will eventually help us to show that we are dealing with a home-base hypergraph.

In Section 3.3, we prove some important properties of home-base hypergraphs, which will be essential for several parts of the rest of the proof.

In Section 3.4, we define and study cromulent and perfectly cromulent triples. A perfectly cromulent triple is a set of vertices such that the rest is a home-base hypergraph that interacts with the rest of the edges in a controlled fashion. This turns out to be precisely the substructure we need so that we can extend the home-base partition given by induction to a home-base partition of the whole hypergraph. Cromulent triples are apparently weaker versions of perfectly cromulent triples, but careful considerations will show that no cromulent triple can actually fail to be perfectly cromulent under the assumption that $\tau = 2\nu$. Therefore, it will be enough to find just a cromulent triple in order to show that we have a home-base hypergraph.

In Section 3.5, we show how to use a good set to find a perfectly cromulent triple and hence conclude that we are dealing with a home-base hypergraph. The rest of Section 3.5 is devoted to exploring how the edges of the link graphs extend to hyperedges under the assumption that there are no good sets and no cromulent triples.

In Section 3.6, we use the information on how the links extend, together with the fact that the links have CP-decompositions to show that the hypergraph must contain a truncated multi-Fano plane that interacts minimally with the rest of the hypergraph, which by induction will have a home-base partition. It is then easy to show that adding the lone F results in a home-base partition of the whole hypergraph.

The proof of Theorem 1.1.2 is assembled from all of the theorems and lemmas of the preceeding four sections in Section 3.7.

In Section 3.8 we prove a couple of facts related to our main theorem, some of them leading to interesting open questions.

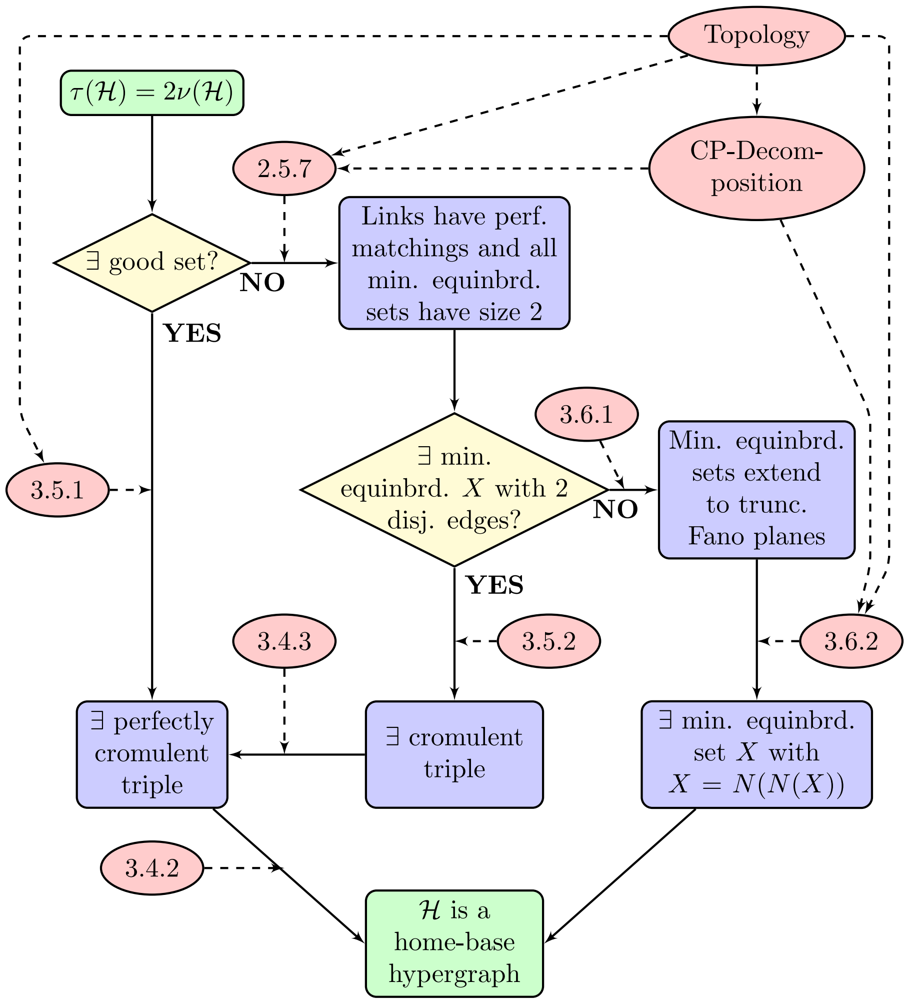

Figure 3.5: A flow-chart describing the logic of the proof with relevant lemmas shown.

3.2 Theorems about the link graph

In this section we collect theorems that will be used in our arguments. For proofs, see Chapter 2.

The *line graph* $L(\mathcal{H})$ of a hypergraph $\mathcal{H}$ is the simple graph $L(\mathcal{H})$ on the vertex set $E(\mathcal{H})$ with $e, f \in V(L(\mathcal{H}))$ adjacent if $e \cap f \neq \emptyset$.

Recall that the *connectedness* of a graph G, denoted $\operatorname{conn}(G)$, is the largest k such that the independence complex of the graph G is k-connected.

Theorem 2.1.1. *Let $\mathcal{G}$ be an r-graph. Then*

$$\operatorname{conn}(L(\mathcal{G})) \geq \frac{\nu(\mathcal{G})}{r} - 2.$$

Definition 2.1.2. Let $\mathcal{H}$ be a 3-partite 3-graph with parts V_1, V_2, and V_3. Let $S \subseteq V_i$ for some $i = 1, 2, 3$. Then the *link graph* $\operatorname{lk}_{\mathcal{H}}(S)$ is the bipartite graph with vertex classes V_j and V_k (where $\{i, j, k\} = \{1, 2, 3\}$) whose edge multiset is $\{e \setminus V_i : e \in E(\mathcal{H}), e \cap V_i \subseteq S\}$.

Proposition 2.3.1. *Let $\mathcal{H}$ be a* 3*-partite* 3*-graph with vertex classes V_1, V_2, and V_3. Then for each $i \in \{1, 2, 3\}$ we have the following:*

(i) For all $S \subseteq V_i$ we have

$$\operatorname{conn}(L(\operatorname{lk}_{\mathcal{H}}(S))) \geq \frac{\tau(\mathcal{H}) - (|V_i| - |S|)}{2} - 2.$$

(ii) There is some $S \subseteq V_i$ such that

$$\operatorname{conn}(L(\operatorname{lk}_{\mathcal{H}}(S))) \leq \nu(\mathcal{H}) - (|V_i| - |S|) - 2.$$

(iii) For every $S \subseteq V_i$ for which the inequality in (ii) holds we have

$$|S| \geq |V_i| - (2\nu(\mathcal{H}) - \tau(\mathcal{H})).$$

Theorem 2.1.3. *If $\mathcal{H}$ is a* 3*-partite* 3*-graph with vertex classes V_1, V_2, and V_3, such that $\tau(\mathcal{H}) = 2\nu(\mathcal{H})$, then for each i we have*

(i) $\operatorname{conn}(L(\operatorname{lk}_{\mathcal{H}}(V_i))) = \nu(\mathcal{H}) - 2$.

(ii) $\nu(\operatorname{lk}_{\mathcal{H}}(V_i)) = \tau(\mathcal{H})$.

In particular

$$\operatorname{conn}(L(\operatorname{lk}_{\mathcal{H}}(V_i))) = \frac{\nu(\operatorname{lk}_{\mathcal{H}}(V_i))}{2} - 2. \tag{3.2.1}$$

Theorem 2.4.3. *Let G be a bipartite graph. Then we have* $\mathrm{conn}(L(G)) = \frac{\nu(G)}{2} - 2$ *if and only if G has a collection of $\nu(G)/2$ pairwise vertex-disjoint subgraphs, each of them a C_4 or a P_4, such that every edge of G is parallel to an edge of one of the C_4's or is incident to an interior vertex of one of the P_4's.*

We refer to such a collection as a *CP-decomposition*. Note that this is just a specialization of the concept of CP-decomposition in Chapter 2 for the entire line graph, which is the only case we will need in this part. As promised in Chapter 2, the "if" direction of this theorem will be proved here. We will postpone the proof until Section 3.8, as it is not necessary for the proof of the main theorem.

For a subset X of the vertices of a graph, we denote the *neighborhood* of X by $N(X)$, meaning the set of vertices adjacent to some vertex in X.

Definition 2.5.1. Let G be a bipartite graph with vertex classes A and B. A subset $X \subseteq B$ is called *decent* if it satisfies the following conditions:

(1) $\nu(G) = |N(X)| + |B \setminus X|$,

(2) For every $x \in X$ and $y \in N(x)$ the edge xy participates in a maximum matching of G.

Definition 2.5.3. Let G be a bipartite graph. A subset X of a vertex class of G is called *equineighbored* if X is nonempty and $|N(X)| = |X|$.

Definition 2.5.6. Let G be a bipartite graph with vertex classes A and B. A subset $X \subseteq B$ is called *good* if it is decent, and if for all $y \in N(X)$ we have $\mathrm{conn}\left(L\left(G - \{yz \in E(G) : z \in B \setminus X\}\right)\right) > \mathrm{conn}(L(G))$.

Note in particular that if X is good, then $\{yz \in E(G) : z \in B \setminus X\} \neq \emptyset$ for all $y \in N(X)$.

Lemma 2.5.7. *Let G be a bipartite graph with vertex classes A and B. Suppose $\nu(G) = 2k$ for some integer k and* $\mathrm{conn}(L(G)) = k - 2$. *If G has no good set in A nor in B, then the following hold:*

(i) G has a perfect matching

(ii) For every minimal equineighbored subset $X \subseteq A$ or $X \subseteq B$ we have $|X| = 2$. In particular, $G[X \cup N(X)]$ is a C_4 (possibly with parallel edges).

Note that the minimality requirement in (ii) is well-defined because by (i) both A and B are equineighbored.

3.3 Properties of Home-Base Hypergraphs

The next couple of sections will establish some basic properties of home-base hypergraphs that we will need in the proof of Theorem 1.1.2.

First is the so-called "monster lemma," which states under which conditions a monster can eat some vertices of a home-base hypergraph without reducing the matching number.

But before we can prove it, we shall need some definitions.

3.3.1 Essential and Superfluous Vertices

Definition 3.3.1. Let G be a bipartite graph with vertex classes X_1 and X_2. A subset $C \subseteq X_i$ is called *essential* if there is a subset $U \subseteq X_{3-i}$ with $|U| = |C|$ and $C = N(U)$.

We remark briefly that non-empty essential subsets are precisely the neighborhoods of equineighbored subsets. We will of course apply this concept to the bipartite graphs B_i from the matchability criterion of FR-partitions.

Let $\mathcal{H}$ be a 3-partite 3-graph on vertex classes V_1, V_2, and V_3 with a matchable FR-partition $(\mathcal{F}, \mathcal{R}, W)$. We call a vertex v in V_i essential if $v \in W$ and $\{v\} \subseteq W \cap V_i$ is essential in B_i. If $R \in \mathcal{R}$ has only $v \in W \cap V_i$ as its neighbor in B_i, then we say v is essential for R.

Lemma 3.3.2. *Let B be a bipartite graph with vertex classes $\mathcal{R}$ and W, which has a matching saturating $\mathcal{R}$. Then W contains a unique maximal essential subset.*

Proof. Let $C_1, C_2 \subseteq W$ be essential. Then we claim $C_1 \cup C_2$ is also essential. Consider $\mathcal{U}_1, \mathcal{U}_2 \subseteq \mathcal{R}$ such that $C_1 = N_B(\mathcal{U}_1)$, $C_2 = N_B(\mathcal{U}_2)$, $|\mathcal{U}_1| = |C_1|$ and $|\mathcal{U}_2| = |C_2|$. Then $N_B(\mathcal{U}_1 \cup \mathcal{U}_2) = C_1 \cup C_2$ and by Hall's Theorem, $|C_1 \cup C_2| \geq |\mathcal{U}_1 \cup \mathcal{U}_2|$. But of course $N_B(\mathcal{U}_1 \cap \mathcal{U}_2) \subseteq C_1 \cap C_2$ and thus again by Hall's Theorem, $|C_1 \cap C_2| \geq |\mathcal{U}_1 \cap \mathcal{U}_2|$. By the inclusion-exclusion principle, we thus have $|C_1| + |C_2| - |C_1 \cup C_2| \geq |\mathcal{U}_1| + |\mathcal{U}_2| - |\mathcal{U}_1 \cup \mathcal{U}_2|$, and since $|\mathcal{U}_1| = |C_1|$ and $|\mathcal{U}_2| = |C_2|$, we find that $|C_1 \cup C_2| \leq |\mathcal{U}_1 \cup \mathcal{U}_2|$, so that in fact there is equality. This proves that $C_1 \cup C_2$ is essential. Therefore the union over all essential subsets of W gives the unique maximal essential set. □

A vertex of W which is not in the maximal essential set is called *superfluous*. Note that any one superfluous vertex can be removed, and the rest of the bipartite graph will still have a matching saturating $\mathcal{R}$. Again, we will apply this to the bipartite graphs B_i from the matchability criterion of FR-partitions.

Let $\mathcal{H}$ be a home-base hypergraph on vertex classes V_1, V_2, and V_3 with a home-base partition $(\mathcal{F}, \mathcal{R}, W)$. Then the auxiliary bipartite graphs B_i have vertex classes $\mathcal{R}$ and $W \cap V_i$ and a matching saturating $\mathcal{R}$. Therefore, each $W \cap V_i$ contains a unique maximum essential subset C_i, and we may call a vertex of V_i superfluous if it is in $W \cap V_i \setminus C_i$. Clearly superfluous vertices are non-essential W-vertices in a stronger form. We can make the following observation:

Observation 3.3.3. *Let $\mathcal{H}$ be a 3-partite 3-graph with a matchable FR-partition $(\mathcal{F}, \mathcal{R}, W)$, and let $S \subseteq W$ be a set of superfluous vertices with at most one vertex in each vertex class. Then $(\mathcal{F}, \mathcal{R}, W \setminus S)$ is a matchable FR-partition of $\mathcal{H} - S$.*

Proof. Since removing any single superfluous vertex s from any of the bipartite graphs B_i leaves a matching saturating $\mathcal{R}$, $(\mathcal{F}, \mathcal{R}, W \setminus \{s\})$ is a matchable FR-partition. Since removing s from one does not change the other graphs B_j at all, we can do this for each vertex class independently. $\square$

We will need the following simple lemma about removing superfluous vertices later in Section 3.5.

Lemma 3.3.4. *Let B be a bipartite graph with vertex classes $\mathcal{R}$ and W that has a matching saturating $\mathcal{R}$, and let $C \subseteq W$ be the maximal essential subset. If $p \in C$ and $s \in W \setminus C$, then p is essential in B if and only if it is essential in $B - s$.*

Proof. If p is essential in B, then it clearly is essential in $B - s$.

Conversely, assume p is essential in $B - s$. Let $\mathcal{U} \subseteq \mathcal{R}$ be such that $N_B(\mathcal{U}) = C$ and $|\mathcal{U}| = |C|$, which exists by the definition of essential subsets. Since p is essential, there is a unique $R \in \mathcal{R}$ such that $N_{B-s}(R) = \{p\}$. We claim that $R \in \mathcal{U}$. Suppose not. Then $N_B(R) \subseteq \{s, p\}$, and hence $N_B(\mathcal{U} \cup \{R\}) \subseteq C \cup \{s\}$. Since $|\mathcal{U} \cup \{R\}| = |\mathcal{U}| + 1 = |C \cup \{s\}|$, this would make $C \cup \{s\}$ an essential set in B, a contradiction, since C is maximal. Hence $R \in \mathcal{U}$, from which follows that $s \notin N_B(R)$, and thus $N_B(R) = \{p\}$, so p is essential in B. $\square$

3.3.2 The Monster Lemma

Lemma 3.3.5. *Let $\mathcal{H}$ be a 3-partite 3-graph that has a matchable FR-partition $(\mathcal{F}, \mathcal{R}, W)$. Let $a, b, c \in V(\mathcal{H})$ be in different vertex classes. Suppose that the following two conditions hold:*

(1) For every $F \in \mathcal{F}$, there is an F-edge avoiding $\{a, b, c\}$,

(2) For every $R \in \mathcal{R}$, there is an R-edge avoiding $\{a, b, c\}$.

Then $\nu(\mathcal{H} - \{a, b, c\}) = \nu(\mathcal{H})$.

Proof. Let V_1, V_2, and V_3 be the vertex classes of $\mathcal{H}$, where $a \in V_1$, $b \in V_2$, and $c \in V_3$. We will select a matching $\mathcal{M} \subseteq E(\mathcal{H})$ of size $\nu(\mathcal{H})$ avoiding $\{a, b, c\}$.

First, for each $F \in \mathcal{F}$ we choose an arbitrary edge from $\mathcal{H}|_F$ avoiding $\{a, b, c\}$ and include it in $\mathcal{M}$. This can be done by condition (1). These edges are all pairwise disjoint, since the members of $\mathcal{F}$ are pairwise disjoint. Furthermore, we will describe a procedure that selects pairwise disjoint $\mathcal{R}$-edges, one for each $R \in \mathcal{R}$, each containing a W-vertex and avoiding $\{a, b, c\}$. Because they contain a W-vertex, these $\mathcal{R}$-edges will all be disjoint from the $\mathcal{F}$-edges we already put into $\mathcal{M}$ (since both W and $V(\mathcal{R})$ are disjoint from $V(\mathcal{F})$). If successful, we will have constructed the required matching $\mathcal{M}$, since $|\mathcal{M}| = |\mathcal{F}| + |\mathcal{R}| = \nu(\mathcal{H})$.

How we choose the $\mathcal{R}$-edges will fall into several cases. We introduce the following convenient notation for talking about $\mathcal{R}$-edges. An $\mathcal{R}$-edge xyz of $\mathcal{H}$ is called a *WRR-edge* if $x \in W \cap V_1$. Analogously, xyz is called an *RWR-edge* or an *RRW-edge* if $y \in W \cap V_2$ or $z \in W \cap V_3$, respectively.

Case 1. At least one of the vertices a, b, or c is in $V(\mathcal{R})$.

We may assume without loss of generality that $a \in V(\mathcal{R})$. First we choose a matching M_1 saturating $\mathcal{R}$ in the auxiliary bipartite graph B_1. Such a matching exists by the matchability of the FR-partition. Each edge $Rw \in M_1$, with $R \in \mathcal{R}$ and $w \in W \cap V_1$ corresponds to a WRR-edge of $\mathcal{H}$ consisting of w and two vertices of R. These edges form a matching $\mathcal{M}'$ of $\mathcal{R}$-edges in $\mathcal{H}$. Each edge in $\mathcal{M}'$ contains a W-vertex in V_1 and hence avoids $a \in V(\mathcal{R}) \cap V_1$. The only problem might be that b or c appear in some of these edges, rendering those edges unsuitable. If b is contained in the R-edge $e_1 \in \mathcal{M}'$ for some $R \in \mathcal{R}$, then replace e_1 in $\mathcal{M}'$ with an arbitrary RWR-edge e_2 for R. Such an edge exists because B_2 has a matching saturating $\mathcal{R}$, and it is disjoint from all other edges in $\mathcal{M}'$ because these are WRR-edges. The vertex of e_2 in V_1 cannot be a, since then all R-edges would intersect $\{a, b\}$, contradicting condition (2). Similarly, the vertex of e_2 in V_3 cannot be c, since then all R-edges would intersect $\{b, c\}$. Finally, if c is contained in the R'-edge $e_3 \in \mathcal{M}'$ for some $R' \in \mathcal{R}$, then replace e_3 in $\mathcal{M}'$ with an arbitrary RRW-edge e_4 for R'. Such an edge exists because B_3 has a matching saturating $\mathcal{R}$, and it is disjoint from all other edges of $\mathcal{M}'$ because they are all WRR- and RWR-edges. The edge e_4 cannot contain a, otherwise all R'-edges would intersect $\{a, c\}$, contradicting (2). The edge e_4 also does not contain b, since otherwise every R'-edge would intersect $\{b, c\}$, again contradicting (2).

Now the vertices of the matching $\mathcal{M}'$ avoid $\{a, b, c\}$ and Case 1 is complete.

Let us assume from now on that none of the vertices a, b, and c are in $V(\mathcal{R})$.

Case 2. None of the vertices a, b, and c are essential.

First we choose a matching M_1 in B_1 saturating $\mathcal{R}$, which exists by the matchability of the FR-partition. This corresponds to a matching $\mathcal{M}'$ in $\mathcal{H}$ consisting of WRR-edges. Clearly, b and c are avoided by the edges of $\mathcal{M}'$ because $b, c \notin V(\mathcal{R})$. If a is contained in an R-edge $e_1 \in \mathcal{M}'$ for some $R \in \mathcal{R}$, then replace e_1 in $\mathcal{M}'$ by an arbitrary RWR-edge e_2 for R that avoids b. This can be done, since b is not essential. The edge e_2 also avoids a and c because $a, c \notin V(\mathcal{R})$, and it is disjoint from all other edges of $\mathcal{M}'$ because they are all WRR-edges.

Hence we have the required matching $\mathcal{M}'$ avoiding $\{a, b, c\}$ and Case 2 is complete.

Case 3. Not all of the vertices a, b, and c are essential W-vertices for the same $R \in \mathcal{R}$.

We may assume without loss of generality that a is essential for $R \in \mathcal{R}$ (If no vertex is essential, we are in Case 2). By assumption, not both b and c are essential for R as well, so assume without loss of generality that b is not essential for R. We choose a matching $M_1 \subseteq E(B_1)$ saturating $\mathcal{R}$. This corresponds to a matching $\mathcal{M}'$ in $\mathcal{H}$ consisting of WRR-edges. Clearly, b and c are avoided by the edges of $\mathcal{M}'$ because $b, c \notin V(\mathcal{R})$. Since a is essential for R, it must be that $Ra \in M_1$ because a is the only neighbor of R in $W \cap V_1$. Let $e_1 \in \mathcal{M}'$ be the edge corresponding to $Ra \in M_1$. We replace e_1 in $\mathcal{M}'$ by an arbitrary RWR-edge e_2 for R that avoids b. This can be done, since b is not essential for R. The edge e_2 also avoids a and c because $a, c \notin V(\mathcal{R})$, and it is disjoint from all other edges of $\mathcal{M}'$ because they are all WRR-edges.

This means that $\mathcal{M}'$ avoids $\{a, b, c\}$, and so Case 3 is complete.

Case 4. The vertices a, b, and c are all essential W-vertices for $R \in \mathcal{R}$.

By condition (2), there must be an R-edge e avoiding a, b, and c. At least two of its vertices must be in R, so assume without loss of generality that $e \cap V_2, e \cap V_3 \subseteq R$. We choose a matching M_1 in B_1 saturating $\mathcal{R}$. It corresponds to a matching $\mathcal{M}'$ of WRR-edges in $\mathcal{H}$. Because a is essential for R, it follows that there is an edge of $\mathcal{M}'$ containing a and two vertices of R. Replace it by e, which avoids a, b, and c and is disjoint from the other edges of $\mathcal{M}'$ because its V_1-vertex is not in W (because a is the only W-vertex in a WRR-edge of R) and its other vertices are in R. The rest of the edges of $\mathcal{M}'$ clearly avoid a, b, and c, since the one edge of $\mathcal{M}'$ containing a has already been replaced, and $b, c \notin V(\mathcal{R})$.

We must be careful because in this case, one of the edges of $\mathcal{M}'$, namely e, is not necessarily contained in $V(\mathcal{R}) \cup W$, as has been true in all other cases. Thus, the V_1-vertex of e may be in some $F \in \mathcal{F}$, and hence could potentially intersect the F-edge which we added to $\mathcal{M}$ in the beginning. However, since $\mathcal{H}|_F$ is a truncated multi-Fano plane, it cannot be covered by one vertex, so there is an F-edge disjoint from e with which we can replace our original choice of edge for $\mathcal{M}$. Note that we do not need to worry about avoiding $\{a, b, c\}$ with this edge, as these are all in W.

Adding the edges in $\mathcal{M}'$ to $\mathcal{M}$ gives us our desired matching avoiding $\{a, b, c\}$. This concludes Case 4.

These cases exhaust all possibilities, so the proof is complete. □

In order to facilitate the use of this lemma, we prove in some specific cases that the conditions are fulfilled.

Corollary 3.3.6. *Let $\mathcal{H}$ be a 3-partite 3-graph with a matchable FR-partition $(\mathcal{F}, \mathcal{R}, W)$. Let $a, b, c \in V(\mathcal{H})$ be in different vertex classes, and let $S \subseteq W$ be a set of superfluous vertices with at most one vertex in each vertex class. Then in any of the following cases we have $\nu(\mathcal{H} - (\{a, b, c\} \cup S)) = \nu(\mathcal{H})$:*

(1) $a \in V(\mathcal{F})$, $b \in W$, and c is arbitrary,

(2) $a \in R \in \mathcal{R}$, $b \notin R$, and $c \notin V(\mathcal{R})$,

(3) $a \in W$ is essential for $R \in \mathcal{R}$, b is not essential for R in $\mathcal{H} - S$, and $c \notin V(\mathcal{R})$,

(4) $a \in W$ is not essential in $\mathcal{H} - S$, $b \notin V(\mathcal{R})$, and c is arbitrary.

Proof. Let V_1, V_2, and V_3 be the vertex classes of $\mathcal{H}$, where $a \in V_1$, $b \in V_2$, and $c \in V_3$. Let $S' = S \setminus \{a, b, c\}$. By Observation 3.3.3, the hypergraph $\mathcal{H}' = \mathcal{H} - S'$ has the matchable FR-partition $(\mathcal{F}, \mathcal{R}, W \setminus S')$, and hence $\nu(\mathcal{H}') = \nu(\mathcal{H})$. We will apply Lemma 3.3.5 to $\mathcal{H}'$ to find a matching in $\mathcal{H}'$ of size $\nu(\mathcal{H}')$ avoiding $\{a, b, c\}$. This constitutes a matching in $\mathcal{H} - (\{a, b, c\} \cup S)$ of size $\nu(\mathcal{H})$, as desired. We must simply check that the two conditions of Lemma 3.3.5 hold.

Case 1. $a \in V(\mathcal{F})$, $b \in W$, and c is arbitrary.

For any $F \in \mathcal{F}$, there is an F-edge avoiding $\{a, b, c\}$, because $b \in W$, and a and c, being in different vertex classes, do not cover every edge of $\mathcal{H}'|_F$ (a truncated multi-Fano plane).

Let $R = \{r_1, r_2, r_3\} \in \mathcal{R}$ (where $r_i \in V_i$). We will find an R-edge avoiding $\{a, b, c\}$. If $c \in R$, then there is an R-edge avoiding $\{a, b, c\}$ because the matchability of B_3 ensures that there is an R-edge $r_1 r_2 w$ with $w \in W \cap V_3$, which clearly avoids $\{a, b, c\}$, because $a, b \notin V(\mathcal{R})$, and $c \in R$. Suppose $c \notin R$. By the matchability of B_1, there is an R-edge $w' r_2 r_3$, where $w' \in W \cap V_1$, and this edge avoids $\{a, b, c\}$ because $a \in V(\mathcal{F})$, $b \in W$, and $c \notin R$.

Therefore Lemma 3.3.5 applies, and we have $\nu(\mathcal{H}' - \{a, b, c\}) = \nu(\mathcal{H})$.

Case 2. $a \in R \in \mathcal{R}$, $b \notin R$, and $c \notin V(\mathcal{R})$.

For any $F \in \mathcal{F}$, there is an F-edge avoiding $\{a, b, c\}$, because $a \in V(\mathcal{R})$, and b and c do not cover every edge of $\mathcal{H}'|_F$ (a truncated multi-Fano plane).

Let $R' = \{r_1, r_2, r_3\} \in \mathcal{R}$ (where $r_i \in V_i$). We will find an R'-edge avoiding $\{a, b, c\}$. If $b \in R'$, then $R' \neq R$, so $a \notin R'$. There is an R'-edge $r_1 w r_3$ with $w \in W \cap V_2$ by matchability applied to B_2. This edge avoids $\{a, b, c\}$ because $a \notin R'$, $b \in R'$, and $c \notin V(\mathcal{R})$. Suppose $b \notin R'$. By the matchability of B_1, there is an R'-edge $w' r_2 r_3$, where $w' \in W \cap V_1$, and this edge avoids $\{a, b, c\}$ because $a \in V(\mathcal{R})$, $b \notin R'$, and $c \notin V(\mathcal{R})$.

Therefore Lemma 3.3.5 applies, and we have $\nu(\mathcal{H}' - \{a, b, c\}) = \nu(\mathcal{H})$.

Case 3. $a \in W$ is essential for $R \in \mathcal{R}$, b is not essential for R in $\mathcal{H} - S$, and $c \notin V(\mathcal{R})$.

Note that if a is essential for R in $\mathcal{H}$, then it is still essential for R in $\mathcal{H}'$, a subgraph of $\mathcal{H}$. Similarly, if b is not essential for R in $\mathcal{H} - S$, then it certainly is not essential for R in $\mathcal{H}'$, since $\mathcal{H} - S$ is a subhypergraph of $\mathcal{H}'$.

For any $F \in \mathcal{F}$, there is an F-edge avoiding $\{a, b, c\}$, because $a \in W$, and b and c do not cover every edge of $\mathcal{H}'|_F$ (a truncated multi-Fano plane).

Let $R' = \{r_1, r_2, r_3\} \in \mathcal{R}$ (where $r_i \in V_i$). We will find an R'-edge avoiding $\{a, b, c\}$. If b is not essential for R', then R' has a neighbor $w \in W \cap V_1$ in B_2 with $w \neq b$. The R' edge $r_1 w r_3$ then avoids $\{a, b, c\}$ because $a \in W$, $b \neq w$, and $c \notin V(\mathcal{R})$. If b is essential for R', then $b \in W$ and $R' \neq R$, so a is not essential for R' (because no vertex can be essential for two different members of $\mathcal{R}$ by matchability). Thus R' has a neighbor $w' \in W \cap V_1$ in B_1 with $w' \neq a$. The R'-edge $w' r_2 r_3$ then avoids $\{a, b, c\}$ because $w' \neq a$ and $b, c \notin V(\mathcal{R})$.

Therefore Lemma 3.3.5 applies, and we have $\nu(\mathcal{H}' - \{a, b, c\}) = \nu(\mathcal{H})$.

Case 4. $a \in W$ is not essential in $\mathcal{H} - S$, $b \notin V(\mathcal{R})$, and c is arbitrary.

Note that if a is not essential in $\mathcal{H} - S$, then it certainly is not essential in $\mathcal{H}'$, since $\mathcal{H} - S$ is a subhypergraph of $\mathcal{H}'$.

For any $F \in \mathcal{F}$, there is an F-edge avoiding $\{a, b, c\}$, because $a \in W$, and b and c do not cover every edge of $\mathcal{H}'|_F$ (a truncated multi-Fano plane).

Let $R = \{r_1, r_2, r_3\} \in \mathcal{R}$ (where $r_i \in V_i$). We will find an R-edge avoiding $\{a, b, c\}$. If $c \in R$, then there is an R-edge avoiding $\{a, b, c\}$ because the matchability of B_3 ensures that there is an R-edge $r_1 r_2 w$ with $w \in W \cap V_3$, which clearly avoids $\{a, b, c\}$, since $a, b \notin V(\mathcal{R})$, and $c \in R$. Suppose $c \notin R$. Since a

is not essential, R has a neighbor $w' \in W \cap V_1$ in B_1 with $w' \neq a$. The R-edge $w'r_2r_3$ then avoids $\{a, b, c\}$ because $w' \neq a$, $b \notin V(\mathcal{R})$, and $c \notin R$.

Therefore Lemma 3.3.5 applies, and we have $\nu(\mathcal{H}' - \{a, b, c\}) = \nu(\mathcal{H})$. □

It is unfortunately necessary in Cases 3 and 4 to make sure that the non-essential W-vertex remains non-essential after removing the superfluous vertices. However, this condition is often very easy to check, since removing superfluous vertices from the hypergraph only affects the status of those W-vertices in their vertex class. This leads to the following observation:

Observation 3.3.7. *Let $\mathcal{H}$ be a 3-partite 3-graph with a matchable FR-partition $(\mathcal{F}, \mathcal{R}, W)$, and let $s \in W$ be a superfluous vertex. Then if $w \in W$ is in a different vertex class from s, it holds that w is non-essential in $\mathcal{H}$ if and only if it is non-essential in $\mathcal{H} - s$.*

3.3.3 Matchability and the Edge-Home Property

One nice consequence of the monster lemma is the following proposition, which will be key to our proof.

Definition 3.3.8. An FR-partition $(\mathcal{F}, \mathcal{R}, W)$ is *proper* if there is no $R \in \mathcal{R}$ and an edge of $\mathcal{H}$ consisting of three vertices of W which together induce a truncated Fano plane. Being proper just means that we have not called anything an R if it could have been part of an F.

Clearly home-base partitions are proper, because they do not contain any edges consisting of W-vertices. It turns out that a converse to this fact is also true.

Proposition 3.3.9. *A proper matchable FR-partition of a 3-partite 3-graph has the edge-home property.*

Proof. Let $\mathcal{H}$ be a 3-partite 3-graph with vertex classes V_1, V_2, V_3, and let $(\mathcal{F}, \mathcal{R}, W)$ be a proper matchable FR-partition of $\mathcal{H}$. Let abc be an edge of $\mathcal{H}$. We aim to show that it is either an $\mathcal{F}$-edge or an $\mathcal{R}$-edge. Suppose it is not. We will aim for a contradiction by applying Lemma 3.3.5 to show $\mathcal{H} - \{a, b, c\}$ has a matching of size $\nu(\mathcal{H})$.

By assumption, abc is not in $\mathcal{H}|_F$ for any $F \in \mathcal{F}$, which means that every $F \in \mathcal{F}$ has an F-edge avoiding $\{a, b, c\}$, since the only way to cover a truncated Fano plane with vertices from different vertex classes is if they form one of its edges. We want to show that it also cannot cover every R-edge for any $R \in \mathcal{R}$.

Since the partition is matchable, each of the auxiliary bipartite graphs B_1, B_2, and B_3 have matchings saturating $\mathcal{R}$, say M_1, M_2, and M_3, respectively. Then each $R = \{r_1, r_2, r_3\} \in \mathcal{R}$ has three W-vertices, $w_i^R \in V_i$ assigned to it, so that $Rw_i^R \in M_i$, which means that $w_i^R r_j r_k$ are edges for each choice of $\{i, j, k\} = \{1, 2, 3\}$. By assumption, abc intersects R in at most one vertex (otherwise, it is an R-edge). If abc intersects R in one vertex, without loss of generality in V_1, then $w_1^R r_2 r_3$ is an R-edge disjoint from abc. If abc does not intersect R in any vertex, then it intersects all the R-edges $w_i^R r_j r_k$ for $\{i,j,k\} = \{1,2,3\}$ only if $abc = w_1^R w_2^R w_3^R$, which would mean that abc, $w_1^R r_2 r_3$, $r_1 w_2^R r_3$, and $r_1 r_2 w_3^R$ form a truncated Fano plane. If this is the case, then we claim that these are in fact the only edges on $\{a,b,c,r_1,r_2,r_3\}$, which would contradict the assumption that $(\mathcal{F},\mathcal{R},W)$ is proper.

Suppose these are not the only edges on $\{a,b,c,r_1,r_2,r_3\}$. Then there are two disjoint edges on $\{a,b,c,r_1,r_2,r_3\}$. Now pick one F-edge for each $F \in \mathcal{F}$, and take the edges $w_1^{R'} r_2' r_3'$ for each $R' \in \mathcal{R} \setminus \{R\}$. These edges form a matching of size $|\mathcal{F}| + |\mathcal{R}| - 1$, and they do not intersect $\{a,b,c,r_1,r_2,r_3\}$. Together with the two disjoint edges on $\{a,b,c,r_1,r_2,r_3\}$, we find a matching of size $|\mathcal{F}| + |\mathcal{R}| + 1 = \nu(\mathcal{H}) + 1$, a contradiction.

Hence a, b, and c fulfill the conditions of Lemma 3.3.5, and $\mathcal{H} \setminus \{a,b,c\}$ would have a matching of size $\nu(\mathcal{H})$, which together with abc would be a matching of size $\nu(\mathcal{H})+1$ in $\mathcal{H}$, a contradiction. Therefore $\mathcal{H}$ has the edge-home property. □

3.4 Cromulent Triples

The aim of this section is to define the appropriate substructure which will facilitate the inductive proof of our main theorem (Theorem 1.1.2). The key definition is that of a cromulent triple.

Definition 3.4.1. Let $\mathcal{H}$ be a 3-partite 3-graph with vertex classes V_1, V_2, and V_3. A triple of nonempty sets (Y_1, Y_2, X) with $Y_1 \subseteq V_i$, $Y_2 \subseteq V_j$ and $X \subseteq V_k$, where $\{i,j,k\} = \{1,2,3\}$ is called a *cromulent* triple if it fulfills the following conditions:

(1) $|Y_1| = |Y_2| \leq |X|$,

(2) $N_{\mathrm{lk}_{\mathcal{H}}(V_i)}(X) = Y_2$,

(3) There is a hypergraph matching in $\mathcal{H}|_{Y_1 \cup Y_2 \cup X}$ of size $|Y_1|$,

(4) The hypergraph $\mathcal{H}_0 = \mathcal{H} - (Y_1 \cup Y_2 \cup X)$ is a home-base hypergraph with $\nu(\mathcal{H}_0) = \nu(\mathcal{H}) - |Y_1|$,

(5) Given any home-base partition $(\mathcal{F},\mathcal{R},W)$ of $\mathcal{H}_0$, we have $N_{\mathrm{lk}_{\mathcal{H}}(V_j)}(X) \subseteq Y_1 \cup V(\mathcal{R}) \cup V(\mathcal{F})$.

Such a triple is called *perfectly cromulent* if it fulfills the following stronger version of condition (5):

(5*) $N_{\mathrm{lk}_{\mathcal{H}}(V_j)}(X) = Y_1$.

The first lemma of this section states that perfectly cromulent triples are the kind of substructure we should look for in order to prove our main theorem.

Lemma 3.4.2. *Let $\mathcal{H}$ be a 3-partite 3-graph with $\tau(\mathcal{H}) = 2\nu(\mathcal{H})$. If $\mathcal{H}$ has a perfectly cromulent triple, then $\mathcal{H}$ is a home-base hypergraph.*

Unfortunately, it is sometimes hard to ensure property (5*), and it will be easier to find just cromulent triples instead. Fortunately, we will be able to prove that this suffices.

Lemma 3.4.3. *If $\mathcal{H}$ is a 3-partite 3-graph with $\tau(\mathcal{H}) = 2\nu(\mathcal{H})$, then every cromulent triple of $\mathcal{H}$ is perfectly cromulent.*

These two lemmas combine to give the main result of this section as an immediate corollary:

Corollary 3.4.4. *Let $\mathcal{H}$ be a 3-partite 3-graph with $\tau(\mathcal{H}) = 2\nu(\mathcal{H})$. If $\mathcal{H}$ has a cromulent triple, then $\mathcal{H}$ is a home-base hypergraph.*

The proofs of the two lemmas follow similar lines, and so they will be handled in parallel. The basic idea is outlined below. We start with Lemma 3.4.2.

Let (Y_1, Y_2, X) be a perfectly cromulent triple, and let $\mathcal{H}_0 = \mathcal{H} - (Y_1 \cup Y_2 \cup X)$ be the hypergraph from the definition of cromulent triples. Let $(\mathcal{F}, \mathcal{R}, W)$ be a home-base partition of $\mathcal{H}_0$. Our goal will be to extend this partition into a home-base partition $(\mathcal{F}', \mathcal{R}', W')$ of $\mathcal{H}$. Fix a maximum hypergraph matching $\mathcal{M}$ in $\mathcal{H}|_{Y_1 \cup Y_2 \cup X}$. Each pair $y \in Y_1$, $y' \in Y_2$ that are together in an edge of $\mathcal{M}$ will participate in a new $R \in \mathcal{R}'$ together with a uniquely determined member of $W \cap V_3$. The vertices in X will be vertices of W', and by virtue of the matching saturating Y_1 and Y_2, they will ensure a matching saturating $\mathcal{R}'$ exists in the bipartite graph B'_3. The rest of the section will be devoted to finding the member of $W \cap V_3$ we can include in our new R's and proving that the resulting partition $(\mathcal{F}', \mathcal{R}', W')$ is indeed a home-base partition. Our fundamental tool in this proof will be Corollary 3.3.6, and we will finish by using Proposition 3.3.9.

If (Y_1, Y_2, X) was simply a cromulent triple, then much of the same proof as above still goes through in a more restricted form, and eventually we will be able to find a contradiction if (Y_1, Y_2, X) violated condition (5*), which will show Lemma 3.4.3.

We first introduce a notion which will be helpful for our upcoming proofs.

3.4.1 Heavy Vertex Covers

Recall the definition of essential subsets and superfluous vertices from Section 3.3.

The following is a particular type of vertex cover for home-base hypergraphs, which will be useful for the proofs in this and the next section.

Definition 3.4.5. Let $\mathcal{H}$ be a home-base hypergraph on vertex classes V_1, V_2, and V_3 with a home-base partition $(\mathcal{F}, \mathcal{R}, W)$, and let $i, j \in \{1, 2, 3\}$ with $i \neq j$. Let $C_i \subseteq W \cap V_i$ be the maximal essential set in B_i and let $\mathcal{U}_i \subseteq \mathcal{R}$ be the set with $|\mathcal{U}_i| = |C_i|$ and $N_{B_i}(\mathcal{U}_i) = C_i$. Then the union of the sets

- $C_i \cup ((V(\mathcal{F}) \cup V(\mathcal{R})) \cap V_i)$
- $\left(\bigcup_{R \in \mathcal{R} \setminus \mathcal{U}_i} R\right) \cap V_j$

is called the *i-heavy (i, j)-cover* of $\mathcal{H}$.

Observation 3.4.6. *Every vertex in V_i which is not in the i-heavy (i, j)-cover is a superfluous vertex in $W \cap V_i$.*

Proposition 3.4.7. *If $\mathcal{H}$ is a home-base hypergraph on vertex classes V_1, V_2, and V_3 with a home-base partition $(\mathcal{F}, \mathcal{R}, W)$, then for every pair $i, j \in \{1, 2, 3\}$ with $i \neq j$, the i-heavy (i, j)-cover is a minimal vertex cover of $\mathcal{H}_0$.*

Proof. Let T be the i-heavy (i, j)-cover of $\mathcal{H}$. Let $e \in E(\mathcal{H})$. Then by the edge-home property, e is at home in some $F \in \mathcal{F}$ or some $R \in \mathcal{R}$. If it is at home in F, then it contains some vertex in $F \cap V_i$, and so it intersects T. If it is at home in $R \in \mathcal{R} \setminus \mathcal{U}_i$, then it contains some vertex in $R \cap (V_i \cup V_j)$, and hence intersects T. The only remaining case is that e is at home in some $R' \in \mathcal{U}_i$. Let $V_i \cap e = \{v\}$. If $v \in V(\mathcal{F}) \cup V(\mathcal{R})$, then e intersects T. If $v \in W \cap V_i$, then vR' is an edge of B_i, and hence $v \in N_{B_i}(\mathcal{U}_i) = C_i$, which shows that e again intersects T. Thus T is a vertex cover of $\mathcal{H}$.

We now calculate the size of T. By the definition of the i-heavy (i, j)-cover, we get $|T| = 2\,|\mathcal{F}| + |\mathcal{R}| + |C_i| + |\mathcal{R}| - |\mathcal{U}_i|$. Since $|C_i| = |\mathcal{U}_i|$, we get $|T| = 2\,|\mathcal{F}| + 2\,|\mathcal{R}| = 2\,|\mathcal{F} \cup \mathcal{R}| = 2\nu(\mathcal{H})$, and because home-base hypergraphs are tight for Ryser's Conjecture by Proposition 3.1.5, we get $|T| = \tau(\mathcal{H})$ as desired. $\square$

3.4.2 Facts About Cromulent Triples

We start with some lemmas about cromulent and perfectly cromulent triples. Note that properties (2) and (5*) make the roles of Y_1 and Y_2 symmetric in perfectly cromulent triples. This gives us the following observation:

Observation 3.4.8. *(Y_1, Y_2, X) is a perfectly cromulent triple if and only if (Y_1, Y_2, X) and (Y_2, Y_1, X) are both cromulent triples.*

Most of the proofs in this section work for cromulent triples, and can be strengthened for perfectly cromulent triples by using Observation 3.4.8.

Assumptions. For the rest of this section, let $\mathcal{H}$ be a 3-partite 3-uniform hypergraph with vertex classes V_1, V_2, and V_3 such that $\tau(\mathcal{H}) = 2\nu(\mathcal{H})$, and assume it has a cromulent triple (Y_1, Y_2, X). We will assume without loss of generality that $Y_1 \subseteq V_1$, $Y_2 \subseteq V_2$, and $X \subseteq V_3$. We also fix a hypergraph matching $M \subseteq E(\mathcal{H}|_{Y_1 \cup Y_2 \cup X})$ of size $|Y_1|$. Let $\mathcal{H}_0 = \mathcal{H} - (Y_1 \cup Y_2 \cup X)$ be the corresponding home-base hypergraph, and fix a home-base partition $(\mathcal{F}, \mathcal{R}, W)$ of $\mathcal{H}_0$.

Lemma 3.4.9. *For every pair $(i, j) \in \{(1, 2), (1, 3), (2, 1)\}$ we have that for every $y \in Y_i$ there is an edge ywu, where $w \in W \cap V_j$, and $u \in V(\mathcal{H}_0) \setminus V(\mathcal{R})$. If (Y_1, Y_2, X) is perfectly cromulent, then this holds also for $(i, j) = (2, 3)$.*

Proof. We will construct a vertex set T of size $\tau(\mathcal{H}) - 1$ which intersects all edges of $\mathcal{H}$ except for the edges of the form in question. Since T cannot be a vertex cover by virtue of its small size, some such edge must exist. Let T be the union of the sets $Y_1 \cup Y_2 \setminus \{y\}$, $(V(\mathcal{F}) \cup V(\mathcal{R})) \cap V_j$, and $V(\mathcal{R}) \cap V_k$, where $k \in \{1,2,3\} \setminus \{i,j\}$. Since we have taken two vertices from each $F \in \mathcal{F}$ and two vertices from each $R \in \mathcal{R}$, and $2|Y_1| - 1$ additional vertices, we get $|T| = 2|\mathcal{F} \cup \mathcal{R}| + 2|Y_1| - 1 = 2\nu(\mathcal{H}_0) + 2|Y_1| - 1 = 2\nu(\mathcal{H}) - 1 = \tau(\mathcal{H}) - 1$, hence T is not a vertex cover of $\mathcal{H}$.

It is clear that T includes a cover of all edges of $\mathcal{H}_0$, so any uncovered edge must contain y or intersect X. It turns out that any edge e intersecting X is also covered by T. If $i = 1$, then e is covered by $N_{\mathrm{lk}_{\mathcal{H}}(V_1)}(X) = Y_2 \subseteq T$. If $i = 2$, then $j = 1$ and e is covered by $N_{\mathrm{lk}_{\mathcal{H}}(V_2)}(X) \subseteq Y_1 \cup (V(\mathcal{F}) \cup V(\mathcal{R})) \cap V_1 \subseteq T$. Therefore, any edge not covered by T must contain y and two vertices of $\mathcal{H}_0$. The V_j-vertex must be a W-vertex because $(V(\mathcal{F}) \cup V(\mathcal{R})) \cap V_j \subseteq T$, and the V_k-vertex cannot be in $V(\mathcal{R})$ because $V(\mathcal{R}) \cap V_k \subseteq T$. □

Lemma 3.4.10. *For every pair* $(i,j) \in \{(1,2),(1,3),(2,1)\}$ *we have that for every* $y \in Y_i$ *there is an edge* ysu*, where* $s \in W \cap V_j$ *is superfluous, and* $u \in V(\mathcal{H}_0)$*. If* (Y_1, Y_2, X) *is perfectly cromulent, then this holds also for* $(i,j) = (2,3)$*.*

Proof. We will construct a vertex set T of size $\tau(\mathcal{H}) - 1$ which intersects all edges of $\mathcal{H}$ except for the edges of the form in question. Since T cannot be a vertex cover by virtue of its small size, some such edge must exist. Let T be the union of $Y_1 \cup Y_2 \setminus \{y\}$ and the j-heavy (j,i)-cover of $\mathcal{H}_0$. Since we have taken $\tau(\mathcal{H}_0)$ vertices from $\mathcal{H}_0$ and $2|Y_1| - 1$ additional vertices, we get $|T| = 2|\mathcal{F} \cup \mathcal{R}| + 2|Y_1| - 1 = \tau(\mathcal{H}) - 1$ (as calculated before). As in the proof of Lemma 3.4.9, the V_i-vertex of any uncovered edge must be y, and the other vertices are in $V(\mathcal{H}_0)$. The V_j-vertex of an uncovered edge must be a superfluous vertex because besides $(V(\mathcal{F}) \cup V(\mathcal{R})) \cap V_j$, the maximal essential subset $C_j \subseteq W \cap V_j$ of B_j is also included in T (and every W-vertex outside of the maximal essential subset is by definition superfluous). □

Lemma 3.4.11. *For* $i = 1$ *and* $j = 3$ *we have that for every* $y \in Y_i$*, if* yvs *is an edge of* $\mathcal{H}$ *with* $v \in V(\mathcal{H}_0)$ *and* $s \in V_j$ *a superfluous vertex, then there is an edge* $yv's$ *with* $v' \in V(\mathcal{H}_0) \setminus V(\mathcal{R})$*. If* (Y_1, Y_2, X) *is perfectly cromulent, then this holds also for* $(i,j) = (2,3)$*.*

Proof. We may assume $v \in V(\mathcal{R})$, otherwise we are done. Let $y' \in Y_2$ be the V_2-vertex of the edge of $\mathcal{M}$ containing y.

By Lemma 3.4.9 (with $(i,j) = (2,1)$ for $y' \in Y_2$), there is an edge $wy'u$ with $w \in W \cap V_1$ and $u \in V(\mathcal{H}_0) \setminus V(\mathcal{R})$. We claim $s = u$.

Suppose not. Then yvs and $wy'u$ are disjoint edges. We can apply Case (2) of Corollary 3.3.6 with $a = v$, $b = w$, $c = u$, and $S = \{s\}$ to find a matching of size $\nu(\mathcal{H}_0)$ in $\mathcal{H}_0 - \{s,u,v,w\}$. This matching together with the edges yvs, $wy'u$, and the rest of $\mathcal{M}$ (besides the edge containing y and y') forms a matching of size $\nu(\mathcal{H}_0) + 2 + |Y_i| - 1 = \nu(\mathcal{H}) + 1$, a contradiction. Hence $s = u$.

By Lemma 3.4.10 (with $(i,j) = (1,2)$ for $y \in Y_1$), there is an edge $yv'u'$ with v' a superfluous vertex in $W \cap V_2$. If $u' \neq s$, then $yv'u'$ and $wy's$ are disjoint edges. We can apply Case (4) of Corollary 3.3.6 with $a = v'$, $b = w$, $c = u'$, and $S = \{s\}$ to find a matching of size $\nu(\mathcal{H}_0)$ in $\mathcal{H}_0 - \{s, v', w, u'\}$. This matching together with the edges $yv'u'$, $wy's$, and the rest of $\mathcal{M}$ (besides the edge containing y and y') forms a matching of size $\nu(\mathcal{H}_0)+2+|Y_i|-1 = \nu(\mathcal{H})+1$, a contradiction.

Therefore $u' = s$, and thus $yv's$ is the edge we are looking for. □

The next lemma is a strengthening of Lemma 3.4.11 in two ways: we can require more of our third vertex, and we can apply it to more combinations of i and j.

Lemma 3.4.12. *For $i = 1$ and for every $j \in \{2, 3\}$ we have that for every $y \in Y_i$, if yvs is an edge of $\mathcal{H}$ with $v \in V(\mathcal{H}_0)$ and $s \in V_j$ a superfluous vertex, then there is an edge $ys's$ with s' also superfluous. If (Y_1, Y_2, X) is perfectly cromulent, then this holds also for $i = 2$ and $j \in \{1, 3\}$.*

Proof. Let yvs be an edge with $v \in V(\mathcal{H}_0)$ and $s \in V_j$ superfluous. Let $y' \in Y_2$ be the V_2-vertex of the edge of $\mathcal{M}$ containing y. There are two cases.

Case 1. $i = 1$, $j = 3$.

By Lemma 3.4.11 (with $(i, j) = (1, 3)$), we may assume $v \in V(\mathcal{H}_0) \setminus V(\mathcal{R})$. By Lemma 3.4.10 (with $(i, j) = (2, 1)$ for $y' \in Y_2$), there is an edge $s''y'u$ with $s'' \in V_i$ a superfluous vertex. If $s \neq u$, then yvs and $s''y'u$ are disjoint edges, and we will reach a contradiction as in the previous lemma. We can apply Case (4) of Corollary 3.3.6 with $a = s''$, $b = v$, $c = u$, and $S = \{s\}$ to find a matching of size $\nu(\mathcal{H}_0)$ in $\mathcal{H}_0 - \{s, s'', u, v\}$. This matching together with the edges yvs, $s''y'u$, and the rest of $\mathcal{M}$ (besides the edge containing y and y') forms a matching of size $\nu(\mathcal{H}_0) + 2 + |Y_i| - 1 = \nu(\mathcal{H}) + 1$, a contradiction.

It follows that $s = u$. Lemma 3.4.10 (with $(i, j) = (1, 2)$ for $y \in Y_1$) tells us that there is an edge $ys'u'$ with $s' \in V_2$ superfluous. It must be the case that $s = u'$ because otherwise $ys'u'$ and $s''y's$ are disjoint edges, and we would reach a similar contradiction. We can apply Case (4) of Corollary 3.3.6 with $a = s''$, $b = s'$, $c = u'$, and $S = \{s\}$ to find a matching of size $\nu(\mathcal{H}_0)$ in $\mathcal{H}_0 - \{s, s', s'', u'\}$. This matching together with the edges $ys'u'$, $s''y's$, and the rest of $\mathcal{M}$ (besides the edge containing y and y') forms a matching of size $\nu(\mathcal{H}_0) + 2 + |Y_i| - 1 = \nu(\mathcal{H}) + 1$, a contradiction.

Therefore there is an edge $ys's$, as required.

Case 2. $i = 1$, $j = 2$.

By Lemma 3.4.10) (with $(i, j) = (1, 3)$ for $y \in Y_1$) there is an edge $yr's'$ with $s' \in V_3$ superfluous, and then by Case 1, above, there is an edge yrs' with $r \in V_2$ and $s' \in V_3$ both superfluous. By Lemma 3.4.10 (with $(i, j) = (2, 1)$ for $y' \in Y_2$), there is an edge $qy'u$ with $q \in V_1$ a superfluous vertex and $u \in V(\mathcal{H}_0)$. If $u \neq s'$, then we will again reach a contradiction. Suppose yrs' and $qy'u$ are disjoint. We can apply Case (4) of Corollary 3.3.6 with $a = q$, $b = r$, $c = u$, and $S = \{s'\}$ to find a matching of size $\nu(\mathcal{H}_0)$ in $\mathcal{H}_0 - \{q, r, s', u\}$. This matching together with the edges yrs', $qy'u$, and the rest of $\mathcal{M}$ (besides the edge containing y and y') forms a matching of size $\nu(\mathcal{H}_0) + 2 + |Y_i| - 1 = \nu(\mathcal{H}) + 1$, a contradiction.

Therefore $u = s'$. A similar contradiction is reached by ysv and $qy's'$ if $v \neq s'$, so that cannot be the case either. Suppose ysv and $qy's'$ are disjoint. We can apply Case (4) of Corollary 3.3.6 with $a = q$, $b = s$, $c = v$, and $S = \{s'\}$ to find a matching of size $\nu(\mathcal{H}_0)$ in $\mathcal{H}_0 - \{q, s, s', v\}$. This matching together with the edges ysv, $qy's'$, and the rest of $\mathcal{M}$ (besides the edge containing y and y') forms a matching of size $\nu(\mathcal{H}_0) + 2 + |Y_i| - 1 = \nu(\mathcal{H}) + 1$, a contradiction.

Therefore we have found our edge yss'. □

Lemma 3.4.13. *Let $y \in Y_1$ and $y' \in Y_2$ be in an edge of $\mathcal{M}$ together. Then there is a unique superfluous vertex $z_{y,y'} \in V_3$ such that*

(i) There are edges $yvz_{y,y'}$ and $uy'z_{y,y'}$ for some vertices $u, v \in V(\mathcal{H}_0)$,

(ii) If $yv's'$ or $u'y's'$ is an edge with s' superfluous, then $s' = z_{y,y'}$.

Proof. By Lemma 3.4.10 (with $(i, j) = (1, 3)$ for $y \in Y_1$) there is an edge yvs with $v \in V(\mathcal{H}_0)$ and $s \in V_3$ superfluous. We claim that s satisfies (i) and (ii).

To see (i), we only need to find $uy's$, since we have yvs. By Lemma 3.4.12 (with $(i, j) = (1, 2)$), we may assume v is superfluous as well. By Lemma 3.4.10 (with $(i, j) = (2, 1)$ for $y' \in Y_2$), we have an edge $s'y'u'$ with $s' \in W \cap V_1$ superfluous. Suppose $u' \neq s$. Then yvs and $s'y'u'$ are disjoint edges. We can apply Case (4) of Corollary 3.3.6 with $a = v$, $b = s'$, $c = u'$, and $S = \{s\}$ to find a matching of size $\nu(\mathcal{H}_0)$ in $\mathcal{H}_0 - \{s, s', u', v\}$. This matching together with the edges yvs, $s'y'u'$, and the rest of $\mathcal{M}$ (besides the edge containing y and y') forms a matching of size $\nu(\mathcal{H}_0) + 2 + |Y_i| - 1 = \nu(\mathcal{H}) + 1$, a contradiction.

Therefore $u' = s$, and we have the desired edge $s'y's$.

We now show (ii). Let $yv's'$ and $u'y's''$ be edges of $\mathcal{H}$ with $s', s'' \in V_3$ both superfluous vertices. By Lemma 3.4.12 (with $(i, j) = (1, 2)$), we may assume v' is superfluous as well. If $s' \neq s''$, then $yv's'$ and $u'y's''$ are disjoint edges. This leads to a contradiction as before. We can apply Case (4) of Corollary 3.3.6 with $a = v'$, $b = s'$, $c = u'$, and $S = \{s''\}$ to find a matching of size $\nu(\mathcal{H}_0)$ in $\mathcal{H}_0 - \{s', s'', u', v'\}$. This matching together with the edges $yv's'$, $u'y's''$, and the rest of $\mathcal{M}$ (besides the edge containing y and y') forms a matching of size $\nu(\mathcal{H}_0) + 2 + |Y_i| - 1 = \nu(\mathcal{H}) + 1$, a contradiction.

Therefore it must be the case that $s' = s''$, which in particular means that $s' = s'' = s$, since we could have substituted yvs or $uy's$ for $yv's'$ or $u'y's''$, respectively. □

Our aim is to make each set $\{y, y', z_{y,y'}\}$ into an R for our home-base partition. We will first show that the $z_{y,y'}$'s are all distinct, and then we will make use of Lemma 3.3.9 to show that combining the new R's with the home-base partition of $\mathcal{H}_0$ forms a home-base partition of $\mathcal{H}$.

Lemma 3.4.14. *For each (y, y')-pair, the associated $z_{y,y'}$ is distinct, and there is a matching saturating $\mathcal{R}$ in the subgraph of B_3 induced by $\mathcal{R} \cup (V_3 \cap W \setminus Z)$, where Z is the set of all $z_{y,y'}$'s.*

Proof. Define the bipartite graph K with parts $\mathcal{R} \cup Y_1$ and $W \cap V_3$, where there is an edge between $R \in \mathcal{R}$ and $w \in W \cap V_3$ precisely when there is an R-edge containing w, and there is an edge between $y \in Y_1$ and $w \in W \cap V_3$ precisely when $w = z_{y,y'}$, where y' is the partner of y in the pairing between Y_1 and Y_2. We claim that K has a matching saturating $\mathcal{R} \cup Y_1$.

We will apply Hall's theorem, so let $\mathcal{R}_0 \subseteq \mathcal{R}$ and $Y_0 \subseteq Y_1$. We construct a vertex cover T of $\mathcal{H}$. Let C_3 be the maximal essential set in the subgraph of K induced by $\mathcal{R}$ and $W \cap V_3$ (this is the graph B_3 associated with $\mathcal{H}_0$), and let $\mathcal{U}_3 \subseteq \mathcal{R}$ be such that $N_K(\mathcal{U}_3) = C_3$, which exists by the definition of essential. Let T be the union of the sets $(Y_1 \cup Y_2) \setminus Y_0$, $N_K(\mathcal{R}_0 \cup Y_0)$, $(V(\mathcal{R}) \cup V(\mathcal{F})) \cap V_3$, C_3, and $\bigcup_{R \in \mathcal{R} \setminus (\mathcal{U}_3 \cup \mathcal{R}_0)} (R \cap V_1)$. Note the similarities to the 3-heavy $(3,1)$-cover of $\mathcal{H}_0$.

We must show that T is indeed a vertex cover. Let $e \in E(\mathcal{H}_0)$. Then e is either an $\mathcal{F}$-edge or an $\mathcal{R}$-edge. If it is an $\mathcal{F}$-edge, it is covered by $V(\mathcal{F}) \cap V_3 \subseteq T$. If it is an $\mathcal{R}$-edge, then it is covered by $(V(\mathcal{F}) \cup V(\mathcal{R}) \cap V_3 \subseteq T$, unless its V_3-vertex is in W, so assume that is the case. Let e be an R-edge. If $R \in \mathcal{R}_0$, then $e \cap V_3 \in N_K(R) \subseteq T$. If $R \in \mathcal{U}_3$, then $e \cap V_3 \in C_3 \subseteq T$. If $R \in \mathcal{R} \setminus (\mathcal{U}_3 \cup \mathcal{R}_0)$, then $e \cap V_1 = R \cap V_1 \subseteq T$. This shows that T covers every edge of $\mathcal{H}_0$. All edges incident to X intersect Y_2, so any uncovered edge must be incident to Y_0 and two vertices of $\mathcal{H}_0$. All such edges whose V_3-vertex is not superfluous intersect T, since $C_3 \cup (V(\mathcal{R}) \cup V(\mathcal{F})) \cap V_3 \subseteq T$. Thus, the only edges we have to worry about are those incident to some $y \in Y_0$ and a superfluous vertex in V_3. Then by Lemma 3.4.13, the V_3-vertices of those edges are the corresponding $z_{y,y'}$, and hence those edges intersect $N_K(Y_0) \subseteq T$. This shows that T is a vertex cover.

We now calculate the size of T. By the definition of T, we calculate $|T| = |Y_1| + |Y_2| - |Y_0| + |N_K(\mathcal{R}_0 \cup Y_0)| + 2|\mathcal{F}| + |\mathcal{R}| + |C_3| - |C_3 \cap N_K(\mathcal{R}_0)| + |\mathcal{R}| - |\mathcal{U}_3 \cup \mathcal{R}_0|$. Because it is a vertex cover, we must have $|T| \geq \tau(\mathcal{H})$. Since $\nu(\mathcal{H}) = \nu(\mathcal{H}_0) + |Y_1|$ by the definition of cromulent triple, and since $\tau(\mathcal{H}) = 2\nu(\mathcal{H})$, we have $\tau(\mathcal{H}) = 2\nu(\mathcal{H}_0) + 2|Y_1| = 2|\mathcal{F} \cup \mathcal{R}| + |Y_1| + |Y_2|$. Combining this with the fact that $\tau(\mathcal{H}) \leq |T|$ yields the inequality $|Y_0| + |\mathcal{U}_3 \cup \mathcal{R}_0| + |C_3 \cap N_K(\mathcal{R}_0)| \leq |N_K(\mathcal{R}_0 \cup Y_0)| + |C_3|$. By the inclusion-exclusion principle we can rewrite this as $|Y_0| + |\mathcal{U}_3| + |\mathcal{R}_0| - |\mathcal{U}_3 \cap \mathcal{R}_0| + |C_3 \cap N_K(\mathcal{R}_0)| \leq |N_K(\mathcal{R}_0 \cup Y_0)| + |C_3|$. Since $C_3 = N_K(\mathcal{U}_3)$, we clearly have $C_3 \cap N_K(\mathcal{R}_0) \supseteq N_K(\mathcal{U}_3 \cap \mathcal{R}_0)$. Since B_3 has a matching saturating $\mathcal{R}$, by Hall's Theorem, we must have $|\mathcal{U}_3 \cap \mathcal{R}_0| \leq |N_K(\mathcal{U}_3 \cap \mathcal{R}_0)|$. Combining this with our previous inequality, we then get $|Y_0| + |\mathcal{U}_3| + |\mathcal{R}_0| - |\mathcal{U}_3 \cap \mathcal{R}_0| + |\mathcal{U}_3 \cap \mathcal{R}_0| \leq |N_K(\mathcal{R}_0 \cup Y_0)| + |C_3|$, which simplifies to $|Y_0| + |\mathcal{R}_0| \leq |N_K(\mathcal{R}_0 \cup Y_0)|$, since $|\mathcal{U}_3| = |C_3|$. This last inequality shows that we can apply Hall's Theorem to find a matching in K saturating $\mathcal{R} \cup Y_0$, which proves the lemma. □

Lemma 3.4.15. *For $i = 2$, let K_i be the bipartite graph with parts $\mathcal{R} \cup Y_{3-i}$ and $W \cap V_i$, where there is an edge between $R \in \mathcal{R}$ and $w \in W \cap V_i$ precisely when there is an R-edge containing w, and there is an edge between $y \in Y_{3-i}$ and $w \in W \cap V_i$ precisely when there is an edge $ywz_{y,y'}$, where y' is the partner of y in the pairing between Y_1 and Y_2. Then K_i has a matching saturating $\mathcal{R} \cup Y_{3-i}$. If (Y_1, Y_2, X) is perfectly cromulent, then this holds also for $i = 1$.*

Proof. We will apply Hall's theorem, so let $\mathcal{R}_0 \subseteq \mathcal{R}$ and $Y_0 \subseteq Y_{3-i}$. We construct a vertex cover T of $\mathcal{H}$. Let C_i be the maximal essential set in the subgraph of K_i induced by $\mathcal{R}$ and $W \cap V_i$ (this is the graph B_i associated with $\mathcal{H}_0$), and let $\mathcal{U}_i \subseteq \mathcal{R}$ be such that $N_{K_i}(\mathcal{U}_i) = C_i$, which exists by the definition of essential. Let T be the union of the sets $(Y_1 \cup Y_2) \setminus Y_0$, $N_{K_i}(\mathcal{R}_0 \cup Y_0)$, $(V(\mathcal{R}) \cup V(\mathcal{F})) \cap V_i$, C_i, and $\bigcup_{R \in \mathcal{R} \setminus (\mathcal{U}_i \cup \mathcal{R}_0)} (R \cap V_3)$. Note the similarities to the i-heavy $(i, 3)$-cover of $\mathcal{H}_0$.

We must show that T is indeed a vertex cover. Let $e \in E(\mathcal{H}_0)$. Then e is either an $\mathcal{F}$-edge or an $\mathcal{R}$-edge. If it is an $\mathcal{F}$-edge, it is covered by $V(\mathcal{F}) \cap V_i \subseteq T$. If it is an $\mathcal{R}$-edge, then it is covered by $(V(\mathcal{F}) \cup V(\mathcal{R}) \cap V_i \subseteq T$, unless its V_i-vertex is in W, so assume that is the case. Let e be an R-edge. If $R \in \mathcal{R}_0$, then $e \cap V_i \in N_K(R) \subseteq T$. If $R \in \mathcal{U}_i$, then $e \cap V_i \in C_i \subseteq T$. If $R \in \mathcal{R} \setminus (\mathcal{U}_i \cup \mathcal{R}_0)$, then $e \cap V_3 = R \cap V_3 \subseteq T$. This shows that T covers every edge of $\mathcal{H}_0$. All edges incident to X intersect Y_2, which if $i = 2$ is part of T, and if $i = 1$, then (Y_1, Y_2, X) is assumed to be perfectly cromulent, in which case all edges incident to X are incident to $Y_1 \subseteq T$. Therefore, any uncovered edge must be incident to Y_0 and two vertices of $\mathcal{H}_0$. All such edges whose V_3-vertex is not superfluous intersect T, since $C_i \cup (V(\mathcal{R}) \cup V(\mathcal{F})) \cap V_i \subseteq T$. Thus, the only edges we have to worry about are those incident to some $y \in Y_0$ and a superfluous vertex $s \in V_i$. By Lemma 3.4.12 (with $(i, j) = (3 - i, i)$), there is an edge containing y and s, whose V_3-vertex is also superfluous. By Lemma 3.4.13, the V_3-vertices of those edges are the corresponding $z_{y,y'}$, and hence their V_2-vertices are in $N_{K_i}(Y_0) \subseteq T$ by the definition of K_i. This shows that T is a vertex cover.

We now calculate the size of T. By the definition of T, we calculate $|T| = |Y_1| + |Y_2| - |Y_0| + |N_{K_i}(\mathcal{R}_0 \cup Y_0)| + 2\,|\mathcal{F}| + |\mathcal{R}| + |C_i| - |C_i \cap N_{K_i}(\mathcal{R}_0)| + |\mathcal{R}| - |\mathcal{U}_i \cup \mathcal{R}_0|$. Because it is a vertex cover, we must have $|T| \geq \tau(\mathcal{H})$. Since $\nu(\mathcal{H}) = \nu(\mathcal{H}_0) + |Y_1|$ by the definition of cromulent triple, and since $\tau(\mathcal{H}) = 2\nu(\mathcal{H})$, we have $\tau(\mathcal{H}) = 2\nu(\mathcal{H}_0) + 2\,|Y_1| = 2\,|\mathcal{F} \cup \mathcal{R}| + |Y_1| + |Y_2|$. Combining this with the fact that $\tau(\mathcal{H}) \leq |T|$ yields the inequality $|Y_0| + |\mathcal{U}_i \cup \mathcal{R}_0| + |C_i \cap N_{K_i}(\mathcal{R}_0)| \leq |N_{K_i}(\mathcal{R}_0 \cup Y_0)| + |C_i|$. By the inclusion-exclusion principle we can rewrite this as $|Y_0| + |\mathcal{U}_i| + |\mathcal{R}_0| - |\mathcal{U}_i \cap \mathcal{R}_0| + |C_i \cap N_{K_i}(\mathcal{R}_0)| \leq |N_{K_i}(\mathcal{R}_0 \cup Y_0)| + |C_i|$. Since $C_i = N_{K_i}(\mathcal{U}_i)$, we clearly have $C_i \cap N_{K_i}(\mathcal{R}_0) \supseteq N_{K_i}(\mathcal{U}_i \cap \mathcal{R}_0)$. Since B_i has a matching saturating $\mathcal{R}$, by Hall's Theorem, we must have $|\mathcal{U}_i \cap \mathcal{R}_0| \leq |N_{K_i}(\mathcal{U}_i \cap \mathcal{R}_0)|$. Combining this with our previous inequality, we then get $|Y_0| + |\mathcal{U}_i| + |\mathcal{R}_0| - |\mathcal{U}_i \cap \mathcal{R}_0| + |\mathcal{U}_i \cap \mathcal{R}_0| \leq |N_{K_i}(\mathcal{R}_0 \cup Y_0)| + |C_i|$, which simplifies to $|Y_0| + |\mathcal{R}_0| \leq |N_{K_i}(\mathcal{R}_0 \cup Y_0)|$, since $|\mathcal{U}_i| = |C_i|$. This last inequality shows that we can apply Hall's Theorem to find a matching in K_i saturating $\mathcal{R} \cup Y_0$, which proves the lemma. □

3.4.3 The Proof of Corollary 3.4.4

It suffices to prove Lemmas 3.4.2 and 3.4.3.

Proof of Lemma 3.4.2. Let (Y_1, Y_2, X) be a perfectly cromulent triple. We set $\mathcal{R}' = \mathcal{R} \cup \{\{y, y', z_{y,y'}\} : y \in Y_1, y' \in Y_2$ in an edge of $\mathcal{M}$ together with $y\}$, and $W' = W \cup X \setminus \{z_{y,y'} : y \in Y_1, y' \in Y_2$ in an edge of $\mathcal{M}$ together with $y\}$, where

$z_{y,y'}$ is the superfluous vertex in V_3 from Lemma 3.4.13. By the application of Lemma 3.4.14, we find that $(\mathcal{F}, \mathcal{R}', W')$ is an FR-partition, since $\nu(\mathcal{H}) = \nu(\mathcal{H}_0) + |Y_1| = |\mathcal{F} \cup \mathcal{R}| + |Y_1| = |\mathcal{F} \cup \mathcal{R}'|$. Applying 3.4.15 for $i = 1, 2$ we get that $(\mathcal{F}, \mathcal{R}', W')$ has a matching in B'_1 and B'_2. We can combine the partial matching in B'_3 that we get from Lemma 3.4.14 with the edges of $\mathcal{M}$ going to X to complete it. Thus $(\mathcal{F}, \mathcal{R}', W')$ is a matchable FR-partition. It is clearly also proper, because there are no edges with three vertices in W' by virtue of the fact that no such edge is in $\mathcal{H}_0$ and all edges going to X have their other vertices in Y_1 and Y_2. Thus, by Proposition 3.3.9, we in fact have a home-base partition. □

Proof of Lemma 3.4.3. Let (Y_1, Y_2, X) be a cromulent triple. We now mean to rule out the possibility that any edge incident to X is also incident to an $\mathcal{F}$- or $\mathcal{R}$-vertex of $\mathcal{H}_0$. Lemma 3.4.15 means that we can find a hypergraph matching $\mathcal{M}'$ of size $|Y_1|$ in $\mathcal{H}$ consisting of edges of the form yss' with $y \in Y_1$, and s, s' superfluous vertices in $\mathcal{H}_0$. Suppose there were an edge $uy'x$ for some $u \in (V(\mathcal{F}) \cup V(\mathcal{R})) \cap V_1$, $y' \in Y_2$, and $x \in X$. By the matchability of B_1, we can choose a matching of WRR-edges for each $R \in \mathcal{R}$, which avoids u, since $u \notin W$. We can also clearly find a matching of $\mathcal{F}$-edges avoiding u. Combining these matchings with $\mathcal{M}'$ yields a hypergraph matching of size $\nu(\mathcal{H})$ which is disjoint from $uy'x$. This is impossible, so such an edge cannot exist. Therefore (Y_1, Y_2, X) is a perfectly cromulent triple. □

Therefore, we have shown that if we have a cromulent triple, we have a home-base hypergraph. The next section is devoted to finding cromulent triples under various assumptions.

3.5 Searching for Cromulent Triples

Let $\mathcal{H}$ be a 3-partite 3-graph with vertex classes V_1, V_2, and V_3, and with $\tau(\mathcal{H}) = 2\nu(\mathcal{H})$. We want to find a home-base partition of $\mathcal{H}$. By Corollary 3.4.4, we are done if $\mathcal{H}$ has a cromulent triple. Therefore, our goal will be to find a cromulent triple inside our hypergraph. We will do this under a few assumptions, and we will later show that if all of these assumptions fail to hold, then we can prove $\mathcal{H}$ is a home-base hypergraph even without cromulent triples.

Finding cromulent triples will entail finding a subgraph which is a home-base hypergraph. We do this by finding a subgraph which is tight for Ryser's Conjecture and has a smaller matching number than $\mathcal{H}$, and then applying induction on Theorem 1.1.2. We would like to pinpoint exactly where in the proof we need to rely on induction. Therefore, we lay out the induction hypothesis here precisely.

Induction Hypothesis (IH(k))**.** If $\mathcal{H}$ is a 3-partite 3-graph with $\nu(\mathcal{H}) \leq k$ and $\tau(\mathcal{H}) = 2\nu(\mathcal{H})$, then $\mathcal{H}$ is a home-base hypergraph.

The first assumption under which we will find a cromulent triple is if we have a good set (see Definition 2.5.6).

3.5.1 Good Subsets Lead to Cromulent Triples

Lemma 3.5.1. *Suppose IH($k-1$) holds. Let $\mathcal{H}$ be a 3-partite 3-graph with vertex classes V_1, V_2, and V_3 such that $\tau(\mathcal{H}) = 2\nu(\mathcal{H}) = 2k$. If $X \subseteq V_3$ is a good set for $\mathrm{lk}_{\mathcal{H}}(V_1)$, then the triple (Y_1, Y_2, X) is perfectly cromulent, where $Y_1 = N_{\mathrm{lk}_{\mathcal{H}}(V_2)}(X)$ and $Y_2 = N_{\mathrm{lk}_{\mathcal{H}}(V_1)}(X)$.*

Proof. Let $X \subseteq V_3$ be a good set, and let $Y_2 = N_{\mathrm{lk}_{\mathcal{H}}(V_1)}(X)$. Let $y \in Y_2$, and let $\mathcal{H}_y = \mathcal{H} - \{vyz \in E(\mathcal{H}) : v \in V_1, z \in V_3 \setminus X\}$. Since the deleted edges can be covered by one vertex (y), we clearly have $\tau(\mathcal{H}_y) \geq \tau(\mathcal{H}) - 1$, and of course $\nu(\mathcal{H}_y) \leq \nu(\mathcal{H})$ as $\mathcal{H}_y \subseteq \mathcal{H}$. It is easy to see that $\mathrm{lk}_{\mathcal{H}_y}(V_1) = \mathrm{lk}_{\mathcal{H}}(V_1) - \{yz \in E(\mathrm{lk}_{\mathcal{H}}(V_1)) : z \in V_3 \setminus X\}$. Therefore, because X is good, we have $\mathrm{conn}(L(\mathrm{lk}_{\mathcal{H}_y}(V_1))) \geq \mathrm{conn}(L(\mathrm{lk}_{\mathcal{H}}(V_1))) + 1$. Recall that by Theorem 2.1.3, we have $\mathrm{conn}(L(\mathrm{lk}_{\mathcal{H}}(V_1))) = \nu(\mathcal{H}) - 2$. Thus, we have $\mathrm{conn}(L(\mathrm{lk}_{\mathcal{H}_y}(V_1))) \geq \nu(\mathcal{H}) - 1$. By Proposition 2.3.1, there is a subset $S \subseteq V_1$ for which we have $\mathrm{conn}(L(\mathrm{lk}_{\mathcal{H}_y}(S))) \leq \nu(\mathcal{H}_y) - (|V_1| - |S|) - 2$ and $|S| \geq |V_1| - (2\nu(\mathcal{H}_y) - \tau(\mathcal{H}_y))$. Plugging in the inequalities for τ and ν, we get

$$\mathrm{conn}(L(\mathrm{lk}_{\mathcal{H}_y}(S))) \leq \nu(\mathcal{H}) - (|V_1| - |S|) - 2$$

and

$$|S| \geq |V_1| - (2\nu(\mathcal{H}) - \tau(\mathcal{H}) + 1) = |V_1| - 1$$

since $\tau(\mathcal{H}) = 2\nu(\mathcal{H})$.

We have seen that V_1 itself does not fulfil the first of these inequalities, so S must be a proper subset of V_1, and thus by the second inequality, $S = V_1 \setminus \{a\}$ for some $a \in V_1$. A priori, we do not know if this a is unique for each $y \in Y_2$, so denote by A_y the set of all V_1-vertices a for which $\mathrm{conn}(L(\mathrm{lk}_{\mathcal{H}_y}(V_1 \setminus \{a\}))) \leq \nu(\mathcal{H}) - 3$.

Let $a \in A_y$ and let $S = V_1 \setminus \{a\}$. By Theorem 2.1.1, we have $\nu(\mathrm{lk}_{\mathcal{H}_y}(S)) \leq 2\,\mathrm{conn}(L(\mathrm{lk}_{\mathcal{H}_y}(S))) + 4 \leq 2\nu(\mathcal{H}) - 2 = \tau(\mathcal{H}) - 2$, which implies that $\nu(\mathrm{lk}_{\mathcal{H}}(S)) \leq \tau(\mathcal{H}) - 1$ because at most one edge of each maximum matching has been erased when passing from $\mathcal{H}$ to $\mathcal{H}_y$ in the link of S. We must have $\tau(\mathcal{H}_y) = \tau(\mathcal{H}) - 1$ because if $\tau(\mathcal{H}_y) = \tau(\mathcal{H})$, then by inequality (i) of Proposition 2.3.1, we would have $\mathrm{conn}(L(\mathrm{lk}_{\mathcal{H}_y}(S))) \geq \tau(\mathcal{H}_y)/2 - 2$ (since $\mathrm{conn}(L(\mathrm{lk}_{\mathcal{H}_y}(S)))$ is an integer and $\tau(\mathcal{H}_y) = \tau(\mathcal{H})$ is even), which is a contradiction. We can in fact show $\nu(\mathrm{lk}_{\mathcal{H}}(S)) = \tau(\mathcal{H}) - 1$, from which $\nu(\mathrm{lk}_{\mathcal{H}_y}(S)) = \tau(\mathcal{H}) - 2$ then follows, by considering the vertex cover T_S of $\mathcal{H}$ consisting of a and a minimum vertex cover of $\mathrm{lk}_{\mathcal{H}}(S)$ (which, by König's Theorem, has size $\nu(\mathrm{lk}_{\mathcal{H}}(S))$).

This means that every maximum matching in $\mathrm{lk}_{\mathcal{H}}(S)$ must contain an edge which is not in $\mathrm{lk}_{\mathcal{H}_y}(S)$. Set $Z = V_3 \setminus X$ and $W = V_2 \setminus Y_2$. We get the following structure for the maximum matchings:

Claim. *For every $y \in Y_2$ and for every $a \in A_y$ every maximum matching in $\mathrm{lk}_{\mathcal{H}}(V_1 \setminus \{a\})$ contains an edge yz for some $z \in Z$, and then saturates $Y_2 \setminus \{y\}$ using (X, Y_2)-edges and saturates $Z \setminus \{z\}$ using (Z, W)-edges.*

Proof. Let $S = V_1 \setminus \{a\}$. As observed, every maximum matching in $\mathrm{lk}_{\mathcal{H}}(S)$ contains an edge from y to Z. Since X is good (hence decent), it satisfies property (1) of Definition 2.5.1, so $\nu(\mathrm{lk}_{\mathcal{H}}(V_1)) = |Y_2| + |Z|$. Then because there are no edges between X and W, it follows that every maximum matching in $\mathrm{lk}_{\mathcal{H}}(V_1)$ saturates Y_2 with edges incident to X and saturates Z with edges incident to W. Since $\nu(\mathrm{lk}_{\mathcal{H}}(S)) = \tau(\mathcal{H}) - 1 = \nu(\mathrm{lk}_{\mathcal{H}}(V_1)) - 1$, we cannot have more than one matching edge between Y_2 and Z. Thus the claim follows. □

This structure immediately implies that the sets A_y are pairwise disjoint.

Claim. *If $y, y' \in Y_2$ with $y \neq y'$, then $A_y \cap A_{y'} = \emptyset$.*

Proof. Let $a \in A_y$, and let $S = V_1 \setminus \{a\}$. Then we know that a maximum matching in $\mathrm{lk}_{\mathcal{H}}(S)$ contains a (y, Z)-edge and the rest of its edges are between X and Y_2 and between Z and W. Thus the only edge between Y_2 and Z in the matching is incident to y. For $a' \in A_{y'}$, the structure of the maximum matchings in $\mathrm{lk}_{\mathcal{H}}(V_1 \setminus \{a'\})$ is different, and thus $a \neq a'$, hence the sets A_y and A'_y must be disjoint. □

Since every A_y is non-empty, we thus clearly have $\left|\bigcup_{y\in Y_2} A_y\right| \geq |Y_2|$.

Claim. *For every $a \in \bigcup_{y\in Y_2} A_y$, every maximum (X, Y_2)-matching in $\mathrm{lk}_{\mathcal{H}}(V_1)$ must have one edge which extends only to a.*

Proof. Suppose there were a maximum (X, Y_2)-matching M' in $\mathrm{lk}_{\mathcal{H}}(V_1)$ in which every edge extended to an element of $S = V_1 \setminus \{a\}$. Then we could take a maximum (V_2, V_3)-matching in $\mathrm{lk}_{\mathcal{H}}(S)$ (which must contain a (y, Z)-edge) and replace the part of the matching which hits Y_2 with M'. Because X has no neighbors outside of Y_2, this modified matching is a matching and is at least as big as the original one and therefore also maximum. This does not use a (y, Z)-edge, so we have a contradiction. Thus M' must contain an edge which does not extend to S, and hence extends only to a. □

From this claim, we see that $\left|\bigcup_{y\in Y_2} A_y\right| = |Y_2|$, since there can be at most as many vertices in $\bigcup_{y\in Y_2} A_y$ as edges in a maximum (X, Y_2)-matching in $\mathrm{lk}_{\mathcal{H}}(V_1)$, of which there are precisely $|Y_2|$.

Claim. *$Y_1 = \bigcup_{y\in Y_2} A_y$ and there is a hypergraph matching in $\mathcal{H}_{Y_1\cup Y_2\cup X}$ saturating Y_1 and Y_2.*

Proof. We clearly have $Y_1 \supseteq \bigcup_{y\in Y_2} A_y$ by the previous claim. We will show the other inclusion as well. Consider any vertex $x \in Y_1$. It follows from the definitions of Y_1 and Y_2 that there is an (X, Y_2)-edge e in $\mathrm{lk}_{\mathcal{H}}(V_1)$ such that $e \cup \{x\} \in E(\mathcal{H})$. Since X is good, e appears in a maximum matching M. For every $y \in Y_2$ and every $a \in A_y$, one edge of the matching between X and Y_2 must extend to a (recall that to be maximum, M must saturate Y_2 using (Y_2, X)-edges and must saturate Z using (Z, W)-edges). Since the A_y's are all disjoint, the matching extends to a hypergraph matching saturating Y_2 and $\bigcup_{y\in Y_2} A_y$. Since e extends to $\bigcup_{y\in Y_2} A_y$, it follows that $x \in \bigcup_{y\in Y_2} A_y$ and hence $Y_1 = \bigcup_{y\in Y_2} A_y$. This proves the claim. □

Now we almost have that (Y_1, Y_2, X) is perfectly cromulent. We just need to show that $\mathcal{H}_0 = \mathcal{H} \setminus (Y_1 \cup Y_2 \cup X)$ is a home-base hypergraph with $\nu(\mathcal{H}_0) = \nu(\mathcal{H}) - |Y_1|$.

Consider the graph $\mathcal{H}_1 = \mathcal{H} \setminus (Y_1 \cup Y_2)$. Since we have removed only $2|Y_1|$ vertices from $\mathcal{H}$, it follows that $\tau(\mathcal{H}_1) \geq \tau(\mathcal{H}) - 2|Y_1|$. We must have $\nu(\mathcal{H}_1) \leq \nu(\mathcal{H}) - |Y_1|$ because to any matching in $\mathcal{H}_1$, we may add the matching of size $|Y_1|$ we just showed exists to it to produce a matching in $\mathcal{H}$ (because no matching edge in the original matching is incident to $Y_1 \cup Y_2 \cup X$). Because $\tau(\mathcal{H}_1) \leq 2\nu(\mathcal{H}_1)$, we must have equality in both cases, whence $\tau(\mathcal{H}_1) = 2\nu(\mathcal{H}_1) = 2\nu(\mathcal{H}) - 2|Y_1|$. Note however that X is a set of isolated vertices in $\mathcal{H}_1$, and so removing them changes neither the matching size nor the covering number. Hence $\mathcal{H}_0 = \mathcal{H}_1 \setminus X$ also has $\tau(\mathcal{H}_0) = 2\nu(\mathcal{H}_0) = 2\nu(\mathcal{H}) - 2|Y_1|$. By IH$(k-1)$, $\mathcal{H}_0$ is a home-base hypergraph. This proves that (Y_1, Y_2, X) is a perfectly cromulent triple. □

This lemma shows that if $\mathrm{lk}_{\mathcal{H}}(V_i)$ has a good set for any i, then we find a perfectly cromulent triple.

3.5.2 No Good Sets

From now on we assume that $\mathrm{lk}_{\mathcal{H}}(V_1)$ has no good set. Recall that by Theorem 2.1.3, we know that $\mathrm{conn}(L(\mathrm{lk}_{\mathcal{H}}(V_1))) = \nu(\mathcal{H}) - 2$, and so by Lemma 2.5.7 $\mathrm{lk}_{\mathcal{H}}(V_1)$ has a perfect matching. Moreover for every minimal equineighbored set $X \subseteq V_3$ both it and its neighborhood $N_{\mathrm{lk}_{\mathcal{H}}(V_1)}(X)$ have size 2 and together induce a C_4 (possibly with parallel edges). Our next assumption will be that there are two disjoint hyperedges incident to some minimal equineighbored set.

Lemma 3.5.2. *Suppose IH$(k-1)$ holds. Let $\mathcal{H}$ be a 3-partite 3-graph with vertex classes V_1, V_2, and V_3 such that $\tau(\mathcal{H}) = 2\nu(\mathcal{H}) = 2k$, and let $\mathrm{lk}_{\mathcal{H}}(V_1)$ have no good sets. Suppose there is a minimal equineighbored set $X \subseteq V_3$ in $\mathrm{lk}_{\mathcal{H}}(V_1)$ such that there are two disjoint hyperedges zyx and $z'y'x'$ of $\mathcal{H}$ with $x, x' \in X$. Let $Y_1 = \{z, z'\} \subseteq V_1$ and $Y_2 = \{y, y'\} \subseteq V_2$. Then (Y_1, Y_2, X) is a cromulent triple.*

Proof. For Condition (1) note that $|Y_1| = |Y_2| = |X| = 2$, since by Lemma 2.5.7 X has size 2.

Then $X = \{x, x'\}$ and because X is equineighbored, the neighborhood of X is also of size 2, that is, $N_{\mathrm{lk}_{\mathcal{H}}(V_1)}(X) = \{y, y'\}$. So Condition (2) is satisfied.

For Condition (3) note that by assumption there are two disjoint hyperedges zyx and $z'y'x'$ in $\mathcal{H}|_{Y_1 \cup Y_2 \cup X}$ and that $|Y_1| = 2$.

For Condition (4) we first prove that $\tau(\mathcal{H}_0) = 2\nu(\mathcal{H}_0) = 2(\nu(\mathcal{H}) - |Y_1|)$. Then we can use IH$(k-1)$ to derive the existence of a home-base partition of $\mathcal{H}_0$. First, consider the graph $\mathcal{H}_1 = \mathcal{H} \setminus (Y_1 \cup Y_2)$. Since we have removed only $2|Y_1|$ vertices from $\mathcal{H}$, it follows that $\tau(\mathcal{H}_1) \geq \tau(\mathcal{H}) - 2|Y_1|$. We must have $\nu(\mathcal{H}_1) \leq \nu(\mathcal{H}) - |Y_1|$ because X consists of isolated vertices in $\mathcal{H}_1$, so we may add zyx and $z'y'x'$ to any matching in $\mathcal{H}_1$ to obtain a matching 2 larger in $\mathcal{H}$. Because $\tau(\mathcal{H}_1) \leq 2\nu(\mathcal{H}_1)$, we must have equality in both cases, whence

$\tau(\mathcal{H}_1) = 2\nu(\mathcal{H}_1) = 2\nu(\mathcal{H}) - 2|Y_1|$. Note however that because X is a set of isolated vertices in $\mathcal{H}_1$, removing them changes neither the matching size nor the covering number. Hence $\mathcal{H}_0 = \mathcal{H}_1 \backslash X$ also has $\tau(\mathcal{H}_0) = 2\nu(\mathcal{H}_0) = 2\nu(\mathcal{H}) - 2|Y_1|$. Thus $\mathcal{H}_0$ has a home-base partition $(\mathcal{F}, \mathcal{R}, W)$.

The proof of Condition (5) is far more involved and will use a number of internal lemmas, so we give a brief overview. Our goal will be to find a contradiction by providing a larger matching than $\nu(\mathcal{H})$ if there is an edge of $\mathcal{H}$ incident to X and a W-vertex of $\mathcal{H}_0$. This matching will consist of a maximum matching in $\mathcal{H}_0$ and a few extra edges whose existence will be guaranteed by the high vertex cover number of $\mathcal{H}$. We utilize the fact that we are quite flexible in choosing a matching for $\mathcal{H}_0$, so that we can usually avoid the vertices of the extra edges when we choose our matching. Recall the definition of superfluous vertices and i-heavy (i, j)-covers from Section 3.4.

Lemma 3.5.3. *There is no edge wyx with $w \in W$. Similarly there is no $wy'x'$.*

Proof. Suppose wyx is an edge. Take the following partial cover of $\mathcal{H}$: y, y', and z' plus the 2-heavy $(2, 3)$-cover of $\mathcal{H}_0$. Since this set of vertices is one too small to be a cover, this implies the existence of an edge zsp avoiding it, where s is superfluous in $\mathcal{H}_0$, and $p \in V(\mathcal{H}_0)$. Indeed, an edge not intersecting the partial cover must avoid Y_2, hence also X, is not in $E(\mathcal{H}_0)$, and by Observation 3.4.6, its V_2-vertex is superfluous. By Case (4) of Corollary 3.3.6 applied to $\mathcal{H}_0$ with $a = s$, $b = w$, $c = p$, and $S = \emptyset$, we can find a matching of size $\nu(\mathcal{H}_0)$ inside $\mathcal{H}_0$ avoiding $\{s, w, p\}$. This matching together with the edges $z'y'x'$, wyx, and zsp gives a matching of size $\nu(\mathcal{H}_0) + 3 = \nu(\mathcal{H}) + 1$, a contradiction. □

Lemma 3.5.4. *If there is an edge of $\mathcal{H}$ incident to X and a vertex of $W \cap V_1$, then there are two disjoint edges of $\mathcal{H}$ whose V_1-vertices are in W, at least one being superfluous, whose V_2-vertices are y and y', and exactly one of whose V_3-vertices are in $V(\mathcal{H}_0)$.*

Proof. Suppose there is an edge incident to $w \in W \cap V_1$ and X. Without loss of generality suppose it is incident to x. Then by Lemma 3.5.3, it is not incident to y, so it must be the edge $wy'x$.

Suppose that w is superfluous in $\mathcal{H}_0$. Then we will show that wyx' is also an edge of $\mathcal{H}$ and that $wy'x$ and wyx' are the only edges extending $y'x$ or yx'.

Since X is a minimal equineighbored of size 2, we have $yx' \in E(\mathrm{lk}_{\mathcal{H}}(V_1))$, and hence there is some edge $vyx' \in E(\mathcal{H})$. Suppose $v \neq w$. Take the partial cover consisting of $\{y, y'\}$ plus the 2-heavy $(2, 3)$-cover of $\mathcal{H}_0$. If $v \in \{z, z'\}$, then add v to the partial cover. If $v \in R_1 \in \mathcal{R}$, then add instead the vertex in $R_1 \cap V_3$ to the partial cover. This leaves an edge of the form (z or z')sp where $s \in V_2$ is superfluous in $\mathcal{H}_0$ and $p \notin R_1$ (in case $v \in V(\mathcal{R})$, hence R_1 exists) which is disjoint from vyx'. Indeed, an edge not intersecting the partial cover must avoid Y_2, hence also X, is not in $E(\mathcal{H}_0)$, and by Observation 3.4.6, its V_2-vertex is superfluous. If $v \in \{z, z'\}$, then we can apply Case (4) of Corollary 3.3.6 to $\mathcal{H}_0$ with $a = w$, $b = s$, $c = p$, and $S = \emptyset$. If $v \in V(\mathcal{R})$, then we can apply Case (2) of Corollary 3.3.6 to $\mathcal{H}_0$ with $a = v$, $b = p$, $c = s$, and $S = \{w\}$. And if $v \in V(\mathcal{H}_0) \setminus V(\mathcal{R})$, then we can apply Case (4) of Corollary 3.3.6 to $\mathcal{H}_0$ with $a = s$, $b = v$, $c = p$, and $S = \{w\}$. In any case, we find a matching in $\mathcal{H}_0$ of size $\nu(\mathcal{H}_0)$ avoiding $\{w, v, s, p\}$. Then this matching together with $wy'x$, vyx', and (z or z')sp gives a matching of size $\nu(\mathcal{H}_0) + 3 = \nu(\mathcal{H}) + 1$, a contradiction.

Therefore the only edge extending yx' is wyx', and because wyx' is an edge, a similar argument shows that $wy'x$ is the only edge extending $y'x$.

Take a partial cover $\{z, z', w\}$ plus the 1-heavy $(1, 2)$-cover of $\mathcal{H}_0$. This leaves an edge w'(y or y')p where w' is superfluous and $w' \neq w$. Indeed, an edge not intersecting the partial cover is not in $E(\mathcal{H}_0)$, and by Observation 3.4.6, its V_1-vertex is superfluous. Also $p \notin \{x, x'\}$, since $w' \neq w$. It is disjoint from one of wyx' and $wy'x$, so w'(y or y')p together with whichever of wyx' and $wy'x$ it is disjoint from are the two disjoint edges we are after.

Suppose on the other hand, that there is no edge incident to $\{x, x'\}$ which extends to a superfluous vertex in V_1. Then in particular w is not superfluous in $\mathcal{H}_0$. Take the partial cover $\{z, z', y'\}$ plus the 1-heavy $(1, 3)$-cover of $\mathcal{H}_0$. This leaves an edge syp where s is superfluous in $\mathcal{H}_0$, and hence $s \neq w$. Indeed, an edge not intersecting the partial cover is not in $E(\mathcal{H}_0)$, and by Observation 3.4.6, its V_1-vertex is superfluous. Also $p \notin \{x, x'\}$, since s is superfluous. Thus $wy'x$ and syp are the two disjoint edges we are after. □

Thus we may suppose that there is an edge incident to $W \cap V_1$ and X. By Lemma 3.5.4, there are two disjoint edges e and f whose vertices intersect $V(\mathcal{H}_0)$ in $s, w \in W \cap V_1$ and $p \in V_3$. At least one of s and w is superfluous in $\mathcal{H}_0$, so suppose without loss of generality that s is the superfluous one. We consider several cases, depending on the location of p. In each case we will reach a contradiction.

Case 1. $p \in V(\mathcal{F})$.

Take the partial cover $\{y, y', z\}$, plus the 3-heavy $(3, 2)$-cover of $\mathcal{H}_0$. This gives an edge $z'p's'$ where s' is superfluous (hence $s' \neq p$). Indeed, an edge not intersecting the partial cover must avoid Y_2, hence also X, is not in $E(\mathcal{H}_0)$, and by Observation 3.4.6, its V_3-vertex is superfluous. We can apply Case (1) of Corollary 3.3.6 with $a = p$, $b = w$, $c = p'$, and $S = \{s, s'\}$ to obtain a matching of size $\nu(\mathcal{H}_0)$ in $\mathcal{H}_0$ avoiding $\{s, s', w, p', p\}$. This matching together with the edges e, f, and $z'p's'$ gives a matching of size $\nu(\mathcal{H}_0) + 3 = \nu(\mathcal{H}) + 1$, a contradiction.

Case 2. $p \in R_1 \in \mathcal{R}$.

Take the partial cover $\{y, y'\}$ together with the vertex in $R_1 \cap V_2$ and the 3-heavy $(3, 2)$-cover of $\mathcal{H}_0$. This gives an edge (z or z')$p's'$ where s' is superfluous (note $s' \neq p$) and p' is not in R_1. Indeed, an edge not intersecting the partial cover must avoid Y_2, hence also X, is not in $E(\mathcal{H}_0)$, and by Observation 3.4.6, its V_3-vertex is superfluous. We can apply Case (2) of Corollary 3.3.6 with $a = p$, $b = p'$, $c = w$, and $S = \{s, s'\}$ to obtain a matching of size $\nu(\mathcal{H}_0)$ in $\mathcal{H}_0$ avoiding $\{s, s', w, p', p\}$. This matching together with the edges e, f, and (z or z')$p's'$ gives a matching of size $\nu(\mathcal{H}_0) + 3 = \nu(\mathcal{H}) + 1$, a contradiction.

Case 3. $p \in W$ is essential for $R_1 \in \mathcal{R}$.

Take the partial cover $\{y, y'\}$, the V_2-vertex essential for R_1 if it exists, plus the 3-heavy $(3, 2)$-cover of $\mathcal{H}_0$. This gives an edge $(z \text{ or } z')p's'$ where s' is superfluous (hence $s' \neq p$) and p' is not essential for R_1. Indeed, an edge not intersecting the partial cover must avoid Y_2, hence also X, is not in $E(\mathcal{H}_0)$, and by Observation 3.4.6, its V_3-vertex is superfluous. We can apply Case (3) of Corollary 3.3.6 with $a = p$, $b = p'$, $c = w$, and $S = \{s, s'\}$ to obtain a matching of size $\nu(\mathcal{H}_0)$ in $\mathcal{H}_0$ avoiding $\{s, s', w, p', p\}$. This matching together with the edges e, f, and $(z \text{ or } z')p's'$ gives a matching of size $\nu(\mathcal{H}_0) + 3 = \nu(\mathcal{H}) + 1$, a contradiction.

Case 4. $p \in W$ is not essential but not superfluous.

Take the partial cover $\{y, y'\}$ plus the 3-heavy $(3, 2)$-cover of $\mathcal{H}_0$. This gives an edge $(z \text{ or } z')p's'$ where s' is superfluous, hence $s' \neq p$. Indeed, an edge not intersecting the partial cover must avoid Y_2, hence also X, is not in $E(\mathcal{H}_0)$, and by Observation 3.4.6, its V_3-vertex is superfluous. By Lemma 3.3.4, p does not become essential after removing a superfluous vertex from V_3. Then we can apply Case (4) of Corollary 3.3.6 with $a = p$, $b = w$, $c = p'$, and $S = \{s, s'\}$ to obtain a matching of size $\nu(\mathcal{H}_0)$ in $\mathcal{H}_0$ avoiding $\{s, s', w, p', p\}$. This matching together with the edges e, f, and $(z \text{ or } z')p's'$ gives a matching of size $\nu(\mathcal{H}_0) + 3 = \nu(\mathcal{H}) + 1$, a contradiction.

Case 5. $p \in W$ is superfluous.

Take the partial cover $\{y, y', p\}$ plus the 2-heavy $(2, 3)$-cover of $\mathcal{H}_0$. This gives an edge $(z \text{ or } z')s'p'$ where s' is superfluous and $p' \neq p$. Indeed, an edge not intersecting the partial cover must avoid Y_2, hence also X, is not in $E(\mathcal{H}_0)$, and by Observation 3.4.6, its V_2-vertex is superfluous. We can apply Case (4) of Corollary 3.3.6 with $a = s'$, $b = w$, $c = p'$, and $S = \{s, p\}$ to obtain a matching of size $\nu(\mathcal{H}_0)$ in $\mathcal{H}_0$ avoiding $\{s, s', w, p', p\}$. This matching together with the edges e, f, and $(z \text{ or } z')s'p'$ gives a matching of size $\nu(\mathcal{H}_0) + 3 = \nu(\mathcal{H}) + 1$, a contradiction.

Thus we conclude that there can be no edge incident to $W \cap V_1$ and X, so Condition (5) must hold, and hence (Y_1, Y_2, X) is a cromulent triple. □

Thus, if we either have a good set, or if we have no good set and there are two disjoint hyperedges incident to a minimal equineighbored subset of some link graph, then we find a cromulent triple, and hence have found a home-base partition by Corollary 3.4.4. Therefore, the only hypergraphs left to check are those which have no good set and where the hyperedges incident to any equineighbored subset of any link graph form intersecting hypergraphs. This case is handled in the next section.

3.6 The End Game

We start with the following easy proposition which will be useful in what is to come:

Proposition 3.6.1. *Let $\mathcal{H}$ be a 3-partite 3-graph with vertex classes V_1, V_2, and V_3 such that each link $\mathrm{lk}_{\mathcal{H}}(V_i)$ has a perfect matching. Suppose $X \subseteq V_j$*

is a minimal equineighbored set of $\mathrm{lk}_{\mathcal{H}}(V_i)$ *with* $|X| = 2$*, and suppose* X *is not incident to two disjoint edges of* $\mathcal{H}$*. Then the edges incident to* X *form a truncated multi-Fano plane.*

Proof. Since X is a minimal equineighbored set of size 2 and $\mathrm{lk}_{\mathcal{H}}(V_i)$ has no isolated vertices, it follows easily that the edges of $\mathrm{lk}_{\mathcal{H}}(V_i)$ incident to X form a C_4 (possibly with parallel edges). By assumption, the edges incident to X form an intersecting hypergraph. Since the hyperedges incident to X all intersect, each pair of opposite edges in the C_4 must extend to one vertex in V_i. If this is the same vertex v for all pairs, then $N_{\mathrm{lk}_{\mathcal{H}}(V_k)}(X) = \{v\}$, where V_k is the third vertex class besides V_i and V_j. This contradicts the fact that $\mathrm{lk}_{\mathcal{H}}(V_k)$ has a perfect matching, so each pair extends to a different vertex, which gives the truncated Fano plane. If there are parallel edges in the C_4, this analysis shows that they also have to extend to the same vertex as the edges to which they are parallel, hence we have a truncated multi-Fano plane. $\square$

We aim to prove the following lemma, which is the missing ingredient in our proof of Theorem 1.1.2.

Lemma 3.6.2. *Suppose IH(*$k-1$*) holds. Let* $\mathcal{H}$ *be a* 3*-partite* 3*-graph with vertex classes* V_1*,* V_2*, and* V_3 *such that* $\tau(\mathcal{H}) = 2\nu(\mathcal{H}) = 2k$*. Suppose that* $\mathcal{H}$ *does not have a cromulent triple. Then there is an* $X \subseteq V_3$*, which is a minimal equineighbored set for* $\mathrm{lk}_{\mathcal{H}}(V_1)$ *such that for its neighborhood* $Y = N_{\mathrm{lk}_{\mathcal{H}}(V_1)}(X)$ *we also have* $N_{\mathrm{lk}_{\mathcal{H}}(V_1)}(Y) = X$*.*

Proof. We have shown in Lemma 3.5.1 that we have a cromulent triple if there is at least one good set, which means we are working under the assumption that $\mathrm{lk}_{\mathcal{H}}(V_1)$ has no good set. By Lemma 2.5.7, we then know that $\mathrm{lk}_{\mathcal{H}}(V_1)$ has a perfect matching and that every minimal equineighbored set is of size 2 and hence is incident to a C_4. Therefore, it is clear that every edge incident to a minimal equineighbored set participates in a perfect matching, so we have shown that every minimal equineighbored set is still decent.

If $X \subseteq V_3$ is a minimal equineighbored set, for $y \in N_{\mathrm{lk}_{\mathcal{H}}(V_1)}(X)$ define the bipartite graph $G_y = \mathrm{lk}_{\mathcal{H}}(V_1) - \{yz \in E(\mathrm{lk}_{\mathcal{H}}(V_1)) : z \in V_3 \setminus X\}$. Since X is decent but not good, it must be that for some $y \in N_{\mathrm{lk}_{\mathcal{H}}(V_1)}(X)$ we have

$$\mathrm{conn}(L(G_y) \leq \mathrm{conn}(L(\mathrm{lk}_{\mathcal{H}}(V_1))).$$

A similar statement holds if $X \subseteq V_2$.

Now suppose for the sake of contradiction to the statement of Lemma 3.6.2 that for every minimal equineighbored subset X in $\mathrm{lk}_{\mathcal{H}}(V_1)$, its neighborhood Y has neighbors outside of X. Again, Theorem 2.1.3 gives that $\mathrm{lk}_{\mathcal{H}}(V_1)$ is extremal for Theorem 2.1.1, and hence it has a CP-decomposition by Theorem 2.4.3. We know that any CP-decomposition of $\mathrm{lk}_{\mathcal{H}}(V_1)$ contains some P_4's, since otherwise the graph would consist entirely of disjoint C_4's, which is not the case if there are edges between Y and $V_3 \setminus X$.

Claim. *The graph* $\mathrm{lk}_{\mathcal{H}}(V_1)$ *contains a minimal equineighbored set* $X \subseteq V_3$ *for which both elements of* $N(X)$ *have neighbors outside* X *in* $\mathrm{lk}_{\mathcal{H}}(V_1)$*.*

Proof. Let Z be the set of endpoints of P_4's in V_3 for some CP-decomposition of $\mathrm{lk}_{\mathcal{H}}(V_1)$ with respect to some perfect matching M. Then Z is equineighbored because the edges incident to the endpoints in V_3 all must contain an interior vertex in V_2 either of the same P_4 or of some other one. The set of interior vertices of P_4's in V_2 is matched by M to the set of endpoints of P_4's in V_3, so these are the same size. Therefore $|Z| = |N(Z)|$. Since Z is equineighbored, it contains a minimal equineighbored subset X.

Since X consists of endpoints of P_4's and $N(X)$ consists of interior vertices of P_4's, the vertices in $N(X)$ all have neighbors outside X: the other interior vertices of their respective P_4's. ☐

Fix a perfect matching M of the link graph $\mathrm{lk}_{\mathcal{H}}(V_1)$. Let $X_3 \subseteq V_3$ be a minimal equineighbored set for which both elements of $N(X_3)$ have neighbors outside X_3, and let $N(X_3) = \{y, y'\}$. Let $X_3 = \{x, x'\}$ so that $yx, y'x' \in M$. Without loss of generality, let y' be a vertex of $N(X_3)$ that witnesses the failure of X_3 to be good; that is, we have

$$\mathrm{conn}(L(G_{y'})) \leq \mathrm{conn}(L(\mathrm{lk}_{\mathcal{H}}(V_1))).$$

Then by Theorem 2.4.3, $G_{y'}$ has a CP-decomposition with respect to M (since no edges of M were erased, and hence $G_{y'}$ is still extremal for Theorem 2.1.1). We claim that in every CP-decomposition of $G_{y'}$, the two vertices of X_3 are together in one of the C_4's of the decomposition. The edge $x'y'$ is an edge of M, so it must be in some C_4 or P_4 of the CP-decomposition. Since $N_{G_{y'}}(y') = X$, and $N_{G_{y'}}(x') = N_{\mathrm{lk}_{\mathcal{H}}(V_1)}(X_3)$, this C_4 or P_4 must be contained in $G_{y'}[X_3 \cup N(X_3)]$. But we know the edges in $G_{y'}[X_3 \cup N(X_3)]$ form a C_4, so $x'y'$ can't be contained in a P_4 of the CP-decomposition (one of the edges xy' and $x'y$ would not be at home anywhere).

Let Z_2 be the set of vertices in V_2 reachable by M-alternating paths in $G_{y'}$ starting at y with an edge not in M (including y itself). Note that $Y \subseteq Z_2$. We have $\left|N_{G_{y'}}(Z_2)\right| = |Z_2|$ because every vertex of V_3 we reach is matched to a vertex of V_2 which is included in Z_2. Then Z_2 contains a minimal equineighbored set X_2. Note that X_2 is disjoint from Y, since $X_2 \backslash Y$ must also be equineighbored (because X_3 is taken out of the neighborhood), and $X_2 \backslash Y$ is not empty because $\left|N_{G_{y'}}(Y)\right| > 2$. This means also that X_2 has exactly the same neighborhood in $G_{y'}$ and in $\mathrm{lk}_{\mathcal{H}}(V_1)$, and so it is also a minimal equineighbored set for $\mathrm{lk}_{\mathcal{H}}(V_1)$. Therefore, $|X_2| = 2$ and the edges incident to X_2 form a C_4.

Lemma 3.6.3. *In any CP-decomposition of $G_{y'}$ all vertices of $Z_2 \setminus N(X_3)$ are endpoints of P_4's, and all vertices of $N(Z_2 \setminus N(X_3))$ are interior vertices of P_4's.*

Proof. Since the $(y', V_3 \setminus X_3)$-edges are erased, any CP-decomposition of $G_{y'}$ must have a C_4 on $X_3 \cup N(X_3)$. So any M-alternating path going out from y (not to X_3) must go first to an interior vertex of a P_4, which is matched to an endpoint of that P_4, and so on, alternating between interior vertices and endpoints. So the neighbors of $Z_2 \setminus N(X_3)$ are interior vertices and the vertices of $Z_2 \setminus N(X_3)$ are endpoints. □

This shows in particular that both vertices of X_2 are endpoints of P_4's, and both vertices of $N(X_2)$ are interior vertices of P_4's, and hence both have neighbors outside of X_2.

Lemma 3.6.4. *If $X \subseteq V_3$ and $X' \subseteq V_2$ are minimal equineighbored subsets of $\mathrm{lk}_{\mathcal{H}}(V_1)$ with $X' \cap N(X) = \emptyset$, and there is an M-alternating path from $N(X)$ to $N(X')$ starting with a non-matching edge, then the edges incident to X and the edges incident to X' extend to the same two vertices $\{z, z'\} \subseteq V_1$.*

Proof. We have seen that each link graph $\mathrm{lk}_{\mathcal{H}}(V_i)$ has a perfect matching, and we know $|X| = 2$ and is not incident to two disjoint hyperedges, so by Proposition 3.6.1, the edges incident to X form a truncated Fano plane.

Let $N(X) = \{y, y'\}$, and let $N(X') = \{w, w'\}$, where without loss of generality y is the last vertex of $N(X)$ visited on the M-alternating path, and w is the first vertex of $N(X')$ visited. Let $G_{y',w'}$ be the graph formed by erasing *both* the $(y', V_3 \setminus X)$-edges and the $(w', V_2 \setminus X_2)$-edges from $\mathrm{lk}_{\mathcal{H}}(V_1)$. We will show that $G_{y',w'}$ does *not* have a CP-decomposition. Suppose it did. Then fix a CP-decomposition of $G_{y',w'}$. Both X and X' would need to consist of vertices of a C_4 in the CP-decomposition of $G_{y',w'}$, as previously observed for $G_{y'}$. However since there is an M-alternating path from y to w starting with a non-matching edge, we will see that this leads to a contradiction. Consider the first edge yv of this path. It is not an edge of a C_4 or P_4 of the CP-decomposition, so it must be at home in some P_4, and since y is not an interior vertex of a P_4 of the CP-decomposition, it follows that v is. The next edge is an edge of M which pairs the interior vertex v with an endpoint. The next edge must be at home in some P_4, hence its other vertex is again an interior vertex of that P_4. Continuing in this manner, one sees that the even vertices of the path (y being the first vertex) are interior vertices of P_4's of the CP-decomposition. However, since w is one of the even vertices, this contradicts the fact that w is a vertex of a C_4 of the CP-decomposition. Therefore no CP-decomposition is possible, and hence by the contrapositive of Theorem 2.4.3, we must have

$$\mathrm{conn}(L(G_{y',w'})) \geq \frac{\nu(G_{y',w'})}{2} - 1 = \frac{\nu(\mathrm{lk}_{\mathcal{H}}(V_1))}{2} - 1 = \nu(\mathcal{H}) - 1, \tag{3.6.1}$$

where the last equality is by Theorem 2.1.3.

Consider the hypergraph $\mathcal{H}_{y',w'}$ that results by removing from $\mathcal{H}$ the edges inducing the $(y', V_3 \setminus X)$-edges and the $(w', V_2 \setminus X_2)$-edges in $\mathrm{lk}_{\mathcal{H}}(V_1)$. Then clearly $\mathrm{lk}_{\mathcal{H}_{y',w'}}(V_1) = G_{y',w'}$. We have $\tau(\mathcal{H}_{y',w'}) \geq \tau(\mathcal{H}) - 2$, since we can cover all of the deleted edges with two vertices, and we clearly have $\nu(\mathcal{H}_{y',w'}) \leq \nu(\mathcal{H})$. Therefore by parts (ii) and (iii) of Proposition 2.3.1, there is some $S \subseteq V_1$ such that $\mathrm{conn}(L(\mathrm{lk}_{\mathcal{H}_{y',w'}}(S))) \leq \nu(\mathcal{H}) - (|V_1| - |S|) - 2$ and $|S| \geq |V_1| - 2$. We know $S \neq V_1$ because the first inequality fails for V_1, as we have just concluded in the preceding paragraph.

Combining the inequality for $\text{conn}(L(\text{lk}_{\mathcal{H}_{y',w'}}(S)))$ with the inequality in Theorem 2.1.1 gives that $\nu(\text{lk}_{\mathcal{H}_{y',w'}}(S)) \leq 2\nu(\mathcal{H}) - 2(|V_1| - |S|)$. Recalling the vertex cover T_S of $\mathcal{H}$ consisting of $V_1 \setminus S$ and a minimal vertex cover of $\text{lk}_{\mathcal{H}}(S)$ gives that $\nu(\text{lk}_{\mathcal{H}}(S)) \geq \tau(\mathcal{H}) - (|V_1| - |S|)$ (by König's Theorem). Thus we have

$$\nu(\text{lk}_{\mathcal{H}_{y',w'}}(S)) \leq \nu(\text{lk}_{\mathcal{H}}(S)) - (|V_1| - |S|). \tag{3.6.2}$$

Therefore, every maximum matching of $\text{lk}_{\mathcal{H}}(S)$ has to contain an edge that gets erased in $\mathcal{H}_{y',w'}$. If xy and $x'y'$ are in $\text{lk}_{\mathcal{H}}(S)$, then we can change any maximum matching to avoid a $(y', V_3 \setminus X)$-edge without changing the cardinality of the matching, and similarly for xy' and $x'y$. Analogously, we can avoid a $(w', V_2 \setminus X')$-edge if either pair of opposite edges of the C_4 incident to X' is contained in $\text{lk}_{\mathcal{H}}(S)$. Therefore for one of the C_4's, no pair of opposite edges is contained in $\text{lk}_{\mathcal{H}}(S)$. This implies that the two vertices of V_1 to which the edges of the C_4 extend are not in S, and hence in fact $|S| = |V_1| - 2$.

This of course means that every maximum matching of $\text{lk}_{\mathcal{H}}(S)$ has to contain *two* edges that get erased in $\mathcal{H}_{y',w'}$, so no pair of opposite edges of either C_4 is contained in $\text{lk}_{\mathcal{H}}(S)$, and hence the vertices of V_1 to which the edges extend are not in S. But each C_4 extends to exactly two vertices, as observed in Lemma 3.6.1, and since $|S| = |V_1| - 2$, they must be the same two vertices for X and X', as claimed. $\square$

Lemma 3.6.4 applied to X_2 and X_3 shows that $\mathcal{H}$ has two truncated Fano planes intersecting in two vertices $\{z, z'\} \subseteq V_1$. We will see that this leads to a contradiction.

Let $X_2 = \{v, v'\}$, and let $N(X_2) = \{w, w'\}$. Assume without loss of generality that the truncated Fano planes consist of the edges $\{zyx, zy'x', z'yx', z'y'x\}$ and $\{zvw, zv'w', z'vw', z'v'w\}$. Consider the hypergraph $\mathcal{H}' = \mathcal{H} - \{y, w, z, z'\}$, and note that X_3 and X_2 consist of isolated vertices in $\mathcal{H}'$, since all edges incident to them are incident to $\{z, z'\}$. Because we have deleted only four vertices, we clearly have $\tau(\mathcal{H}') \geq \tau(\mathcal{H}) - 4$. To any matching in $\mathcal{H}'$ we may add zyx and $z'vw$ to get a matching two larger in $\mathcal{H}$, so we must have $\nu(\mathcal{H}') \leq \nu(\mathcal{H}) - 2$. Combining this with the assumption that $\tau(\mathcal{H}) = 2\nu(\mathcal{H})$ and the fact that Ryser's Conjecture is true for 3-partite hypergraphs we get the following sequence of inequalities:

$$\tau(\mathcal{H}') \leq 2\nu(\mathcal{H}') \leq 2\nu(\mathcal{H}) - 4 = \tau(\mathcal{H}) - 4 \leq \tau(\mathcal{H}').$$

Since the first and last expressions are the same, all inequalities are actually equalities, and hence $\mathcal{H}'$ is also extremal for Ryser's Conjecture, with $\nu(\mathcal{H}') = k - 2$. Therefore, by the inductive hypothesis IH$(k-1)$, $\mathcal{H}'$ has a home-base partition $(\mathcal{F}, \mathcal{R}, W)$.

We will find either a vertex cover of size $\tau(\mathcal{H}) - 1$, or a matching of size $\nu(\mathcal{H}) + 1$ in $\mathcal{H}$, either of which gives our desired contradiction.

Consider the minimal vertex cover of $\mathcal{H}'$ consisting of $V(\mathcal{F}) \cap V_1$ and $V(\mathcal{R}) \cap (V_1 \cup V_3)$. If adding the three vertices z, z', and w to this set would form a vertex cover T of $\mathcal{H}$, we would have a contradiction and be done, so we may assume that there is some edge $e \in E(\mathcal{H})$ which avoids T. Its V_1-vertex must be in W, since $(V(\mathcal{F}) \cup V(\mathcal{R})) \cap V_1 \cup \{z, z'\} \subseteq T$. Its V_3-vertex must be in $V(\mathcal{F}) \cup W$, since $V(\mathcal{R}) \cap V_3 \cup \{w\} \subseteq T$ and any edge incident to X_3 intersects T in $\{z, z'\}$. Its V_2-vertex cannot be in $V(\mathcal{H}')$, since otherwise e would be an edge of $\mathcal{H}'$ and hence intersect T, and its V_2-vertex also cannot be in X_2, since all edges incident to X_2 intersect T in $\{z, z'\}$. Therefore e must go through y, so it is of the form ayb for some vertices $a \in W \cap V_1$ and $b \in (V(\mathcal{F}) \cup W) \cap V_3$.

Suppose we can find a maximum matching in $\mathcal{H}'$ avoiding a, y', and b. Then this matching plus the three disjoint edges $zy'x'$, $z'v'w$, and ayb would form a matching of size $\nu(\mathcal{H}) + 1$ in $\mathcal{H}$, a contradiction.

By the monster lemma (Lemma 3.3.5), we can find a matching of size $\nu(\mathcal{H}')$ in $\mathcal{H}' - \{a, y', b\}$ if there is an F-edge avoiding $\{a, y', b\}$ for each $F \in \mathcal{F}$, and an R-edge avoiding $\{a, y', b\}$ for each $R \in \mathcal{R}$. Since $a \in W$, and y' and b are in different vertex classes, we do not cover all F-edges for any $F \in \mathcal{F}$. Since $a, b \notin V(\mathcal{R})$, we could pick an RWR-edge for any $R \in \mathcal{R}$ avoiding $\{a, y', b\}$ *unless* y' is a W-vertex essential for some $R \in \mathcal{R}$. This means that if $y' \notin W$, we have the desired contradictory matching, and hence we may assume $y' \in W$.

Consider the 1-heavy $(1, 3)$-cover of $\mathcal{H}'$ (see Section 3.4 for the definition), which is a minimal vertex cover of $\mathcal{H}'$. If adding the three vertices z, z', and w to this set would form a vertex cover T' of $\mathcal{H}$, we would again have a contradiction, so we may assume that some edge $e' \in E(\mathcal{H})$ avoids T'. Its V_1-vertex must be a superfluous W-vertex, since all other V_1-vertices are in T'. Its V_3-vertex must be in $V(\mathcal{H}')$, since $w \in T'$ and any edge incident to X_3 intersects T' in $\{z, z'\}$. Its V_2-vertex cannot be in $V(\mathcal{H}')$, since otherwise e' would be an edge of $\mathcal{H}'$ and hence intersect T', and its V_2-vertex also cannot be in X_2, since all edges incident to X_2 intersect T' in $\{z, z'\}$. Therefore e' must go through y, so it is of the form $a'yb'$ for some superfluous vertex $a' \in W \cap V_1$ and some vertex $b' \in V(\mathcal{H}') \cap V_3$.

By part (4) of Corollary 3.3.6 of the monster lemma applied to $\mathcal{H}'$ with $a = a'$, $b = y'$, and $c = b'$, there is a matching of size $\nu(\mathcal{H}')$ in $\mathcal{H}'$ avoiding a', y', and b'. Combining this matching with the three disjoint edges $zy'x'$, $z'v'w$, and $a'yb'$ yields a matching of size $\nu(\mathcal{H}) + 1$, a contradiction.

Therefore, in all cases we have found a contradiction, and since we have assumed the negation of the statement of Lemma 3.6.2, we have proven the lemma. $\square$

3.7 The Proof of Theorem 1.1.2

Proof of Theorem 1.1.2. The proof is by induction. IH(0) holds trivially: Let $\mathcal{H}$ be a 3-partite 3-graph with $\nu(\mathcal{H}) = 0$. Then $\mathcal{H}$ has no edges, so $(\emptyset, \emptyset, V(\mathcal{H}))$ is a home-base partition of $\mathcal{H}$ as can easily be seen. Now assume IH$(k-1)$ holds. We will show IH(k).

Let $\mathcal{H}$ be a 3-partite 3-graph with vertex classes V_1, V_2, and V_3 such that $\tau(\mathcal{H}) = 2\nu(\mathcal{H}) = 2k$. If it has a cromulent triple, then by Corollary 3.4.4, it is a home-base hypergraph, and we are done.

Therefore, assume there is no cromulent triple. Then by Lemma 3.6.2 there is a minimal equineighbored $X \subseteq V_3$ such that for $Y = N_{\mathrm{lk}_{\mathcal{H}}(V_1)}(X)$ we also have $N_{\mathrm{lk}_{\mathcal{H}}(V_1)}(Y) = X$. By Proposition 3.6.1, the edges incident to X form a truncated Fano plane F. Let A be the set of V_1-vertices of the hyperedges of F. Set $\mathcal{H}_1 = \mathcal{H} \setminus A$. Since we have removed two vertices, we have $\tau(\mathcal{H}_1) \geq \tau(\mathcal{H}) - 2$, and since any matching in $\mathcal{H}_1$ can be enlarged by adding an edge of F (as no edge of $\mathcal{H}_1$ is incident to X or Y), we have $\nu(\mathcal{H}_1) \leq \nu(\mathcal{H}) - 1$. Combining these inequalities with the fact that $\tau(\mathcal{H}_1) \leq 2\nu(\mathcal{H}_1)$ yields that all three inequalities are actually equalities. Since X and Y consist of isolated vertices, the same holds true for $\mathcal{H}_0 = \mathcal{H}_1 \setminus (Y \cup X)$. Thus, we can apply induction to get a home-base partition of $\mathcal{H}_0$ and add the F to it to get a proper matchable FR-partition of $\mathcal{H}$, which by Lemma 3.3.9 is a home-base partition.

Thus in all cases, $\mathcal{H}$ is a home-base hypergraph, so IH(k) holds.

Therefore Theorem 1.1.2 holds by induction. □

For interest, we can directly show also that IH(1) holds.

Proposition 3.7.1. *Let $\mathcal{H}$ be a 3-partite 3-graph with $\nu(\mathcal{H}) = 1$ and $\tau(\mathcal{H}) = 2$. Then $\mathcal{H}$ is a home-base hypergraph.*

Proof. Suppose $\mathcal{H}$ is an intersecting 3-partite 3-graph with $\tau(\mathcal{H}) = 2$. If every pair of edges intersect in two vertices, then it is easy to see that there must then be two vertices which are in every edge, and thus $\mathcal{H}$ would in fact have a vertex cover of size 1 (pick any one of the two vertices). Therefore there must be two edges which intersect in one vertex. Label these edges abc and ade. Since a alone does not form a vertex cover, there must be an edge which misses a, but it must intersect both of these edges, each in a different vertex class of $\mathcal{H}$. Thus WLOG, we have the edge fbe. If fdc is also an edge of $\mathcal{H}$, then we have an F. In this case, no further edge can be present unless it is parallel to one of the existing edges, since no other edge can intersect all four of these edges. Therefore in this case, $\mathcal{H}$ is indeed a home-base hypergraph which consists of a single F.

If fdc is not an edge of $\mathcal{H}$, then we let $R = \{a, b, e\}$, and we claim that every edge of $\mathcal{H}$ contains at least two of the vertices a, b, or e. If an edge misses any two of these vertices, then its third vertex must be the vertex outside of R of the edge among abc, ade, and fbe that contains those two vertices (since $\mathcal{H}$ is intersecting). Since this vertex is not in R either, by symmetry the same is true of each of the other edges we have given. Thus the edge must in fact be fdc, which is not the case by assumption. Thus $(\emptyset, \{R\}, V(\mathcal{H}) \setminus R)$ forms an FR-partition of $\mathcal{H}$ with the edge-home property. It is matchable because the graphs B_1, B_2, and B_3 contain edges Rf, Rd, and Rc, respectively, which obviously form matchings saturating $\{R\}$. Thus in this case, $\mathcal{H}$ is a home-base hypergraph consisting of a single R and at least three W-vertices. This proves the case $\nu(\mathcal{H}) = 1$. □

3.8 Concluding Remarks and Open Questions

3.8.1 Proof of the Reverse Implication for Theorem 2.4.3

As promised, we prove here the "if" direction of Theorem 2.4.3.

Proof of Theorem 2.4.3, (⇐). Let G be a bipartite graph with a collection of $\nu(G)/2$ pairwise vertex-disjoint subgraphs, each of them a C_4 or a P_4, such that every edge of G is either an edge of one of the C_4's or is incident to an interior vertex of one of the P_4's. We will construct a home-base hypergraph $\mathcal{H}$ with G as one of its links.

Let V_1 and V_2 be the vertex classes of G. Let V_3 be a set of sufficiently many new vertices ($\nu(G)$ suffice). Let $\mathcal{H}$ be the empty 3-graph. Then $(\mathcal{F}, \mathcal{R}, W) = (\emptyset, \emptyset, \emptyset)$ is a home-base partition of $\mathcal{H}$. We will add edges to $\mathcal{H}$, maintaining a home-base partition $(\mathcal{F}, \mathcal{R}, W)$.

For each C_4 in the collection we do the following. Let $\{a, b, c, d\}$ be the vertices of the C_4, so that $a, c \in V_1$, $b, d \in V_2$, and $ab, bc, cd, da \in E(G)$. Take two unused vertices $e, f \in V_3 \setminus V(\mathcal{H})$, and add the edges abe, adf, cbf, and cde to $\mathcal{H}$. These edges form a truncated Fano plane. For each edge parallel to an edge of the C_4, add an edge parallel to the corresponding one of these edges to $\mathcal{H}$, forming a truncated multi-Fano plane. We can then add the set $F = \{a, b, c, d, e, f\}$ to $\mathcal{F}$, maintaining that $(\mathcal{F}, \mathcal{R}, W)$ is a home-base partition of $\mathcal{H}$. Clearly, the C_4 is now present in the link $\mathrm{lk}_{\mathcal{H}}(V_3)$ together with all its parallel edges.

Then, for each P_4 in the collection we do the following. Let $\{a, b, c, d\}$ be the vertices of the P_4, so that $a, c \in V_1$, $b, d \in V_2$, and $ab, bc, cd \in E(G)$. Take two unused vertices $e, f \in V_3 \setminus V(\mathcal{H})$, and add the edges abe, cbf, and cde to $\mathcal{H}$. For each edge parallel to an edge of the P_4, add an edge parallel to the corresponding one of these edges to $\mathcal{H}$. Add the set $R = \{b, c, e\}$ to $\mathcal{R}$, and add the vertices a, d, and f to W. The edges abe, cbf, and cde are R-edges with a W-vertex in V_1, V_3, and V_2, respectively. Thus a, d, and f can be matched to R in B_1, B_3, and B_2, respectively, without disturbing matchability, since the W-vertices are new. Clearly the P_4 is now present in the link $\mathrm{lk}_{\mathcal{H}}(V_3)$ along with all parallel edges, and note especially that its interior vertices are members of R.

Once we've processed all the C_4's and P_4's, any edges of G not yet present in the link $\mathrm{lk}_{\mathcal{H}}(V_3)$ are incident to an interior vertex of one of the P_4's. Let $xy \in E(G)$ be such an edge, and suppose y is an interior vertex of one of the P_4's. Then $y \in R$ for some $R \in \mathcal{R}$. Let $z \in R \cap V_3$. Then, we add the edge xyz to $\mathcal{H}$. If x was not previously a vertex of $\mathcal{H}$, we add it to W, otherwise, we leave it where it is. Since xyz is an R-edge, $\mathcal{H}$ is still a home-base hypergraph with home-base partition $(\mathcal{F}, \mathcal{R}, W)$. After this addition, xy is present in the link $\mathrm{lk}_{\mathcal{H}}(V_3)$. We process every remaining edge this way.

If G has any isolated vertices, we add them to $\mathcal{H}$, putting them in W (these clearly do not disturb the home-base partition of $\mathcal{H}$). Now $\mathcal{H}$ is a home-base hypergraph with $\mathrm{lk}_{\mathcal{H}}(V_3) = G$. We know $\mathcal{H}$ satisfies $\tau(\mathcal{H}) = 2\nu(\mathcal{H})$ by Proposition 3.1.5, and hence by equation (3.2.1) we have $\mathrm{conn}(L(G)) = \frac{\nu(G)}{2} - 2$, as desired. $\square$

3.8.2 The Connectedness of the Line Graphs of Home-Base Hypergraphs

For 3-graphs $\mathcal{H}$, Theorem 2.1.1 gives

$$\mathrm{conn}(L(\mathcal{H})) \geq \frac{\nu(\mathcal{H})}{3} - 2.$$

Using our characterization, we can show that the Ryser-extremal 3-graphs are far from tight for this theorem. For a Ryser-extremal 3-partite 3-graph we can improve the bound to the following:

Proposition 3.8.1. *If $\mathcal{H}$ is a home-base hypergraph, then*

$$\mathrm{conn}(L(\mathcal{H})) \geq \frac{2}{3}\nu(\mathcal{H}) - 2.$$

Proof. Let $\mathcal{H}$ be a home-base hypergraph with vertex classes V_1, V_2, and V_3, and let $(\mathcal{F}, \mathcal{R}, W)$ be the home-base partition of $\mathcal{H}$. For each auxiliary bipartite graph B_i, let M_i be a matching saturating $\mathcal{R}$. For each $R \in \mathcal{R}$, let R^+ be the three edges corresponding to it in the respective matchings M_i. For an edge $e \in E(\mathcal{H})$, let $\mathrm{home}(e)$ be the member of $\mathcal{F} \cup \mathcal{R}$ where e is at home. Call an edge $e \in E(\mathcal{H})$ *crossing* if $\mathrm{home}(e) \in \mathcal{R}$ but $e \notin \mathrm{home}(e)^+$, and call it a *home-edge* otherwise. We will prove a slightly stronger statement, so that we can use induction.

Claim. *Let $k \in \mathbb{N}$, and let $J \subseteq L(\mathcal{H})$ such that $V(J)$ contains all the home-edges of at least k members of $\mathcal{F} \cup \mathcal{R}$. Then*

$$\mathrm{conn}(J) \geq \frac{2k}{3} - 2.$$

Proof of claim. We prove this by induction on $|E(J)|$. If no home-edge is adjacent in J to any crossing edge, then J contains at least k connected components, and so $\mathrm{conn}(J) \geq k-2$, since in this case $\mathcal{I}(J)$ is the join of at least k complexes that are (-1)-connected.

Thus, we may assume we have a crossing edge e which is J-adjacent to a home-edge f. We know e is J-adjacent to home-edges of at most two members of $\mathcal{F} \cup \mathcal{R}$. If e is not J-adjacent to home-edges of both, or one of those members is not among the k, then we are done, since by induction $\mathrm{conn}(J - ef) \geq 2k/3 - 2$ and $\mathrm{conn}(J \divideontimes ef) \geq 2(k-1)/3 - 2 > 2k/3 - 3$ (because $J \divideontimes ef$ contains all the home-edges of the $k-1$ members of $\mathcal{F} \cup \mathcal{R}$ which J contained, except $\mathrm{home}(f)$), and thus by Theorem 2.1.5, $\mathrm{conn}(J) \geq 2k/3 - 2$.

Thus assume e is J-adjacent to a home-edge f with $\mathrm{home}(e) = \mathrm{home}(f)$. Again, by induction we have $\mathrm{conn}(J - ef) \geq 2k/3 - 2$, so we just need $\mathrm{conn}(J \divideontimes ef) \geq 2k/3 - 3$ in order to be able to finish the proof using Theorem 2.1.5.

We can assume e is also J-adjacent to a home-edge f' with $\text{home}(f') \neq \text{home}(e)$. Unfortunately, since $J \divideontimes ef$ therefore does not contain all the home-edges of at least $k-1$ members of $\mathcal{F} \cup \mathcal{R}$, we cannot directly use induction to show the bound we need. Consider the other home-edges with home $\text{home}(f')$ remaining in $J' = J \divideontimes ef$. We know this set is non-empty because there is at least one $\text{home}(f')$-edge disjoint from e.

If there is a crossing edge e_1 which is J'-adjacent to a home-edge f_1 with $\text{home}(f_1) = \text{home}(f')$, then by induction

$$\text{conn}(J' \divideontimes e_1 f_1) \geq \frac{2(k-3)}{3} - 2 = \frac{2k}{3} - 4,$$

since $J' \divideontimes e_1 f_1$ still contains all the home-edges of the members of $\mathcal{F} \cup \mathcal{R}$ which J contains except those of $\text{home}(e)$, $\text{home}(f')$, and $\text{home}(e_1)$ (so still at least $k-3$). Therefore we only need $\text{conn}(J' - e_1 f_1) \geq 2k/3 - 3$ in order to be able to apply Thoerem 2.1.5. We can show this holds by iteratively deleting all of the adjacencies between home-edges of $\text{home}(f')$ and crossing edges so that we get a sequence $e_1, f_1, \ldots, e_r, f_r$, where the e_i are crossing edges, the f_j are home-edges of $\text{home}(f')$, and e_i is J'-adjacent to f_i for all i. Then we know by induction that

$$\text{conn}(J' - e_1 f_1 - \cdots - e_i f_i \divideontimes e_{i+1} f_{i+1}) \geq \frac{2(k-3)}{3} - 2 = \frac{2k}{3} - 4$$

for every $i < r$. We claim that $\text{conn}(J' - e_1 f_1 - \cdots - e_r f_r) \geq 2k/3 - 3$, since the home-edges of $\text{home}(f')$ are separated from the rest of the graph, so that $\mathcal{I}(J' - e_1 f_1 - \cdots - e_r f_r)$ is the join of the independence complex of those edges with the independence complex of the rest of J. Since the rest of J has connectedness $2(k-3)/3 - 2 = 2k/3 - 4$ by induction, and since the join with a non-empty complex increases the connectedness by at least one, we have $\text{conn}(J' - e_1 f_1 - \cdots - e_r f_r) \geq 2k/3 - 3$ as promised.

With this, we see that $\text{conn}(J' - e_1 f_1 - \cdots - e_i f_i) \geq 2k/3 - 3$ for every i by Theorem 2.1.5, and so $\text{conn}(J' - e_1 f_1) \geq 2k/3 - 3$ as desired. Then by Theorem 2.1.5, $\text{conn}(J \divideontimes ef) \geq 2k/3 - 3$, and thus again by Theorem 2.1.5, we have $\text{conn}(J) \geq 2k/3 - 2$. □

Therefore, since for the whole line graph we have all of the home-edges of $\nu(\mathcal{H})$ members of $\mathcal{F} \cup \mathcal{R}$, the inequality falls out of the claim. □

It is also not difficult to show that this bound is tight. For instance, disjoint copies of the following home-base hypergraph give tight examples:

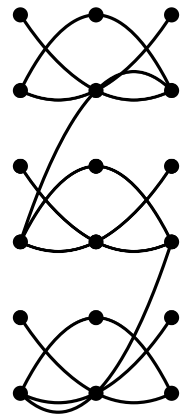

Figure 3.6: A 3-partite 3-graph $\mathcal{H}$ with $\tau(\mathcal{H}) = 6$, $\nu(\mathcal{H}) = 3$, and $\mathrm{conn}(L(\mathcal{H})) = 0$.

Since Proposition 3.8.1 is a strengthening of Theorem 2.1.1 when $\tau(\mathcal{H}) = 2\nu(\mathcal{H})$, one could ask for the best possible extension of it when the ratio τ/ν is different from 2. To make this precise, let us define the function $f : [1,2] \to \mathbb{R}$ by

$$f(x) = \inf\left\{\frac{\mathrm{conn}(L(\mathcal{H}))+2}{\nu(\mathcal{H})} : \mathcal{H} \text{ is a 3-partite 3-graph}, \tau(\mathcal{H}) \geq x\nu(\mathcal{H})\right\}.$$

We then have that for any 3-partite 3-graph $\mathcal{H}$ with $\tau(\mathcal{H}) = x\nu(\mathcal{H})$ it holds that

$$\mathrm{conn}(L(\mathcal{H})) \geq f(x)\nu(\mathcal{H}) - 2.$$

Clearly f is monotone increasing and bounded below by $1/3$, by Theorem 2.1.1. Since Proposition 3.8.1 is tight, we have $f(2) = 2/3$, while there are easy examples showing $f(1) = 1/3$. One could speculate whether there is a linear lower bound on f interpolating these two extremes, so that $f(x) \geq x/3$. This would be very interesting, as it would imply Ryser's Conjecture for 4-partite 4-graphs by a straightforward generalization of Aharoni's argument for 3-partite 3-graphs. Unfortunately, this does not turn out to be the case, as there is a violation of this bound for $x = 4/3$, as we'll see in detail in the next chapter:

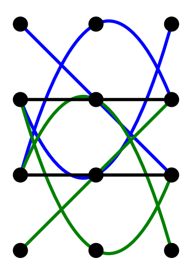

Figure 3.7: The 3-partite 3-graph $\mathcal{F}_4^{(3)}$.

The 3-partite 3-graph $\mathcal{F}_4^{(3)}$, pictured above, has $\tau(\mathcal{F}_4^{(3)}) = 4$, $\nu(\mathcal{F}_4^{(3)}) = 3$, and $\mathrm{conn}(L(\mathcal{F}_4^{(3)})) = -1$. This shows that $f(x) = 1/3$ for $x \in [1, 4/3]$. It can also be shown that $f(x) \geq x/5$ for every $x \in [1, 2]$, but this only represents an improvement when $x \in (\frac{5}{3}, 2)$ (see Chapter 4). We conjecture that $f(x) \geq x/4$ for every $x \in [1, 2]$.

To approach Ryser's Conjecture for 4-graphs, we seem to need a much better understanding of the potential link 3-graphs, in particular those with $\tau(\mathcal{H}) > \nu(\mathcal{H})$. We believe the function f will be a useful tool for this purpose, even though the extension of Aharoni's argument, at least in its most straightforward version, does not succeed due to the fact that $f(4/3) = 1/3$.

Chapter 4

Tau and the Connectedness of Line Graphs

4.1 Introduction

Combinatorial lower bounds on the connectedness of independence complexes can be quite useful. For instance, recall the following bound from Chapter 2, which relates the connectedness of line graphs to the matching number of their underlying hypergraphs:

Theorem 2.1.1. *If $\mathcal{H}$ is an r-graph, then*

$$\text{conn}(L(\mathcal{H})) \geq \frac{\nu(\mathcal{H})}{r} - 2.$$

This bound was sufficient to prove Ryser's Conjecture for 3-partite 3-graphs, and it played an integral role in our characterization of Ryser-extremal 3-graphs in the previous two chapters. The aim of this chapter is to investigate bounds on the connectedness of line graphs in terms of the vertex cover number of their underlying hypergraphs. One such bound is the following one for general hypergraphs, which may be proven via Meshulam's Theorem (Theorem 2.1.5):

Theorem 4.1.1. *If $\mathcal{H}$ is an r-graph, then*

$$\text{conn}(L(\mathcal{H})) \geq \frac{\tau(\mathcal{H})}{2r-1} - 2.$$

This bound is tight for general r-graphs, but the extremal hypergraphs we know of are not r-partite, so there is hope for an improved bound for r-partite r-graphs. In this light, we offer the following conjecture:

Conjecture 2. *If $\mathcal{H}$ is an r-partite r-graph, then*

$$\text{conn}(L(\mathcal{H})) \geq \frac{\tau(\mathcal{H})}{2r-2} - 2.$$

This conjecture, if true, would be tight. The main result of this chapter is that Conjecture 2 holds for 3-partite 3-graphs with vertex cover number at most 12.

4.2 Theorem 4.1.1 and its Tightness

As mentioned, we will use Meshulam's Theorem to prove Theorem 4.1.1, so we quote it here for convenience.

Theorem 2.1.5. *Let G be a graph and let $e \in E(G)$. Then we have*

$$\operatorname{conn}(G) \geq \min\{\operatorname{conn}(G - e), \operatorname{conn}(G \divideontimes e) + 1\}.$$

If $J \subseteq L(\mathcal{H})$ is a subgraph of the line graph, let $\mathcal{H}_J \subseteq \mathcal{H}$ denote the subhypergraph whose vertices are the vertices of $\mathcal{H}$, and whose edges are the vertices of J.

Theorem 4.1.1 is a special case of the following more general theorem:

Theorem 4.2.1. *If $J \subseteq L(\mathcal{H})$ is a subgraph of the line graph of an r-graph $\mathcal{H}$, then*

$$\operatorname{conn}(J) \geq \frac{\tau(\mathcal{H}_J)}{2r-1} - 2.$$

Proof. We prove this by induction on $|E(J)|$. Let $J \subseteq L(\mathcal{H})$ be a subgraph of the line graph. If J is empty, then $\tau(\mathcal{H}_J) = 0$, so the bound we want to prove is $\operatorname{conn}(J) \geq -2$, which is always true. If J is not empty but has no edges, then $\operatorname{conn}(J) = \infty$, so the bound is satisfied.

Otherwise, assume J has an edge ef, where $e, f \in E(\mathcal{H})$. Since $\mathcal{H}_J = \mathcal{H}_{J-ef}$, by induction we have

$$\operatorname{conn}(J - ef) \geq \frac{\tau(\mathcal{H}_{J-ef})}{2r-1} - 2 = \frac{\tau(\mathcal{H}_J)}{2r-1} - 2.$$

We will also need a bound on $\operatorname{conn}(J \divideontimes ef)$. Taking a minimum vertex cover of $\mathcal{H}_{J \divideontimes ef}$ plus the vertices in e and f forms a vertex cover of $\mathcal{H}_J$, since all of the edges removed by exploding ef intersect e or f because they are neighbors of one of these edges in the line graph of $\mathcal{H}$. Since e and f must intersect by virtue of the fact that $ef \in E(L(\mathcal{H}))$, we have $|e \cup f| \leq 2r - 1$, so we have

$$\tau(\mathcal{H}_{J \divideontimes ef}) + 2r - 1 \geq \tau(\mathcal{H}_{J \divideontimes ef}) + |e \cup f| \geq \tau(\mathcal{H}_J),$$

which we may rearrange to get $\tau(\mathcal{H}_{J \divideontimes ef}) \geq \tau(\mathcal{H}_J) - 2r + 1$. By induction, we then have

$$\operatorname{conn}(J \divideontimes ef) \geq \frac{\tau(\mathcal{H}_{J \divideontimes ef})}{2r-1} - 2 \geq \frac{\tau(\mathcal{H}_J) - 2r + 1}{2r-1} - 2 = \frac{\tau(\mathcal{H}_J)}{2r-1} - 3.$$

Therefore, by Meshulam's Theorem (Theorem 2.1.5), we have

$$\operatorname{conn}(J) \geq \min(\operatorname{conn}(J - ef), \operatorname{conn}(J \divideontimes ef) + 1) \geq \frac{\tau(\mathcal{H}_J)}{2r-1} - 2.$$

Thus by induction, the theorem holds. □

This immediately implies Theorem 4.1.1, since $\mathcal{H}_{L(\mathcal{H})} = \mathcal{H}$.

We note that together with Theorem 2.1.7, this theorem implies an old theorem of Haxell[13, Theorem 3]:

Theorem 4.2.2 (Haxell)**.** *Let $\mathcal{H}$ be an r-graph whose vertices are partitioned into two sets A and B, such that every edge of $\mathcal{H}$ has exactly one vertex from A. If for every subset $S \subseteq A$ the $(r-1)$-graph $\mathcal{H}_S$ on B with edges $\{e \subseteq B : e \cup s \in E(\mathcal{H})$ for some $s \in S\}$ satisfies $\tau(\mathcal{H}_S) > (2r-3)(|S|-1)$, then $\nu(\mathcal{H}) = |A|$.*

Proof. We plan to apply Theorem 2.1.7 with $d = 0$ to $\mathcal{I}(\mathcal{H}_A)$. Indeed, if we color the edges of $\mathcal{H}_A$ according to which member of A they extend, a rainbow matching in $\mathcal{H}_A$ corresponds to a matching in $\mathcal{H}$. This induces a coloring on the vertices of $\mathcal{I}(L(\mathcal{H}_A))$ that we claim satisfies the conditions of Thoerem 2.1.7 for $d = 0$. Clearly, for any subset $S \subseteq A$, we have $\mathcal{I}(L(\mathcal{H}_A))|_S = \mathcal{I}(L(\mathcal{H}_S))$, and since $\tau(\mathcal{H}_S) > (2r-3)(|S|-1)$ by assumption, we have by Theorem 4.1.1 that

$$\operatorname{conn}(L(\mathcal{H}_S)) > \frac{(2r-3)(|S|-1)}{2r-3} - 2 = |S| - 3,$$

as $\mathcal{H}_S$ is an $(r-1)$-graph. Since the connectedness is an integer, this implies $\operatorname{conn}(L(\mathcal{H}_S)) \geq |S|-2$, which is the condition of Theorem 2.1.7 for $d = 0$. Thus, $\mathcal{H}_A$ has a rainbow matching of size $|A|$, so $\mathcal{H}$ has a matching of size $|A|$, and since A is a cover for $\mathcal{H}$, there clearly is no larger matching, meaning $\nu(\mathcal{H}) = |A|$, as promised. □

Next, we will show that Theorem 4.1.1 is tight. To be precise, we prove the following:

Proposition 4.2.3. *For every integer $r \geq 2$, and every integer $k \geq 0$, there is an r-graph $\mathcal{H}$ with $\tau(\mathcal{H}) = k$ and*

$$\operatorname{conn}(L(\mathcal{H})) = \left\lceil \frac{\tau(\mathcal{H})}{2r-1} \right\rceil - 2.$$

To show this we will need an easy lemma about the connectedness of joins, which is an easy consequence of Proposition 2.2.1 and the Künneth formula for joins [23]:

Lemma 4.2.4. *If $X_1, \ldots, X_n$ are topological spaces with $\operatorname{conn}(X_i) = -1$ for all $i = 1, \ldots, n$, then*

$$\operatorname{conn}(X_1 * \cdots * X_n) = n - 2.$$

Armed with this lemma, we need only show tightness for $\tau(\mathcal{H}) \leq 2r-1$, and then we can build larger tight examples out of disjoint unions.

For $k = 1, \ldots, 2r-1$, we define the r-graph $\mathcal{G}_k^{(r)}$ to have vertex set $V(\mathcal{G}_k^{(r)}) = [r]^2$ and the following edge set:

- If $k \le r$, we set $E(\mathcal{G}_k^{(r)}) = \{\{(i,j) : j \in [r]\} : i \in [k]\} \cup \{\{(i,1) : i \in [r]\}\}$.
- If $k \ge r+1$, we set

$$\begin{aligned} E(\mathcal{G}_k^{(r)}) = &\{\{(i,j) : j \in [r]\} : i \in [r]\} \\ &\cup \{\{(i,j) : i \in [r]\} : j \in [k-r+1]\} \\ &\cup \{\{(i,\sigma(i)) : i \in [r]\} : \sigma \in S_r\}, \end{aligned}$$

  where S_r denotes the set of permutations of $[r]$.

In words, the vertices of $\mathcal{G}_k^{(r)}$ form an $r \times r$ grid; for $k \le r$, the edges consist of k rows and one column, while for $k \ge r+1$, the edges consist of r rows, $k-r+1$ columns, and all transversals.

In order to clear up any confusion in visualizing these hypergraphs, we note here that our coordinates are laid out in the style of matrix indices, so that (i,j) is the vertex in row i and column j, where the rows are numbered from top to bottom, and the columns from left to right.

Proposition 4.2.5. *For $k = 1, \dots, 2r-1$ we have* $\operatorname{conn}(L(\mathcal{G}_k^{(r)})) = -1$.

Proof. We claim that the set of columns (edges of the form $\{(i,j) : i \in [r]\}$ for some j) forms a connected component of $\mathcal{I}(L(\mathcal{G}_k^{(r)}))$. Clearly, the columns are all disjoint, so they form a simplex in $\mathcal{I}(L(\mathcal{G}_k^{(r)}))$, and every other edge intersects all of the columns, so no other edge is in a simplex with a column. Hence they form a path component. Since there is always at least one edge that is not a column, $\mathcal{I}(L(\mathcal{G}_k^{(r)}))$ has at least two components, showing $\operatorname{conn}(L(\mathcal{G}_k^{(r)})) = -1$, as desired. □

We now show that $\tau(\mathcal{G}_k^{(r)}) = k$ for $k = 1, \dots 2r-1$. This is easy for $k \le r$, since for these, $\mathcal{G}_k^{(r)}$ has a matching of size k and clearly has a vertex cover of size k as well. For $k \ge r+1$, things get trickier. We start with a lemma:

Lemma 4.2.6. $\tau(\mathcal{G}_{2r-1}^{(r)}) = 2r-1$.

Proof. Clearly, taking any row together with any column forms a vertex cover of $\mathcal{G}_{2r-1}^{(r)}$ of size $2r-1$, so we have $\tau(\mathcal{G}_{2r-1}^{(r)}) \le 2r-1$. It remains only to see that $\tau(\mathcal{G}_{2r-1}^{(r)}) > 2r-2$. Suppose $T \subseteq [r]^2$ is a vertex cover of size $2r-2$. Rearranging the columns does not change $\mathcal{G}_{2r-1}^{(r)}$, so we may assume that the columns are sorted so that the number of elements of T in each column is monotone decreasing. For $j = 1, \dots, r$, let t_j be the number of elements of T in column j, so that $t_1 \ge \cdots \ge t_r$. Now in order to be a vertex cover, there must be an element in each column, so $t_j \ge 1$ for all j. This means that for each k we have

$$|T| = \sum_{j=1}^{r} t_j \ge k \cdot t_k + (r-k) \cdot 1.$$

Since $|T| = 2r - 2$, we get

$$t_k \le \frac{r-2}{k} + 1.$$

We claim that for $1 \le k \le r-1$, we have $t_k \le r-k$. Indeed, for such k we have that $(k-1)(k-r+1) \le 0$, hence by rearranging we get

$$r - k \ge \frac{r-1}{k}.$$

From this it follows that

$$t_k \le \frac{r-2}{k} + 1 < \frac{r-1}{k} + 1 \le r - k + 1,$$

and since t_k is an integer, this means that $t_k \le r - k$. This means that for $k = 1, \ldots, r-1$, we have at least k elements in column k that are not in T, which makes it easy to choose a partial transversal S avoiding T containing one element from each of the first $r-1$ columns: In the first column, there is an element v_1 not in T, and in general in the k-th column there is an element v_k not in T, which is in a different row from $v_1, \ldots, v_{k-1}$. We can thus find $S = \{v_1, \ldots, v_{r-1}\}$ avoiding T with one element from each of the first $r-1$ columns and one element from each of $r-1$ different rows. Consider now the single row i that does not intersect S. If T does not contain the last element in row i, then S together with that element would form a transversal avoiding T. Thus we may assume that $(i,r) \in T$. Then this is the only element of the last column in T, since $t_r < 2$. There must be an element of (i,j) of row i that is not in T, otherwise T must miss some other row, since $|T| = 2r-2$. If $v_j = (\ell, j)$, then $S \setminus \{v_j\} \cup \{(i,j), (\ell, r)\}$ is a transversal avoiding T, concluding the proof that $\tau(\mathcal{G}_{2r-1}^{(r)}) > 2r-2$. Therefore $\tau(\mathcal{G}_{2r-1}^{(r)}) = 2r-1$, as claimed. □

This will allow us to find $\tau(\mathcal{G}_k^{(r)})$.

Proposition 4.2.7. *For $k = 1, \ldots, 2r-1$, we have $\tau(\mathcal{G}_k^{(r)}) = k$.*

Proof. For $k = 1, \ldots, r$, we have already seen that $\tau(\mathcal{G}_k^{(r)}) = k$. For $k = r+1, \ldots, 2r-1$, we must show that there is a vertex cover of size k and no vertex cover of size $k-1$.

Let $k \in \{r+1, \ldots, 2r-1\}$. Then there is a vertex cover T of $\mathcal{G}_k^{(r)}$ given by

$$T = \{(i,1) : i \in [r]\} \cup \{(1,j) : j \in [k-r+1]\}.$$

Indeed, every transversal and every row intersects $\{(i,1) : i \in [r]\}$, and every column in $\mathcal{G}_k^{(r)}$ intersects $\{(1,j) : j \in [k-r+1]\}$. It is clear that $|T| = k$, so $\tau(\mathcal{G}_k^{(r)}) \le k$.

Now let $T \subseteq [r]^2$ with $|T| = k-1$. We claim that T is not a vertex cover. Indeed, if it were, then T would cover every edge of $\mathcal{G}_{2r-1}^{(r)}$ except possibly the last $2r-1-k$ columns. Adding one vertex of each of these columns to T yields a set of size $k-1+2r-1-k=2r-2$, which would cover every edge of $\mathcal{G}^{(r)}_{2r-1}$, a contradiction to Lemma 4.2.6. Thus T is not a vertex cover, so $\tau(\mathcal{G}^{(r)}_k) \geq k$. This completes the proof. □

Now we are ready to prove Proposition 4.2.3.

Proof of Proposition 4.2.3. If $k=0$, the empty r-graph will do. Otherwise, let $k = q(2r-1)+p$ with p and q integers, such that $1 \leq p \leq 2r-1$. We construct $\mathcal{H}$ as follows:

For $i = 1, \ldots, q$, let $\mathcal{H}_i$ be a copy of $\mathcal{G}^{(r)}_{2r-1}$, and let $\mathcal{H}_{q+1}$ be a copy of $\mathcal{G}^{(r)}_p$. Then let $\mathcal{H}$ be the disjoint union of the r-graphs $\mathcal{H}_1, \ldots, \mathcal{H}_{q+1}$. We claim $\mathrm{conn}(L(\mathcal{H})) = q-1$. Indeed $\mathrm{conn}(L(\mathcal{H}_i)) = -1$ by Proposition 4.2.5. Since the independence complex of the disjoint union of graphs is the join of the independence complexes of the graphs, we get by Lemma 4.2.4 that

$$\mathrm{conn}(L(\mathcal{H})) = \mathrm{conn}(\mathcal{I}(L(\mathcal{H}_1)) * \cdots * \mathcal{I}(L(\mathcal{H}_{q+1}))) = q-1,$$

as promised. We have $\tau(\mathcal{H}) = \tau(\mathcal{H}_1) + \cdots + \tau(\mathcal{H}_{q+1}) = q(2r-1)+p = k$ by Proposition 4.2.7, and since $q-1 = \lceil k/(2r-1) \rceil - 2$, we have constructed the desired r-graph $\mathcal{H}$. □

4.3 Towards Conjecture 2

The tight examples we constructed to prove Proposition 4.2.3 are not r-partite, so it leaves room to hope that this bound could be strengthened for r-partite r-graphs, which leads us to Conjecture 2. Let us start by showing that it would be tight.

Proposition 4.3.1. *For every integer $r \geq 2$, and every integer $k \geq 0$, there is an r-partite r-graph $\mathcal{H}$ with $\tau(\mathcal{H}) = k$ and*

$$\mathrm{conn}(L(\mathcal{H})) = \left\lceil \frac{\tau(\mathcal{H})}{2r-2} \right\rceil - 2.$$

For $k = 1, \ldots, 2r-2$, we define the r-partite r-graph $\mathcal{F}^{(r)}_k$ as follows:

- If $k \leq r-1$, we set $V(\mathcal{F}^{(r)}_k) = ([k] \times [r]) \cup \{(i,i) : i \in [r]\}$ and $E(\mathcal{F}^{(r)}_k) = \{\{(i,j) : j \in [r]\} : i \in [k]\} \cup \{\{(i,i) : i \in [r]\}\}$.
- If $k \geq r$, we set $V(\mathcal{F}^{(r)}_k) = ([r-1] \cup S_{r-1}) \times [r]$ and

$$\begin{aligned} E(\mathcal{F}^{(r)}_k) = &\{\{(i,j) : j \in [r]\} : i \in [k-r+1]\} \\ &\cup \{e_{\sigma,j} : \sigma \in S_{r-1}, j \in [r]\}, \end{aligned}$$

  where S_{r-1} denotes the set of permutations of $[r-1]$, and where for a permutation $\sigma \in S_{r-1}$ and an integer $j \in [r]$, the edge $e_{\sigma,j}$ is given by

$$e_{\sigma,j} = \{(\sigma,j)\} \cup \{(\sigma(i), [i+j]_r) : i \in [r-1]\},$$

  where $[p]_r$ is the remainder of p mod r that belongs to $[r]$.

In words, for $k \leq r-1$, $\mathcal{F}_k^{(r)}$ consists of a matching of size k and one diagonal edge that crosses all edges of the matching, while for $k \geq r$, $\mathcal{F}_k^{(r)}$ lives on an $(r-1) \times r$ grid, plus $(r-1)!$ additional vertices in each vertex class, one for each permutation of $[r-1]$. The edges include $k-r+1$ rows of the grid, plus an edge for every vertex class (column) j and permutation σ, which restricts to the transversal corresponding to σ on the $(r-1) \times (r-1)$ subgrid obtained by removing the column j (and permuting the columns cyclically, so that the j'th would be at the end) and passes through the vertex corresponding to σ in column j (and being the only edge incident to that vertex).

We will prove that these hypergraphs combine to form tight examples for every k in the same fashion as we did for Proposition 4.2.3 in the previous section.

Proposition 4.3.2. *For $k = 1, \ldots, 2r-2$ we have* $\operatorname{conn}(L(\mathcal{F}_k^{(r)})) = -1$.

Proof. We claim that the set of rows (edges of the form $\{(i,j) : j \in [r]\}$ for some i) forms a connected component of $\mathcal{I}(L(\mathcal{F}_k^{(r)}))$. Clearly, the rows are all disjoint, so they form a simplex in $\mathcal{I}(L(\mathcal{F}_k^{(r)}))$, and every other edge intersects all of the rows, so no other edge is in a simplex with a row. Hence they form a path component. Since there is always at least one edge that is not a row, $\mathcal{I}(L(\mathcal{F}_k^{(r)}))$ has at least two components, showing $\operatorname{conn}(L(\mathcal{G}_k^{(r)})) = -1$, as desired. □

We now show that $\tau(\mathcal{F}_k^{(r)}) = k$ for $k = 1, \ldots, 2r-2$. This is easy for $k \leq r-1$, since for these, $\mathcal{F}_k^{(r)}$ has a matching of size k and clearly has a vertex cover of size k as well. For $k \geq r$, things get trickier. We start with a lemma:

Lemma 4.3.3. $\tau(\mathcal{F}_{2r-2}^{(r)}) = 2r-2$.

Proof. Clearly, taking any row together with one vertex from the remaining rows forms a vertex cover of $\mathcal{F}_{2r-2}^{(r)}$ of size $2r-2$, so we have $\tau(\mathcal{F}_{2r-2}^{(r)}) \leq 2r-2$. It remains only to see that $\tau(\mathcal{F}_{2r-2}^{(r)}) > 2r-3$. Suppose $T \subseteq V(\mathcal{F}_{2r-2}^{(r)})$ is a vertex cover of size $2r-3$. We may assume that $T \subseteq [r-1] \times [r]$, since if T contains any vertex outside this grid, that vertex only covers one edge, hence we may substitute it by any other vertex of that edge, and all edges would still be covered. Rearranging the rows does not change $\mathcal{F}_{2r-2}^{(r)}$, so we may assume that the rows are sorted so that the number of elements of T in each row is monotone decreasing. For $i = 1, \ldots, r-1$, let t_i be the number of elements of T in row i, so that $t_1 \geq \cdots \geq t_{r-1}$. Now in order to be a vertex cover, there must be an element in each row, so $t_i \geq 1$ for all i. This means that for each k we have

$$|T| = \sum_{i=1}^{r} t_i \geq k \cdot t_k + (r-1-k) \cdot 1.$$

Since $|T| = 2r-3$, we get

$$t_k \leq \frac{r-2}{k} + 1.$$

We claim that for $1 \leq k \leq r-1$, we have $t_k \leq r-k$. Indeed, for such k we have that $(k-1)(k-r+1) \leq 0$, hence by rearranging we get

$$r - k \geq \frac{r-1}{k}.$$

From this it follows that

$$t_k \leq \frac{r-2}{k} + 1 < \frac{r-1}{k} + 1 \leq r - k + 1,$$

and since t_k is an integer, this means that $t_k \leq r-k$. This means that for $k = 1, \ldots, r-1$, we have at least k elements in row k that are not in T, which makes it easy to choose a transversal S avoiding T containing one element from each row: In the first row, there is an element v_1 not in T, and in general in the k-th row there is an element v_k not in T, which is in a different column from $v_1, \ldots, v_{k-1}$. We can thus find $S = \{v_1, \ldots, v_{r-1}\}$ avoiding T with one element from each row and one element from each of $r-1$ different columns. Consider now the single column j that does not intersect S. There is an edge of $\mathcal{F}_{2r-2}^{(r)}$ consisting of the transversal S in the $(r-1) \times (r-1)$ subgrid obtained by removing column j, and whose vertex in column j corresponds to the permutation corresponding to S. This edge is disjoint from T, contradicting the fact that T is a vertex cover. Thus $\tau(\mathcal{F}_{2r-2}^{(r)}) > 2r-3$, and so $\tau(\mathcal{F}_{2r-2}^{(r)}) = 2r-2$, as claimed. □

This will allow us to find $\tau(\mathcal{F}_k^{(r)})$.

Proposition 4.3.4. *For $k = 1, \ldots 2r-2$, we have $\tau(\mathcal{F}_k^{(r)}) = k$.*

Proof. For $k = 1, \ldots, r-1$, we have already seen that $\tau(\mathcal{F}_k^{(r)}) = k$. For $k = r, \ldots, 2r-2$, we must show that there is a vertex cover of size k and no vertex cover of size $k-1$.

Let $k \in \{r, \ldots, 2r-2\}$. Then there is a vertex cover T of $\mathcal{F}_k^{(r)}$ given by

$$T = \{(i,1) : i \in [k-r+1]\} \cup \{(1,j) : j \in [r]\}.$$

Indeed, every edge $e_{\sigma,j}$ intersects $\{(1,j) : j \in [r]\}$, and every row in $\mathcal{F}_k^{(r)}$ intersects $\{(i,1) : i \in [k-r+1]\}$. It is clear that $|T| = k$, so $\tau(\mathcal{F}_k^{(r)}) \leq k$.

Now let $T \subseteq V(\mathcal{F}_k^{(r)})$ with $|T| = k-1$. We claim that T is not a vertex cover. Indeed, if it were, then T would cover every edge of $\mathcal{F}_{2r-2}^{(r)}$ except possibly that last $2r-2-k$ rows. Adding one vertex of each of these rows to T yields a set of size $k-1+2r-2-k = 2r-3$, which would cover every edge of $\mathcal{F}_{2r-2}^{(r)}$, a contradiction to Lemma 4.3.3. Thus T is not a vertex cover, so $\tau(\mathcal{F}_k^{(r)}) \geq k$. This completes the proof. □

Now we are ready to prove Proposition 4.3.1.

Proof of Proposition 4.3.1. If $k = 0$, the empty r-graph will do. Otherwise, let $k = q(2r-2)+p$ with p and q integers, such that $1 \le p \le 2r-2$. We construct $\mathcal{H}$ as follows:

For $i = 1, \dots, q$, let $\mathcal{H}_i$ be a copy of $\mathcal{F}^{(r)}_{2r-2}$, and let $\mathcal{H}_{q+1}$ be a copy of $\mathcal{F}^{(r)}_p$. Then let $\mathcal{H}$ be the disjoint union of the r-partite r-graphs $\mathcal{H}_1, \dots, \mathcal{H}_{q+1}$. We claim $\operatorname{conn}(L(\mathcal{H})) = q-1$. Indeed $\operatorname{conn}(L(\mathcal{H}_i)) = -1$ by Proposition 4.3.2. Since the independence complexes of the disjoint union of graphs is the join of the independence complexes of the graphs, we get by Lemma 4.2.4 that

$$\operatorname{conn}(L(\mathcal{H})) = \operatorname{conn}(\mathcal{I}(L(\mathcal{H}_1)) * \cdots * \mathcal{I}(L(\mathcal{H}_{q+1}))) = q-1,$$

as promised. We have $\tau(\mathcal{H}) = \tau(\mathcal{H}_1) + \cdots + \tau(\mathcal{H}_{q+1}) = q(2r-2)+p = k$ by Proposition 4.3.4, and since $q-1 = \lceil k/(2r-2)\rceil - 2$, we have constructed the desired r-graph $\mathcal{H}$. □

Theorem 2.1.1 shows that Conjecture 2 holds for $r = 2$, since in bipartite graphs, $\tau = \nu$ by König's Theorem. The goal of the rest of the section is to show that it holds for $r = 3$ when τ is small.

4.3.1 Conjecture 2 for $r = 3$

The first value for which Conjecture 2 offers an improvement over Theorem 4.1.1 is for $\tau = 5$, where the conjecture states that the independence complex is path-connected. We could show directly that this is the case, but in order to go further, we will characterize the tight examples for $\tau = 4$, which will imply the bound for $\tau = 5$.

Pictured below is the 3-partite 3-graph $\mathcal{F}^{(3)}_4$, which was used to show the tightness of Conjecture 2:

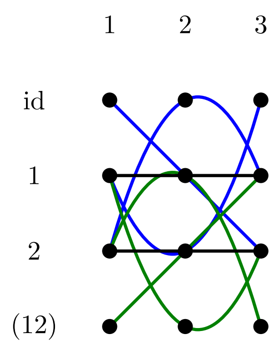

Figure 4.1: The 3-partite 3-graph $\mathcal{F}^{(3)}_4$.

We will call the two black horizontal edges (the edges $\{(1,1),(1,2),(1,3)\}$ and $\{(2,1),(2,2),(2,3)\}$) the *central edges* of $\mathcal{F}^{(3)}_4$. We then have the following characterization theorem, which states that $\mathcal{F}^{(3)}_4$ is the unique minimal tight example, and will lead us to easily be able to infer that Conjecture 2 is true for $r = 3$ when $\tau \le 8$:

Theorem 4.3.5. *If $\mathcal{H}$ is a 3-partite 3-graph and $J \subseteq L(\mathcal{H})$ is a subgraph of its line graph with $\tau(\mathcal{H}_J) \geq 4$ and $\operatorname{conn}(J) \leq -1$, then $\mathcal{H}_J$ contains a copy of $\mathcal{F}_4^{(3)}$, and every edge outside of that copy intersects both central edges of the copy.*

To prove this, we will need the stronger formulation of Theorem 2.1.1 as given in Chapter 2, as well as a similar formulation of Proposition 3.8.1 (which easily follows from the proof given in Chapter 3). We state them here for convenience, and for consistency of notation:

Lemma 4.3.6. *If $\mathcal{H}$ is an r-graph and $J \subseteq L(\mathcal{H})$ is a subgraph of its line graph, then*

$$\operatorname{conn}(J) \geq \frac{\nu(\mathcal{H}_J)}{r} - 2.$$

Lemma 4.3.7. *If $\mathcal{H}$ is a 3-partite 3-graph and $J \subseteq L(\mathcal{H})$ is a subgraph of its line graph with $\tau(\mathcal{H}_J) = 2\nu(\mathcal{H}_J)$, then*

$$\operatorname{conn}(J) \geq \frac{2\nu(\mathcal{H}_J)}{3} - 2.$$

With these lemmas in mind, we are ready to proceed with the proof.

Proof of Theorem 4.3.5. Let $\mathcal{H}$ be a 3-partite 3-graph with vertex classes V_1, V_2, and V_3; and let $J \subseteq L(\mathcal{H})$ be a subgraph of its line graph with $\tau(\mathcal{H}_J) \geq 4$ and $\operatorname{conn}(J) \leq -1$. We remark that $\tau(\mathcal{H}_J) \geq 4$ implies J is not empty, hence $\operatorname{conn}(J) > -2$, so we in fact have $\operatorname{conn}(J) = -1$.

As a first step, we show that $\nu(\mathcal{H}_J) = 3$. Suppose that $\nu(\mathcal{H}_J) \geq 4$. Then $\operatorname{conn}(J) \geq \lceil 4/3 \rceil - 2 = 0$ by Lemma 4.3.6, a contradiction. Now suppose that $\nu(\mathcal{H}_J) \leq 2$. Then since $\tau(\mathcal{H}_J) \geq 4$, and since Ryser's Conjecture holds for 3-graphs, we must have $\tau(\mathcal{H}_J) = 4$ and $\nu(\mathcal{H}_J) = 2$. But then Lemma 4.3.7 implies $\operatorname{conn}(J) \geq \lceil 4/3 \rceil - 2 = 0$, again a contradiction. Thus we have $\nu(\mathcal{H}_J) = 3$.

Fix a matching $M \subseteq V(J)$ of size 3. It forms a simplex in the independence complex $\mathcal{I}(J)$, so it is part of one connected component. By assumption, $\operatorname{conn}(J) = -1$, hence $\mathcal{I}(J)$ is not path-connected, and thus has more than one connected component. Therefore there must be some other component $C \subseteq V(J) \subseteq E(\mathcal{H})$ not containing M. Since there are no simplices of $\mathcal{I}(J)$ with vertices in both components, it must be that every element of C is adjacent in J to every element of M. Since J is a subgraph of $L(\mathcal{H})$, it follows that every edge (of $\mathcal{H}$) in C intersects every edge of M. Now consider the size of the largest matching among edges in C. If this is 1, then any edge of C forms a vertex cover of size 3, since every edge in C intersects it by assumption, and every edge outside of C intersects it because the edges must be adjacent in J, a subgraph of the line graph. Thus there is a matching of size at least 2 in C. If on the other hand there were a matching of size 3 in C, then because every one of these edges must intersect every edge of M, one in each vertex class, it follows that the V_1-vertices of these edges coincide with the V_1-vertices of the edges in M, and would form a vertex cover of size 3. This is because every edge of C must intersect every edge of M, at least one intersection occurring in V_1, and every edge outside C must intersect every edge of the supposed matching in C, again one intersection occurring in V_1. Therefore the largest matching that can be found among the edges of C is exactly 2. So let e and e' be two disjoint edges in C. These will be the central edges of the copy of $\mathcal{F}_4^{(3)}$ we are looking for.

We will now find an explicit isomorphism of a subhypergraph of $\mathcal{H}_J$ with $\mathcal{F}_4^{(3)}$. For $j = 1, 2, 3$, let $m_j \in M$ be the edge whose V_j-vertex is in neither of e and e'. Without loss of generality, we may assume m_1 intersects e in V_2 and e' in V_3 (otherwise exchange the labels e and e'). Then label the vertices of e by $(1,1)$, $(1,2)$, $(1,3)$, and the vertices of e' by $(2,1)$, $(2,2)$, $(2,3)$, so that every vertex (i,j) is in vertex class V_j. We now have that m_j intersects e in $(1, j+1)$ and e' in $(2, j+2)$ (all arithmetic is done modulo 3). Let $m_{j,j}$ denote the remaining vertex of m_j (the one not listed above).

Now for each vertex class V_j we apply the following procedure:

Consider the set T_j consisting of the V_j-vertices of the edges in M. Since this set is too small to be a vertex cover, there is an edge $g_j \in V(J)$ that avoids T_j. Now g_j cannot intersect every edge of M, so it must be in the same component as M, in particular it intersects both e and e', one in V_{j+1} and the other in V_{j+2}. If there is any edge g'_j avoiding T_j which intersects e in $(1, j+2)$, and e' in $(2, j+1)$, then label its V_j-vertex by $((12), j)$, and $m_{j,j}$ by (id, j), so that we have the edges $m_j = e_{\mathrm{id},j}$ and $g'_j = e_{(12),j}$, and we can proceed to the next vertex class. If on the other hand there is no such edge, then consider the set $U_j = \{(1,j), (2,j), (1,j+1)\}$, which is also too small to be a vertex cover. Thus there is an edge $h_j \in V(J)$ that avoids it. If h_j is in C, then it must intersect every edge of M, and the only possibilities avoiding U_j are $h_j = \{m_{j,j}, m_{j+1,j+1}, m_{j+2,j+2}\}$ and $h_j = \{m_{j,j}, (2, j+1), (1, j+2)\}$. The first of these possibilities is impossible, since it would mean C would contain a matching of size 3, which we have shown to be false. Thus the latter possibility holds in this case. If h_j is not in C, then it must intersect both e and e', and the only way to do this avoiding U_j is by containing $(2, j+1)$ and $(1, j+2)$. By assumption, the V_j-vertex of h_j must be $m_{j,j}$, hence in all cases we have the edge $\{m_{j,j}, (2, j+1), (1, j+2)\}$. In this case, we label $m_{j,j}$ by $((12), j)$, and the V_1-vertex of g_j by (id, j), so that we have the edges $g_j = e_{\mathrm{id},j}$ and $h_j = e_{(12),j}$, and we can proceed to the next vertex class.

After each vertex class is processed, we will have found a subhypergraph of $\mathcal{H}_J$ with an explicit isomorphism with $\mathcal{F}_4^{(3)}$. Now we must show that every edge of $\mathcal{H}_J$ outside of this copy intersects both e and e'. Every edge outside of the component C has to intersect both e and e', so the only edges we need to worry about are in the component C. The edges of C must cross every edge in the component of M, in particular, they must intersect $e_{\sigma,j}$ for every σ and every j, and a simple case analysis shows that e and e' are the only 3-sets that do this (consult Figure 4.1). Hence C must consist solely of e and e', and thus we have proven the claim. $\square$

As a corollary we have the following:

Corollary 4.3.8. *If $\mathcal{H}$ is a 3-partite 3-graph and $J \subseteq L(\mathcal{H})$ is a subgraph of its line graph with $\tau(\mathcal{H}_J) \geq 5$, then $\operatorname{conn}(J) \geq 0$.*

Proof. Let $\mathcal{H}$ be a 3-partite 3-graph, and let $J \subseteq L(\mathcal{H})$ be a subgraph of its line graph with $\tau(\mathcal{H}_J) \geq 5$. Now suppose that $\operatorname{conn}(J) \leq -1$. Then by Theorem 4.3.5, $\mathcal{H}_J$ contains a copy of $\mathcal{F}_4^{(3)}$ and every edge outside of it intersects both its central edges. Then the set T consisting of the vertices of one of the central edges and the V_1-vertex of the other is a vertex cover of $\mathcal{H}_J$ of size 4, a contradiction. Indeed, T is a vertex cover of the copy of $\mathcal{F}_4^{(3)}$, and since every outside edge intersects both central edges, these also intersect T. Thus, we must have $\operatorname{conn}(J) \geq 0$. $\square$

In particular, this shows that Conjecture 2 holds for $r = 3$ when $\tau = 5$. Now note that since the connectedness is an integer, Conjecture 2 is only stronger than Theorem 4.1.1 for $r = 3$ when $\tau \in \{5, 9, 10, 13, 14, 15\}$ or if $\tau \geq 17$. Therefore, Corollary 4.3.8 settles the conjecture up to $\tau = 8$. Our next task is to verify it for $\tau = 9$, which will prove it for $\tau \leq 12$:

Theorem 4.3.9. *If $\mathcal{H}$ is a 3-partite 3-graph and $J \subseteq L(\mathcal{H})$ is a subgraph of its line graph with $\tau(\mathcal{H}_J) \geq 9$, then $\operatorname{conn}(J) \geq 1$.*

Proof. Our proof is via contradiction. As it is quite involved, let us give a short overview to start. We suppose that we have a minimal counterexample, and aim to use Meshulam's Theorem together with Theorem 4.3.5 to find a contradiction. We find that the explosion of any edge of J results in a hypergraph satisfying the conditions of Theorem 4.3.5, hence contains a copy of $\mathcal{F}_4^{(3)}$. We make use of the high vertex cover number to prove the existence of various types of edges, which we show must intersect in certain vertices. We use these vertices and various covers to eventually construct a set of edges that cannot all intersect both central edges of a copy of $\mathcal{F}_4^{(3)}$, but yet must by Theorem 4.3.5, which will be our contradiction.

Suppose there were a counterexample to the statement of the theorem. Then we can choose one with a minimal number of edges. So let $\mathcal{H}$ be a 3-partite 3-graph with vertex classes V_1, V_2, and V_3, and let $J \subseteq L(\mathcal{H})$ be a subgraph of its line graph with $\tau(\mathcal{H}_J) \geq 9$ and $\operatorname{conn}(J) \leq 0$. We may assume that $|E(J)|$ is minimal among such graphs. Note that $\tau(\mathcal{H}_J) \geq 9$ implies that J is not empty, so if J has no edges, then $\operatorname{conn}(J) = \infty$, a contradiction. Thus J must have some edges. We start with some facts about these edges.

Claim. *For all edges $e, f \in V(J)$ with $ef \in E(J)$, we have $|e \cap f| = 1$, $\tau(\mathcal{H}_{J \circledast ef}) \geq \tau(\mathcal{H}_J - (e \cup f)) \geq 4$, and $\operatorname{conn}(J \circledast ef) \leq -1$.*

Proof of claim. Let $e, f \in V(J)$ with $ef \in E(J)$. Then e and f intersect, since J is a subgraph of the line graph. Since $\mathcal{H}_J = \mathcal{H}_{J-ef}$, we have $\tau(\mathcal{H}_{J-ef}) \geq 9$, and since $|E(J - ef)| < |E(J)|$, we have $\operatorname{conn}(J - ef) \geq 1$ because J was minimal. If $\operatorname{conn}(J \circledast ef) \geq 0$, then Meshulam's Theorem would imply $\operatorname{conn}(J) \geq 1$, so we must have $\operatorname{conn}(J \circledast ef) \leq -1$.

It is clear that any vertex cover of $\mathcal{H}_J - (e \cup f)$ together with $e \cup f$ forms a vertex cover of $\mathcal{H}_J$, so we have $\tau(\mathcal{H}_J - (e \cup f)) + |e \cup f| \geq \tau(\mathcal{H}_J) \geq 9$. Since $|e \cup f| = 6 - |e \cap f|$, we have $\tau(\mathcal{H}_J - (e \cup f)) \geq 3 + |e \cap f| \geq 4$.

Since the set $e \cup f$ intersects every edge we delete when passing from $\mathcal{H}_J$ to $\mathcal{H}_{J \divideontimes ef}$, it follows that $\mathcal{H}_J - (e \cup f)$ is a subhypergraph of $\mathcal{H}_{J \divideontimes ef}$, so $\tau(\mathcal{H}_{J \divideontimes ef}) \geq \tau(\mathcal{H}_J - (e \cup f))$.

Now if $|e \cap f| \geq 2$, then we would have $\tau(\mathcal{H}_{J \divideontimes ef}) \geq 5$, and so by Corollary 4.3.8 we would have $\operatorname{conn}(J \divideontimes ef) \geq 0$, a contradiction. Thus we must have $|e \cap f| = 1$, as claimed. $\square$

Claim. *For all edges $e, f \in V(J)$ with $ef \in E(J)$, we have that $\mathcal{H}_J - (e \cup f)$ contains a copy of $\mathcal{F}_4^{(3)}$ and every edge in $\mathcal{H}_{J \divideontimes ef}$ is either a central edge of the copy or intersects both central edges.*

Proof of claim. By the previous claim, we know $\tau(\mathcal{H}_{J \divideontimes ef}) \geq 4$ and $\operatorname{conn}(J \divideontimes ef) \leq -1$. By Theorem 4.3.5 we then have that $\mathcal{H}_{J \divideontimes ef}$ contains a copy of $\mathcal{F}_4^{(3)}$ and every edge in $\mathcal{H}_{J \divideontimes ef}$ is either a central edge of the copy or intersects both central edges. Let c and c' denote the central edges in question.

First, we claim that c and c' are both disjoint from e and f. Indeed, suppose $c \cup c'$ and $e \cup f$ shared a vertex v. Without loss of generality, assume $v \in c$. Then consider the set $T \subseteq V(\mathcal{H}_J)$ given by $T = e \cup f \cup c'$, which is of size at most 8, since $|e \cup f| = 5$. We claim T is a vertex cover of $\mathcal{H}_J$. We know by Theorem 4.3.5 that every edge of $\mathcal{H}_{J \divideontimes ef}$ except c intersects c', and c intersects T in $v \in e \cup f$. Furthermore, every edge we removed in the explosion intersects $e \cup f$, hence T is a vertex cover. This is a contradiction, since $\tau(\mathcal{H}_J) \geq 9$. Therefore, c and c' must indeed be disjoint from e and f.

Consider the subgraph $J \cap L(\mathcal{H}_J - (e \cup f))$ of $J \divideontimes ef$. Since $\{c, c'\}$ forms a path-component of $\mathcal{I}(J \divideontimes ef)$, every edge between $\{c, c'\}$ and the rest of $V(J \divideontimes ef)$ is present in $J \divideontimes ef$. Because we may obtain $J \cap L(\mathcal{H}_J - (e \cup f))$ by deleting the vertices of $J \divideontimes ef$ that correspond to those edges of $\mathcal{H}_{J \divideontimes ef}$ that intersect $e \cup f$, we have that every edge between $\{c, c'\}$ and the rest of $V(J \cap L(\mathcal{H}_J - (e \cup f)))$ is present in $J \cap L(\mathcal{H}_J - (e \cup f))$ as well. It only remains to be seen that c and c' are not the only edges present in $\mathcal{H}_J - (e \cup f)$. This is easily seen to be true, since $\tau(\mathcal{H}_J - (e \cup f)) \geq 4$ by the previous claim. Therefore, $\mathcal{I}(J \cap L(\mathcal{H}_J - (e \cup f)))$ has at least two path-components, showing that $\operatorname{conn}(J \cap L(\mathcal{H}_J - (e \cup f))) \leq -1$. Therefore by Theorem 4.3.5, there is a copy of $\mathcal{F}_4^{(3)}$ in $\mathcal{H}_J - (e \cup f)$. It is easy to see that its central edges must be c and c', since all other edges in $\mathcal{H}_J - (e \cup f)$ are in a path-component of $\mathcal{I}(J \cap L(\mathcal{H}_J - (e \cup f)))$ of size greater than 2. Thus the claim holds. $\square$

Since J has an edge, let us fix $e, f \in V(J)$ with $ef \in E(J)$, and let c and c' be the central edges of the copy of $\mathcal{F}_4^{(3)}$ in $\mathcal{H}_J - (e \cup f)$. Let e and f intersect in the vertex $v \in V_j$, and let $i \in \{1, 2\}$. Consider the sets $E_i = (e \cup f) \setminus (e \cap V_{j+i}) \cup ((c \cup c') \cap (V_j \cup V_{j+3-i}))$ and $F_i = (e \cup f) \setminus (f \cap V_{j+i}) \cup ((c \cup c') \cap (V_j \cup V_{j+3-i}))$ (all arithmetic is done modulo 3). These consist of a minimal vertex cover of $\mathcal{H}_{J \divideontimes ef}$ and all but one vertex of $e \cup f$, hence are of size 8. Since these are too small to be vertex covers, there must be edges avoiding E_i and edges avoiding F_i for $i = 1, 2$.

To help classify these edges, we will seek the help of the following lemma:

Lemma 4.3.10. *For every edge $g \in \{e, f\}$, and every edge $h \in V(J)$ with $gh \in E(J)$ that is disjoint from c and c', there is **no** edge $d \in V(J \circledast gh)$ such that d is disjoint from c and c', and intersects at most one of e and f.*

Proof of Lemma 4.3.10. Let $g \in \{e, f\}$, and let $h \in V(J)$ with $gh \in E(J)$ be disjoint from c and c'. Suppose there were an edge $d \in V(J \circledast gh)$ such that d is disjoint from c and c', and intersects at most one of e and f. Since d is disjoint from one of e and f, d does not intersect the other in V_j (since this is the class of the common vertex of e and f). Since d is disjoint from c and c', it is not in $V(J \circledast ef)$, so it is adjacent in J to one of e and f. Thus, h intersects e or f in one vertex, and let V_i be the vertex class of that vertex. Since $gh \in E(J)$, we have seen above that $\mathcal{H}_{J \circledast gh}$ contains a copy of $\mathcal{F}_4^{(3)}$ and every edge in $\mathcal{H}_{J \circledast gh}$ is either a central edge of the copy or intersects both central edges. Let $\hat{c}$ and $\hat{c}'$ be the central edges of the copy.

By assumption, c, c', and d form a matching of size 3 in $\mathcal{H}_{J \circledast gh}$. This implies that none of them are central edges of $\mathcal{H}_{J \circledast gh}$, hence they all intersect $\hat{c}$ and $\hat{c}'$. There are two cases to consider.

Case 1. Neither $\hat{c}$ nor $\hat{c}'$ intersect d in V_i.

In this case, consider the set $S = g \cup h \cup ((c \cup c' \cup d) \cap V_j)$. This has size at most 8, hence is not a vertex cover of $\mathcal{H}_J$. Thus there is some edge m that avoids S. Since m does not intersect either of g and h, we have $m \in E(\mathcal{H}_{J \circledast gh})$. It thus must intersect $\hat{c}$ and $\hat{c}'$, and clearly does not do so in V_j, since the V_j-vertex of $\hat{c}$ and $\hat{c}'$ are both in S. Without loss of generality, suppose it meets $\hat{c}$ in V_i. Let V_k be the third vertex class, besides V_j and V_i. Then m meets $\hat{c}'$ in V_k. We claim that m is disjoint from one of c and c'. Indeed, if it meets both, it meets one of them in V_i, and the other in V_k. Since $\hat{c}'$ does not meet d in V_i by assumption, and does not meet m in V_i either, it follows that $\hat{c}'$ must meet the same member of $\{c, c'\}$ in V_i as it does in V_k, which cannot be the case. Therefore, m must be disjoint from one of c and c'. But m is also disjoint from e and f, since its V_j vertex is not in g, and its V_i- and V_k-vertices are in $\hat{c} \cup \hat{c}'$, which are disjoint from e and f because we are in Case 1. This is a contradiction, because every edge disjoint from e and f must intersect both c and c' (except c and c' themselves). Therefore, this case is impossible.

Case 2. Without loss of generality, $\hat{c}$ intersects d in V_i.

In this case, consider the set $S = g \cup h \cup ((c \cup c; \cup d) \cap V_i))$. This has size at most 8, hence is not a vertex cover of $\mathcal{H}_J$. Thus there is some edge m that avoids S. Since m does not intersect either of g and h, we have $m \in E(\mathcal{H}_{J \circledast gh})$. It thus must intersect $\hat{c}$ and $\hat{c}'$, and clearly does not do so in V_i, since the V_i-vertex of $\hat{c}$ and $\hat{c}'$ are both in S. Let V_k be the third vertex class, besides V_j and V_i. We claim that again m is disjoint from one of c and c'. Indeed, if it meets both, it meets one of them in V_j, and the other in V_k. But $\hat{c}$ also meets one of c and c' in V_j, and the other in V_k, so either $\hat{c}$ misses m, which is a contradiction, or it hits m twice, which would mean m does not intersect $\hat{c}'$, also a contradiction. Therefore, m must be disjoint from one of c and c'. But m is also disjoint from e and f, since its V_j vertex is not in g, and its V_i- and V_k-vertices are in $\hat{c} \cup \hat{c}'$, which are disjoint from e and f because we are in Case 1. This is a contradiction, because every edge disjoint from e and f must intersect both c and c' (except c and c' themselves). Therefore, this case is also impossible.

Since these cases cover all of the possibilites, the existence of such an edge d leads to a contradiction, thus proving the lemma. □

Lemma 4.3.10 implies the following claim:

Claim. *For $i, k \in \{1, 2\}$ it holds that for every edge $a \in V(J)$ avoiding E_i and every edge $b \in V(J)$ avoiding F_k we have $ab \in E(J)$.*

Proof of claim. Suppose there were edges a avoiding E_i and b avoiding F_k with $ab \notin E(J)$. We must have $ae \in E(J)$, since otherwise we would have $a \in V(J \circledast ef)$, which is a contradiction, as a is disjoint from the central edges c and c' of $\mathcal{H}_{J \circledast ef}$. But then Lemma 4.3.10 applied with $g = e$ and $h = a$ gives us a contradiction, since $b \in V(J \circledast ae)$ is disjoint from c and c', intersects only f, and does so in only one vertex. Hence we must have $ab \in E(J)$. □

This immediately implies the following:

Claim. *There is a vertex $x_1 \in V_j$ such that every edge avoiding E_1 and every edge avoiding F_2 contain it, and similarly there is a vertex $x_2 \in V_j$ such that every edge avoiding E_2 and every edge avoiding F_1 contain it.*

Proof of claim. For $i = 1, 2$ we have that any edge avoiding E_i has its V_{j+i}-vertex in e and its V_{j+3-i}-vertex outside $e \cup f$, while any edge avoiding F_{3-i} has its V_{j+i}-vertex outside $e \cup f$ and its V_{j+3-i}-vertex in f. But since any two such edges are adjacent in J, they must intersect, and they can do so only in V_j. The existance of edges avoiding E_i as well as edges avoiding F_{3-i} together with transitivity implies the existence of a vertex $x_i \in V_j$ contained in all such edges. □

We then also have the following:

Claim. *Either $x_1 = x_2$, or there are vertices $y_1 \in V_{j+1}$ and $y_2 \in V_{j+2}$ such that every edge avoiding E_1 and every edge avoiding F_1 contain y_2, and similarly every edge avoiding E_2 and every edge avoiding F_2 contain y_1.*

Proof of claim. Suppose $x_1 \neq x_2$. For $i = 1, 2$ we have that any edge avoiding E_i has its V_j-vertex equal to x_i and its V_{j+i}-vertex in e, while any edge avoiding F_i has its V_j-vertex equal to x_{3-i} and its V_{j+i}-vertex in f. But since any two such edges are adjacent in J, they must intersect, and they can do so only in V_{j+3-i}. The existance of edges avoiding E_i as well as edges avoiding F_i together with transitivity implies the existence of a vertex $y_{3-i} \in V_{j+3-i}$ contained in all such edges. □

Claim. *We have $x_1 \neq x_2$.*

Proof of claim. Suppose $x_1 = x_2$. We claim that the set $T = \{v, x_1\} \cup c \cup c'$ is a vertex cover of $\mathcal{H}_J$ of size 8. Indeed, suppose there were an edge $m \in V(J)$ avoiding T. We must have $em \in E(J)$ or $fm \in E(J)$, since otherwise we would have $m \in V(J \circledast ef)$, which is a contradiction, since m is disjoint from the two central edges c and c' of $\mathcal{H}_{J \circledast ef}$. If m intersects $e \cup f$ in only one vertex, then m avoids one of the sets E_i or F_i for some $i \in \{1, 2\}$. But this cannot be the case, since m does not contain $x_1 = x_2$. Thus m intersects $e \cup f$ in two vertices (it does not intersect in 3 vertices, since it does not contain v). It does not intersect either of them in two vertices because it is adjacent to one of them in J. Thus m intersects e in V_{j+i} and f in V_{j+3-i} for some $i \in \{1, 2\}$.

If $em \in E(J)$, then let $b \in V(J)$ be an edge avoiding F_i. Now b is disjoint from e and m, so $b \in V(J \circledast em)$, and as b is disjoint from c and c', we will get a contradiction by applying Lemma 4.3.10 with $g = e$ and $h = m$. Similarly, if $fm \in E(J)$, we will get a contradiction in the same way from any edge $a \in V(J)$ avoiding E_{3-i} by applying Lemma 4.3.10 with $g = f$ and $h = m$. Thus, there is no such edge m, hence T is a vertex cover of $\mathcal{H}_J$. This is a contradiction, since $\tau(\mathcal{H}_J) \geq 9$ by assumption. Thus, we must have $x_1 \neq x_2$. $\square$

Therefore the previous claim gives the existence of the vertices $y_1 \in V_{j+1}$ and $y_2 \in V_{j+2}$ satisfying the conditions laid out in the claim.

Consider the copy $\mathcal{H}' \subseteq \mathcal{H}_J - (e \cup f)$ of $\mathcal{F}_4^{(3)}$ with central edges c and c'. There is an edge $g \in E(\mathcal{H}')$ whose V_{j+1}-vertex is not in $c \cup c'$ and is not y_1. There are also distinct edges $h, h' \in E(\mathcal{H}')$ whose V_{j+2}-vertices are not in $c \cup c'$. One of these edges is disjoint from g, and we may assume without loss of generality that this edge is h. Now h intersects one of c and c' in V_{j+1}, and since the roles of c and c' have been entirely symmetrical so far, we may assume without loss of generality that h intersects c in V_{j+1}.

Now consider the set $S = \{v, x_1, x_2\} \cup c \cup (c' \setminus V_{j+1})$. This is a set of size 8, so it is too small to be a vertex cover of $\mathcal{H}_J$. Thus there exists an edge $m \in V(J)$ avoiding S. Clearly m must be adjacent in J to one of e and f, since it fails to intersect c (thus it is not in $V(J \circledast ef)$). If m contains two vertices of $e \cup f$, we may proceed as in the proof of the claim showing $x_1 \neq x_2$ to reach a contradiction. Also, if m avoids both c and c', then it avoids E_1 or F_1, which would mean that it contains x_1 or x_2, also a contradiction. Thus, we may assume that $m \cap V_{j+1} = c' \cap V_{j+1}$.

Now if m intersects e, then let $b \in V(J)$ be an edge avoiding F_2 and set $\mathcal{H}^* = \mathcal{H}_J - (b \cup f)$. If on the other hand m intersects f, then let $a \in V(J)$ be an edge avoiding E_2 and set $\mathcal{H}^* = \mathcal{H}_J - (a \cup e)$. Now both a and b have y_1 as their V_{j+1}-vertex, and have x_1 or x_2 as their V_j-vertex, and thus g, h, h', c, and c' are disjoint from them. Thus these edges along with m are in $\mathcal{H}^*$. We know by an earlier claim that $\mathcal{H}^*$ contains a copy of $\mathcal{F}_4^{(3)}$, such that every edge of $\mathcal{H}^*$ is a central edge of the copy or intersects both central edges of the copy. Now m, g, and h form a matching of size 3 in $\mathcal{H}^*$, so none of these can be a central edge in the copy. Also, m and c are disjoint, so c cannot be a central edge, which implies c' is not a central edge either, since it is disjoint from c. Similarly, h' cannot be a central edge, since it is disjoint from the non-central edge h. Now a simple case analysis shows that there is no way to find two disjoint edges, each intersecting all of m, g, h, h', c and c'. Thus, we have reached a contradiction. This means that there can be no $J \subseteq L(\mathcal{H})$ with $\tau(\mathcal{H}_J) \geq 9$ and $\mathrm{conn}(J) \leq 0$, proving the proposition. $\square$

Chapter 5

Triangulations

5.1 Introduction

In this chapter, we aim to provide a solid foundation for the topological machinery we used in the previous chapters. It may be seen as a sort of appendix. Along the way, we fix an oversight that was recently discovered in certain proofs involving triangulations.

Triangulations of spheres and balls have been an object of study for a long time. Specifically relevant to the topic of this thesis was the paper of Aharoni and Haxell [6], from which we implicitly get Theorem 2.1.6. Aharoni, Chudnovsky, and Kotlov [5] gave techniques for extending triangulations of spheres to special triangulations of balls. Szabó and Tardos [27] expanded on these techniques in order to prove some degree conditions on the existence of transversals with various properties. Both of these latter papers relied on the supposed fact that the links of simplices in the interior of triangulated balls are triangulations of spheres, which, as we will come to see, is not necessarily true for general triangulations. One way to guarantee that this fact holds is by considering so-called PL-triangulations. To be fair, Szabó and Tardos do mention this in [27], but one of their construction involves iteratively replacing parts of the ball with a different triangulation, hence one must check that this replacement preserves the PL property.

In the first part of this chapter, we will provide the necessary definitions and known results regarding triangulations of spheres and balls, culminating in the proof that the replacement technique of Szabó and Tardos is sound. In the second part, we will use this same technique to give a triangulation proof of Meshulam's Theorem (Theorem 2.1.5), which is one of the most heavily used tool in the previous chapters of this thesis.

5.2 Topological Definitions and Theorems

Simplicial complexes give discrete descriptions of topological spaces. Indeed, a triangulation of a topological space is a simplicial complex whose polyhedron is homeomorphic to the space. In applications, we often deal with abstract simplicial complexes, which encode only combinatorial information about which simplices are incident to which simplices. However, in order to define such things as subdivisions and PL-triangulations, we will need to deal with geometric simplicial complexes, which come with a concrete embedding of its simplices into $\mathbb{R}^d$. Both concepts have infinite versions, but here we shall only consider finite simplicial complexes.

Note that many of the definitions that follow were given in Chapter 2. They are repeated here for the convenience of the reader.

For a general reference, we refer the reader to [10] and [18].

5.2.1 Abstract Simplicial Complexes

Definition 5.2.1. An *abstract simplicial complex* $\mathcal{C}$ is a finite collection of finite sets that is closed under taking subsets. The set of *vertices* of $\mathcal{C}$ is $V(\mathcal{C}) = \bigcup_{\sigma\in\mathcal{C}} \sigma$. The elements $\sigma \in \mathcal{C}$ are called *simplices*, and the subsets of a simplex are called its *faces*.

Definition 5.2.2. If σ is a simplex of an abstract simplicial complex, then the *dimension* of σ is one less than the number of elements in σ and is denoted $\dim(\sigma) = |\sigma| - 1$. The dimension of an abstract simplicial complex $\mathcal{C}$ is $\dim(\mathcal{C}) = \max_{\sigma\in\mathcal{C}} \dim(\sigma)$.

Definition 5.2.3. A *simplicial map* between abstract simplicial complexes $\mathcal{C}$ and $\mathcal{D}$ is a map $f : V(\mathcal{C}) \to V(\mathcal{D})$ such that $f(\sigma) \in \mathcal{D}$ for every simplex $\sigma \in \mathcal{C}$.

Definition 5.2.4. An *isomorphism* between abstract simplicial complexes $\mathcal{C}$ and $\mathcal{D}$ is a simplicial bijection $f : V(\mathcal{C}) \to V(\mathcal{D})$ whose inverse is simplicial. If an isomorphism between $\mathcal{C}$ and $\mathcal{D}$ exists, we say $\mathcal{C}$ and $\mathcal{D}$ are *isomorphic*, and we write $\mathcal{C} \cong \mathcal{D}$.

Definition 5.2.5. The *join* of two abstract simplicial complexes $\mathcal{C}$ and $\mathcal{D}$ is the abstract simplicial complex $\mathcal{C} * \mathcal{D} = \{(\sigma \times \{0\}) \cup (\tau \times \{1\}) : \sigma \in \mathcal{C}, \tau \in \mathcal{D}\}$.

Definition 5.2.6. Let $\mathcal{C}$ be an abstract simplicial complex, and let $\sigma \in \mathcal{C}$.

The *open star* of σ is $\mathrm{star}_{\mathcal{C}}(\sigma) = \{\tau \in \mathcal{C} : \sigma \subseteq \tau\}$.

The *link* of σ is $\mathrm{lk}_{\mathcal{C}}(\sigma) = \{\tau \in \mathcal{C} : \tau \cup \sigma \in \mathcal{C} \text{ and } \tau \cap \sigma = \emptyset\}$.

5.2.2 Geometric Simplicial Complexes

Definition 5.2.7. A *geometric simplex* $\sigma \subseteq \mathbb{R}^d$ is the convex hull of a set of affinely independent points. These points are its *vertices*, denoted by $V(\sigma)$.

Definition 5.2.8. A *face* of a geometric simplex is the convex hull of a subset of its vertices.

Definition 5.2.9. A *geometric simplicial complex* $\mathcal{K}$ is a finite collection of geometric simplices such that for any geometric simplex $\sigma \in \mathcal{K}$ every face of σ is also in $\mathcal{K}$, and if $\sigma, \tau \in \mathcal{K}$, then $\sigma \cap \tau$ is a common face of σ and τ. The *vertex set* of $\mathcal{K}$, is the set $V(\mathcal{K}) = \bigcup_{\sigma \in \mathcal{K}} V(\sigma)$.

Definition 5.2.10. The *polyhedron* of a geometric simplicial complex Δ is the space $\|\mathcal{K}\| = \bigcup_{\sigma \in \mathcal{K}} \sigma$.

Definition 5.2.11. The *boundary* of a geometric simplex σ is the geometric simplicial complex $\partial\sigma = \{\operatorname{conv}(U) : U \subsetneq V(\sigma)\}$.

Definition 5.2.12. The *interior* of a geometric simplex σ is the convex set $\operatorname{int}(\sigma) = \sigma \setminus \|\partial\sigma\|$.

Note that a geometric simplex $\mathcal{K}$ is the disjoint union of the interiors of its faces. In particular, for any $x \in \|\mathcal{K}\|$, there is a unique $\sigma_x \in \mathcal{K}$ such that $x \in \operatorname{int}(\sigma_x)$.

Definition 5.2.13. A *simplicial map* between geometric simplicial complexes $\mathcal{K}$ and $\mathcal{L}$ is a map $f : V(\mathcal{K}) \to V(\mathcal{L})$ such that $\operatorname{conv}(f(V(\sigma))) \in \mathcal{L}$ for every simplex $\sigma \in \mathcal{K}$.

Definition 5.2.14. Let $\mathcal{K}$ and $\mathcal{L}$ be geometric simplicial complexes, and let $f : V(\mathcal{K}) \to V(\mathcal{L})$ be a simplicial map. Then the *polyhedron* of f is the map $\|f\| : \|\mathcal{K}\| \to \|\mathcal{L}\|$ given by linear extension of f on each of the simplices of $\mathcal{K}$. Concretely, if $x \in \operatorname{int}(\sigma)$ for some $\sigma \in \mathcal{K}$, then $x = \sum_{v \in V(\sigma)} \lambda_v v$ for uniquely determined λ_v, and we define $\|f\|(x) = \sum_{v \in V(\sigma)} \lambda_v f(v)$.

It is clear that if $f : V(\mathcal{K}) \to V(\mathcal{L})$ is simplicial, then $\|f\| : \|\mathcal{K}\| \to \|\mathcal{L}\|$ is continuous.

Definition 5.2.15. An *isomorphism* between geometric simplicial complexes $\mathcal{K}$ and $\mathcal{L}$ is a simplicial bijection $f : V(\mathcal{K}) \to V(\mathcal{L})$ whose inverse is simplicial. If an isomorphism between $\mathcal{K}$ and $\mathcal{L}$ exists, we say $\mathcal{K}$ and $\mathcal{L}$ are *isomorphic*, and we write $\mathcal{K} \cong \mathcal{L}$.

Definition 5.2.16. Two geometric simplices σ and τ in $\mathcal{R}^d$ are called *joinable* if they are disjoint and the union of their vertex sets is affinely independent. If σ and τ are joinable, then their *join* is the geometric simplex $\sigma * \tau = \operatorname{conv}(\sigma \cup \tau) = \operatorname{conv}(V(\sigma) \cup V(\tau))$.

Definition 5.2.17. Two geometric simplicial complexes $\mathcal{K}$ and $\mathcal{L}$ are called *joinable* if for every pair of simplices $\sigma \in \mathcal{K}$ and $\tau \in \mathcal{L}$ we have the following:

(1) σ and τ are joinable,

(2) If $\sigma' \in \mathcal{K}$ and $\tau' \in \mathcal{L}$, then $(\sigma * \tau) \cap (\sigma' * \tau')$ is a common face of $\sigma * \tau$ and $\sigma' * \tau'$.

If $\mathcal{K}$ and $\mathcal{L}$ are joinable, then their *join* is the geometric simplicial complex $\mathcal{K} * \mathcal{L} = \{\sigma * \tau : \sigma \in \mathcal{K}, \tau \in \mathcal{L}\}$.

Definition 5.2.18. Let $\mathcal{K}$ be a geometric simplicial complex, and let $\sigma \in \mathcal{K}$.
The *open star* of σ is $\mathrm{star}_{\mathcal{K}}(\sigma) = \{\tau \in \mathcal{K} : \sigma \subseteq \tau\}$.
The *link* of σ is $\mathrm{lk}_{\mathcal{K}}(\sigma) = \{\tau \in \mathcal{K} : \sigma * \tau \in \mathcal{K} \text{ and } \tau \cap \sigma = \emptyset\}$.

Note that the link of a simplex is a geometric simplicial complex, while the open star is not necessarily one.

In the case when σ is 0-dimensional, i.e. consists of a single vertex v, we usually write $\mathrm{star}(v)$ instead of $\mathrm{star}(\{v\})$, and similarly for the the link.

Definition 5.2.19. If $\mathcal{K}$ is a geometric simplicial complex, then a *subdivision* of $\mathcal{K}$ is a geometric simplicial complex $\mathcal{K}'$ with $\|\mathcal{K}'\| = \|\mathcal{K}\|$ such that every simplex in $\mathcal{K}'$ is contained in a simplex in $\mathcal{K}$.

5.2.3 Vertex Schemes and Realizations

Of course one can translate between abstract and geometric simplicial complexes.

Definition 5.2.20. If $\mathcal{K}$ is a geometric simplicial complex, then the *vertex scheme* of $\mathcal{K}$ is the abstract complex $\mathrm{vs}(\mathcal{K}) = \{U \subseteq V(\mathcal{K}) : \mathrm{conv}(U) \in \mathcal{K}\}$.

It is clear that this does in fact produce an abstract simplicial complex, since the faces of simplices in a geometric simplicial complex are themselves in the complex. Thus every geometric simplicial complex corresponds to an abstract simplicial complex. The correspondence also goes the other way.

Definition 5.2.21. If $\mathcal{C}$ is an abstract simplicial complex, then a *geometric realization* of $\mathcal{C}$, also called an *embedding*, is a geometric simplicial complex $\mathcal{K}$ such that $\mathrm{vs}(\mathcal{K}) \cong \mathcal{C}$.

It is not at first glance clear that every abstract simplicial complex has a geometric realization. The next theorem of Menger and Nöbeling gives us such a result.

Theorem 5.2.22. *If $\mathcal{C}$ is a d-dimensional abstract simplicial complex, then $\mathcal{C}$ has a geometric realization in $\mathbb{R}^{2d+1}$.*

In this way, we can carry over definitions and constructions from one type of simplicial complex to the other. For instance, the definitions of joins, stars, and links in geometric and abstract simplicial complexes translate into each other in this way. More fundamentally, the notion of isomorphism translates between the abstract and geometric setting, so any two geometric realizations of isomorphic abstract simplicial complexes are isomorphic.

We can use this correspondence to define subdivisions of abstract simplicial complexes.

Definition 5.2.23. If $\mathcal{C}$ is an abstract simplicial complex, then an *(abstract) subdivision* of $\mathcal{C}$ is an abstract simplicial complex $\mathcal{C}'$ such that there are geometric realizations $\mathcal{K}$ of $\mathcal{C}$ and $\mathcal{K}'$ of $\mathcal{C}'$ with $\mathcal{K}'$ a subdivision of $\mathcal{K}$.

A related observation is that subdivisions can be carried over between isomorphic geometric simplicial complexes. Let $\mathcal{K}$ and $\mathcal{L}$ be isomorphic geometric simplicial complexes, with $\phi : V(\mathcal{K}) \to V(\mathcal{L})$ a simplicial isomorphism. Now suppose $\mathcal{K}'$ is a subdivision of $\mathcal{K}$. Then ϕ induces an isomorphic subdivision $\mathcal{L}' = \|\phi\| (\mathcal{K}')$ of $\mathcal{L}$ by mapping a simplex $\sigma \in \mathcal{K}'$ to the simplex $\|\phi\| (\sigma)$. Since $\|\phi\|$ is linear on each simplex of $\mathcal{K}$ and every simplex of $\mathcal{K}'$ is contained in a simplex of $\mathcal{K}$, this clearly produces a subdivision of $\mathcal{L}$.

Another useful definition is that of the k-skeleton, which also applies to both abstract and geometric complexes.

Definition 5.2.24. If $\mathcal{C}$ is a simplicial complex, then the *k-skeleton* $\mathcal{C}^{(k)}$ of $\mathcal{C}$ is the subcomplex of $\mathcal{C}$ consisting of all simplices of dimension at most k.

5.2.4 Connectedness

We define the d-sphere concretely by $S^d = \left\{x \in \mathbb{R}^{d+1} : |x| = 1\right\}$, and the d-ball by $B^d = \left\{x \in \mathbb{R}^d : |x| \leq 1\right\}$.

Definition 5.2.25. Let $k \geq -1$ be an integer. A topological space X is said to be *k-connected* if for any integer j with $-1 \leq j \leq k$, any continuous map from the j-dimensional sphere S^j into the space X can be extended to a continuous map from the $(j+1)$-dimensional ball B^{j+1} to X. The *connectedness* of X, denoted $\mathrm{conn}(X)$ is the largest k for which X is k-connected. Note that this may be ∞, which is the case if the space is contractible, i.e. can be shrunk continuously to a single point.

The -1-sphere is the empty set and the 0-ball is a single point, so a space is -1-connected if and only if it is non-empty. 0-connected means path-connected, and 1-connected means simply connected.

A geometric simplicial complex is said to be k-connected if its polyhedron is, and an abstract simplicial complex is k-connected if its geometric realization is (and observe that this does not depend on the choice of geometric realization, as the geometric realizations are all homeomorphic).

A useful fact relating connectedness to joins is the following:

Proposition 5.2.26 (Lemma 2.3 in [23])**.** *If $\mathcal{C}$ and $\mathcal{D}$ are abstract simplicial complexes, then*

$$\mathrm{conn}(\mathcal{C} * \mathcal{D}) \geq \mathrm{conn}(\mathcal{C}) + \mathrm{conn}(\mathcal{D}) + 2$$

5.2.5 Triangulations

Definition 5.2.27. A *triangulation* of a topological space X is an abstract simplicial complex $\mathcal{C}$ for which the polyhedron of its geometric realization is homeomorphic to X.

We are interested in triangulations of balls and spheres. One important fact about balls is that their boundaries are spheres. To translate this notion to the setting of triangulations, note that in a triangulation of a d-ball, there are two kinds of $(d-1)$-dimensional simplices: those which are in two d-dimensional simplices, and those which are in only one.

Definition 5.2.28. If $\mathcal{B}$ is a triangulation of a ball B^d, then the *boundary* of $\mathcal{B}$ is the subcomplex whose maximal simplices are the $(d-1)$-dimensional simplices of $\mathcal{B}$ which are in only one d-simplex of $\mathcal{B}$.

This gives that the boundary of a triangulated d-ball $\mathcal{B}$ is a triangulated $(d-1)$-sphere $\mathcal{S}$, and in fact for any homeomorphism between a geometric realization of $\mathcal{B}$ and B^d, the image of the boundary is S^{d-1}.

As the notion of connectedness calls for extending maps from spheres to the ball they are the boundary of, we will want a simplicial version for triangulations of spheres. In order to be able to freely apply certain gluing procedures, we will require a bit more of our "filling" than one might initially expect.

Definition 5.2.29. If $\mathcal{S}$ is a triangulation of a sphere S^d, then a *filling* of $\mathcal{S}$ is a triangulation $\mathcal{B}$ of the ball B^{d+1} whose boundary is $\mathcal{S}$, and such that if $\sigma \in \mathcal{B}$ with $\sigma \subseteq V(\mathcal{S})$, then $\sigma \in \mathcal{S}$.

The more restricted definition of filling ensures us that if we have a triangulation of a ball and we remove a triangulated ball from the interior, leaving a shell, then adding any filling of the interior boundary of the shell again results in a ball (as long as we avoid using vertices from the shell in the filling, apart from the inner boundary). We will always assume that any filling uses its own distinct set of vertices in the interior, so that there are never any unfortunate coincidences with vertices from other complexes. The following lemma makes this gluing precise.

Lemma 5.2.30. *Let $\mathcal{C}$, $\mathcal{D}$ and $\mathcal{D}'$ be abstract simplicial complexes with $\mathcal{C} \cap \mathcal{D} = \mathcal{C} \cap \mathcal{D}'$ and with $\mathcal{D}$ homeomorphic to $\mathcal{D}'$ via a homeomorphism that is the identity on $\mathcal{C} \cap \mathcal{D}$. Then $\mathcal{C} \cup \mathcal{D}$ is homeomorphic to $\mathcal{C} \cup \mathcal{D}'$.*

Proof. Let $\phi : \|\mathcal{D}\| \to \|\mathcal{D}'\|$ be a homeomorphism that is the identity on $\mathcal{C} \cap \mathcal{D}$. Then define $\psi : \|\mathcal{C} \cup \mathcal{D}\| \to \|\mathcal{C} \cup \mathcal{D}'\|$ to be the identity on $\|\mathcal{C}\|$ and to be ϕ on $\|\mathcal{D}\|$. Since the pieces agree on the intersection, and both $\|\mathcal{C}\|$ and $\|\mathcal{D}\|$ are closed subsets of $\|\mathcal{C} \cup \mathcal{D}\|$, it follows that ψ is continuous. Its inverse is continuous by the same reasoning, hence ψ is a homeomorphism. $\square$

In order to apply certain proof techniques, we will need our triangulations of spheres and balls to be "piecewise linear," or "PL" for short.

Definition 5.2.31. A triangulated d-ball or $(d-1)$-sphere is called a *PL-ball* or *PL-sphere*, respectively, if it has a subdivision which is isomorphic to a subdivision of the abstract d-simplex or its boundary, respectively.

A *geometric* PL-ball or PL-sphere is a geometric realization of a PL-ball or PL-sphere.

This technical property is needed to ensure the following:

Proposition 5.2.32 (Corollary 1.16 in [18])**.** *If $\mathcal{B}$ is a PL-d-ball, and σ is a k-dimensional simplex not contained in its boundary, then $lk_{\mathcal{B}}(\sigma)$ is a PL-$(d-k-1)$-sphere.*

This nice property of PL-balls is unfortunately not true in general. The classic counterexample is the double-suspension of a homology sphere, which by the double-suspension theorem [11] is homeomorphic to a sphere. Removing a maximal simplex of a triangulation creates a ball that fails the conclusion of Proposition 5.2.32.

5.2.6 Simplicial Approximation

A useful fact is that we can check for connectedness using fillings of PL-spheres:

Proposition 5.2.33. *A simplicial complex $\mathcal{C}$ is k-connected if and only if for every j with $-1 \leq j \leq k$ and for every simplicial map $f : V(\mathcal{S}) \to V(\mathcal{C})$, where $\mathcal{S}$ is a PL-j-sphere, there is a filling of $\mathcal{S}$ by a PL-$(j+1)$-ball $\mathcal{B}$, and a simplicial map $\hat{f} : V(\mathcal{B}) \to V(\mathcal{C})$ extending f.*

It will also be important that we can do this even with subdivisions of simplices:

Proposition 5.2.34. *A simplicial complex $\mathcal{C}$ is k-connected if and only if for every j with $-1 \leq j \leq k$ and for every simplicial map $f : V(\mathcal{S}) \to V(\mathcal{C})$, where $\mathcal{S}$ is a subdivision of the boundary of a $(j+1)$-simplex, there is a subdivision $\mathcal{B}$ of a $(j+1)$-simplex with $\mathcal{S}$ as its boundary, and a simplicial map $\hat{f} : V(\mathcal{B}) \to V(\mathcal{C})$ extending f.*

The proof of these follows along the lines of Proposition 2.8 in [27], using the simplicial approximation theorem. We give a more explicit proof.

Definition 5.2.35. If $f : \|\mathcal{K}\| \to \|\mathcal{L}\|$ is a continuous map, then a *simplicial approximation* of f is a simplicial map $g : V(\mathcal{K}) \to V(\mathcal{L})$ such that $f(\mathrm{star}_{\mathcal{K}}(v)) \subseteq \mathrm{star}_{\mathcal{L}}(g(v))$ for every vertex $v \in V(\mathcal{K})$.

In order to cover both cases, we will just refer to a triangulated sphere and a filling of it. One must take these to be the appropriate type for the particular lemma. We note that both classes of spheres and balls are closed under taking subdivisions and under taking cones.

Proof of Lemmas 5.2.33 and 5.2.34. We prove both directions.

($\Leftarrow$) Suppose that for every integer j with $-1 \leq j \leq k$ and every triangulated j-sphere $\mathcal{S}$, every simplicial map $f : V(\mathcal{S}) \to V(\mathcal{C})$ has a simplicial extension $\hat{f} : V(\mathcal{B}) \to V(\mathcal{C})$, where $\mathcal{B}$ is a filling of $\mathcal{S}$.

Let $-1 \leq j \leq k$, and let $f : S^j \to \|\mathcal{C}\|$ be a continuous map. Our goal is to extend f continuously to the ball B^{j+1}. Let $\mathcal{S}$ be a triangulated j-sphere, which means $\|\mathcal{S}\|$ is homeomorphic to S^j, so let $\phi : \|\mathcal{S}\| \to S^j$ be a homeomorphism. Then the composition $f \circ \phi : \|\mathcal{S}\| \to \|\mathcal{C}\|$ is a continuous map between polyhedra of simplicial complexes, so it has a simplicial approximation $s : V(\mathcal{S}') \to V(\mathcal{C})$, where $\mathcal{S}'$ is a subdivision of $\mathcal{S}$. Then $\|s\|$ is homotopic to $f \circ \phi$, so let $H : \|\mathcal{S}\| \times [0,1] \to \|\mathcal{C}\|$ be a homotopy with $H(x,0) = \|s\|(x)$ and $H(x,1) = f \circ \phi(x)$ for all $x \in \|\mathcal{S}\|$. By our supposition, s has a simplicial extension $\hat{s} : V(\mathcal{B}) \to V(\mathcal{C})$, where $\mathcal{B}$ is a filling of $\mathcal{S}'$. We then have that $\|\hat{s}\|$ is a continuous extension of $\|s\|$ to the polyhedron $\|\mathcal{B}\|$. We know $\|\mathcal{B}\|$ is homeomorphic to the ball B^{j+1} and has boundary $\|\mathcal{S}'\| = \|\mathcal{S}\|$, so let $\psi : \|\mathcal{B}\| \to B^{j+1}$ be a homeomorphism extending ϕ.

Define $g : B^{j+1} \to \|\mathcal{C}\|$ by

$$g(x) = \begin{cases} \|\hat{s}\| \circ \psi^{-1}(2x) & \text{if } |x| \leq \frac{1}{2}, \\ H(\phi^{-1}(\frac{x}{|x|}), 2|x| - 1) & \text{if } |x| \geq \frac{1}{2}. \end{cases}$$

We will show that g is a continuous extension of f, which will prove the if direction of our proposition. Both pieces of g are clearly continuous. We need to show that they agree on the boundary. Let $|x| = \frac{1}{2}$. Then $2x$ is in S^j, so $\psi^{-1}(2x)$ is in $\|\mathcal{S}\|$, and $\psi^{-1}(2x) = \phi^{-1}(2x)$, since ψ and ϕ agree on the boundary. Thus $\|\hat{s}\|(\psi^{-1}(2x)) = \|s\|(\phi^{-1}(2x))$.

($\Rightarrow$) Now suppose that $\|\mathcal{C}\|$ is k-connected.

Let $-1 \leq j \leq k$, and let $f : V(\mathcal{S}) \to V(\mathcal{C})$ be a simplicial map from some triangulated j-sphere $\mathcal{S}$ to $\mathcal{C}$. We will find a filling of $\mathcal{S}$ and a simplicial extension of f to the filling. We start by taking the cone $\mathcal{B} = p * \mathcal{S}$. We will assume for convenience that $\mathcal{S}$ lies in $\mathbb{R}^d \times 0$ for some d, and that $p = (0, \ldots, 0, 1)$ in $\mathbb{R}^d \times \mathbb{R}$, but make special note that the construction does not rely on this fact. Consider $\|\mathcal{B}\| \cap (\mathbb{R}^d \times [1/2, 1])$. This is also the polyhedron of a cone over $\mathcal{S}$, so let $\mathcal{B}'$ be the corresponding simplicial complex, and let $\mathcal{S}' \subseteq \mathcal{B}'$ be the subcomplex corresponding to $\|\mathcal{B}\| \cap (\mathbb{R}^d \times 1/2)$, which is isomorphic to $\mathcal{S}$ in the obvious way. Let $f' : V(\mathcal{S}') \to V(\mathcal{C})$ be the map corresponding to f via this isomorphism. Now let $\phi : \|\mathcal{S}'\| \to S^j$ be a homeomorphism, and let $\psi : \|\mathcal{B}'\| \to B^{j+1}$ be a homeomorphism extending ϕ. Then $\|f\| \circ \phi^{-1}$ is a continuous map from S^j to $\|\mathcal{C}\|$, and since $\|\mathcal{C}\|$ is k-connected, it can be extended to a continuous map $g : B^{j+1} \to \|\mathcal{C}\|$. Then by the simplicial approximation theorem, there is a subdivision $\mathcal{B}''$ of $\mathcal{B}$ such that there is a simplicial approximation $h : V(\mathcal{B}'') \to V(\mathcal{C})$ of $g \circ \psi$.

We will now extend the subdivision $\mathcal{B}''$ to a subdivision of $\mathcal{B}$. We will define a chain of complexes $\mathcal{B}_0, \ldots, \mathcal{B}_n$ with $\mathcal{B}_0 = \mathcal{B}''$ and $\mathcal{B}_n$ a subdivision of $\mathcal{B}$, such that each is a subcomplex of the next. For each j-simplex σ of $\mathcal{S}$, let σ' be the corresponding simplex of $\mathcal{S}'$, and let $\mathcal{K}_0(\sigma)$ be the subcomplex of $\mathcal{B}''$ that is a subdivision of σ'. Let $v_1, \ldots, v_n$ be the vertices of $\mathcal{S}$, and $v'_1, \ldots, v'_n$ the corresponding vertices of $\mathcal{S}'$. Supposing $\mathcal{B}_{i-1}$ has been defined, and that $\mathcal{K}_{i-1}(\sigma)$ is a subdivision of a j-simplex for each j-simplex σ of $\mathcal{S}$, let $\mathcal{B}_i$ be the union of $\mathcal{B}_{i-1}$ and the joins of v_i with $\mathcal{K}_{i-1}(\sigma)$ as σ ranges over the j-simplices of $\mathcal{S}$ containing v_i. We also define $\mathcal{K}_i(\sigma)$ for each j-simplex σ of $\mathcal{S}$. If σ does not contain v_i, let $\mathcal{K}_i(\sigma) = \mathcal{K}_{i-1}(\sigma)$, otherwise let $\tau = \|\mathcal{K}_{i-1}(\sigma)\|$ be a simplex, and let $\mathcal{L}_i(\sigma)$ be the subdivision of the facet of τ disjoint from the vertex of $\mathcal{S}'$ corresponding to v_i. Define $\mathcal{K}_i(\sigma) = v_i * \mathcal{L}_i(\sigma)$, which is again a subdivision of a j-simplex. In each step, the $\mathcal{K}_i(\sigma)$ make up the boundary of the ball $\mathcal{B}_i$. In the end, $\mathcal{B}_n$ is a subdivision of $\mathcal{B}$.

We set $\hat{f} : V(\mathcal{B}_n) \to V(\mathcal{C})$ to be equal to h on $V(\mathcal{B}'')$ and equal to f on $V(\mathcal{S})$ (these are all of the vertices of $\mathcal{B}_n$). We must check that it is simplicial. Consider a simplex $\sigma \in \mathcal{B}_n$. If it is contained in $\mathcal{S}$, f maps it to a simplex of $\mathcal{C}$, and if it is contained in $\mathcal{B}''$, h maps it to a simplex of $\mathcal{C}$. Otherwise, by constuction $\sigma = \sigma_1 * \sigma_2$, where σ_1 is a simplex of $\mathcal{S}$ and σ_2 is a simplex of $\mathcal{B}''$ that is part of the subdivision of a j-simplex τ of $\mathcal{S}'$, which has σ_1' as a face. Since $g \circ \psi$ is linear on τ, σ_2 is mapped by $g \circ \psi$ into the simplex τ is mapped to by f'. Since h is a simplicial approximation of $g \circ \psi$, h must maps the vertices of σ_2 to the vertices of the simplex τ is mapped to by f'. Therefore $\hat{f}$ maps σ to a face of that simplex, hence to a simplex. Therefore, $\hat{f}$ is simplicial, and the lemmas follow. $\square$

5.2.7 Star Replacement

One technique for proving the connectedness of a simplicial complex using Proposition 5.2.33 is the following: Take a PL-sphere of the desired dimension together with a simplicial map from it to the complex. Find an initial PL-filling and some extension of the map to the filling, which may fail to be simplicial. Then fix the filling and the map by iteratively replacing "bad" simplices with good ones. One method of doing this involves replacing the open star of a simplex with a filling of its link. This makes sense, since by Proposition 5.2.32, the link is a sphere.

Definition 5.2.36. Let $\mathcal{B}$ be a PL-ball, let $\sigma \in \mathcal{B}$ be a simplex not contained in its boundary, and let $\mathcal{F}$ be a filling of $\mathrm{lk}_{\mathcal{B}}(\sigma)$. Then the *star-replacement* of σ by $\mathcal{F}$ is the complex $\mathrm{starrep}_{\mathcal{B}}(\sigma, \mathcal{F}) = \mathcal{B} \setminus \mathrm{star}_{\mathcal{B}}(\sigma) \cup (\partial\sigma * \mathcal{F})$.

It is important to note that the star-replacement leaves the boundary of the ball unchanged, and does in fact produce another ball, which is due to our extra requirement of fillings.

For example, this approach was used in [27]. We will also apply this technique in the proof of Meshulam's Theorem. If we want to perform this operation more than once, then we had better make sure that the result again produces a PL-ball. The fact that this does in fact happen is stated in the following theorem, whose proof will be the topic of the rest of the subsection.

Theorem 5.2.37. *Let $\mathcal{B}$ be a PL-ball, and let $\sigma \in \mathcal{B}$ be a simplex not contained in the boundary of $\mathcal{B}$. If $\mathcal{F}$ is a PL-filling of* $\mathrm{lk}_{\mathcal{B}}(\sigma)$*, then* $\mathrm{starrep}_{\mathcal{B}}(\sigma, \mathcal{F})$ *is a PL-ball.*

Lemma 5.2.38 (Lemma 1.13 in [18])**.** *Let $\mathcal{S}$ and $\mathcal{S}'$ be PL-spheres, and let $\mathcal{B}$ and $\mathcal{B}'$ be PL-balls. Then $\mathcal{S} * \mathcal{S}'$ is a PL-sphere and $\mathcal{S} * \mathcal{B}$ and $\mathcal{B} * \mathcal{B}'$ are PL-balls.*

The following lemma we will prove later:

Lemma 5.2.39. *If $\mathcal{B}$ and $\mathcal{B}'$ are geometric PL-balls with $\phi : V(\partial\mathcal{B}) \to V(\partial\mathcal{B}')$ an isomorphism between their boundaries, then there are subdivisions $\mathcal{K}$ of $\mathcal{B}$ and $\mathcal{K}'$ of $\mathcal{B}'$ with an isomorphism $\psi : V(\mathcal{K}) \to V(\mathcal{K}')$ such that $\|\psi|_{\partial\mathcal{K}}\| = \|\phi\|$.*

Lemma 5.2.40 (Lemma 1.3 in [18])**.** *Let $\mathcal{K}$ be a geometric simplicial complex, and let $\mathcal{L} \subseteq \mathcal{K}$ be a subcomplex. Then the following holds:*

(1) If $\mathcal{K}'$ is a subdivision of $\mathcal{K}$, then there is a subcomplex $\mathcal{L}' \subseteq \mathcal{K}$ that is a subdivision of $\mathcal{L}$.

(2) If $\mathcal{L}'$ is a subdivision of $\mathcal{L}$, then there is a subdivision $\mathcal{K}'$ of $\mathcal{K}$ that has $\mathcal{L}'$ as a subcomplex.

Proof of Theorem 5.2.37. Since σ is a simplex, hence a PL-ball, and $\mathrm{lk}_{\mathcal{B}}(\sigma)$ is a PL-sphere, $\sigma * \mathrm{lk}_{\mathcal{B}}(\sigma)$ is a PL-ball by Lemma 5.2.38. Moreover, since $\partial\sigma$ is a PL-sphere, and $\mathcal{F}$ is a PL-ball, $\partial\sigma * \mathcal{F}$ is also a PL-ball by Lemma 5.2.38. Furthermore, their boundaries are both $\partial\sigma * \mathrm{lk}_{\mathcal{B}}(\sigma)$, so we will apply Lemma 5.2.39. To make this precise, we take geometric realizations $\mathcal{K}$ of $\mathcal{B}$ and $\mathcal{K}'$ of $\mathrm{starrep}_{\mathcal{B}}(\sigma, \mathcal{F})$. Since both $\mathcal{K}$ and $\mathcal{K}'$ contain geometric realizations L and L', respectively, of $\mathcal{B} \setminus \mathrm{star}_{\mathcal{B}}(\sigma)$, there is a natural isomorphism $\phi : V(L) \to V(L')$. This restricts to an isomorphism between the realizations $\mathcal{M}$ and $\mathcal{M}'$ of $\partial\sigma * \mathrm{lk}_{\mathcal{B}}(\sigma)$, which form the common boundary of the PL-balls we are replacing with one another. By Lemma 5.2.39, there are isomorphic subdivisions $\mathcal{C}$ of $\mathcal{M}$ and $\mathcal{C}'$ of $\mathcal{M}'$, with the isomorphism induced by $\|\phi\|$ on the boundary. By Lemma 5.2.40, there is a subdivision $\mathcal{B}'$ of $\mathcal{K}$ which contains $\mathcal{C}$ as a subcomplex. $\mathcal{B}'$ of course contains a subdivision $\mathcal{E}$ of $\mathcal{L}$, which corresponds via $\|\phi\|$ to a subdivision $\mathcal{E}'$ of $\mathcal{L}'$. This yields that $\mathcal{B}' = \mathcal{E} \cup \mathcal{C}$ is isomorphic to $\mathcal{E}' \cup \mathcal{C}'$. Hence $\mathrm{starrep}_{\mathcal{B}}(\sigma, \mathcal{F})$ has a subdivision isomorphic to a subdivision of a PL-ball, which implies that it must itself be a PL-ball, as desired. □

All that is left is to prove Lemma 5.2.39. To do this, we will need one more Lemma.

Lemma 5.2.41 (Corollary 1.6 in [18])**.** *If two geometric simplicial complexes have the same polyhedron, then they have a common subdivision.*

Proof of Lemma 5.2.39. Since the boundaries of $\mathcal{B}$ and $\mathcal{B}'$ are isomorphic, $\mathcal{B}$ and $\mathcal{B}'$ must have the same dimension d. Let Δ be a geometric realization of the abstract d-simplex. Since $\mathcal{B}$ and $\mathcal{B}'$ are PL-d-balls, there are geometric subdivisions $\mathcal{L}$ of $\mathcal{B}$ and $\mathcal{L}'$ of $\mathcal{B}'$, which are isomorphic to geometric subdivisions $\mathcal{D}$ and $\mathcal{D}'$ of Δ via isomorphisms $\eta : V(\mathcal{L}) \to V(\mathcal{D})$ and $\eta' : V(\mathcal{L}') \to V(\mathcal{D}')$, respectively. Then $\mathcal{L}$ and $\mathcal{L}'$ induce geometric subdivisions $\|\phi\|(\partial\mathcal{L})$ and $\partial\mathcal{L}'$ of $\partial\mathcal{B}'$, and by Lemma 5.2.41, these have a common subdivision $\mathcal{T}$.

Now let v be a point in the interior of Δ. Consider the subdivisions $\mathcal{E} = \|\eta\| \circ \|\phi^{-1}\|(\mathcal{T})$ and $\mathcal{E}' = \|\eta'\|(\mathcal{T})$ of $\partial\Delta$. Then $\mathcal{E} * v$ and $\mathcal{E}' * v$ are subdivisions of Δ. Since both $\mathcal{E}$ and $\mathcal{E}'$ are isomorphic to $\mathcal{T}$, there is a natural isomorphism $\xi : V(\mathcal{E}) \to V(\mathcal{E}')$ induced by the isomorphisms with $\mathcal{T}$, such that $\|\xi\| = \|\eta'\| \circ \|\phi\| \circ \|\eta^{-1}\|$. Extend this to an isomorphism $\hat{\xi}$ of $\mathcal{E} * v$ and $\mathcal{E}' * v$ by fixing

v. By Lemma 5.2.41, there is a common subdivision $\mathcal{F}$ of $\mathcal{E} * v$ and $\mathcal{D}$, and this corresponds to a subdivision $\mathcal{F}' = \left\|\hat{\xi}\right\|(\mathcal{F})$ of $\mathcal{E}' * v$. Now $\mathcal{F}'$ and $\mathcal{D}'$ have a common subdivision $\mathcal{G}'$ by Lemma 5.2.41, and this corresponds to a subdivision $\mathcal{G} = \left\|\hat{\xi}^{-1}\right\|(\mathcal{G}')$ of $\mathcal{F}$. Thus we have found isomorphic subdivisions $\mathcal{G}$ and $\mathcal{G}'$ of Δ, which are subdivisions of $\mathcal{D}$ and $\mathcal{D}'$, respectively. Hence $\mathcal{K} = \left\|\eta^{-1}\right\|(\mathcal{G})$ and $\mathcal{K}' = \left\|\eta'^{-1}\right\|(\mathcal{G}')$ are isomorphic subdivisions of $\mathcal{B}$ and $\mathcal{B}'$, respectively via an isomorphism $\psi : V(\mathcal{K}) \to V(\mathcal{K}')$ induced by $\left\|\eta'^{-1}\right\| \circ \left\|\hat{\xi}\right\| \circ \|\eta\|$. We claim that $\|\psi|_{\partial\mathcal{K}}\| = \|\phi\|$, which would complete the proof of the lemma. This is straightforward, as $\|\psi|_{\partial\mathcal{K}}\| = \left\|\eta'^{-1}|_{\partial\mathcal{D}'}\right\| \circ \left\|\hat{\xi}|_{\mathcal{E}}\right\| \circ \|\eta|_{\partial\mathcal{L}}\|$, and since $\hat{\xi}$ restricts to ξ on the boundary, we thus have $\|\psi|_{\partial\mathcal{K}}\| = \left\|\eta'^{-1}|_{\partial\mathcal{D}'}\right\| \circ \|\xi\| \circ \|\eta|_{\partial\mathcal{L}}\| = \|\phi\|$, as desired. □

5.3 Meshulam's Theorem

Now we're ready to give the proof of Meshulam's Theorem via triangulations. For convenience, here is the statement of the theorem:

Theorem 2.1.5. *Let G be a graph and let $e \in E(G)$. Then we have*

$$\operatorname{conn}(G) \geq \min(\operatorname{conn}(G - e), \operatorname{conn}(G \divideontimes e) + 1).$$

Proof. Let $k = \min(\operatorname{conn}(G - e), \operatorname{conn}(G \divideontimes e) + 1)$. Since G has an edge, it is nonempty, hence $G - e$ is nonempty, and thus $k \geq -1$ (since also $\operatorname{conn}(G \divideontimes e) \geq -2$). The theorem is trivial for $k = -1$, since G is nonempty by assumption, hence $\operatorname{conn}(G) \geq -1$. Therefore, assume $k \geq 0$.

We want to show that $\mathcal{I}(G)$ is k-connected, so we aim to apply Proposition 5.2.33. Therefore, consider a PL-j-sphere $\mathcal{S}$ for some integer j with $-1 \leq j \leq k$ and a simplicial map $f : V(\mathcal{S}) \to V(\mathcal{I}(G))$. If we can find a PL-filling $\mathcal{B}$ of $\mathcal{S}$ and a simplicial map $\hat{f} : V(\mathcal{B}) \to V(\mathcal{I}(G))$ extending f, then by Proposition 5.2.33 this would show that $\mathcal{I}(G)$ is k-connected.

We briefly outline how we will proceed. We start by using the fact that $\mathcal{I}(G - e)$ is k-connected to find a filling of $\mathcal{S}$ and a simplicial extension of f which maps to $\mathcal{I}(G - e)$. This extension might not be simplicial as a map into $\mathcal{I}(G)$, however, so call any simplex of the filling "ruined," if its image under the extension is e (since e is a simplex of $\mathcal{I}(G - e)$, but not a simplex of $\mathcal{I}(G)$). We replace ruined simplices one by one, starting with the highest dimensional ones and working our way down by utilizing the star-replacement operation referred to in Theorem 5.2.37. In the end, we will have a PL-filling of $\mathcal{S}$ and a simplicial extension of f with no ruined simplices, which will mean that the extension is also a simplicial map to $\mathcal{I}(G)$.

Since $V(\mathcal{I}(G - e)) = V(\mathcal{I}(G)) = V(G)$, and since $\mathcal{I}(G)$ is in fact a subcomplex of $\mathcal{I}(G - e)$ (every independent set of G is an independent set in $G - e$),

f is also a simplicial map from $\mathcal{S}$ to $\mathcal{I}(G-e)$. Since $\mathrm{conn}(\mathcal{I}(G-e)) \geq k$ by assumption, by Proposition 5.2.33, there is a PL-filling $\mathcal{B}$ of $\mathcal{S}$ and a simplicial map $\hat{f} : V(\mathcal{B}) \to V(\mathcal{I}(G-e))$ extending f. Call a simplex $\sigma \in \mathcal{B}$ "ruined," if $\hat{f}(\sigma) = e$. Clearly, any simplex of $\mathcal{B}$ is witness to the fact that $\hat{f}$ is not a simplicial map into $\mathcal{I}(G)$ if and only if it contains a ruined simplex. We will change the triangulation and the map until there are no more ruined simplices. Let $\sigma_1, \ldots, \sigma_n$ be the set of ruined simplices in order of decreasing dimension. Note that the dimension of any ruined simplex is at least 1, since it must have at least two vertices, one mapped to each endpoint of e. We will define a sequence of fillings $\mathcal{B}_0, \ldots, \mathcal{B}_n$, and simplicial maps $\hat{f}_i : V(\mathcal{B}_i) \to V(\mathcal{I}(G-e))$, with $\mathcal{B}_0 = \mathcal{B}$, $\hat{f}_0 = \hat{f}$, and $\mathcal{B}_n$ having no ruined simplices under $\hat{f}_n$ as follows.

Suppose $\mathcal{B}_i$ has already been defined and has ruined simplices $\sigma_{i+1}, \ldots, \sigma_n$. Let $d = \dim(\sigma_{i+1})$. By Proposition 5.2.32, $\mathrm{lk}_{\mathcal{B}_i}(\sigma_{i+1})$ is a PL-$(j-d)$-sphere, since σ_{i+1} is of course not contained in the boundary $\mathcal{S}$ of $\mathcal{B}_i$, because that part of $\hat{f}_i$ is equal to f, and hence simplicial into $\mathcal{I}(G)$. Note that $\hat{f}_i$ maps the vertices of the link to $V(\mathcal{I}(G \divideontimes e))$, because the vertices of the link are by definition in simplices together with every vertex of σ_{i+1} and these simplices are not ruined as σ_{i+1} is a maximal ruined simplex, and hence the images must not be adjacent to either endpoint of e in $G-e$. Because $\mathrm{conn}(\mathcal{I}(G \divideontimes e)) \geq k-1$ by assumption, and since $d \geq 1$ (hence $j-d \leq k-1$), by Proposition 5.2.33, there is a PL-filling $\mathcal{K}$ of $\mathrm{lk}_{\mathcal{B}_i}(\sigma_{i+1})$ together with a simplicial map $g : V(\mathcal{K}) \to V(\mathcal{I}(G \divideontimes e))$ extending the restriction of $\hat{f}_i$ to the link. Then let $\mathcal{B}_{i+1} = \mathcal{B}_i \setminus \mathrm{star}_{\mathcal{B}_i}(\sigma_{i+1}) \cup (\partial\sigma_{i+1} * \mathcal{K})$ and let $\hat{f}_{i+1}$ equal $\hat{f}_i$ on $V(\mathcal{B}_i \setminus \mathrm{star}_{\mathcal{B}_i}(\sigma_{i+1}))$ and equal g on the vertices of $\mathcal{K}$ ($\hat{f}_i$ is equal to g on the intersection). By Proposition 5.2.37, this is a PL-ball if $\mathcal{B}_i$ was. We claim that its only ruined simplices are $\sigma_{i+2}, \ldots, \sigma_n$. To see this, note that σ_{i+1} has been removed, $\sigma_{i+2}, \ldots, \sigma_n$ have been untouched (as their dimensions are at most d, and hence are not in the open star), and no new ruined simplices have been added, since all the new simplices include vertices from $\mathcal{K}$, which are all mapped to $V(G \divideontimes e)$, and hence are not ruined (even though they may contain a ruined simplex). Since $\mathcal{I}(G \divideontimes e)$ is a subcomplex of $\mathcal{I}(G)$, $\hat{f}_{i+1}$ is a simplicial map from $\mathcal{B}_{i+1}$ to $\mathcal{I}(G-e)$.

In the end, $\mathcal{B}_n$ has no ruined simplex, so $\hat{f}_n$ will be a simplicial map from $\mathcal{B}_n$ to $\mathcal{I}(G)$, which is what was wanted. Therefore, $\mathrm{conn}(G) \geq k$ and the theorem follows. $\square$

Zusammenfassung

Rysers Vermutung aus dem Jahre 1971 besagt, dass für einen r-partiten r-uniformen Hypergraphen $\mathcal{H}$ die Ungleichung $\tau(\mathcal{H}) \leq (r-1)\nu(\mathcal{H})$ erfüllt ist, wobei $\tau(\mathcal{H})$ die Knotenüberdeckungzahl und $\nu(\mathcal{H})$ die Matchingzahl bezeichnet. Diese Vermutung ist im Allgemeinen weiterhin offen. Fortschritte in verschiedenen Richtungen gab es unter anderem von Aharoni, Berger, Füredi, Haxell, Lovász, Mansour, Scott, Song, Tuza, Yuster, und Ziv. Im Spezialfall $r = 3$ hat Aharoni die Vermutung im Jahre 1999 bewiesen.

Das Hauptthema dieser Dissertation ist die Charakterisierung aller tripartiten 3-uniformen Hypergraphen $\mathcal{H}$, die $\tau(\mathcal{H}) = 2\nu(\mathcal{H})$ erfüllen, also der extremalen Hypergraphen für Rysers Vermutung für $r = 3$. Diese haben alle eine besondere Form, die wir "Home-Base" Hypergraphen nennen. Sie bestehen im Grunde aus $\nu(\mathcal{H})$ Teilhypergraphen mit $\tau = 2$ und $\nu = 1$, zusammen mit möglicherweise extra Hyperkanten, die diese Teile nur auf bestimmte Weise schneiden. Auf dem Weg zu einem Beweis dieser Charakterisierung finden wir auch eine Charakterisierung von bipartiten Graphen, die extremal für ein bestimmtes topologisches Problem sind.

Für beide Charakterisierungen benutzen wir Kenntnisse über die Topologie des sogenannten "Independence Complex" $\mathcal{I}$ von Kantengraphen. Deshalb untersuchen wir zunächst eine untere Schranke des Zusammenhangs von $\mathcal{I}(L(\mathcal{H}))$ in Abhängigkeit von $\tau(\mathcal{H})$. Wir vermuten, dass diese Schranke verbessert werden kann für r-partite r-uniforme Hypergraphen, und bestätigen diese Vermutung für den Spezialfall $r = 3$ und $\tau(\mathcal{H}) \leq 12$.

Ein Satz von Meshulam, welcher eine Aussage über den Zusammenhang von dem "Independence Complex" eines Graphen macht, spielt eine wichtige Rolle in unseren Beweisen. Der Beweis dieses Satzes den man in der Literatur findet ist algebraisch geprägt. Wir geben einen eher geometrischen Beweis, in dem wir bestimmte Triangulierungsmethoden benutzen. Die Richtigkeit dieser Methoden, die unter anderem von Szabó und Tardos benutzt werden, wurde vor ein paar Jahren in Frage gestellt. Im letzten Teil dieser Dissertation liefern wir einen ausführlichen Beweis für die Richtigkeit dieser Methoden.

Eidesstattliche Erklärung

Gemäß §7 (4) der Promotionsordung des Fachbereichs Mathematik und Informatik der Freien Universität Berlin versichere ich hiermit, dass ich alle Hilfsmittel und Hilfen angegeben und auf dieser Grundlage die Arbeit selbständig verfasst habe. Des Weiteren versichere ich, dass ich diese Arbeit nicht schon einmal zu einem früheren Promotionsverfahren eingereicht habe.

Berlin, den

Lothar Narins

Bibliography

[1] M. Adamaszek and J. A. Barmak, *On a lower bound for the connectivity of the independence complex of a graph*, Discrete Mathematics **311** (2011), 2566-2569.

[2] R. Aharoni, *Ryser's conjecture for tri-partite* 3*-graphs*, Combinatorica **21** (2001), no. 1, 1-4.

[3] R. Aharoni and E. Berger, *The intersection of a matroid and a simplicial complex*, Trans. Amer. Math. Soc. **358** (2006), 4895-4917.

[4] R. Aharoni, E. Berger, and R. Ziv, *Independent systems of representatives in weighted graphs*, Combinatorica **27** (2007), no. 3, 253-267.

[5] R. Aharoni, M. Chudnovsky, and A. Kotlov, *Triangulated spheres and colored cliques*, Discrete and Computational Geometry **28** (2) (2002), 223-229.

[6] R. Aharoni and P. Haxell, *Hall's theorem for hypergraphs*, J. Graph Theory **35** (2000), no. 2, 83-88.

[7] L. Babai and P. Frankl, *Linear Algebra Methods in Combinatorics*, (September 1992), University of Chicago.

[8] E. Berger and R. Ziv, *A note on the cover number and independence number in hypergraphs*, manuscript.

[9] E. R. Berlekamp, *On subsets with intersections of even cardinality*, Canad. Math. Bull. **12** (1969), 471-474.

[10] A. Björner, *Topological methods*, Handbook of Combinatorics, Vol. 2 (eds. R. L. Graham, M. Grötschel and L. Lovász), North-Holland, Amsterdam, 1995, 1819-1872.

[11] J. W. Cannon, *Shrinking cell-like decompositions of manifolds. Codimension three*, Annals of Mathematics. Second Series **110** (1) (1979), 83112.

[12] Z. Füredi, *Maximum degree and fractional matchings in uniform hypergraphs*, Combinatorica **1** (1981), 155-162.

[13] P. Haxell, *A condition for matchability in hypergraphs*, Graphs and Combinatorics **11** (1995), 245-248.

[14] P. Haxell, *A topology-free topological method*, submitted.

[15] P. Haxell, L. Narins, T. Szabó, *Extremal Hypergraphs for Ryser's Conjecture I: Connectedness of line graphs of bipartite graphs*, submitted, http://arxiv.org/abs/1401.0169.

[16] P. Haxell, L. Narins, T. Szabó, *Extremal Hypergraphs for Ryser's Conjecture II: Home-base Hypergraphs*, submitted, http://arxiv.org/abs/1401.0171.

[17] P. Haxell and A. D. Scott, *On Ryser's conjecture*, The Electronic Journal of Combinatorics 19.1 (2012), paper 23.

[18] J. F. P. Hudson, *Piecewise Linear Topology* (Benjamin, New York).

[19] J. L. Xiang Li and B. Szegedy, *On the logarithimic calculus and Sidorenko's conjecture*, http://arxiv.org/pdf/1107.1153v1.pdf.

[20] L. Lovász, *On minimax theorems of combinatorics*, Matematikai Lapok **26** (1975), 209-264 (in Hungarian).

[21] T. Mansour, C. Song, and R. Yuster, *A comment on Ryser's conjecture for intersecting hypergraphs*, Graphs and Combinatorics **25** (2009), 101-109.

[22] R. Meshulam, *Domination numbers and homology*, J. Combin. Theory Ser. A 102 (2003), 321-330.

[23] J. Milnor, *Construction of Universal Bundles, II*, Ann. of Math **63**, 430-436.

[24] A.F. Sidorenko, *A correlation inequality for bipartite graphs*, Graphs and Combinatorics **9** (1993), 201-204.

[25] M. Simonovits, *Extremal graph problems, degenerate extremal problems and super-saturated graphs*, Progress in Graph Theory (Waterloo, Ont., 1982), Academic Press, Toronto, ON (1984), 419-437.

[26] E. Sperner, *Neuer Beweis für die Invarianz der Dimensionszahl und des Gebietes*, Abh. Math. Sem. Univ. Hamburg **6** (1928), 265-272.

[27] T. Szabó and G. Tardos, *Extremal problems for transversals in graphs with bounded degree*, Combinatorica **26** (2006), 333-351.

[28] Zs. Tuza, *Some special cases of Ryser's conjecture*, manuscript, 1979.

[29] Zs. Tuza, *Ryser's conjecture on transversals of r-partite hypergraphs*, Ars Combinatoria **16** (1983), 201-209.

Curriculum Vitae

For reasons of data protection,
the curriculum vitae is not included in the online version.