Games on Infinite Graphs and Structural Properties of Directed and Bidirected Graphs

Dissertation

zur Erlangung des Doktorgrades an der Fakultät für Mathematik, Informatik und Naturwissenschaften der

Universität Hamburg

vorgelegt am Fachbereich Mathematik von

Florian Gut

Hamburg, September 2023

Universität Hamburg
DER FORSCHUNG | DER LEHRE | DER BILDUNG

Betreuer
Prof. Dr. Nathan Bowler

Gutachter
Prof. Dr. Nathan Bowler (Erstgutachter)
Dr. Attila Joó (Zweitgutachter)
Dr. Jan Kurkofka (Drittgutachter)

Mitglieder der Prüfungskommission
Prof. Dr. Nathan Bowler
PD Dr. Stefan Geschke (Schriftführung)
Dr. Attila Joó
Prof. Dr. Janko Latschev (Vorsitz)
Prof. Dr. Natalie Neumeyer (Stellv. Vorsitz)

Die Disputation fand am 23. Januar 2024 statt.

Contents

Chapter 1

Introduction

In this thesis we will touch on various interesting topics in structural graph theory, divided into two parts. In Part I we delve into a fascinating field with many recent results: two player games on infinite graphs. We present winning strategies in different games that all fall under this category. In Part II we concern ourselves with structural properties of finite as well as infinite directed graphs. We give an overview of the different topics in the following.

1.1 Games on Graphs

Games have been of interest to mathematicians for centuries. The field was sparked by the analysis of historic games such as tic-tac-toe. This is a finite game and naturally, this was the first class of games analysed. In recent decades infinite games have increasingly drawn the attention of researchers. An intuitive starting point is just moving the rules of a finite game to an infinite board. Consider the aforementioned Tic-tac-toe. Its counterpart on an infinite board became known as *unrestricted* 3*-in-a-row* and was further generalised to *unrestricted* n*-in-a-row* (see e.g. Beck [Bec08]). This often yields interesting insights, as some results in the finite version simply rely on the fact that there are only finitely many possible plays. These types of games are called *semi-infinite*; a comprehensive overview of the known results about such games can be found in [Bec08].

Games also give rise to an interesting field in set theory, which is e.g. described in a book by Moschovakis [Mos09, Chapter 6]: one may assume as an axiom that every game is determined, i.e. that for any two player game with complete information at least one of the players has a winning strategy. It is known that this is not consistent with the axiom of choice, but it nevertheless provides a good framework to study descriptive set theory, itself an essential field of research in present-day mathematical logic.

One classic game is the following two player game, which we will call the *strong* (G, H)*-game*. The players alternately colour exactly one uncoloured edge of a sufficiently large complete graph G, the *board*, in their respective colour. They agree upon some graph H as the objective of the game. The game ends as soon as H is contained as a subgraph in the graph induced by some player's claimed edges. The word *strong* in the name of the game indicates that both players have the same objective (to build a copy of H), thus the two players' roles differ only in who plays first. This is to distinguish it from the *Maker-Breaker* variant of the game, in which only one of the players (Maker) is trying to build a copy of H and the other (Breaker) is simply trying to prevent this. In particular, Breaker does not win simply by virtue of building his own copy of H first.

Finite strong games have been extensively studied, and in particular the strong (G, H)-game is fairly well understood. Indeed, as long as the board is sufficiently large, any such game is a first player win: due to Ramsey's Theorem (see e.g. Diestel [Die17, Theorem 9.1.1]), after all edges of G have been claimed, there will always be a copy of H contained in one of the players' graphs if $|G| \geq R(|H|)$. So one of the players must have a winning strategy, and using a technique called *strategy stealing* (see Hefetz, Krivelevich, Stojaković and Szabó [HKSS14, Theorem 1.3.1]) it cannot be the second player. The argument can be roughly sketched as follows: suppose the second player has a winning strategy. Then the first player can make an arbitrary move in her first turn and from then on play according to the winning strategy as if she is the second player.

Note that even though we now know that there is a winning strategy for the first player if $|G| \geq R(|H|)$, we only have an abstract argument and can not deduce an actual strategy according to which the first player should play. In particular, this argument does not show the existence of an upper bound only depending on H for the number of moves the first player needs to win.

The problem of establishing such bounds is in general very hard, we resolve the game where H is the complete graph on four vertices in Chapter 3, which answers a question by Beck, [Bec08, Open Problem 4.6]. The strategy and the proof for this game is already very complex so that we utilise a computer algorithm that enables us to better check every possible course of the game. By contrast, such bounds are well known for the Maker-Breaker (G, H)-game (see [HKSS14, Chapter 2]). Indeed, Maker-Breaker variants are usually easier to analyse, and are often investigated as a preliminary step before analysing the strong game.

In contrast to the fertility of the study of finite combinatorial games, i.e. games played on an infinite board with perfect information where the players do moves sequentially, infinite combinatorial games have thus far proved barren ground, in that the strong games are too hard to analyse and the Maker-Breaker games are too easy. Let us consider first the strong (G, H)-game, but now played on an infinite complete graph. Since there are now infinitely many edges on the board to choose from these edges need not be exhausted in the course of the game, and there is no guarantee that the players will ever even claim all edges of a finite complete subgraph K^n between them. Thus Ramsey's Theorem is no longer applicable. Play could continue forever without either player ever winning!

So even though we can still use a strategy-stealing argument to rule out the existence of a winning strategy for the second player, it is possible that he has a strategy to force a draw. Although it might seem implausible that this could really happen, in fact for the analogous games played on 5-regular hypergraphs rather than graphs cases are known in which the second player can force a draw. Such an example was constructed by Hefetz, Kusch, Narins, Pokrovskiy, Requilé and Sarid in [HKN+17].

In fact, it is folklore that the existence of a winning strategy for the first player in the infinite strong (G, H)-game is equivalent to the existence of a finite upper bound on the number of moves the first player needs to win in the finite strong (G, H)-games, a straightforward proof by a compactness argument can be found in (Leader [Lea08, Proposition 4]), and we saw above that such problems currently seem intractable.

Then again, the Maker-Breaker (G, H)-game is trivial; the winning strategies for the finite variants also work in the infinite variant. However, Erde recently noticed that there are interesting Maker-Breaker (G, H)-games on infinite graphs; one just has to take H itself to be infinite [Erd19].

We believe that the investigation of such infinite Maker-Breaker games will prove very fruitful. In Chapter 4 we begin this investigation by considering a few simple variants of the $(K^{\aleph_0}, K^{\aleph_0})$-game. We will present a winning strategy for Maker in the basic version of the game in Section 4.1.

A natural variation arises if we colour the vertices of the board beforehand and demand that Maker respects this colouring in such a way that her $K^{\aleph_0}$ again contains infinitely many vertices of every colour. In fact some variations introduce so much complexity that there is a winning strategy for Breaker. We investigate these types of games in Section 4.2 and Section 4.3.

Another variation emerges by introducing additional structural elements such as an order of the vertex set. In Section 4.4 we consider the complete graph with the rational numbers $\mathbb{Q}$ as the vertex set. We call this graph $K^{\mathbb{Q}}$ and investigate the $(K^{\mathbb{Q}}, K^{\mathbb{Q}})$-game.

If one moves beyond countable graphs and allows G and H to be uncountable, the set-theoretic framework becomes much more important. We consider Maker-Breaker games on uncountable boards in Chapter 5.

1.2 Directed and Bidirected Graphs

In Part II we study different structural properties of directed graphs, a generalisation of graphs, and even of bidirected graphs, a generalisation of directed graphs.

1.2.1 Ubiquity

A (di)graph H is called $\trianglelefteq$-*ubiquitous* for a binary (di)graph relation $\trianglelefteq$ if any (di)graph G that contains k disjoint copies of H for every $k \in \mathbb{N}$ also contains infinitely many disjoint copies of H with respect to $\trianglelefteq$. Possible relations for $\trianglelefteq$ are e.g. the subgraph, topological minor or minor relation for graphs or the subdigraph relation for digraphs.

Halin started the investigation of ubiquity in graphs with his landmark result that rays are subgraph-ubiquitous in [Hal65]. Andreae conjectured that every locally finite connected graph is minor-ubiquitous after studying minor-ubiquity in [And02, And13]. Noteworthy progress towards this conjecture was recently achieved by Bowler, Elbracht, Erde, Heuer, Pitz and Teegen in a series of papers [BEE+22, BEE+18, BEE+20], in which they proved, among several other results, that all trees are topological minor-ubiquitous. Throughout the years several results proving and disproving the ubiquity of certain graphs have been published, including results concerning different notions of ubiquity as in [BCP15, Kur21].

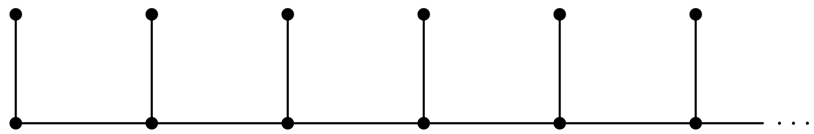

Figure 1.1: The comb, a graph that is not subgraph-ubiquitous.

An example for a graph that is not subgraph-ubiquitous is the comb [And02] (see Figure 1.1). For graphs that are not topological minor-ubiquitous, see [And13]. Very recently, Carmesin provided an example of a locally finite graph that is not minor-ubiquitous in [Car23]. We remark that this does not contradict Andreae's conjecture as the graph is not connected.

In [BCP15] Bowler, Carmesin and Pott first suggested the topic of ubiquity in digraphs by asking whether any digraph containing arbitrarily many edge-disjoint directed double rays also contains infinitely many of them.

In Chapter 6 we take on the quest of investigating ubiquity in digraphs by examining two simple classes of directed graphs. These classes are *oriented rays* which are digraphs whose underlying undirected graphs are rays and *oriented double rays* which are digraphs whose underlying undirected graphs are double rays. In both of these classes a *turn* is a vertex that is incident with either two heads or two tails. We prove that oriented rays are ubiquitous if and only if they have finitely many turns in Section 6.1. In Section 6.2 we prove that any oriented double ray other than the directed double ray (that is the double ray in which every vertex has in- and out-degree one) is ubiquitous if and only if it has an odd number of turns.

1.2.2 Flames

The study of minimal subgraphs witnessing a connectivity property is an important field in graph theory. A key early theorem was proved by Lovász: let $D = (V, E)$ be a finite digraph and let $r \in V$. The local connectivity $\kappa_D(r, v)$ from r to v is defined to be the maximal number of internally disjoint r–v paths in D. A spanning subdigraph

L of D with $\kappa_L(r,v) = \kappa_D(r,v)$ for every $v \in V - r$ must have at least $\sum_{v\in V-r} \kappa_D(r,v)$ edges. Lovász proved that, maybe surprisingly, this lower bound is always sharp for finite digraphs.

Let us call a spanning subdigraph L of a finite 'r-rooted' digraph $D = (V,E)$ large (with respect to D) if L preserves the local connectivity from the root, i.e. $\kappa_D(r,v) = \kappa_L(r,v)$ for every $v \in V - r$. Furthermore, a finite r-rooted digraph $D = (V,E)$ is defined to be a vertex-flame if $\kappa_D(r,v) = |\mathsf{in}_D(v)|$ for every $v \in V - r$, where $\mathsf{in}_D(v)$ is the set of incoming edges of v. Using this terminology, Lovász' Theorem says that every finite rooted digraph admits a large vertex-flame. It was shown by Calvillo Vives in [Cal78], that in every finite r-rooted digraph D every vertex-flame subgraph (with respect to root r) can be extended to a large vertex-flame of D. This was further generalised by Joó in [Joó21]. He proved that the edge sets of the vertex-flame subdigraphs of a finite rooted digraph $D = (V,E)$ are the feasible sets of a greedoid on E whose bases are exactly the large vertex-flames in Lovász' Theorem.

There are many results in infinite graph theory that were first proved only for finite graphs and a deeper understanding of the underlying concept and more complex arguments were necessary to generalise them to infinite ones. Sometimes even the appropriate formulation of the problem for infinite graphs is already non-trivial because the equivalent forms in the finite case might not be equivalent in general. For example it is well-known and easy to prove that the edge set of a finite graph can be partitioned into cycles if and only if there is no vertex with odd degree. The condition can be rephrased as the non-existence of odd cuts. A deep theorem of Nash-Williams [NW60, p. 235 Theorem 3] says that the reformulated condition is actually sufficient to partition the edges of a graph of any size into cycles, whereas the original condition is insufficient which is for example witnessed by the double ray. Other classical results fail at some cardinalities; for example every countable graph admits a normal spanning tree but an uncountable complete graph does not.

For one of the most influential theorems in infinite graph theory the necessity of choosing the 'right' formulation was also true. The result in question is the generalisation of Menger's Theorem for arbitrary graphs which play an important role in the main result of Chapter 7. Erdős observed that the maximal size of a system of pairwise disjoint paths in a graph between two prescribed vertex sets and the minimal size of a vertex set meeting all the paths between these two sets is the same regardless of the size of the graph. He realised that considering this min-max formulation of Menger's Theorem does not lead to a really strong infinite generalisation. Indeed, choosing the path-system inclusion-wise maximal and taking all the vertices of these paths as a separator is suitable whenever the path-system in question is infinite, although this separator is clearly way too 'big' in a structural sense. Erdős conjectured a structural infinite generalisation of Menger's Theorem, (see Theorem 2.3, it was known as the Erdős-Menger conjecture) which was eventually proved after several partial results by Aharoni and Berger [AB09, Theorem 1.6].

As in the case of the Erdős-Menger conjecture, quantities are not appropriate to obtain the right infinite generalisation of Lovász Theorem, thus we need to look at the structural properties of L.

In Chapter 7 we extend the definition of vertex-flames to rooted digraphs of any size by demanding for every $v \neq r$ the existence of internally disjoint r–v paths covering *all* incoming edges of v instead of just the equality of the in-degree of v and the local connectivity from r to v. The condition $\kappa_D(v) = \kappa_L(v)$ translates to the existence of a maximal-sized internally disjoint r–v path-system $\mathcal{P}$ of D that lives in L. We strengthen this by asking $\mathcal{P}$ to be 'big' not just cardinality-wise but in the structural sense from the Erdős-Menger conjecture. Namely, we demand the existence of a separation of v from r in D that can be obtained by choosing exactly one edge or one internal vertex from each path in $\mathcal{P}$. A spanning subdigraph is called large if there is such a $\mathcal{P}$ for every $v \neq r$. We will see that in a large vertex-flame for each v the path-system witnessing largeness and the path-system covering the incoming edges of v can be chosen to be the same.

1.2.3 Directed Separations

Separations in (di-)graphs are a basic structure that has long been of interest for mathematicians. In undirected graphs for small k, there are simple and canonical combinatorial structures displaying all k-separations of a given k-connected undirected graph and the relationships between them. For $k = 0$ this is trivial; it is simply the decomposition of the graph into its connected components. For $k = 1$ this is the block-cut decomposition [Die17, Chapter 3.1] and for $k = 2$ it is the Tutte decomposition [Tut66]. Recent work of Carmesin and Kurkofka has developed an analogous such structure with $k = 3$ [CK23].

For digraphs the situation is very different, with almost nothing being known. In the case $k = 0$ we again have the decomposition of the digraph into its strongly connected components, together with the partial order on those components given by the reachability relation. However, already for $k = 1$ previous progress on this problem was limited to a partial result by Lovász [Lov87][1].

In order to explain what Lovász' result means for digraphs, we need to consider an operation on strongly connected digraphs. Let D be a digraph, and let (A, B) be a directed 1-separation of D with separator v. Let D_A be the digraph obtained from D by contracting all of B onto v and let D_B be the digraph obtained by contracting all of A onto v. We say that D_A and D_B are obtained from D by *pulling it apart* along (A, B).

Now imagine taking a strongly connected digraph, pulling it apart along some directed 1-separation, then continuing to pull those parts apart for as long as possible until you are left with a list of strongly 2-connected digraphs. In this context, Lovász'

[1] In fact Lovász, like us, works in the more general framework of matching covered graphs.

result says that it does not matter in what order you carry out this process or which 1-separations you pick, you will always end up with the same list up to rearrangement and digraph isomorphism.

To understand the limitations of this result, let's compare it with the Tutte decomposition. An analogous decomposition procedure would be the following. Given a 2-connected *undirected* graph G and a 2-separation (A, B) of G with separator $\{x, y\}$, we could define G_A and G_B to be the graphs obtained from $G[A]$ and $G[B]$ respectively by adding the edge xy to each of them. We could once again consider a procedure of repeatedly pulling a graph apart along such separations until we are left with a list of 3-connected graphs, and it follows from the Tutte decomposition that you will always end up with the same list of 3-connected graphs up to list rearrangement and graph isomorphism.

However, this adaption to 3-connected graphs yields a weaker result than the original Tutte decomposition in a number of ways. First, the elements of the list are only given up to isomorphism, but in the Tutte decomposition they are given as the torsos of a canonical tree-decomposition of G. Second, the adapted decomposition procedure cuts up the 2-connected graph too much. In the Tutte decomposition, some of the torsos are cycles, and these are sensibly not cut up any further because there is no way to do so canonically. But the adapted decomposition procedure will happily cut up an n-cycle into $n - 2$ triangles, and in so doing lose the possibility of finding canonical representatives in the original graph. Third, the Tutte decomposition provides a global tree structure along which the parts are arranged, and this information is not contained in the list of 3-connected graphs obtained by the adapted decomposition procedure.

Lovász' result has the same 3 limitations. The elements of the list are only given up to isomorphism, not as canonical structures within the original digraph. It cuts any directed n-cycle up into $n - 1$ directed 2-cycles, although there is no hope of finding canonical representatives for these in the original graph. Finally, it does not provide any global structure along which the parts are arranged.

In Chapter 8, we resolve the first two limitations of Lovász' result, but not the third. More precisely, we find canonical structures, which we call torsoids, within the digraph corresponding to the parts into which it would be cut by Lovász' procedure. However, we do not cut up directed cycles any further, since there would be no way to find canonical structures representing the parts. As an illustration of the canonicity of our results and the deeper structural understanding they provide, we are able to extend Lovász' results to infinite digraphs.

Having discussed the limitations of Lovász' result, it is worth noting one major advantage. He was able to prove his result in the more general context of *matching covered graphs* (that is, connected graphs such that every edge is contained in at least one perfect matching). There is a well-understood correspondence between strongly connected digraphs and bipartite matching-covered graphs [McC00, RST99], under which directed 1-separations correspond to tight cuts. Lovász' result deals with the process of cutting up (not necessarily bipartite) matching covered graphs along their tight cuts.

Since all our arguments also work just as easily in this more general context, and since there has been a renewed interest in the structure of matching covered graphs in recent years [NT06, NT07, HRW19, GW21], we also phrase all of our results in Chapter 8 in these more general terms.

1.2.4 Menger's Theorem in Bidirected Graphs

A fundamental result in the field of graph theory is due to Karl Menger [Men27] which is nowadays just known as Menger's Theorem. It gives an important insight into the connectivity of two sets of vertices in digraphs. In its common form, Menger's Theorem asserts that the minimum number of vertices separating two sets of vertices X and Y in a digraph D is the same as the maximum number of disjoint X–Y paths in D, see e.g. [BJG08, Theorem 7.3.1] in the book of Bang-Jensen and Gutin.

A slight structural strengthening of Menger's Theorem was proved by Böhme, Göring and Harant [BGH01]: for a set $\mathcal{P}$ of k disjoint X–Y paths, there is either a set of k vertices separating X from Y, one vertex on each path in $\mathcal{P}$, or a set of $k+1$ disjoint X–Y paths where k of them use the same start vertices in X as the paths of $\mathcal{P}$.

In Chapter 9 we make this stronger form of Menger's Theorem available in the realm of bidirected graphs. Bidirected graphs were first introduced by Kotzig in a series of papers [Kot59a, Kot59b, Kot60]. These objects can be understood as a generalisation of undirected and directed graphs. They can be obtained from undirected graphs by assigning to each endpoint of every edge one of two signs, $+$ or $-$. Directed graphs can then be envisioned as bidirected graphs in which each edge has the sign $+$ at one endpoint and $-$ at the other.

Bidirected graphs have recently received increased attention by the graph theoretic community with both structural [AK06, Wie22] and algorithmic results [AFN96, BZ06]. Especially, Wiederrecht proposed bidirected graphs as another angle of attack to investigate the structure of undirected graphs with perfect matchings. This approach is inspired by a recent proof of Norin's Matching Grid Conjecture for bipartite graphs with perfect matchings, which was given by Hatzel, Rabinovich and Wiederrecht [HRW19]. Their proof applies the Directed Grid Theorem [KK15] using the structural relation of digraphs and bipartite graphs with perfect matchings. In the same structural sense, bidirected graphs are in a one-to-one correspondence with general undirected graphs with perfect matchings; for a detailed introduction, we refer the reader to [Wie22].

Unfortunately, Menger's Theorem does not hold if we simply replace 'directed' by 'bidirected'. The main reason for this is an intricate complication in the structure of bidirected graphs: unlike in directed graphs, a walk between two fixed vertices need not contain a path between them. This property of bidirected graphs in particular prevents a direct transfer of the usual proofs such as the one given in [BGH01].

To overcome this complication, we introduce a sufficient condition on the set $\mathcal{X}$ of 'signed' start vertices of our paths, where a *signed vertex* is a pair of a vertex and a sign $+$ or $-$: we require $\mathcal{X}$ to be 'clean', that is, we forbid the existence of non-trivial $\mathcal{X}$–$\mathcal{X}$ paths. In particular, this condition is satisfied by any set of start vertices in a digraph.

Chapter 2

Tools and Terminology

Throughout this thesis we assume familiarity with the basic notions of graph theory. For a thorough introduction to graph theory we refer to the book ‘Graph Theory’ by Diestel [Die17] and for notions regarding directed graphs we refer to the book ‘Digraphs: Theory, Algorithms and Applications’ by Bang-Jensen and Gutin [BJG08]. Any definition not introduced here can be found in these books.

2.1 Basic Concepts

In the entire thesis we assume $\mathbb{N} = \{0, 1, 2, \ldots\}$. For numbers $k \leq \ell$ we write $[k, \ell] := \{n \in \mathbb{N} \colon k \leq n \leq \ell\}$, and we abbreviate $[k] := [1, k]$. For a set X and κ a cardinal we denote by $[X]^\kappa$ the set of κ-element subsets of X. For a finite sequence $(a_1, a_2, \ldots, a_n)$ in X and $a \in X$ we define $(a_1, a_2, \ldots, a_n)^\frown a := (a_1, a_2, \ldots, a_n, a)$. Let Y also be a set and $X' \subseteq X$. For a map $f : X \to Y$ we define $f[X'] := \{f(x) \colon x \in X'\}$. If X and Y are disjoint we write $X \sqcup Y$ for their union.

A classic result of graph theory is due to Karl Menger [Men27] and is nowadays just known as Menger’s Theorem. We use it repeatedly throughout this thesis.

Theorem 2.1 ([Die17, Theorem 3.3.1])**.** *Let G be a graph and X and Y be subsets of the vertex set of G. Then the minimum number of vertices separating X from Y is the same as the maximum number of disjoint X–Y-paths.*

2.2 Games on Graphs

In Part I we study two closely connected versions of a two player game, called the *(G, H)-game*. We will investigate the *strong* version and the *Maker-Breaker* version of the (G, H)-game, both of which we will explain below. In either version the players alternately claim exactly one unclaimed edge of a graph G, which we call the *board*. They claim an edge by choosing an uncoloured edge of the board and colouring it in their respective colour. The players agree upon some graph H before the start of the game. In the strong (G, H)-game we call the players *first player* and *second player* and the respective aim of the players is to be the first to have a copy of H contained as a subgraph in the graph induced by their coloured edges. We assume that the first player chooses the colour red and refer to that player as *she* or *her*. Likewise we assume that the second player chooses the colour blue and refer to that player as *he* or *him*. In case the board G is infinitely large the game could continue indefinitely, in which case the second player is declared the winner of the game. In the Maker-Breaker (G, H)-game we call the players *Maker* and *Breaker*. Akin to the strong version we refer to Maker as *she* or *her* and to Breaker as *he* or *him*. We will assume that Maker colours her edges in magenta and that Breaker uses the colour blue. Maker has the same objective as in the strong version, i.e. to have a copy of H contained in the subgraph of the board induced by her coloured edges. Breaker has the task to stop Maker from achieving her goal. In case the board G is infinitely large Breaker is declared the winner of the game if the game continues indefinitely. We also call the (strong/Maker-Breaker) $(K^{\aleph_0}, H)$-game the (strong/Maker-Breaker) *H-building game* respectively.

For the rest of this section when we talk about a *player* this stands for a player in either version of the (G, H)-game, that is it can stand for the first player, the second player, Maker or Breaker. Let $c \in \{FP, SP, M, B\}$ where FP stands for the first player, SP for the second player, M for Maker and B for Breaker. When we say that a player *c-connects* a vertex v to a vertex w, we mean that the respective player claims the edge vw. We may also shorten this to simply *connect* if the player is clear from the context and we mean the same when we say that a player *plays from v to w*. At any point of a game we assume $E(G_c)$ to be the edges that the respective player has claimed up to that point and $V(G_c)$ to be the vertices of the board that are incident with at least one edge of $E(G_c)$. With this we define $G_c := (V(G_c), E(G_c))$. $N_c(v)$ are the *neighbours* of a vertex $v \in V(G_c)$ in G_c and the *c-degree* $\deg_c(v)$ of a vertex $v \in G$ is $|N_c(v)|$ if $v \in G_c$ and 0 otherwise. Accordingly, we will say that two vertices v and w are *c-connected* or *c-adjacent*, if $vw \in G_c$. If we talk about a *fresh* vertex we mean a vertex $v \in V(G) \setminus V(G_{\mathrm{FP}} \cup G_{\mathrm{SP}})$ in the strong (G, H)-game and a vertex $v \in V(G) \setminus V(G_M \cup G_B)$ in the Maker-Breaker (G, H)-game.

By a *turn* of a player we mean that that player chooses an edge of the board that has not been coloured by either of the players and colours that edge in her or his respective colour. We alternatively also call a turn a *move* interchangeably and also say that that player *claims* an edge. A sequence of one move by (the first player/Maker) and then one by (the second player/Breaker) respectively is called a *round*. The turns are indexed by ordinals by their chronological order and so are the rounds. We make use of that and assign indices to the vertices accordingly, i.e. v_k is the k-th vertex that (the first player/Maker) adds to her subgraph. In case she claims an edge incident with two fresh vertices, we indicate a choice how the vertices are enumerated, which will only happen on her first turn in our constructions.

When we want to prove that a game is a win for Breaker, we shall always do this by means of a *pairing strategy*. That is, we will define a family of disjoint pairs of edges from $E(G)$ with the intention that whenever Maker claims one edge from a pair Breaker claims the other one in his following turn. It then suffices to verify, for the (G, H)-game in question, that any subgraph of the board isomorphic to H must include both edges of at least one such pair.

2.2.1 Games on Uncountable Boards

In Chapter 5 we concern ourselves with Maker-Breaker games on uncountable boards. For these games we alter the rules such that the game terminates only when all edges are claimed. The objectives for Maker and Breaker stay the same. For a cardinal κ the complete graph on κ is $K^\kappa := (\kappa, [\kappa]^2)$. For two cardinals the complete bipartite graph with vertex classes of size λ and κ is denoted by $K^{\lambda,\kappa}$, its vertex set is $(\lambda\times\{0\})\cup(\kappa\times\{1\})$ and its edge set is $\{(\alpha, 0),(\beta, 1)\colon \alpha < \lambda, \beta < \kappa\}$. Note that the vertex sets of these graphs are already well-ordered, and so we generally do not need to invoke the axiom of choice. Furthermore, in Subsection 5.2.3 we consider the Maker-Breaker $(K^\kappa, K^{\mathsf{club}})$-game, that is the Maker-Breaker game on K^κ where it is the goal for Maker to build a complete graph whose vertex set is a closed unbounded subset of κ which we denote by K^{club}.

2.3 Directed and Bidirected Graphs

For a digraph D and vertices $v, w \in D$, we write $d_D(v, w)$ for the distance of v and w in the underlying undirected graph. Let D' also be a digraph. We write $D \cong D'$ if D is isomorphic to D' and $D \leq D'$ if D is isomorphic to a subdigraph of D'.

We denote the set of incoming edges of a vertex set X by $\mathsf{in}_D(X)$ and $\mathsf{out}_D(X)$ stands for the set of the outgoing edges. For the set of in-neighbours of X (i.e. the tails of the edges in $\mathsf{in}_D(X)$) we write $N^-_D(X)$ and the out-neighbours $N^+_D(X)$ are defined analogously. The subdigraph induced by a vertex set U is $D[U]$ and $H \subseteq D$ expresses that H is a subgraph of D. We define $D_0 \cap D_1 := (V_0 \cap V_1, E_0 \cap E_1)$ if $D_i = (V_i, E_i)$ are digraphs.

Paths in Digraphs

An *oriented path* is a digraph whose underlying undirected graph is a path and we call an oriented path a *dipath* or simply *path* if all its edges are consistently oriented, i.e. each vertex has in- and out-degree at most 1. A dipath is *trivial* if it consists of a single vertex.

For vertex sets X and Y we say that a path is an X–Y path if exactly its first vertex is in X and exactly its last vertex is in Y. In case X or Y is a singleton, we replace the vertex set by the vertex. For paths P and Q with $v \in V(P) \cap V(Q)$, let PvQ be the digraph consisting of the initial segment of P up to v and the terminal segment of Q from v. We call a set of paths $\mathcal{P}$ a *path-system* and say that it is *disjoint* if the paths in it are pairwise vertex-disjoint. We define *internally disjoint* similarly except that the first and last vertices are allowed to be shared. We denote the united vertex set and edge set of the paths in $\mathcal{P}$ by $V(\mathcal{P})$ and by $E(\mathcal{P})$ respectively. Let us write $V^-(\mathcal{P})$ and $V^+(\mathcal{P})$ for the respective set of the first and last vertices of the paths in $\mathcal{P}$. We define $E^-(\mathcal{P})$ and $E^+(\mathcal{P})$ similarly with edges but only for path-systems without trivial paths. We write simply $\mathsf{in}_{\mathcal{P}}(v)$ for the set of the incoming edges of v in the digraph $(\bigcup_{P\in\mathcal{P}} V(P), \bigcup_{P\in\mathcal{P}} E(P))$. A *$v$-fan* is a system of paths sharing only their initial vertex v. A *v-infan* is what we obtain by reversing the edges of a v-fan. A set $X \subseteq V - v$ is *linked from v* in D if there is a v-fan $\mathcal{P}$ in D with $V^+(\mathcal{P}) = X$. Similarly, X is *linked to v* if there is a v-infan $\mathcal{P}$ with $V^-(\mathcal{P}) = X$.

Menger's Theorem, Theorem 2.1, is also true in digraphs. A slight structural strengthening of Menger's Theorem was proved by Perfect [Per68].

Theorem 2.2 ([Per68]). *Let X and Y be sets of vertices of a digraph D, and let $P_1, \ldots, P_k$ be vertex-disjoint X–Y paths in D where P_i starts in $v_i \in X$ for $i \in [k]$. Then precisely one of the following is true:*

(1) *There is a set S of k vertices of D such that $D - S$ contains no X–Y path.*

(2) *There are $k+1$ vertex-disjoint X–Y paths $P'_1, \ldots, P'_{k+1}$ in D where P'_i starts in v_i for $i \in [k]$.*

A structural generalisation to infinite digraphs was conjectured by Erdős, proved by Aharoni and Berger in [AB09] and reads as follows:

Theorem 2.3 (Aharoni & Berger [AB09, Theorem 1.6]). *For every digraph D and $X, Y \subseteq V(D)$, there is a system $\mathcal{P}$ of pairwise disjoint X–Y paths in D such that one can choose exactly one vertex from each path in $\mathcal{P}$ in such a way that the resulting vertex set S separates Y from X in D.*

2.3.1 Bidirected Graphs

A *bidirected graph* $B = (G, \sigma)$ consists of a graph $G = (V, E)$, a corresponding set of *half-edges* defined as

$$\mathbf{E}(B) := \{(u, e) \colon e \in E \text{ and } u \in e\}$$

and a *signing* $\sigma \colon \mathbf{E} \to \{+, -\}$ assigning to each half-edge (u, e) its *sign* $\sigma(u, e) := \sigma((u, e))$; we say that e *has sign* $\sigma(u, e)$ *at* u. Then $V(B) := V$ is the *vertex set* of B and $E(B) := E$ is its *edge set*. We refer to the elements of $V(B)$ and $E(B)$ as the *vertices* and the *edges* of B, respectively.

For technical simplification, we do not allow bidirected graphs to have distinct edges e and f that have the same endvertices and the same signs at them.

A *signed vertex* of B is a pair (v, α) of a vertex v of B and a sign $\alpha \in \{+, -\}$. Given a set $\mathcal{X}$ of signed vertices, we write $V(\mathcal{X}) := \{v \in V(B) \colon \exists\, \alpha \in \{+, -\} \colon (v, \alpha) \in \mathcal{X}\}$.

An *oriented edge* $\vec{e}$ of a bidirected graph B is formally defined as a triple (e, u, v) where e is an edge of B with incident vertices $u, v \in V(B)$; we call e its *underlying edge*, u its *startvertex* and v its *endvertex*, and think of $\vec{e}$ as orienting e from u to v. The edge e has precisely two *orientations*, one with startvertex u and endvertex v and the other one with startvertex v and endvertex u. We denote the two orientations of e as $\vec{e}$ and $\overleftarrow{e}$; there is no default orientation of e, but if we are given one of them as $\vec{e}$, say, then the other one is written as $\overleftarrow{e}$. Given an oriented edge $\vec{e}$ we conversely write e for its underlying edge. Given a set A of edges of B, we write $\overleftrightarrow{A}$ for the set of all orientations of edges in A.

Walks, Trails and Paths in Bidirected Graphs

Let $B = (G, \sigma)$ be a bidirected graph. A sequence

$$W = v_0 \vec{e}_1 v_1 \vec{e}_2 v_2 \dots v_{\ell-1} \vec{e}_\ell v_\ell$$

of vertices $v_i \in V(B)$ and oriented edges $\vec{e}_j \in \overleftrightarrow{E}(B)$ is a *walk* W of *length* ℓ in B if the oriented edge $\vec{e}_j$ has startvertex v_{j-1} and endvertex v_j for $j \in [\ell]$ and we have $\sigma(v_i, e_i) \neq \sigma(v_i, e_{i+1})$ for $i \in [\ell - 1]$. A *subwalk* of W is a walk of the form $v_i \vec{e}_i \dots \vec{e}_{j-1} v_j$ for some $i \leq j$. The *inverse of* W is $W^- := v_\ell \overleftarrow{e}_\ell v_{\ell-1} \overleftarrow{e}_{\ell-1} v_{\ell-2} \dots v_1 \overleftarrow{e}_1 v_0$. A walk is *trivial* if it has length 0 and *non-trivial* otherwise.

The walk W *starts* in its *startvertex* v_0 and *ends* in its *endvertex* v_ℓ. If W is non-trivial, we say that it starts with sign $\sigma(e_1, v_0)$ in its *signed startvertex* $(v_0, \sigma(e_1, v_0))$. Likewise we say that it ends with sign $\sigma(e_\ell, v_\ell)$ in its *signed endvertex* $(v_\ell, \sigma(e_\ell, v_\ell))$. All other vertices of W are *internal vertices* of W. The set of all vertices in W is denoted as $V(W) := \{v_0, \dots, v_\ell\}$. The non-trivial walk W *arrives in* v_i with sign $\sigma(e_i, v_i)$ for $i \in [1, \ell]$ and *leaves* v_i with sign $\sigma(e_{i+1}, v_i)$ for $i \in [0, \ell - 1]$. Analogously, we say that a non-trivial walk W *starts* in $\vec{e}_1$ and *ends* in $\vec{e}_\ell$.

The set of all oriented edges in a walk W is denoted by $\vec{E}(W) := \{\vec{e}_1, \dots, \vec{e}_\ell\}$, the set of edges underlying $\vec{E}(W)$ is $E(W) := \{e_1, \dots, e_\ell\}$, and we write $\overleftrightarrow{E}(W) := \overleftrightarrow{E(W)}$. If $E(W) \subseteq A$ for some set A of edges of B, then W is a walk *in* A.

A non-trivial walk W is an *x–y walk* in B if W starts in x and ends in y, where x and y can be oriented edges or vertices or signed vertices of B. Given two sets $\mathcal{X}$ and $\mathcal{Y}$ of signed vertices of B we define an *$\mathcal{X}$–$\mathcal{Y}$ walk* as follows: a trivial walk $W = v$ is an $\mathcal{X}$–$\mathcal{Y}$ walk if there is $\alpha \in \{+, -\}$ such that $(v, \alpha) \in \mathcal{X}$ and $(v, -\alpha) \in \mathcal{Y}$. A non-trivial walk W is an $\mathcal{X}$–$\mathcal{Y}$ walk if it is an x–y walk for some $x \in \mathcal{X}$ and $y \in \mathcal{Y}$, no internal vertex of W is in $V(\mathcal{X}) \cup V(\mathcal{Y})$ and neither the start- nor the endvertex of W forms a trivial $\mathcal{X}$–$\mathcal{Y}$ walk. Note that no proper subwalk of an $\mathcal{X}$–$\mathcal{Y}$ walk is again an $\mathcal{X}$–$\mathcal{Y}$ walk.

Two walks W and W' in B are *vertex-disjoint* if $V(W) \cap V(W') = \emptyset$. Similarly, W and W' are *edge-disjoint* if $E(W) \cap E(W') = \emptyset$. A set of walks is *vertex-disjoint* (respectively *edge-disjoint*) if its elements are pairwise vertex-disjoint (respectively edge-disjoint).

A walk W in B is a *trail* if all edges underlying $\vec{E}(W)$ are distinct. If not only all edges, but also all vertices of W are distinct, then W is a *path*. For both trails and paths we transfer all notions and notation defined for walks (such as e.g. subtrail or v–w path).

For a given trail $Q = v_0\vec{e}_1v_1\vec{e}_2v_2\dots v_{\ell-1}\vec{e}_\ell v_\ell$ and $0 \le i \le j \le \ell$, we write $\vec{e}_iQ := v_{i-1}\vec{e}_i\dots v_\ell$, $Q\vec{e}_i := v_0\dots\vec{e}_iv_i$ and $\vec{e}_iQ\vec{e}_j := v_{i-1}\vec{e}_i\dots\vec{e}_jv_j$ for the appropriate subtrails of Q. Analogously, given a path $P = v_0\vec{e}_1v_1\vec{e}_2v_2\dots v_{\ell-1}\vec{e}_\ell v_\ell$ and $0 \le i \le j \le \ell$, we write $v_iP := v_i\dots v_\ell$, $Pv_i := v_0\dots v_i$ and $v_iPv_j := v_i\dots v_j$ for the appropriate subpaths of P, and let $\vec{e}_iPv_j := v_{i-1}Pv_j$ as well as $v_iP\vec{e}_j := v_iPv_j$. In particular, these two notions combine for paths to the respective notations $\vec{e}_iPv_j$ and $v_iP\vec{e}_j$.

2.3.2 Rays and Double Rays

In this thesis rays and double rays are digraphs together with linear orders on their sets of vertices: a *ray* or *double ray* is a digraph R together with a linear order $\le_R$ on $V(R)$ isomorphic to $\mathbb{N}$ or $\mathbb{Z}$, respectively, such that for all vertices v, w of R:

- if v and w are consecutive in the linear order, then either $vw \in E(R)$ or $wv \in E(R)$ but not both, and
- if v and w are not consecutive in the linear order, then $vw, wv \notin E(R)$.

Hence, rays and double rays are oriented versions of the more frequently considered undirected rays and undirected double rays, together with linear orders.

We say that a (double) ray is a *subdigraph* of another (double) ray or is *isomorphic* to another (double) ray if this is true for the digraphs in the usual sense, regardless of the linear orders.

Let R be a ray or a double ray and let $v \in V(R)$. We write vR for the subdigraph of R induced by all vertices w of R with $w \ge_R v$ and Rv for the subdigraph of R induced by all vertices w of R with $w \le_R v$. Similarly, for $e = vw \in E(R)$ with $v < w$ we define $eR := vR$ and $Re := Rw$. An infinite subdigraph of this form is called a *tail* of R. For $w \in D$ we say that $v \in R$ lies *beyond w on R* if $w \notin V(vR)$. Two rays or double rays R_1, R_2 *traverse an edge $uv \in E(R_1) \cap E(R_2)$ in the same direction* if either v lies beyond u on R_1 and R_2 or u lies beyond v on R_1 and R_2. Otherwise R_1 and R_2 *traverse uv in opposite directions*. We call R *periodic* if R has a non-trivial $\le_R$-preserving endomorphism. (Note that when R is a ray, any endomorphism must preserve $\le_R$. However, the same is not true for double rays.) Let $v \in V(R)$ and choose a non-trivial $\le_R$-preserving endomorphism f of R such that the distance d of v and $f(v)$ in R is minimal. Then we say that R has *periodicity* d (and d is independent of the choice of v).

A vertex of a ray or double ray which is incident with two outgoing or two incoming edges is called a *turn* and a maximal (possibly infinite) dipath contained in a ray or double ray is called a *phase*.

The next definitions are concerned only with rays rather than double rays. We call the unique $\leq_R$-smallest vertex of a ray R the *root* of R. A ray without turns is an *in-ray* if the first edge points towards the root and an *out-ray* otherwise. Furthermore we call an edge of an oriented ray R *in-oriented* if it is directed towards the root of R and *out-oriented* otherwise. Similarly we call a phase of a ray *in-oriented* if its edges are in-oriented and *out-oriented* otherwise.

A ray that has infinitely many in-oriented and infinitely many out-oriented edges can be represented by a sequence of natural numbers where the n-th term of the sequence represents the length of the n-th phase. We call this the *representing sequence* of the ray. The representing sequence is *bounded* if there is $b \in \mathbb{N}$ such that all elements of this sequence are contained in $[b]$. Otherwise the representing sequence is called *unbounded*.

2.3.3 Connection of Digraphs and Matching Covered Graphs

Every strongly connected (possibly infinite) digraph corresponds to a bipartite graph with a perfect matching. We can obtain this bipartite graph with partition classes V_0 and V_1 from a digraph D by splitting every vertex v into two vertices v_0, v_1, one of each partition class, that are connected by an edge, such that all in-edges of v become in-edges of v_0 and all out-edges of v become out-edges of v_1. Then we remove all directions from the edges. Clearly, this yields a bipartite graph and the new edges form a perfect matching in it. We refer to this graph as G_D and to the canonical perfect matching in it as M_D.

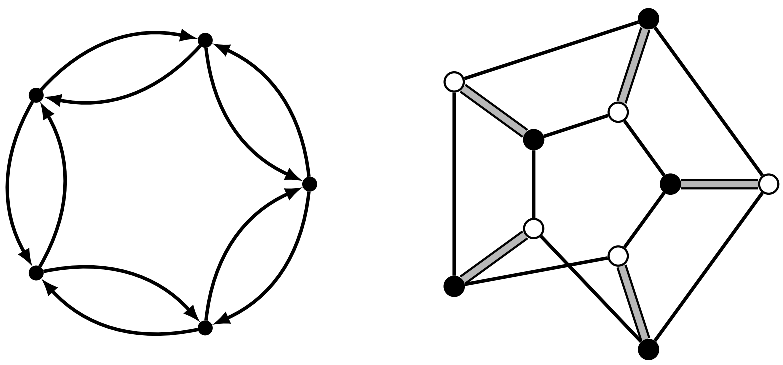

Figure 2.1: A strongly connected digraph D and the corresponding bipartite graph G_D with the canonical perfect matching M_D indicated with light gray edges.

Conversely, we can also obtain a digraph from any given (possibly infinite) bipartite graph with a perfect matching, by directing all non-matching edges from V_1 to V_0 and then contracting all matching-edges into single vertices. For a given graph G with perfect matching M, we call this the M-direction of G and denote it $D(G, M)$.

For an undirected graph G we refer to the set of all perfect matchings in G by $\mathcal{M}(G)$. An undirected graph G is called *matching covered with respect to a perfect matching* $M \in \mathcal{M}(G)$ if every edge lies in a perfect matching that has finite symmetric difference to the matching M. We write $\mathcal{M}(G, M)$ for the set of all perfect matchings that have finite symmetric difference to M. We simply say G is *matching covered*, when it is matching covered with respect to some matching. Note that on finite graphs this is equivalent to saying that every edge lies in a perfect matching as $\mathcal{M}(G, M) = \mathcal{M}(G)$ for all perfect matchings M in G.

Observation 2.4. Given an undirected graph G that is matching covered with respect to a perfect matching M, the symmetric difference between any two perfect matchings in $\mathcal{M}(G, M)$ is a collection of disjoint cycles.

For a graph G and a set $X \subseteq V(G)$ we define the *cut induced by* X, $\partial_G(X) := \{e \in E(G) \colon |e \cap X| = 1\}$, to be the set of all edges with exactly one endpoint in X.

Let G be a graph that is matching covered with respect to a perfect matching M. A subset $X \subseteq V(G)$ is a *tight set* in G if every $M' \in \mathcal{M}(G, M)$ has exactly one edge in $\partial(X)$ and we refer to $\partial(X)$ as a *tight cut*.

A tight set X and its corresponding tight cut $\partial(X)$ are called *trivial* if $|X| = 1$ or $|V(G) \setminus X| = 1$.

Tight cuts C_1, C_2 are *nested* if there are $X_1, X_2 \subset V(G)$ with $X_1 \cap X_2 = \emptyset$ such that $C_i = \partial(X_i)$ for $i \in [2]$. Tight sets X_1 and X_2 are *nested* if $\partial(X_1)$ and $\partial(X_2)$ are nested tight cuts.

A bipartite graph with no non-trivial tight sets is called a *brace*, while a non-bipartite graph with no non-trivial tight sets is called a *brick*. Since thus every graph with no non-trivial tight sets is either a brick or a brace, we refer to such a graph as a *BoB* (Brick or Brace) for short.

For a graph G and a maximal family $\mathcal{C}$ of nested tight cuts in G, we define $\vec{\mathcal{C}} := \{X \subseteq V(G) \colon \partial(X) \in \mathcal{C}\}$, i.e. the set of tight sets corresponding to tight cuts in $\mathcal{C}$. Note that a tight set X is contained in $\vec{\mathcal{C}}$ if $\partial(X)$ is nested with $\mathcal{C}$ since $\mathcal{C}$ is maximal. Furthermore, any two elements of $\vec{\mathcal{C}}$ are nested.

Unless stated otherwise, we consider tight cuts and tight sets with respect to an arbitrary but fixed perfect matching of G, without explicitly mentioning that perfect matching.

Proposition 2.5. *Let G be a connected, matching covered graph. For any tight set $X \subset V(G)$ the subgraphs $G[X]$ and $G[V(G) \setminus X]$ are connected.*

Proof. For any two distinct edges e and f in $\partial(X)$, the symmetric difference of any matching containing e with any matching containing f is a disjoint union of cycles. The cycle containing e must also contain f, and so contains paths joining e to f through both X and $V(G) \setminus X$.

Now for any two vertices v and w of X we know there is a path P joining them in G, and we can find a walk joining them in $G[X]$ by replacing each segment where P leaves $G[X]$ with a path in $G[X]$ as above. $\square$

2.3.4 Ubiquity

We call a digraph H *ubiquitous* if any digraph D that contains k disjoint copies of H as a subdigraph for every $k \in \mathbb{N}$ also contains infinitely many disjoint copies of H as a subdigraph.

Let $\overleftarrow{D}$ be the digraph obtained from a digraph D by inverting the orientation of every edge. To warm up with the definition of ubiquity, we prove the following proposition:

Proposition 2.6. *A digraph H is ubiquitous if and only if $\overleftarrow{H}$ is ubiquitous.*

Proof. It suffices to show that $\overleftarrow{H}$ is ubiquitous if H is ubiquitous. Let D be any digraph containing arbitrarily many disjoint copies of $\overleftarrow{H}$. Hence $\overleftarrow{D}$ contains arbitrarily many disjoint copies of $\overleftarrow{\overleftarrow{H}} = H$. Then $\overleftarrow{D}$ also contains infinitely many disjoint copies of H since H is ubiquitous. Therefore $\overleftarrow{\overleftarrow{D}} = D$ contains infinitely many disjoint copies of $\overleftarrow{H}$, which proves that $\overleftarrow{H}$ is ubiquitous. $\square$

Tribes and Ubiquity

In accordance with [BEE$^+$22, Definition 5.1], for digraphs D and R we call a collection $\mathcal{F}$ of finite sets of disjoint copies of R in D an *R-tribe in D*. If R is clear from the context we may also just say *tribe* instead, similarly for the surrounding graph D. Furthermore, for an R-tribe $\mathcal{F}$ in D we call $F \in \mathcal{F}$ a *layer of* $\mathcal{F}$, any element of F a *member of* $\mathcal{F}$ and say that $\mathcal{F}$ is *thick* if for each $n \in \mathbb{N}$ there is a layer F of $\mathcal{F}$ with $|F| \geq n$. A tribe $\mathcal{F}'$ in D is an *(R-)subtribe*[1] of an R-tribe $\mathcal{F}$ in D if every layer of $\mathcal{F}'$ is a subset of a layer of $\mathcal{F}$. Whenever we consider a copy R' of a digraph R we implicitly fix an isomorphism $\varphi\colon R \to R'$ and for a subdigraph $\hat{R} \subseteq R$ we write in short $\hat{R}'$ for $\varphi(\hat{R})$. With this, we say that an R-tribe $\mathcal{F}$ is *forked at* $\hat{R}$ if $\hat{R}' \cap R'' = \emptyset$ for any two distinct members R', R'' of $\mathcal{F}$.

Note that we may also use tribes to define ubiquity. A digraph H is ubiquitous if and only if every digraph D that contains a thick H-tribe also contains infinitely many disjoint copies of H as a subdigraph.

2.3.5 Flames

Let V be some fixed vertex set with a prescribed 'root vertex' $r \in V$. For an r-rooted digraph D, $v \in V - r$ and an arbitrary set I we write $D \restriction_v I$ for the subdigraph we obtain from D by deleting those incoming edges of v that are not in $I + rv$. For

[1]Note that even though this definition of subtribe differs from the one given in [BEE$^+$22], every subtribe in our sense is also a subtribe according to [BEE$^+$22, Definition 5.1].

$v \in V - r$ we denote by $\mathcal{G}_D(v)$ the set of those $I \subseteq \mathsf{in}_D(v)$ for which there exists an internally disjoint r–v path-system $\mathcal{P}$ in D with $E^+(\mathcal{P}) = I$. We say that D has the *vertex-flame property* at $v \in V - r$ if $\mathsf{in}_D(v) \in \mathcal{G}_D(v)$ and we call D a *vertex-flame* if it has the vertex-flame property at every $v \in V - r$. The *quasi-vertex-flame property* at v means that all finite subsets of $\mathsf{in}_D(v)$ are in $\mathcal{G}_D(v)$ and D is a *quasi-vertex-flame* if it has the quasi-vertex-flame property at every $v \in V - r$.

Erdős-Menger Separations and Path-Systems

For $S \subseteq V - r - v$ let $\mathfrak{P}_D(v, S)$ be the set of those internally disjoint r–v path-systems $\mathcal{P}$ in D that are *orthogonal* to S, i.e. for which S can be obtained by choosing exactly one internal vertex from each $P \in \mathcal{P}$ (observe that a path consisting of a single edge cannot be in $\mathfrak{P}_D(v, S)$). For $v \in V - r$, we define $\mathfrak{S}_D(v)$ to be the set of *Erdős-Menger separations* between r and v, i.e. the set of those $S \subseteq V - r - v$ that separate r from v in $D - rv$ (separation means that every r–v path in $D - rv$ meets S) and for which $\mathfrak{P}_D(v, S) \neq \emptyset$. We call $\mathfrak{P}_D(v) := \bigcup\{\mathfrak{P}_D(v, S) \colon S \in \mathfrak{S}_D(v)\}$ the set of the *Erdős-Menger paths-systems*. Note that the infinite version of Menger's Theorem (i.e. Theorem 2.3) applied to $X = N^+_{D-rv}(r)$ and $Y = N^-_{D-rv}(v)$ in $D - rv$ ensures that $\mathfrak{S}_D(v) \neq \emptyset$ and therefore $\mathfrak{P}_D(v) \neq \emptyset$. Observe that an $S \in \mathfrak{S}_D(v)$ is always a minimal set separating r from v in $D - rv$ since for every $s \in S$ each $\mathcal{P} \in \mathfrak{P}_D(v, S)$ contains some r–v path P_s that meets S only at s. One can show (see [Joó19a, Theorem 3.5]) that $\mathfrak{S}_D(v)$ is a complete lattice with respect to the partial order in which $S \leq T$ if S separates T from r (equivalently T separates S from v) in $D - rv$. We denote the smallest and the largest element of $\mathfrak{S}_D(v)$ by $S_{D,v}$ and $T_{D,v}$, respectively.

Large Spanning Subdigraphs

A system $\mathcal{P}$ of internally disjoint r–v paths in D is called *strongly maximal* w.r.t. D if for every internally disjoint r–v path-system $\mathcal{Q}$ we have $|\mathcal{Q} \setminus \mathcal{P}| \leq |\mathcal{P} \setminus \mathcal{Q}|$. In a finite digraph D strongly maximal simply means 'maximal-sized' but in general digraphs it is a stronger assumption than cardinality-wise maximality. The set of the strongly maximal internally disjoint r–v path-systems in D is exactly $\mathfrak{P}_D(v)$ if $rv \notin E(D)$ and the extensions of the elements of $\mathfrak{P}_D(v)$ with the single-edge path rv if $rv \in E(D)$, see for example [Joó19b, Proposition 3.4] for a proof. For a fixed D and $v \in V(D) - r$ we call a spanning subdigraph L of D *v-large* w.r.t. D if there is a strongly maximal internally disjoint r–v path-system of D that lies in L, moreover, L is *D-large* (or just *large* if D is fixed) if it is v-large w.r.t. D for every $v \in V - r$. For a finite D the largeness of $L \subseteq D$ is equivalent with the preservation of the local connectivities from the root, i.e. with $\kappa_L(v) = \kappa_D(v)$ for every $v \in V - r$ but it has a stronger structural meaning for general digraphs. Largeness of L can be rephrased as: $\mathfrak{P}_D(v) \cap \mathfrak{P}_L(v) \neq \emptyset$ (equivalently $\mathfrak{S}_D(v) \cap \mathfrak{S}_L(v) \neq \emptyset$) for every $v \in V - r$ and $\mathsf{out}_D(r) \subseteq L$.

2.3.6 Directed Separations

Let D be a digraph. A tuple (A, B) with $A, B \subseteq V(D)$ and $A \cup B = V(D)$ is a *(directed) separation* if there is no edge with tail in $B \setminus A$ and head in $A \setminus B$ or there is no edge with tail in $A \setminus B$ and head in $B \setminus A$. The integer $k := |A \cap B|$ is called the *order* of the separation and we also refer to (A, B) as a (directed) k-separation.

Tight Set Partitions

A partition $\mathcal{P}$ of the vertex set of a matching covered graph G is a *tight set partition* of G if every $P \in \mathcal{P}$ is a tight set. For every tight set partition $\mathcal{P}$ of a matching covered graph G we define the collapse $\mathsf{coll}(\mathcal{P})$ of $\mathcal{P}$ to be the graph with vertex set $\mathcal{P}$ and an edge between P and Q if and only if there are $p \in P$ and $q \in Q$ such that $pq \in E(G)$.

Part I

Games on Graphs

Chapter 3

The Strong K^4-Building Game

In this chapter we analyse the $(K^{\aleph_0}, K^4)$-game and prove that there is a winning strategy for the first player. We present it in Section 3.2. The strategy is also a winning strategy for the first player in the (K^n, K^4)-game if n is at least 17, thus in particular there is a winning strategy for the first player in the (K^{17}, K^4)-game. Therefore there is a winning strategy on a board that is smaller than suggested by Ramsey theory, as $R(4) = 18$ (see [GG55, Section 3] by Greenwood and Gleason).

The winning strategy presented in Section 3.2 draws on ideas first introduced by Beck in [Bec02, Section 5] in 2002. More precisely, our strategy broadly follows the one presented in [Bec02]. Unfortunately, the latter pools together some cases that must be considered separately and, more importantly, claims that some cases must not be investigated further for brevity but in fact there are ways to win for the second player in these cases. We examine these issues in Section 3.1. Beck noticed these shortcomings and therefore posed finding an explicit strategy for the $(K^{R(4)}, K^4)$-game as a question in [Bec08, Open Problem 4.6] in 2008, which we answer in this chapter. Our proof deals with all cases separately, which means there are a large number of cases, thus we use a computer algorithm to validate the first player's strategy in all of them. This allows us to also deduce that it takes the first player at most 21 turns to win the game.

The $(K^{\aleph_0}, K^4)$-game is also the subject of the bachelor thesis of the author [Gut17], thus the strategy presented in Section 3.2 is largely akin to the one in [Gut17]. However, there are some improvements upon the strategy and the proof technique in this chapter. Utilising a computer algorithm to do the case checking is completely novel in comparison to the bachelor thesis.

3.1 Problems with an Earlier Strategy

We begin by examining the proof of [Bec02, Theorem A.1]. We introduce the following terms that we use throughout this chapter: we denote by K^4_- the unique graph on 4 vertices with 5 edges. A *threat* is a monochromatic induced subgraph of $G_{\mathrm{FP}} \cup G_{\mathrm{SP}}$ that is isomorphic to K^4_-. Based on this we call a graph H a *threat seed graph* or simply a *threat seed* if it is a monochromatic induced subgraph of $G_{\mathrm{FP}} \cup G_{\mathrm{SP}}$ and it is a graph with precisely four vertices and four edges. That is, it is isomorphic to either a cycle of length four or a triangle with an attached edge.

In the proof of [Bec02, Theorem A.1] the discussion of case 'L(3)' has a flaw, the assertion 'The graph of the first five blue edges always forms a path or two separated paths with some red edges between them.' does not hold true. Note that in [Bec02] the colour blue refers to the second player and red refers to the first player, as in this thesis. We present a course of the game that falls under L(3) and in which after five turns G_{SP} is a tree with three leaves. This has a large enough impact that if the first player follows the strategy given in [Bec02] then the second player can win the game, as illustrated in Figure 3.1. Note that a similar problem arises if the first player claims the

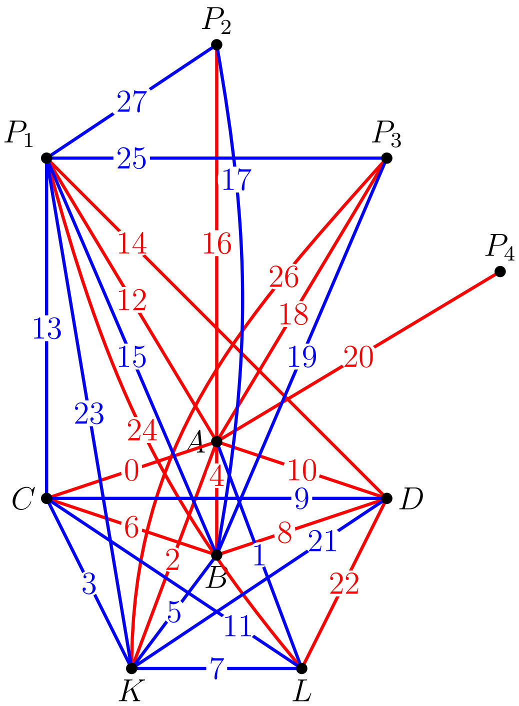

Figure 3.1: A course of the game in case 'L(3)' in the proof of [Bec02, Theorem A.1] in which the second player can win the game. Note that the edges labelled 22, 24 and 26 are forced moves for the first player in that they are reactions to threats by the second player. After claiming the edge P_1P_2, the second player can win by either claiming P_2P_3 or KP_2 with his next move.

edges BP_x instead of AP_x in her turns. Thus, in case $G_{\mathrm{FP}} \cup G_{\mathrm{SP}}$ is isomorphic to the graph mentioned in L(3) or a similar one (see Figure 3.6), the first player must instead divert from her usual course of play. We present a possible strategy for this special case in Stage 6 of the K^4 building strategy presented in Section 3.2.

The case 'L(4)' also has an issue. The second to last paragraph of the proof of [Bec02, Theorem A.1] asserts that in case L(4) if there is no blue triangle either containing the vertex B or the vertex C after six turns, then the first player can continue with the standard strategy of the proof and win. This is not true. In Figure 3.2 we present a possible course of the game falling under case L(4) where after six turns there is no blue triangle at all but the second player wins that game if the first player adheres to her regular strategy. Thus in that case the first player must use a different strategy for any of the board states depicted in Figure 3.5. We define a possible strategy in Stage 5 of the K^4 building strategy.

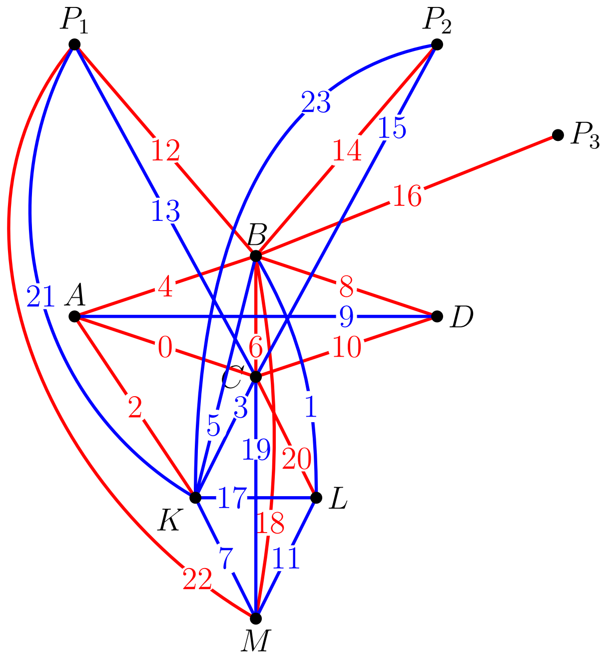

Figure 3.2: A course of the game in case 'L(4)' in the proof of [Bec02, Theroem A.1] in which the second player can win the game. Note that the edges labelled 18, 20 and 22 are moves in which the first player reacts to threats by the second player. In his next move the second player can win by claiming either P_1P_2 or MP_2.

Finally, there is another inaccuracy in the case 'L(4)' In the subcase where there are two blue triangles, one containing the vertex B and one containing the vertex C of case 'L(4)', the proof says that the first player should divert from her usual strategy and instead claim BD_1, BD_2 and BD_3 in her sixth, seventh and eighth turn respectively for suitable fresh vertices D_1, D_2 and D_3. The only course of the game that is considered is the one where the second player reacts by claiming AD_x in each of these three turns, claiming that if he does not do this then the first player can claim AD_x instead and the board is then isomorphic to the board in one of the other cases. Unfortunately this is not true. Suppose the second player does not claim AD_3. Then the first player claims AD_3 and the second player must claim CD_3 in his next turn since that is a threat by the first player. Under a desired isomorphism each D_x would then need to be mapped to a P_y, except for D_3 which would be mapped to D (which in turn must also be mapped to a P_y. Suppose the second player had claimed D_1D_2 instead of AD_3, then that maps to an edge P_iP_j and the proof never considers a situation where there is precisely one blue edge and no red edge on the subgraph of the board induced by $\{P_i : i \in \{1, 2, 3, 4, 5\}\}$. Moreover, if the first player adheres to the described strategy then there is a way for the second player to win the game. We deal with this special case in Stage 5.

3.2 The K^4-Building Game

We begin by investigating the $(K^{\aleph_0}, K^3)$-game. The result, stated in Lemma 3.1, is folklore, we give a possible proof in the following. We define the K^3*-building strategy* as follows. In her first two turns, the first player claims a path of length 2: she first claims some edge and calls it ab. If the second player claimed an edge incident with ab, by renaming the vertices if necessary, we can assume that it is incident with b. In any case the first player chooses a fresh vertex, calls it c and claims the edge bc. If the second player does not claim ac in his second turn, the first player does so in her third turn and thereby finishes a monochromatic triangle on a, b and c. Otherwise, if the second player's first edge uv was disjoint from her first edge ab and she cannot finish her triangle in three turns, then she claims bu in her third turn. If the second player's first edge was not disjoint from the first player's first edge, then it is incident with b by assumption. In this case the first player chooses a fresh vertex u and colours bu in her third turn. In her fourth turn at least one of the edges au, cu are unclaimed and so she can complete a K^3 on at least one of a, b and u or a, c and u.

Note that the first player needs at most five vertices to execute her strategy with the only condition that there are no claimed edges in between these vertices. Note further that the dependence of the third move of the first player on the first move of the second player is not necessary for Lemma 3.1 but it is for Theorem 3.3, thus we include it here already.

Lemma 3.1. *The K^3 building strategy is a winning strategy for the first player in the K^3-building game. Moreover, the first player needs at most 4 turns to win the K^3 game with the K^3 building strategy.*

Proof. The moreover part follows from the definition of the strategy. If it takes the first player only three turns to construct her K^3, she wins the game, as the second player has only had two turns up to that point. If it takes her four turns, then the second player colours ac in his second turn. Let b_1 denote the edge the second player claimed in his first turn. Either b_1 is disjoint from the first player's first edge ab or it is incident with b. Thus it is not incident with a. As the first player chose c as a fresh vertex in her second turn, b_1 is also not incident with c. So ac and b_1 have no vertices in common. Thus, the first two blue edges already use four vertices and therefore cannot be two edges of a triangle. □

Before we give the complete strategy for the first player for the K^4-building game, we need to make some adaptions to the K^3 building strategy, which we call the *modified*

K^3 *building strategy.* This strategy will be used on multiple occasions by the first player in her strategy for the K^4-building game to build new triangles in particular regions of the board.

In principle, the modified K^3 building strategy consists of the same steps as the K^3 building strategy but in the details the first player is more cautious, since there may be more coloured edges on the board at later points of the game. We define $P := \{p_0, p_1, p_2, p_3, p_4\}$.

- the first player's first edge is an edge that is incident with two vertices of P that have smallest possible SP-degree.
- For the first player's second edge, which we call f,
  - If the second player's previous edge is incident with a vertex of f then the first player's next edge is incident with that vertex and a vertex of P that is fresh in $G[P]$ with the smallest possible SP-degree.
  - If the second player's previous edge is incident with two vertices of P and not with a vertex of f then the first player's next edge is incident with a vertex of f with smallest possible SP-degree and a vertex of the second player's edge with smallest possible SP-degree.
  - Otherwise, if not all vertices of $P \setminus f$ have the same SP-degree, the first player claims an edge that is incident with a vertex of f and a vertex of $P \setminus f$ of minimum SP-degree that is not incident with the second player's previous edge.
  - If all vertices of $P \setminus f$ have the same SP-degree and there is a K^4 on the board of which 5 edges have been claimed by the first player and one has been claimed by the second player, let $\hat{K}$ be the subgraph of the board with this property that has been the first to emerge and call the two vertices of SP-degree 1 in $\hat{K}$ ℓ and r. Then the first player claims an edge that is incident with a vertex of f and a vertex of $P \setminus f$ that is also neighbouring as many vertices of ℓ and r as possible in the coloured subgraph.
- For the first player's third edge,
  - if the first player can finish a triangle with her third edge then she does so.
  - Otherwise, let v be the unique vertex of FP-degree 2 on the subgraph induced by P and let u and w be the vertices of FP-degree 1 in the subgraph induced by P and let a and b be the remaining two vertices of P. Then for one $x \in \{a, b\}$ the edges xv, xu and xw are unclaimed. The first player claims xv for such a vertex x.
- If there is a fourth turn, then in that turn the first player can finish a monochromatic triangle on P.

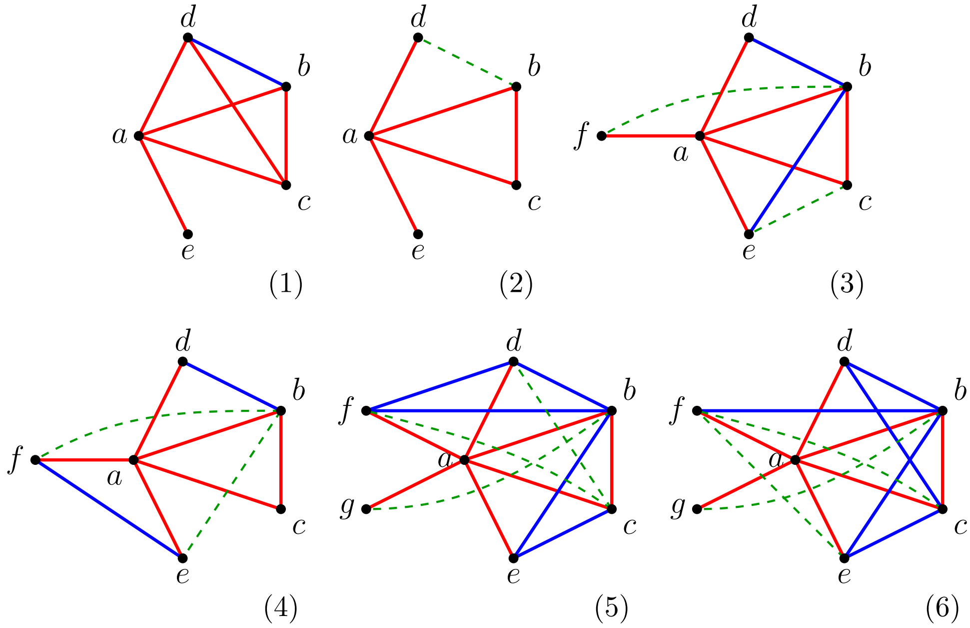

Figure 3.3: The six types of promising graphs. In each of the graphs the set of dashed edges represents the set of vulnerable edges. The straight edges represent the edges claimed by the first player and the second player. The non-edges are unclaimed edges.

We call a complete graph whose claimed edges induce a graph isomorphic to any of the graphs depicted in Figure 3.3 such that the dashed edges are either unclaimed or claimed by the second player, a *promising graph*. We call the respective dashed edge(s) the *vulnerable edge(s)* of the promising graph of that type. Additionally, we call a complete graph H that is isomorphic to a promising graph where the vulnerable edges are unclaimed a *realised promising graph* and say that H is *realised* in G. Sometimes we also refer to the intersection of H with $G_{\mathrm{FP}} \cup G_{\mathrm{SP}}$ as realised promising graph to simplify notation.

Lemma 3.2. *Let G be a board, H be a promising graph and suppose that H is realised in G. Suppose that it is the first player's turn and further that there is no monochromatic K^4 and no threat by the second player on the board even if the vulnerable edges are added as edges claimed by the second player. Then there is a sequence of moves for the first player, each one a threat, the last of which creates two threats, thus the first player wins the game.*

Proof. In (1), by claiming ce the first player immediately creates two threats. In (2), by claiming cd the first player creates a threat. If the second player reacts to that threat by claiming bd then the resulting subgraph is isomorphic to (1), thus she can create two threats by ce. In (3), the first player can create a threat by claiming cf. If the second player reacts to the threat by claiming bf then the induced subgraph on $\{a, c, d, e, f\}$ is isomorphic to (2), thus the first player can win with the corresponding moves presented in (2). Similarly in (4), the first player can create a threat by claiming ce and if the second player reacts by claiming be, in her next move the first player creates a threat by claiming cf. If the second player reacts to the threat by claiming bf then the induced subgraph on $\{a, c, d, e, f\}$ is isomorphic to the graph in (1), that is the first player can create two threats by claiming cd. In (5), by claiming cg the first player creates a threat and if the second player reacts by claiming bg then the subgraph induced by $\{a, c, d, e, f, g\}$ is isomorphic to the one in (4). Likewise in (6), by claiming cg the first player creates a threat. If the second player reacts with bg, then the subgraph induced by $\{a, c, d, e, f, g\}$ is isomorphic to the one in (3). In each case if the second player does not claim the respective edge mentioned then the first player could claim it in her next turn and finish a monochromatic K^4.

Since there is no threat or monochromatic K^4 by the second player on the board by assumption, he cannot create a threat or a monochromatic K^4 by claiming any of the vulnerable edges by assumption, all of his moves are forced moves and in the last move there are two threats by the first player, the first player wins the game as claimed. □

Note that for any of the promising graphs, the given sequence of moves for the first player also leads to a win for the first player if any of the edges of the second player indicated in blue in Figure 3.3 is unclaimed instead. In particular this is true for the edge bd in (6), which we use in the computer algorithm.

In the following, we define the K^4 *building strategy.*

Standing assumption: At all points we assume that the first player will need to react to any immediate threat by the second player. This overrides anything else suggested below. She also checks at any point whether she can just finish a monochromatic K^4 in her turn. If so, she does this and wins.

Stage 1: The first player constructs a triangle using the K^3 building strategy.

Stage 2: Throughout this stage the first player wants to build a monochromatic K^4_-, incorporating the K^3 from Stage 1. This is the point of the game where the first player must check whether the second player played in such a way that she cannot adhere to her usual strategy, as mentioned in Section 3.1. This is the first thing she checks at the start of each turn in Stage 2.

- The first player checks whether the graph induced by the claimed edges is isomorphic to one of the graphs in Figure 3.5 or Figure 3.6. If that is the case, she switches to Stage 5 or Stage 6 respectively.
- Otherwise, she checks whether she can claim an edge such that there is a monochromatic K^4_- in her colour on the board. If that is the case, she claims such an edge and switches to Stage 3.
- Otherwise, let c be the vertex of the monochromatic triangle of the first player with maximum SP-degree. the first player chooses a new vertex g and claims cg. Then she will be able to connect g to a or b in her following turn.

Stage 3: Generally in this stage the first player wants to claim five edges, all incident with the same particular vertex of the K^4_- constructed in Stage 2. More specifically, the first player does this as follows:

- when the first player is in Stage 3 for the first time and she cannot immediately win, then there is a subraph of the board that is a coloured K^4 of which five edges are claimed by the first player and one is claimed by the second player. The first player assigns roles to the four vertices of that subgraph and sticks to them throughout Stage 3 and Stage 4, see Figure 3.4 for an illustration.

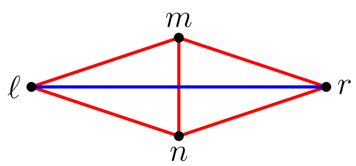

Figure 3.4: The K^4_- and the vertex names which the first player assigns in Stage 3.

  - We call the two vertices of the K^4 that are incident with the second player's edge ℓ and r, with arbitrary assignment.
  - Of the remaining two vertices, if one of them is contained within a monochromatic triangle of the second player, call this vertex m. If both are not contained in a triangle of the second player, call one of the two with the biggest SP-degree m.
  - We call the remaining vertex n.

  Then the first player chooses a fresh vertex p_0 and claims the edge from m to p_0.

- In a later turn during Stage 3, let p_i, $i \in \{0, 1, 2, 3, 4\}$ be the most recent fresh vertex that the first player chose.
  - If the second player claimed the edge between n and p_i,
    * if $i < 4$, then the first player chooses a fresh vertex p_{i+1} and claims the edge between m and p_{i+1},
    * if $i = 4$, then the first player switches to Stage 4.
  - If the second player claimed an edge between p_i and one of ℓ, r, then the first player claims the respective other edge.
  - If the second player claimed none of the edges between p_i and one of n, ℓ, r, then the first player claims the edge between p_i and n.

Stage 4: In this stage the first player constructs a triangle on P with the modified K^3 building strategy. At the end of Stage 4 that triangle together with m is a monochromatic K^4.

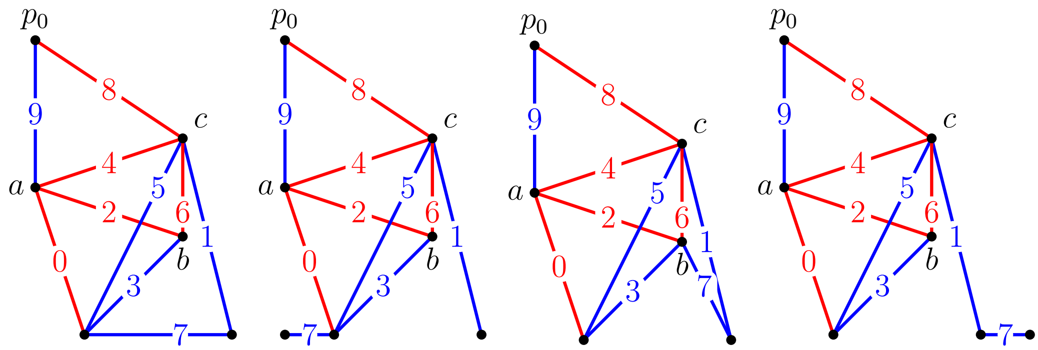

Figure 3.5: The four cases in which the first player applies her strategy of Stage 5.

Stage 5: If after 5 turns $G_{\mathrm{FP}} \cup G_{\mathrm{SP}}$ is isomorphic to any of the four graphs depicted in Figure 3.5, then the first player diverts from Stage 2 of her regular strategy and instead plays as follows. Her main goal is to either have a promising graph contained in the board as an induced subgraph or an independent set of five vertices in $G_{\mathrm{FP}} \cup G_{\mathrm{SP}}$ all with a common neighbour in G_{FP}. Let us call the vertices of the monochromatic red triangle in the board positions a, b and c, where c is the vertex of FP-degree three and SP-degree two, a is the other vertex of FP-degree 3 and b is the last vertex of the triangle. Furthermore let us give the name p_0 to the vertex of FP-degree and SP-degree one.

- In her sixth turn the first player chooses a fresh vertex p_1 and claims cp_1.
- In her seventh turn,
  - if the second player did not colour ap_1 or bp_1 in his sixth turn, then the first player claims ap_1 (which is a threat to the second player in this case) and switches to Stage 3 with the K_-^4 induced by $\{a, b, c, p_1\}$ and the p_i as assigned in this stage for $i \in \{0, 1\}$.
  - Otherwise the second player claimed either ap_1 or bp_1 in his sixth turn. Then the first player chooses a fresh vertex p_2 and claims cp_2.
- In the first player's eighth turn,
  - if the second player did not colour ap_2 or bp_2 or p_1p_2 in her seventh turn, then the first player claims ap_2 in her eighth turn and again switches to Stage 3 with the K_-^4 induced by $\{a, b, c, p_2\}$ and the p_i as assigned in this stage for $i \in \{0, 1, 2\}$.
  - Otherwise the second player claimed either ap_2, bp_2 or p_1p_2 in his seventh turn. Then the first player chooses a fresh vertex p_3, claims cp_3 in her eighth turn and switches to Stage 6.

Stage 6: The first player may enter this stage in two different ways. The first is that she enters this stage after exiting Stage 5. The second is that if after four turns $G_{\mathrm{FP}} \cup G_{\mathrm{SP}}$ is isomorphic to one of the two graphs depicted in Figure 3.6, then

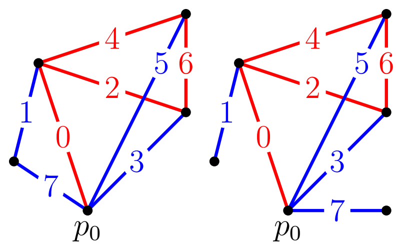

Figure 3.6: The two cases in which the first player applies her strategy of Stage 6.

in Stage 2 the first player diverts from her regular strategy and instead plays as follows. Only in the latter case in the fifth turn she gives the name p_0 to the vertex of FP-degree one and a to the vertex of FP-degree three. In either case, from turn eight or turn six onwards respectively, the first player goes through the list of conditions below and executes the strategy of the first item from the list whose conditions are met.

(a) The first player checks whether she secured a way to win in a previous turn, defined from one of the options.

(b) The first player checks, whether there is a realised promising graph present as a subgraph of the board such that the corresponding promising graph fulfils the conditions of Lemma 3.2. If there is one of those present, then this lets her continue with a threat each turn until she finally wins as proved in Lemma 3.2. Thus she saves the information on how to continue for the next turns.

(c) The first player checks, whether there are five p_i on the board and also checks whether there are no edges in between any two of them. Then she can build a triangle on the five vertices with the modified K^3 building strategy. If this applies, she saves the information that she wants to continue to build the triangle on these five vertices.

(d) The first player checks, whether there are fewer than five vertices p_i, $i \in \{0, 1, 2, 3, 4\}$. If this applies, she finds a fresh vertex, calls it p_j with $j := \max\{i + 1 \colon p_i \text{ is defined}\}$ and claims cp_j.

Note that one could get the impression that there can be courses of the game where the first player uses her strategy from Stage 6 but the prerequisites are not met for any of the considered situations. In practice this never occurs, as one is able to deduce from the proof of the following theorem.

Theorem 3.3. *The K^4 building strategy is a winning strategy for the first player in the K^4 game. Moreover, the first player needs at most 21 moves to win the K^4 game with the K^4 building strategy.*

Proof. We prove this theorem with the help of the algorithm given in `K_4_game.py` and `functions.py`, see Chapter 10 for the provision of the algorithm. There we implement the K^4 building strategy in the function `FP_edge()`, while `move()` recursively goes through every possible course of the game. For the exact implementation, we refer to the comments in the code. We do however make two modifications that are purely made for runtime considerations.[1] First, we make a list of possible non-isomorphic board states after Stage 2.

Claim. There are no additional courses of the game by considering other, isomorphic board states after Stage 2. ■

Second, in Stage 3 we only consider that the second player claims the edge incident with n and p_i for $i \in \{0, 1, 2\}$ when the first player claimed the edge incident with m and p_i in her turn and no other possible edge. Additionally, for p_3 we only allow that the second player claims an edge incident with p_3 and one of n, ℓ, r.

Claim. There are no additional courses of the game if the second player does not react to the p_i in the implemented way in Stage 3 up to isomorphism.

Proof. Suppose that it is the second player's turn and the first player chose a fresh vertex p_i and claimed $p_i m$ in her previous turn. In each of the mentioned turns, if the second player can not create a threat somewhere on the board himself, he must claim an edge incident with p_i and one of n, ℓ, r because otherwise the first player can win in two turns by claiming the edge $p_i n$. If he can create a threat somewhere else on the board, he can also do so at the beginning of Stage 4 and since the first player claims edges incident with vertices that are fresh at some point during Stage 3, edges claimed by the first player during Stage 3 will not mitigate that threat. Thus, if anything the second player has more opportunities to create threats at the beginning of Stage 4.

Now suppose that at some turn during Stage 3 the second player claims $p_i\ell$. Then the first player claims $p_i r$ in her following turn. This creates a threat to which the second player must react by claiming $p_i n$. For a $j \neq i$ the second player cannot claim $p_j\ell$ any more because by then claiming $p_j r$, the first player creates two threats and thus wins the game.

For a similar reason the second player can claim an edge incident with a p_i and r at most once.

Thus the second player can claim an edge incident with a p_i and ℓ at most once and only an edge incident with a different p_j and r at most once. After that he must always claim the edge $p_k n$ right after the first player claimed $p_k m$. Thus any possible board state after Stage 3 will be isomorphic to one we implemented in the algorithm, apart from threats by the second player. This, together with the fact that this does not infringe on the possibilities for the second player to create threats proves the claim. ■

This finishes the proof. □

[1]The runtime for the program with the adaptions but without multiprocessing on a MacBookPro with an Intel i5 1.4GHz Quad-Core processor is about 12h.

Chapter 4

Maker-Breaker $K^{\aleph_0}$-Building Games

In this chapter we investigate different variants of the $(K^{\aleph_0}, K^{\aleph_0})$-game. We begin by investigating the basic version in Section 4.1. As there is a winning strategy for Maker in this game, we introduce additional properties that Maker's graph should fulfil such as colouring the vertex set in different ways in Section 4.2 and Section 4.3 or enumerating the vertices with the rational numbers in Section 4.4. Parts of this chapter are based upon the master thesis of the author [Gut20]. In particular Section 4.1, Section 4.2 and Section 4.3 deal with variations of the $(K^{\aleph_0}, K^{\aleph_0})$-game that are also investigated in the thesis.

4.1 The Basic $K^{\aleph_0}$-Building Game

In this section we will prove that Maker can win the $K^{\aleph_0}$-building game. We will achieve this by first describing a strategy according to which Maker should play and then verifying in the proof of Theorem 4.1 that, in fact, $K^{\aleph_0} \subseteq G_M$ holds true if Maker adheres to the given strategy.

Our focus will be on two different kinds of activity by Maker. On the one hand, she will regularly want to add fresh vertices to her subgraph G_M. On the other hand, she must ensure that G_M is as interconnected as possible and thus contains large complete graphs. The same interplay between making G_M highly interconnected and regularly moving on to fresh vertices will also provide the basic rhythm for our strategies for Maker in later sections.

We will call the following strategy for Maker the *structured greedy strategy*. In her first turn, she picks some edge v_1v_2. In case Breaker was the first player, she picks one that only uses fresh vertices. In a later turn, suppose v_n is the last vertex that was added to Maker's subgraph. Now for $1 \leq i < n$, if there is some v_i such that

(□1) v_iv_n has not yet been claimed in either colour,

(□2) $N_M(v_n) \subseteq N_M(v_i)$, and

(□3) i is minimal subject to (□1) and (□2), then Maker claims $v_i v_n$. If there is no such v_i, she picks a fresh vertex v_{n+1} and claims $v_1 v_{n+1}$.

Theorem 4.1. *The structured greedy strategy is a winning strategy for Maker in the basic version of the $K^{\aleph_0}$-building game.*

Proof. We consider an arbitrary play of the game in which Maker follows the structured greedy strategy. We must show that at the end of the game G_M includes a $K^{\aleph_0}$. We will recursively construct a complete graph $K^n \subseteq G_M$ as well as a set of vertices $W_n \subseteq V(G_M)$ for every $n \in \mathbb{N} \setminus \{0\}$ such that:

(■1) $|K^n| = n$, $|W_n| = \aleph_0$,

(■2) $K^n \subseteq K^{n+1}$ for $n > 1$, and

(■3) for any $w \in W_n$ the first n vertices to which w was M-connected were $V(K^n)$.

If we successfully construct such a sequence $K^1 \subset K^2 \subset K^3 \subset \dots$, the claim follows immediately for

$$\bigcup_{i \in \mathbb{N}} K^i = K^{\aleph_0}.$$

The purpose of the sets W_n is to ensure that there will be a suitable candidate to enlarge the complete graph at each step.

Initial step: We can set $K^1 = (\{v_1\}, \emptyset)$ and $W_1 = V(G_M) \setminus \{v_1\}$. This immediately satisfies (■1) and (■2). (■3) holds true because every vertex in G_M other than v_1 got M-connected to v_1 right after it was chosen as a fresh vertex.

Recursion step: Now suppose K^n and W_n subject to (■1) to (■3) are given for some fixed $n \in \mathbb{N}$. Consider the first $n+1$ vertices that are completely M-adjacent to every $v \in V(K^n)$. Such vertices exist since every w in the infinite set W_n has this property by (■3). Let us call this set of vertices F.

Now consider any vertex $w \in W_n \setminus F$. At the point in the game n turns after w was chosen as a fresh vertex by Maker, it was already M-connected to $V(K^n)$, and at that point w was B-adjacent to at most n other vertices, so at least one vertex $v' \in F$ was still available. Let $\hat{v}$ be the vertex in F with this property and the smallest possible index. Then Maker claimed $w\hat{v}$ in her $(n+1)$-st move of M-connecting w. As w was arbitrary, this is true for every one of the infinitely many vertices of $W_n \setminus F$ and so, as F is finite, at least one vertex of F gets chosen in this way for infinitely many vertices from W_n. We denote the smallest such vertex in F by v^* and set

$$\begin{aligned} K^{n+1} &:= (V(K^n) \cup \{v^*\}, E(K^n) \cup \{vv^* : v \in V(K^n)\}), \text{ as well as} \\ W_{n+1} &:= \{w \in W_n : v^* \text{ was } M\text{-connected in the } (n+1)\text{-st turn} \\ &\qquad \text{after being picked as a fresh vertex}\}. \end{aligned}$$

This takes care of (■2) and additionally, we have $|K^{n+1}| = n+1$ and $|W_{n+1}| = \aleph_0$, thus (■1) is satisfied. (■3) follows from the recursion assumption together with the choice of W_{n+1}. □

As we know now that the $K^{\aleph_0}$-building game is a Maker win, we will go on to consider some variants in which we make life a little harder for her. The natural way to do so might be to allow Breaker to claim more than one edge for any edge that Maker claims. This kind of variant has also been studied in the finite case, see [HKSS14, Chapter 3] and is called a *biased* game. In our setting it doesn't make much difference, at least as long as Breaker is only allowed to claim the same finite number k of edges on each turn. Maker does not even need to adapt her strategy. In the verification we need to take F to be of size $kn + 1$ rather than $n + 1$, and the argument works just as before.

What if Breaker is allowed to claim a monotone increasing number of edges on his turns? It turns out that regardless of how slow the increment actually is, as long as the number of edges he claims tends to infinity he has a winning strategy: at the beginning of the game, he picks an enumeration $e_1 = x_1y_1, e_2 = x_2y_2, e_3 = x_3y_3, \ldots$ of the edges of the board. For any $n \in \mathbb{N}$ there is an $N \in \mathbb{N}$ such that from the N-th turn on, Breaker is allowed to claim n edges in each of his turns, for any edge $e = \{x, y\}$ that Maker claims. Beginning at $i = 1$, for every $i \leq n$, whenever $G[\{x, y, x_i, y_i\}] \not\subseteq G_M \cup G_B$, Breaker claims one of the available edges from $G[\{x, y, x_i, y_i\}]$ in his i-th turn. This strategy ensures that e_n can be part of a complete graph in G_M of at most some finite size dependent on N. As this holds for any $n \in \mathbb{N}$, Maker cannot construct a $K^{\aleph_0} \subseteq G_M$.

In the following sections we will take a closer look at other, more challenging variations. One could also consider biased versions of the games considered later in this chapter, but the theory of such biased games is always just the same as that outlined above and so we will not discuss it further.

4.2 The Finitely Coloured $K^{\aleph_0}$-Building Game

An interesting way to make Maker's objective more demanding can be obtained by assigning a colour to each vertex of the board and demand that Maker's graph reflects the colouring in a given way. For this purpose let $k \in \mathbb{N} \setminus \{0\}$. We say that a map

$$c\colon V(G) \longrightarrow [k]$$

is a *colouring* of $V(G)$ if $|c^{-1}(i)| = \aleph_0$ for every colour $i \in [k]$. Moreover, for a set $W \subseteq V(G)$ we define $c[W] := \{c(v)\colon v \in W\}$. Let c be a colouring and $j \in \mathsf{im}(c)$. We call

$$c^{-1}(j) \subseteq V(G)$$

the *colour class* of j. In this section it is Maker's objective to build a $K^{\aleph_0} \subseteq G_M$ as before but with the additional property that it includes infinitely many vertices from every colour class. We will call this version of the game the finitely coloured $K^{\aleph_0}$-building game. This is again a win for Maker with the following strategy.

Let $k \in \mathbb{N} \setminus \{0\}$, suppose that the board is coloured by a colouring $c\colon V(G) \longrightarrow [k]$ and let v_n be the vertex added to G_M most recently. Now suppose $\deg_M(v_n) \equiv \ell \mod k$

and $v_n \in c^{-1}(h)$ for $h, \ell \in [k]$ not necessarily distinct. Then, if Maker connects v_n to a vertex of colour ℓ in the following fashion, we say that she plays according to the *finite colour balanced greedy strategy.*

Let $F \subseteq V(G_M)$ be the set of the first $k \cdot \deg_M(v_n) + 1$ many vertices such that for all $v_m \in F$:

- $N_M(v_n) \subseteq N_M(v_m)$,
- $v_m \in c^{-1}(\ell)$,
- $m < n$, and
- $a < m$ for all $v_a \in N_M(v_n)$.

If there are fewer than $k \cdot \deg_M(v_n) + 1$ vertices satisfying these conditions, Maker chooses a fresh vertex v_{n+1} of colour m, where $n + 1 \equiv m \mod k$, and claims $v_1 v_{n+1}$. Otherwise, she considers the set $K \subseteq V(G_M)$ of all vertices v_i satisfying:

- $j < i$ for all $v_j \in F$,
- $N_M(v_n) \supseteq N_M(v_i)$, and
- $v_i \in c^{-1}(h)$.

Maker assigns a tuple in $\mathbb{N} \times \mathbb{N}$ to every $v_i \in F$ via the injective map

$$\begin{aligned} f \colon F &\longrightarrow \mathbb{N} \times \mathbb{N}, \\ v_i &\longmapsto (|N_M(v_i) \cap K|, i) \end{aligned}$$

and then she orders $f(F)$ lexicographically, which results in an ordered set

$$(f(F), \leq). \tag{4.1}$$

Maker determines the smallest tuple $(|N_M(v_\delta) \cap K|, \delta) \in f(F)$ such that $v_\delta v_n \notin E(G_B)$ and claims this edge. By the size of F it is clear that there will be a vertex v_δ available, as Breaker had only $\deg_M(v_n) < k \cdot \deg_M(v_n) + 1$ many moves where he could have claimed edges that are incident with v_n.

Let us shed some light on two aspects of this strategy, namely the size of F and the purpose of the order on F induced by f.

Our verification that this strategy works will be similar to that in Section 4.1, in that we will again recursively build a nested sequence of complete graphs K^n for every n and in every step make sure that there is an infinite set $W_n \subseteq V(G_M)$ such that for every vertex $v \in W_n$ the entire K^n is contained in its neighbourhood, i.e. the induced subgraph on $V(K^n) \cup \{v\}$ is a potential candidate to continue the sequence. Then the crucial part is to carefully pick a vertex such that there still is an infinite set $W_{n+1} \subseteq W_n$ left. Note that the set K for some vertex v in the strategy will be contained in the corresponding set W_n in the proof if v is considered as a potential next vertex in the recursion. Because of the role the sets W_n and therefore the sets K play, we will informally refer to them as *reservoir*. In contrast to the proof in Section 4.1, we need to also make sure that the sets W_n also contain infinitely many vertices of every colour. This is precisely the motivation for the map f introduced in the strategy above: if Maker just chose to play to the vertex from F with the smallest possible index as she does in the basic version, Breaker could ensure that all elements of the reservoir of colour a are joined to some vertex v_a, but that all elements of the reservoir of colour b are joined to some other vertex v_b. Then there would be no vertex that has infinitely many neighbours of both colours. Thus, instead of designating one vertex that has infinitely many neighbours of every colour, Maker instead ensures that for any colour Breaker can bar at most $\deg_M(v_n)$ vertices of F from having infinitely many neighbours of that colour. This excludes at most $k \cdot m$ vertices of F (recall that k is the number of colours and m the current M-degree of v_n). Maker wants to utilise this fact in order to ensure that there is at least one suitable vertex, i.e. a vertex with infinitely many neighbours of every colour, in the recursion step of the proof. She can achieve this by ensuring that the connection from vertices in the reservoir are spread as evenly as possible across F. The tool to do this is the function f and the lexicographic ordering.

Picking v_δ minimally in (4.1) makes the choice of the vertex unique for Maker, this is ensured by the second entry of the ordering. More importantly, as we have argued above, the vertices of F must be M-connected in a balanced fashion and this is achieved by choosing v such that $|N_M(v) \cap K|$ is smallest possible. To illustrate what we mean by that, one may think of the vertices in K as being the set of vertices in G_M that are identical to v_n in the following sense: they were added to G_M later than all vertices in F, the vertices got M-connected to G_M during their first $\deg_M(v_n)$ many turns in the same manner as v_n, and they have the same colour as v_n, namely h. Via f, Maker finds the elements in F that have the fewest neighbours in K and out of these she chooses the one that has the smallest index.

Therefore, by ensuring that F has size $k \cdot m + 1$ playing to vertices of F as evenly as possible via f Maker ensures that there is a suitable vertex in the recursion step of the proof.

Theorem 4.2. *The finite colour balanced greedy strategy is a winning strategy for Maker in the finitely coloured version of the $K^{\aleph_0}$-building game.*

Proof. We want to show that at the end of the game, if Maker plays according to the finite colour balanced greedy strategy, there is a $K^{\aleph_0} \subseteq G_M$ that uses infinitely many vertices of each colour class.

Recursive construction: For every $n \in \mathbb{N}$ we will construct a complete graph $K^n \subseteq G_M$ together with a set of vertices $W_n \subseteq V(G_M)$ with the properties

(▲1) $K^n \subset K^{n+1}$,

(▲2) $|W_n \cap c^{-1}(i)| = \aleph_0$ for all $i \in [k]$,

(▲3) for each $w \in W_n$ we have $N_M(w) \supseteq V(K^n)$ and the vertices of K^n were the first n to become M-connected to w, and

(▲4) $|K^n| = n$ and there is an enumeration $\{v'_1, v'_2, \dots, v'_n\}$ of $V(K^n)$ such that v'_j is coloured in m and $j \equiv m \mod k$ for every $1 \leq j \leq n$.

Note that by (▲2) we have in particular $|W_n| = \aleph_0$. As in the proof of Theorem 4.1, we can get the desired $K^{\aleph_0}$ from properties (▲1) and (▲4) by considering

$$\bigcup_{n \in \mathbb{N}} K^n = K^{\aleph_0}.$$

Here (▲1) ensures that there is the sequence $K^1 \subset K^2 \subset K^3 \subset \dots$ of complete graphs. (▲4) ensures that there are infinitely many vertices of each colour. (▲2) and (▲3) are needed to ensure that there always is a next vertex that can be added to K^n to form K^{n+1}. It remains to show that the conditions above can be preserved in every step.

Initial step: Again, we can set $K^1 := (\{v_1\}, \emptyset)$ and $W_1 = V(G_M) \setminus \{v_1\}$. As this is the initial step, (▲1) holds true. Since Maker repeatedly added vertices of all colours to G_M, (▲2) is true as well. (▲3) holds, since v_1 was the first vertex to be joined to each v_i with $i \geq 2$. Finally, as $V(K^1) = \{v_1\}$ and v_1 is coloured with colour $1 \in [k]$, (▲4) is also true. This concludes the base case.

Recursion step: Let $n \in \mathbb{N}$, $1 \leq i \leq k$, let K^n and W_n subject to (▲1) and (▲4) be given and let $i \in [k]$ such that $n + 1 \equiv i \mod k$. We want to construct K^{n+1} and W_{n+1} with the required properties. In order to do so, let F be the set of the first $kn+1$ vertices of colour i that have a common magenta edge with every vertex of K^n. Such a set exists, since all the infinitely many vertices of colour i in W_n have this property.

Let j be the largest index of a vertex in F, fix an arbitrary colour $\ell \in [k]$ and let $w \in (W_n \setminus F) \cap c^{-1}(\ell)$ be a vertex with an index larger than j. After Maker claimed wv' for every $v' \in V(K^n)$ in her first n moves of connecting w to G_M, the statement

$$N_M(w) \subseteq N_M(v)$$

held for all $v \in F$. Since these are the first $kn + 1$ such vertices, Maker chose the smallest available one of them with respect to the ordering derived from f as defined in (4.1). Breaker could block at most n edges vw for $v \in F$, thus there are at least $(k-1)n + 1$ possible edges for Maker to choose from. Therefore, at most k vertices of F individually have only finitely many vertices in $W_n \cap c^{-1}(\ell)$ that chose them in the $(n+1)$-st move of connecting to G_M.

As the colour ℓ was arbitrary, the argument above holds true for each of the k different colours. Therefore there is at least one vertex u in F that has infinitely many neighbours in every colour class in W_n. We choose the vertex $u^* \in F$ of these with the smallest index and let K^{n+1} be the graph obtained from K^n by adding u^* and all edges from it to K^n. As W_{n+1} we take the set of vertices in W_n such that u^* was the $(n+1)$-st vertex to which they were M-connected. This ensures (▲1). Moreover, it means that W_{n+1} contains infinitely many vertices of every colour class by the choice of u^*, therefore ensuring (▲2). The first $n+1$ vertices to be joined to any $w \in W_{n+1}$ were those of the K^{n+1} by the induction hypothesis and the construction of u^*, so we have (▲3). Lastly, the fact that we considered only vertices of colour $n+1 \equiv i \mod k$ for F, together with the assumption that (▲4) holds for n, ensures (▲4) for step $n+1$. □

4.3 The Infinitely Coloured $K^{\aleph_0}$-Building Game

As we now know that the finitely coloured $K^{\aleph_0}$-building game is a win for Maker, we now extend the definition of *colouring* to also include maps

$$c\colon V(K^{\aleph_0}) \longrightarrow \mathbb{N}$$

that fulfil $|c^{-1}(i)| = K^{\aleph_0}$ for every $i \in \mathbb{N}$. Let $j \in \mathbb{N}$ and $W \subseteq V(K^{\aleph_0})$. We define *colour class* and $c[W]$ analogously to the respective definitions for colourings with finitely many colours.

In this section we consider what happens if the board is coloured with infinitely many colours. If it is Maker's aim to construct a $K^{\aleph_0}$ that uses every colour class infinitely often, she is doomed to fail, as there is a strategy for Breaker with which he can keep Maker from doing so. However, if it is Maker's aim to construct a $K^{\aleph_0}$ that only uses infinitely many different colour classes, then there is a strategy with which she can secure this. We will first present Breaker's strategy for the first variant and after that we will give the strategy according to which Maker should play to win the second variant.

4.3.1 Using All Colours of the Board

Let a colouring $c\colon V(G) \longrightarrow \mathbb{N}$ of the board be given. Our aim is to define a pairing strategy such that for every edge e of the board there is some colour class i of which Maker may use no vertices together with e. First note that there are countably infinitely many edges in a $K^{\aleph_0}$ as $V(K^{\aleph_0})$ is countably infinite and the edges correspond to the two element subsets of $V(K^{\aleph_0})$.

Before the beginning of the game, Breaker picks an enumeration $e_1, e_2, e_3, \ldots$ of all edges of the board. He then recursively finds an enumeration $c_1, c_2, c_3, \ldots$ of infinitely many colours that are present on the board such that for all $i \in \mathbb{N}$:

(★1) $c_i \neq c_j$ for all $j < i$ and

(★2) $c_i \notin c\left[\bigcup_{j\leq i} e_j\right]$.

Fix some $m \in \mathbb{N}$. We set $V_m := c^{-1}(c_m)$ and according to (★2) we have $V_m \cap e_m = \emptyset$. Suppose $e_m = xy$. Then for any $v \in c^{-1}(c_m)$ there are exactly two edges from v to e_m, namely vx and vy. Whenever Maker claims one of those edges, Breaker claims the other one in his following turn.

Lemma 4.3. *The all infinite colour class pairing strategy is a pairing strategy and furthermore a winning strategy for Breaker in the infinitely coloured version of the $K^{\aleph_0}$-building game where Maker must have all colours contained in her $K^{\aleph_0}$.*

Proof. We first check whether the defined strategy actually is a pairing strategy, i.e. that any edge lies in at most one of the pairs of edges. Indeed, suppose for a contradiction, that there is an edge $e = uw$, that lies in two different pairs of edges.

Case 1: e lies in the pair of edges for two distinct edges e_i, e_j that are incident with the same vertex of e, u say: then $w \in c^{-1}(c_i) \cap c^{-1}(c_j)$. Thus we have $c_i = c_j$, a contradiction to (★1), as either $i < j$ or $j < i$.

Case 2: e lies in the pair of edges for two non-adjacent edges e_i, e_j: without loss of generality we may assume $u \in e_i$ and $w \in e_j$. This can only happen, if $u \in V_j$ and $w \in V_i$, a contradiction to (★2), as either $i < j$ or $j < i$.

Thus the pairs used for the strategy are disjoint. Therefore the given strategy actually is a pairing strategy for Breaker.

Furthermore Maker cannot build a $K^{\aleph_0}$ that uses all colours present on the board, if Breaker plays according to the defined strategy: for any edge that she wants to incorporate into her $K^{\aleph_0}$, there is a colour class that corresponds to it according to the construction which she therefore cannot use in her $K^{\aleph_0}$. □

Note that this result can be strengthened in the following sense: as every edge in a $K^{\aleph_0}$ of Maker has a different colour assigned to it and the $K^{\aleph_0}$ contains infinitely many edges, there are infinitely many colours of which Maker cannot incorporate infinitely many into her $K^{\aleph_0}$, this means that Breaker can even stop Maker from using cofinitely many colours, each infinitely often, in a $K^{\aleph_0}$.

4.3.2 Using Infinitely Many Colours of the Board

Let us now investigate how Maker should play in order to ensure that G_M contains infinitely many vertices from infinitely many different colour classes. As before, she needs to add fresh vertices to G_M, making sure that she keeps track of the colours as well as taking care of the vertices that are already part of G_M, while also ensuring that they are as interconnected as possible in general and paying attention to the colours of these fresh vertices in particular.

We first introduce one additional definition. For a finite subset $U \subseteq V(G_M)$ we let

$$\varphi_U \colon [|U|] \longrightarrow \{i \in \mathbb{N} \colon v_i \in U\} \subseteq \mathbb{N}$$

be the unique order preserving bijection. For an infinite subset $W \subseteq V(G_M)$ we consider the unique order preserving bijection

$$\varphi_W \colon \mathbb{N} \longrightarrow \{i \in \mathbb{N} \colon \; v_i \in W\} \subseteq \mathbb{N}\,.$$

In this variant of the game, Maker cannot rotate through the colour classes in the same fashion as with finitely many colours and thus she has to work on her objective in a diagonal fashion. We further specify this in the strategy.

At the beginning of the game, Maker chooses a sequence $s_1, s_2, s_3, \ldots$ of all the colours appearing on the board such that each individual colour appears infinitely often. Let us call this sequence S.

We call the following strategy for Maker the *infinite colour balanced greedy strategy.* In her first turn, she picks two fresh vertices of colours s_1 and s_2, calls them v_1 and v_2 respectively and claims the edge v_1v_2 for herself. When Maker adds vertices to her subgraph in later stages of the game, letting $|G_M| = n-1$, she adds a fresh vertex v_n of colour s_n to G_M.

After M-connecting a fresh vertex v_n of colour s_n by claiming v_1v_n, on the next few turns Maker determines which edge v_iv_n to claim by considering the set $U \subseteq V(G_M)$ of all vertices v_u that satisfy

- $N_M(v_n) \subseteq N_M(v_u)$, moreover in the first $\deg_M(v_n)$ turns of M-connecting v_u Maker M-connected v_u to the same vertices as v_n, in the same order, and
- $i < u$ for all $v_i \in N_M(v_n)$.

Then Maker considers the subset

$$U' := \{v \in U \colon c(v) = c(v_n)\} \tag{4.2}$$

and with φ_U as defined above she determines that the vertex she plays to next should be of colour

$$j = c\left(v_{\varphi_U(|U'|)}\right). \tag{4.3}$$

Note that $v_n \in U$, thus $\varphi_U\left(|U'|\right)$ is well-defined. Next, Maker lets $F \subseteq V(G_M)$ be the set of the first $(|c[N_M(v_n)]| + 2) \cdot \deg_M(v_n) + 1$ vertices, such that

- $k < i$ for all $v_i \in F$ and all $v_k \in N_M(v_n)$,
- $i < n$ for all $v_i \in F$,
- $v_i \in c^{-1}(j)$ for all $v_i \in F$, and
- all $v_i \in F$ satisfy $N_M(v_n) \subseteq N_M(v_i)$.

If there are fewer than $(|c[N_M(v_n)]|+2)\cdot\deg_M(v_n)+1$ vertices satisfying these conditions, Maker instead chooses a fresh vertex v_{n+1} as described above. If there is such a set, she chooses its subset of the first $(|c[N_M(v_n)]|+2)\cdot\deg_M(v_n)+1$ many and calls this F. Maker wants to M-connect a vertex from F to v_n analogously to the finite colour balanced greedy strategy: she considers the set $K \subseteq V(G_M)$ of all vertices v_k that satisfy

- for all $v_i \in F$ we have $i < k$,
- for all $v_\ell \in N_M(v_n)$ we have $\ell < k$,
- $N_M(v_k) \supseteq N_M(v_n)$, and
- $c(v_k) = c(v_n)$.

Maker assigns a tuple to every $v_i \in F$ as follows:

$$\begin{aligned} g\colon F &\longrightarrow \mathbb{N}\times\mathbb{N}, \\ v_i &\longmapsto \left(|N_M(v_i)\cap K|, i\right), \end{aligned}$$

and then orders $g(F)$ lexicographically, which results in an ordered set

$$(g(F),\leq)\,. \tag{4.4}$$

Maker determines the smallest $v_\delta \in F$ such that $v_\delta v_n \notin E(G_B)$ and claims this edge.

Note that as $|F| = (|c[N_M(v_n)]|+2)\cdot\deg_M(v_n)+1$, there will be a vertex v_δ available, as Breaker has only had $\deg_M(v_n)$ many moves where he could have B-connected vertices from F with v_n. As in Section 4.2, considering the ordering $(g(F),\leq)$ ensures that Maker plays from vertices similar to v_n to the vertices in F in a 'balanced fashion', which is crucial in the verification step.

Let us investigate why the size of F should be $(|c[N_M(v_n)]|+2)\cdot\deg_M(v_n)+1$. Recall that the size of the corresponding set in the finite colour balanced greedy strategy was '(number of colours on the board $\cdot$ degree of the active vertex)+1'. The sizes thus only differ by 'number of colours on the board' vs 'number of colours in the neighbourhood of the active vertex+2'. It is clear that 'number of colours on the board' cannot be used in the infinitely coloured game, as the size of F must be finite. It becomes clear, why the chosen number gives a good compromise for the following reason, when we suppose that the vertex v_n will be considered as an element of the set W_m of potential future vertices for some $m \in \mathbb{N}$. In the proof we must make sure that for any colour d already present on the K^m there are infinitely many vertices of colour d in W_m. This is ensured by the 'number of colours in the neighbourhood of the active vertex'-part. On top of that, in order to eventually have infinitely many colours present on the $K^{\aleph_0}$, we need to (a) allow for one additional colour in case we want to add a new colour to the K^{m+1} and (b) ensure that there will still be infinitely many other potential colours present in W_{m+1} to add in the future. This gives rise to the '+2'-part.

Lastly, let us elucidate Makers' choice of the colour j given in (4.3) before we move on. Breaker might render some colours unusable for Maker, but which these will be will not be clear until after the game. Thus, Maker needs a method to ensure that for each two such colours j and k she infinitely often tries to connect from a vertex of colour k down to one of colour j. The given function fulfils this purpose, which we will prove later (see (4.5)).

Theorem 4.4. *The infinite colour balanced greedy strategy is a winning strategy for Maker in the infinitely coloured version of the $K^{\aleph_0}$-building game where Maker must have infinitely many colours contained infinitely often in her $K^{\aleph_0}$.*

Proof. We want to prove that, after infinitely many turns, there is a $K^{\aleph_0} \subseteq G_M$ that uses infinitely many vertices of infinitely many different colours, if Maker plays according to the strategy above. Before we begin with the recursion, we pick a sequence $\hat{C} = c_1, c_2, c_3, \ldots$ of colours of $c[V]$ which contains every element of $c[V]$ infinitely often. We just require that $c_1 = s_1 = c(v_1)$.

Recursive construction: For every $n \in \mathbb{N} \setminus \{0\}$ we will construct a complete graph $K^n \subseteq G_M$ together with a set of vertices $W_n \subseteq V$, and a set of colours $C_n \subseteq c[V]$ with the properties

(♦1) $K^n \subset K^{n+1}$,

(♦2) $|W_n \cap c^{-1}(i)| = \aleph_0$ for every $i \in C_n$,

(♦3) for each $w \in W_n$ the first n moves of connecting w to G_M by Maker were claiming the edges that join w to the K^n,

(♦4) $|K^n| = n$ and there is an enumeration $\{v'_1, v'_2, \ldots, v'_n\}$ of $V(K^n)$ such that $c(v'_i) = c(v'_j)$ if and only if $c_i = c_j$ for $1 \leq i \leq j \leq n$, and

(♦5) $|C_n| = \aleph_0$ and $c[V(K^n)] \subseteq C_n$.

Note that in (♦4) we do not require $c(v_i) = c_i$. This is indeed impossible to achieve. But it secures that any colour that appears really appears infinitely often and together with (♦5) it furthermore secures that infinitely many different colours appear infinitely often in the inclusive chain

$$K^1 \subset K^2 \subset K^3 \subset K^4 \subset K^5 \subset \ldots \ ,$$

thus

$$\bigcup_{n \in \mathbb{N}} K^n = K^{\aleph_0}$$

is the desired complete subgraph of G_M.

Initial step: Set $K^1 := (\{v_1\}, \emptyset)$, $W_1 := V(G_M) \setminus \{v_1\}$ and $C_1 = c[W_1]$. (♦1) holds true, since this is the initial step. As $|K^1| = 1$, (♦4) holds true as well. $|C_1| = \aleph_0$ is true and $c[V(K^1)] \subseteq C_1$ is satisfied because S contains every colour infinitely often, thus C_1 satisfies (♦5). Moreover, as every vertex of $G_M \setminus \{v_1\}$ was first M-connected to v_1, (♦3) is true. Finally, as S contains every colour infinitely often, this ensures (♦2) for C_1. This concludes the base case.

Recursion step: Let $n \geq 1$ and K^n, W_n and C_n subject to (♦1) to (♦5) be given and suppose c_m was the entry of $\hat{C}$ we worked with in the previous step (which means in particular that $c(v_n) = c_m$). We want to construct K^{n+1}, W_{n+1} and C_{n+1} with the required properties. As before, we will need to make sure that we add vertices of colours that are already present in the K^n but we will sometimes also need to add vertices of colours that are not. If there is $i \leq n$ such that $c_i = c_{n+1}$, we set $c_p := c(v_i)$ and otherwise we choose $c_p \in C_n \setminus c[V(K^n)]$ arbitrarily. We want to add a vertex of colour c_p next and let F be the set of the first $(|c[V(K^n)]| + 2) \cdot n + 1$ vertices of colour c_p that have a common magenta edge with every vertex of the K^n. This set exists since there are infinitely many vertices of colour c_p in W_n by (♦2). Moreover, we need to restrict W_n to only contain vertices whose $(n+1)$-st turn of connecting it to G_M was a vertex of colour c_p and we want to ensure that (♦2) holds for this restriction of W_n as well. Note that since (♦1), (♦4) and (♦5) are independent of W_n, they still hold and (♦3) will hold for the restriction as it is a subset of W_n.

Fix an order preserving map $\psi\colon \mathbb{N} \setminus \{0\} \to I$ such that

$$\left\{v_{\psi(i)} \colon i \in \mathbb{N} \setminus \{0\}\right\} = W_n \cap c^{-1}(c_p)\,. \tag{4.5}$$

Then, for every $m \in \mathbb{N} \setminus \{0\}$ and every $d \in C_n$ the $(\psi(m))$-th vertex of colour d in W_n got M-connected to a vertex of colour c_p in the $(n+1)$-st move of connecting it to G_M. Thus, there are infinitely many vertices of colour d in W_n whose $(n+1)$-st neighbour in G_M (according to the order in which they were connected to it) was a vertex of colour c_p. As d was arbitrary, this is true for every colour in C_n. We can thus restrict W_n to these vertices and work with this set W_n' from here on.

Let $\ell \in c[V(K^n)] \cup \{c_p\}$. Since Maker played to the vertices of F in a balanced fashion, there are at most n vertices $v \in F$ such that only finitely many vertices $w \in W_n' \cap c^{-1}(\ell)$ got M-connected to v in their $(n+1)$-st move of connecting them to G_M. As ℓ was arbitrary, this is true for every colour in $c[V(K^n)] \cup \{c_p\}$ and thus there are at least $n+1$ vertices that have infinitely many such vertices in $W_n' \cap c^{-1}(p)$ for every $p \in c[V(K^n)] \cup \{c_p\}$. Conversely, regarding the infinitely many colours in $C_n \setminus (c[V(K^n)] \cup c_p)$, since Breaker can block at most n vertices for any of them, there are at most n vertices in F that are chosen by only finitely many vertices of cofinitely many colours not yet occurring in the K^n in the $(n+1)$-st move of connecting them to G_M. Combining this means that there is at least one vertex $u' \in F$ that got chosen by infinitely many vertices of every colour in $c[V(K^n)] \cup \{c_p\}$ as well as infinitely many vertices of infinitely many distinct colours in C_n in their $(n+1)$-st move of connecting them to G_M. We choose the smallest such vertex and call it v_{n+1}'. We set

- $K^{n+1} := G\left[V(K^n) \cup \{v_{n+1}'\}\right]$,

- C_{n+1} the set of colours $i \in C_n$ for which infinitely many vertices that lie in $c^{-1}(i) \cap W_n$ got M-connected to v'_{n+1} in their $(n+1)$-st move of connecting them to G_M, and

- $W_{n+1} \subseteq W'_n$ as the vertices in W_n that got M-connected to v'_{n+1} in their $(n+1)$-st move of connecting them to G_M and that are coloured with a colour in C_{n+1}.

This ensures (♦1), (♦2) and (♦5). (♦4) holds true by the choice of c_p and the definition of v'_{n+1}. All vertices of W_{n+1} are completely M-adjacent to the K^n by the induction hypothesis and to v'_{n+1} according to the construction, so all vertices of W_{n+1} are completely (M-)adjacent to the K^{n+1}. It follows from the induction hypothesis and the choice of v'_{n+1} that those were the first $n+1$ moves that Maker made for each element of W_{n+1}. This verifies (♦3).

This shows that all required properties are preserved throughout the induction and thus the claim is proved. □

This concludes our investigation into vertex colourings.

4.4 The Ordered $K^{\aleph_0}$-Building Game

Given the close relationship of Ramsey theory and Ramsey games, it is natural to ask how close the relationship is exactly. For this purpose consider a generalisation of the H-building game which we call the *structural H-building game*: One may introduce further structural information into the board G and then require that H satisfies an additional property with regards to the structural information of G. For example one may consider an order on the vertices of G and require that the subset of vertices that form Maker's copy of H is order-isomorphic to $V(G)$. With this we return to the relationship of Ramsey theory and Ramsey games. There are two sides to the question of how closely the two fields are related.

Question 4.5. Let G be a complete graph with a structural property and H be a graph with a structural property of its vertex set that is compatible with that of $V(G)$. Suppose that for any 2-colouring of the edges of G there is a monochromatic copy of H contained in G as a subgraph. Is there a winning strategy for Maker in the structural H-building game on G?

So far all research seems to be supportive of this assertion, the author is not aware of any example disproving Question 4.5.

Question 4.6. Let G be a complete graph with a structural property and H be a graph with a structural property of its vertex set that is compatible with that of $V(G)$. Suppose that there is a winning strategy for Maker in the structural H-building game on G. Does this imply that for any 2-colouring of the edges of G there is a monochromatic copy of H contained in G as a subgraph?

A good example for this is the K^4-building game investigated in Chapter 3. We saw there that for $n \geq 17$ the K^4-building game on K^n is a first player win and noted that this implies that there is a winning strategy for the first player for every board size of $n \geq R(4)$. Notably there is a small discrepancy between the sufficient board size of 17 and the Ramsey number of 18. Note that in Chapter 3 there is no structural property considered but one could e.g. assume that the vertices of the board $K^{\aleph_0}$ and the vertices of K^4 are totally ordered, which gives the same result.

Similarly in the Maker-Breaker version of the $K^{\aleph_0}$-building game studied in Section 4.1 there is a winning strategy for Maker and for any 2-colouring of the countably infinite complete graph there is a monochromatic $K^{\aleph_0}$, see e.g. Dushnik and Miller [DM41, Theorem 5.22] for a proof of Ramsey's theorem for infinite graphs.

In this section we answer Question 4.6 in the negative. We present a game and a winning strategy for Maker in that game where the corresponding statement in Ramsey theory is not true.

Before stating the main result of this section let us shed some light on another disparity of finite and infinite graphs which has to do with total orders. While any two total orders on a given finite set are isomorphic, this is not true for infinite total orders. For example there are two non-isomorphic total orders on the rational numbers: The usual order and one induced by an enumeration of the rationals. We carry this over to the setting of structured H-building games, and make use of this discrepancy: let $K^{\mathbb{Q}}$ be the complete graph with the rational numbers $\mathbb{Q}$ as the vertices. In the *$K^{\mathbb{Q}}$-building* game we call $K^{\mathbb{Q}}$ the *board*, of which the two players Maker and Breaker alternately claim edges. The aim of Maker is to have contained in the subgraph induced by her claimed edges a copy of the board. That is, a complete graph such that its vertex set (which is a subset of $\mathbb{Q}$) with the order induced by $\mathbb{Q}$ is order isomorphic to $\mathbb{Q}$. It is Breaker's goal to stop Maker from achieving this. We will prove that there is a winning strategy for Maker.

Theorem 4.7. *There is a winning strategy for Maker in the $K^{\mathbb{Q}}$-building game.*

The corresponding statement in Ramsey theory is false: There is a colouring of the graph $K^{\mathbb{Q}}$ with two colours such that there is no monochromatic copy of $K^{\mathbb{Q}}$ contained in either of the colour classes, which is shown in Proposition 4.10.

In Subsection 4.4.3 we show that the result of Theorem 4.7 cannot be made stronger in the following sense: consider the *dense $K^{\mathbb{Q}}$-building game*, the Maker-Breaker game played on $K^{\mathbb{Q}}$ where it is Maker's goal to finish a copy of the board as in the $K^{\mathbb{Q}}$-building game with the additional property that the vertex set of Maker's copy is dense in the vertex set of the board. We prove that there is a winning strategy for Breaker in the dense $K^{\mathbb{Q}}$-building game.

4.4.1 Ramsey Properties of $\mathbb{Q}$ and $\mathbb{Q}^2$

In preparation for the investigation of the $K^{\mathbb{Q}}$-building game we need three folklore results. While these were known before, we give proofs for each of them here.

Proposition 4.8. *There is a partition $\mathcal{P}$ of $\mathbb{Q}$ into pairwise disjoint open intervals and an order $\leq_{\mathcal{P}}$ of $\mathcal{P}$ defined by $P \leq_{\mathcal{P}} Q$ if and only if either $P = Q$ or $p \leq q$ for every $p \in P$ and $q \in Q$ such that $(\mathcal{P}, \leq_{\mathcal{P}})$ is order isomorphic to $(\mathbb{Q}, \leq)$.*

Note that since any open interval of $\mathbb{Q}$ is again isomorphic to $\mathbb{Q}$ (since any two countable dense total orders without smallest and largest element are isomorphic, see [BMMN06, Theorem 9.3]), there also exists such a collection for any open interval of $\mathbb{Q}$.

Proof. Let $\preceq$ be the lexicographic order on $\mathbb{Q}^2$ and $f\colon (\mathbb{Q}^2, \preceq) \to (\mathbb{Q}, \leq)$ be an order isomorphism. Then $\mathcal{P} := \{f[\{q\} \times \mathbb{Q}] \colon q \in \mathbb{Q}\}$ is as desired as $\leq_{\mathcal{P}}$ inherits its properties from $\leq$. □

Proposition 4.9. *For every 2-colouring of the rational numbers there is a monochromatic subset that is isomorphic to the rational numbers.*

Proof. Choose a collection $(\mathcal{P}, \leq_{\mathcal{P}})$ as in Proposition 4.8 and let a 2-colouring of $\mathbb{Q}$ be given. Call the colours red and blue. First suppose that there is a set $P \in \mathcal{P}$ that contains no red element. Then P is an open blue interval, which is isomorphic to $(\mathbb{Q}, \leq)$ by Cantor's Isomorphism Theorem [BMMN06, Theorem 9.3]. Now suppose that every $P \in \mathcal{P}$ contains a red element. Choose a red element from each P. The union of these elements together with the order $p \leq q \Leftrightarrow P \leq Q$ is then isomorphic to $(\mathbb{Q}, \leq)$. □

Proposition 4.10. *There is a 2-colouring of $K^{\mathbb{Q}}$ such that there is no monochromatic subgraph $K^{\mathbb{Q}}$.*

Proof. Let $(q_i)_{i \in \mathbb{N}}$ be an enumeration of $\mathbb{Q}$. We colour $q_i q_j \in E(K^{\mathbb{Q}})$ blue if either $q_i \leq q_j$ and $i \leq j$ or $q_j \leq q_i$ and $j \leq i$. Otherwise we colour it red. Now suppose that there is a subset Q of $\mathbb{Q}$ that is isomorphic to $(\mathbb{Q}, \leq)$ such that the complete graph induced by Q is monochromatic. This implies that Q is also order-isomorphic to $\mathbb{N}$ or its reversal, a contradiction, as $\mathbb{N}$ and $\mathbb{Q}$ are not order isomorphic, since $\mathbb{N}$ has a smallest element, but $\mathbb{Q}$ does not. □

4.4.2 The $K^{\mathbb{Q}}$-Building Game

In this subsection we present a winning strategy for Maker in the $K^{\mathbb{Q}}$-game. As in the previous sections we will do this by first describing a strategy according to which Maker should play and then proving that there actually is a subgraph with the desired property if Maker adheres to the strategy in the proof of Theorem 4.11.

In the following we define the $\mathbb{Q}$-*game strategy*. At the beginning of the game Maker chooses

- a partition $\mathcal{P}$ of the interval $(0,1) \subseteq \mathbb{Q}$ into open intervals such that there is an order $\leq_{\mathcal{P}}$ of $\mathcal{P}$ induced by the regular order on $\mathbb{Q}$ such that $(\mathcal{P}, \leq_{\mathcal{P}})$ is order isomorphic to $(\mathbb{Q}, <)$ where "$<$" is the usual ordering (this is possible by Proposition 4.8),
- an enumeration $(P_i)_{i\in\mathbb{N}}$ of $\mathcal{P}$, and
- a sequence $(n_i)_{i\in\mathbb{N}}$ of natural numbers such that every finite sequence of natural numbers appears infinitely often as a subsequence.

Informally speaking, by containing vertices of sufficiently many different intervals in $\mathcal{P}$, Maker can ensure that a subset of $V(G_M)$ actually is isomorphic to $\mathbb{Q}$, as $(\mathcal{P}, \leq_{\mathcal{P}})$ is isomorphic to $(\mathbb{Q}, \leq)$. The enumeration $(P_i)_{i\in\mathbb{N}}$ together with the sequence $(n_i)_{i\in\mathbb{N}}$ will be used to ensure that different intervals P_i, P_j are well-connected to ensure that there actually is a complete graph using vertices of many different P_i.

In her first turn, Maker claims the edge $\{0, 1\}$ for herself and sets $v_1 = 0$ and $v_2 = 1$. In a later turn, suppose $V(G_M) = \{v_1, \ldots, v_k\}$ where v_i is the i-th vertex that Maker added to G_M. For every vertex $v \in V(G_M)$ we define a finite sequence $S_v = (v_{i_1}, \ldots, v_{i_\ell})$ of the vertices of $N_M(v)$ that were added to G_M before v that represents the order in which Maker claimed the edges vv_{i_j}. In particular ℓ was the degree of v in G_M when Maker first chose a new fresh vertex after v. Maker uses the sequence S_{v_k} to determine from which interval $P \in \mathcal{P}$ to choose the vertex that she plays to next: suppose that $v_k \in P \in \mathcal{P}$ and $S_{v_k} = (v_{i_1}, \ldots, v_{i_\ell})$. Suppose further that $|\{v \in V(G_M) \colon S_v|_\ell = S_{v_k} \wedge v \in P\}| = m$, i.e. there are m previous vertices in P that Maker connected to the same vertices as v_k in the same order in the first ℓ moves of connecting them to G_M. Now consider the $(m+1)$-st time that $(i_1, i_2, \ldots, i_\ell)$ appears in $(n_i)_{i\in\mathbb{N}}$ as a subsequence and let n be the number appearing next. Then Maker wants to play to a vertex of P_n next.

For that purpose she considers the vertices v_i such that

(1) $v_i \in P_n \cap V(G_M)$,

(2) $S_{v_i}|_\ell = S_{v_k}$, and

(3) there are at most $\ell \cdot (\ell+1)$ vertices $v_j \in V(G_M)$ with $j < i$ satisfying (1) and (2).

Let us call this set F. If F contains fewer than $\ell(\ell+1)+1$ vertices, Maker chooses a fresh vertex $v_{k+1} \in P_{n_{k+1}}$, claims v_1v_{k+1} and begins the aforementioned process for v_{k+1}. Otherwise, F has size precisely $\ell(\ell+1)+1$ and Maker continues as follows: Define $L := \{v \in P \cap V(G_M) \colon S_v|_\ell = S_{v_k}\}$, i.e. L is the set of vertices of G_M that come from the same partition class of $\mathcal{P}$ as v_k and in their first ℓ moves of being connected to G_M Maker connected them to G_M in the same manner as v_k. We define a total order on F via

$$v_i < v_j :\Leftrightarrow \begin{cases} |N_M(v_i) \cap L| < |N_M(v_j) \cap L| & \text{, or} \\ |N_M(v_i) \cap L| = |N_M(v_j) \cap L| \text{ and } i < j & \end{cases}. \tag{4.6}$$

Maker claims v_iv_k such that v_i is minimal with respect to that order. Note that this choice is unique. In fact the second clause is only there to make this true.

Maker will use this order to play to the vertices of F in a 'balanced fashion', that is, Maker connects v_k to an available vertex with the smallest possible number of neighbours v_i that behave like v_k in the sense that $S_{v_i}|_\ell = S_{v_k}|_\ell$. This is useful because this ensures that for an infinite set of vertices of which each vertex behaves like v_k up to M-degree ℓ, for all but at most ℓ many vertices of F there are infinitely many vertices that get connected to that vertex.

Theorem 4.11. *The $\mathbb{Q}$-game strategy is a winning strategy for Maker in the complete rational number game.*

Note that Theorem 4.11 implies Theorem 4.7.

Proof. For the proof we again fix a sequence $(a_i)_{i\in\mathbb{N}}$ of natural numbers. This sequence should contain every number infinitely often and each a_i should be at most i. Moreover, we reuse the partition $\mathcal{P}$ and its enumeration $(P_i)_{i\in\mathbb{N}}$ from the $\mathbb{Q}$-game strategy. We recursively build for every $m \in \mathbb{N}$ a complete graph K^m and a set W_m such that

(a) $K^{m-1} \subseteq K^m \subseteq G_M$ with $V(K^m) = \{u_1, \dots, u_m\}$, and

(b) $W_m = \{w \in V(G_M)\colon S_w|_m = (u_1, \dots, u_m)\}$.

For $i \in [m]$ we denote by P^i the partition class of $\mathcal{P}$ that contains u_i. We define

$$\begin{aligned} \mathcal{Q}^i_m := \{P \in \mathcal{P}\colon |W_m \cap P| = \aleph_0 \text{ and } P^i <_\mathcal{P} P \text{ and}\\ \text{there is no } j \in [m] \text{ with } P^i <_\mathcal{P} P^j <_\mathcal{P} P\}\,. \end{aligned} \tag{4.7}$$

With these definitions we also require W_m to satisfy that

(c) $\mathcal{Q}^i_m$ contains a subset that, together with the order induced by $<_\mathcal{P}$, is order isomorphic to $(\mathbb{Q}, <)$ for every $i \in [m]$, and

(d) $u_m \in \bigcup \mathcal{Q}^{a_{m-1}}_{m-1}$, if $m > 1$.

Clearly, (a) implies that $\bigcup_{m\in\mathbb{N}} K^m = K^{\aleph_0}$ is a complete graph contained in Maker's subgraph G_M. By (d) we add a vertex between any pair of vertices $u_i, u_j \in K^{\aleph_0}$, thus $V(\bigcup_{m\in\mathbb{N}} K^m)$ contains a subset that is order isomorphic to $(\mathbb{Q}, \leq)$. The subgraph induced by these is the desired $K^\mathbb{Q}$. Properties (b) and (c) are needed to ensure that the recursion can be continued indefinitely.

Recursion start: Recall that $v_1 = 0$ was added to G_M as the first vertex according to the $\mathbb{Q}$-game strategy. We set $K^1 := (\{v_1\}, \emptyset)$ and $W_1 := V(G_M) \setminus \{v_1\}$. The requirements of (a) and (d) are empty in the initial step and therefore satisfied. Property (b) is true for W_1, as according to the $\mathbb{Q}$-game strategy every other vertex of G_M is connected to v_1 first. For every $P \in \mathcal{P}$ Maker added infinitely many vertices to G_M, as the i-th vertex she adds is a vertex of P_{n_i} and $(n_i)_{i\in\mathbb{N}}$ contains every number infinitely often. Thus $\mathcal{Q}_1^1 = \mathcal{P}$ and therefore $\mathcal{Q}_1^1$ satisfies (c).

Recursion step: Let $k \geq 2$ and let K^k and W_k satisfying (a) to (d) be given. We demonstrate how we can find a vertex u_{k+1} and a set W_{k+1} such that $K^{k+1} := G[V(K^k) \cup \{u_{k+1}\}]$ and W_{k+1} comply with (a) to (d).

By (c), $\mathcal{Q}_k^{a_k}$ contains a subset that is order isomorphic to $(\mathbb{Q}, <)$. Let $\mathcal{Q}$ be such a subset and x be the first element of $(a_i)_{i\in\mathbb{N}}$ such that P_x is an element of $\mathcal{Q}$ but contains no vertex of $V(K^k)$. We choose x in such a way because this ensures that there is a partition of $\mathcal{Q} \setminus \{P_x\}$ into two $\mathbb{Q}$-isomorphic subsets $\mathcal{Q}^-$ and $\mathcal{Q}^+$ such that all $Q^- \in \mathcal{Q}^-$ and $Q^+ \in \mathcal{Q}^+$ fulfil $Q^- <_{\mathcal{P}} P_x <_{\mathcal{P}} Q^+$. We will choose a vertex of P_x as u_{k+1}. This choice ensures (d).

Next we want to suitably restrict W_k to a set of vertices that were connected to a vertex of P_x, but we also need to make sure that we can preserve (c) for the subsequent steps. For this purpose we let $W' := \{w \in W_k \colon \text{there is } v \in P_x \text{ such that } S_w|_{k+1} = (u_1, \dots, u_k, v)\}$ and use this to define $\hat{\mathcal{Q}}^i$ similarly to the definition in (4.7): for $i \in [k]$ let P^i be the partition class of $\mathcal{P}$ that contains u_i and set $P^{k+1} := P_x$. We define

$$\hat{\mathcal{Q}}^i := \{P \in \mathcal{P} \colon |W' \cap P| = \aleph_0 \text{ and } P^i <_{\mathcal{P}} P \text{ and}$$
$$\text{there is no } j \in [k] \text{ with } P^i <_{\mathcal{P}} P^j <_{\mathcal{P}} P\}\,.$$

Claim 1. $\hat{\mathcal{Q}}^i$ contains a subset that is order isomorphic to $(\mathbb{Q}, <)$ for every $i \in [k+1]$.

Proof. **Case 1:** $i \in [k+1] \setminus \{a_k, k+1\}$**.** We prove this case by showing that any element of $\mathcal{Q}_k^i$ is also an element of $\hat{\mathcal{Q}}^i$. The claim then follows from (c).
Consider any $P \in \mathcal{Q}_k^i$. By definition we have $|P \cap W_k| = \aleph_0$ and by (b) every $w \in P \cap W_k$ was connected precisely to $V(K^k)$ in its first k moves of being connected to G_M. According to the $\mathbb{Q}$-game strategy, for infinitely many of them Maker played to a vertex of P_x next, thus $|P \cap W'| = \aleph_0$. As this is true for any $P \in \mathcal{P}$ with $|P \cap W_k| = \aleph_0$, this proves the claim for $i \in [k+1] \setminus \{a_k, k+1\}$.

Case 2: $i \in \{a_k, k+1\}$**.** With the same reasoning as in Case 1 any $P \in \mathcal{Q}^-$ is also in $\hat{\mathcal{Q}}^{a_k}$ and any $P \in \mathcal{Q}^+$ is also in $\hat{\mathcal{Q}}^{k+1}$. Since both, $\mathcal{Q}^-$ and $\mathcal{Q}^+$, are isomorphic to $\mathbb{Q}$, this implies the claim for $i \in \{a_k, k+1\}$. ■

To continue, consider the set F of the first $k(k+1)+1$ many vertices $u \in P^{k+1}$ with $S_u|_k = (u_1, \dots, u_k)$. There is such a set as any vertex $u \in P^{k+1} \cap W_k$ fulfils this property and there are infinitely many such vertices by the choice of P^{k+1}. Further, fix $i \in [k+1]$ and consider $Q \in \hat{\mathcal{Q}}^i$. For any vertex $q \in Q \cap W'$, Maker considered the set F in the $(k+1)$-st move of connecting q to G_M. As Maker played from $Q \cap W'$ to F in a balanced fashion, that is she played to the smallest vertex of F with respect to the order defined in (4.6), there are at most k vertices in F that have only finitely many neighbours in $Q \cap W'$. By choosing some superset of k vertices if necessary, we obtain a colouring of $\hat{\mathcal{Q}}^i$ indicating for every $Q \in \hat{\mathcal{Q}}^i$ a subset of F of size k such that all other vertices of F have infinitely many neighbours in $Q \cap W'$. By taking the complements of the sets of size k, the same colouring indicates for each Q for which subset $F' \subseteq F$ every vertex in F' was picked by Maker for infinitely many vertices v of $Q \cap W'$ to be played to in the $(k+1)$-st move of connecting v to G_M. There are $\binom{k(k+1)+1}{k}$ colours in this colouring of $\hat{\mathcal{Q}}^i$, since every k element subset of F gets assigned a colour and F has $k(k+1)+1$ elements. In particular, this is a finite number. Thus by Proposition 4.9 and Claim 1, there is a colour class that again contains a subset that is order isomorphic to $\mathbb{Q}$. We fix a suitable colour class $C^i \subseteq \hat{\mathcal{Q}}^i$ for every $i \in [k+1]$. As seen, any of the $k+1$ fixed colour classes excludes k vertices of F as the next candidate and since F has size $k(k+1)+1$, there is at least one vertex in F that is met by the fixed colour class C^i for every $i \in \mathbb{N}$. We choose the smallest such vertex as u_{k+1}. Then (b) and (c) are fulfilled by construction and (d) is fulfilled by the choice of x as mentioned above. Lastly, we set $K^{k+1} := G[V(K^k) \cup \{u_{k+1}\}]$, thus also (a) is ensured. $\square$

4.4.3 The Dense $K^{\mathbb{Q}}$-Building Game

Since Maker can always win in the $K^{\mathbb{Q}}$-building game, we now consider a variant of the game in which it is harder for Maker to achieve her goal, the dense $K^{\mathbb{Q}}$-building game. In this game Maker is doomed to fail. We prove this by providing a winning strategy for Breaker in this variant of the game.

In the following we define the *dense $\mathbb{Q}$-building pairing strategy.* At the beginning of the game Breaker picks

- an infinite sequence $\mathcal{I} := (I_j)_{j\in\mathbb{N}}$ of pairwise disjoint intervals of $\mathbb{Q}$,
- an enumeration $\mathcal{Q}$ of $\mathbb{Q}$, and
- an enumeration $(\{p_j, q_j\})_{j\in\mathbb{N}}$ of the 2-element subsets of $\mathbb{Q}$.

For any $s \in I_j \setminus \{p_j, q_j\}$ where s appears later in the enumeration $\mathcal{Q}$ than both p_j and q_j he pairs the edges $\{p_j s, q_j s\}$. Then, whenever Maker claims either $p_j s$ or $q_j s$ in one of her turns, Breaker claims the other in his following turn.

Lemma 4.12. *The dense $K^{\mathbb{Q}}$-building pairing strategy is a pairing strategy that is a winning strategy for Breaker in the dense rational number game.*

Proof. To verify that the strategy is a pairing strategy we need to verify that the pairs of edges $\{p_j s, q_j s\}$ are pairwise disjoint. For fixed j and $s \neq t \in I_j \setminus \{p_j, q_j\}$ it is true that $\{p_j s, q_j s\} \cap \{p_j t, q_j t\} = \emptyset$. For $i \neq j$ with $s \in I_i$ and $t \in I_j$, the intersection of $\{p_i s, q_i s\}$ and $\{p_j t, q_j t\}$ can only be non-empty if both $s \in \{p_j, q_j\}$ and $t \in \{p_i, q_i\}$. But this cannot happen because p_i, q_i, p_j and q_j appear in a unique order in the enumeration $\mathcal{Q}$ which ensures that Breaker considers only one of the pairs $\{p_i s, q_i s\}$ and $\{p_j t, q_j t\}$ in this case.

To see that the dense $\mathbb{Q}$-game strategy is a winning strategy for Breaker, suppose for a contradiction that Maker finishes a $K^{\mathbb{Q}}$ whose vertices are dense in $\mathbb{Q}$. Then there is $n \in \mathbb{N}$ such that $\{p_n, q_n\} \subseteq V(K^{\mathbb{Q}})$. But according to the dense $K^{\mathbb{Q}}$-building pairing strategy only finitely many vertices of I_n can be in $K^{\mathbb{Q}}$: since at some point both p_n and q_n have appeared in $\mathcal{Q}$, only finitely many other elements have appeared before and for every vertex $v \in I_n$ that appears later in $\mathcal{Q}$ the pair $\{p_n v, q_n v\}$ is considered in the dense $\mathbb{Q}$-game strategy. Thus, $V(K^{\mathbb{Q}})$ cannot be dense in I_n, which implies that the set of vertices is not dense in $\mathbb{Q}$. $\square$

4.5 Open Problems

One obvious variation of the game proposed in Section 4.1 immediately comes to mind: while in Section 4.2 and Section 4.3 the basic version of the game was altered by colouring the vertices, one could instead colour the edges of the board and again demand that G_M contains an isomorphic copy of the board as a subgraph. We will call this game the $K^{\aleph_0}$ *edge colouring game.* While in the vertex case the colour classes are very symmetric, this is different for edge colourings: it could happen that the subgraph induced by some colour class is locally finite, while the one for another is not.

Question 4.13. Let $n \in \mathbb{N}$. For which colourings $c\colon E(K^{\aleph_0}) \longrightarrow [n]$ is there a winning strategy for one of the players in the $K^{\aleph_0}$ edge colouring game?

Certainly one may consider a colouring with infinitely many colours as well.

Question 4.14. For which colourings $c\colon E(K^{\aleph_0}) \longrightarrow \mathbb{N}$ is there a winning strategy for one of the players in the $K^{\aleph_0}$ edge colouring game?

Another possible generalisation is to adapt the game to hypergraphs.

Question 4.15. Let H be a complete infinite k-regular hypergraph. In the Maker-Breaker game on H where it is Makers aim to have an isomorphic copy of H be contained in G_M, is there a winning strategy for Maker?

Naturally, we can consider a $\leq k$-regular hypergraph or even an infinite complete hypergraph instead of a k-regular one. Furthermore, we could also apply a vertex or an edge colouring to the board in any of these variants, just as in the $K^{\aleph_0}$-building game. Since the k-regular hypergraph variant is already very advanced, we will not state these as questions here and rather highlight one particularly intriguing question.

Question 4.16. Consider the Maker-Breaker game in which the players alternately claim finite subsets of $\mathbb{N}$ and it is Makers aim to claim an infinite set $\mathcal{F}$ of pairwise disjoint finite subsets of $\mathbb{N}$ as well as the union of every finite subset of $\mathcal{F}$. Is there a winning strategy for one of the players?

Chapter 5

Maker-Breaker (G, H)-Games on Uncountable Boards

Recall that we defined in Subsection 2.2.1 that for Maker-Breaker games on uncountable boards we assume that a game ends only after all edges of the board have been claimed. Thus, each outcome of the game defines a 2-colouring of $E(G)$. This suggests an even deeper connection to Ramsey type problems as in Chapter 4, although still the colourings in question are not arbitrary but are produced by players with particular goals in mind. Since in the context of uncountable boards the set-theoretic framework in which we work becomes very important we fix some notation: we write CH, GCH, DC, AD and $\mathfrak{p}$ for the continuum hypothesis, generalised continuum hypothesis, axiom of dependent choice, axiom of determinacy and the pseudo-intersection number respectively. See e.g. [Kun11] for an introduction to these concepts. There are colourings of the edges of K^{ω_1} with two colours without any monochromatic K^{ω_1} in ZFC (see [Sie33]), but if instead of the axiom of choice one assumes DC+AD, then there is always a monochromatic K^{ω_1} because ω_1 becomes measurable (see [Kan08, Theorem 28.2]) and hence weakly compact.

The existence of a monochromatic K^{ω,ω_1} when colouring the edges of K^{ω,ω_1} with two colours is even more dependent on the set-theoretic framework. While there is a colouring without a monochromatic copy in ZFC+CH, there is no such colouring in ZFC+$\omega_1 < \mathfrak{p}$. Since we could not find these particular statements formulated anywhere in the literature on infinite Ramsey theory, for the sake of completeness we include them here as Corollary 5.6 and Corollary 5.11.

These Ramsey-type results compare well to the corresponding results about the existence of a winning strategy for either player. Our main results are as follows:

Theorem 5.1. *It is independent of ZFC if Breaker has a winning strategy in the Maker-Breaker $(K^{\omega,\omega_1}, K^{\omega,\omega_1})$-game. He has one under ZFC+GCH*[1]*, while Maker has one under ZFC+$\omega_1 < \mathfrak{p}$.*

[1] A closer analysis shows that only CH is needed here, but we have chosen a simpler exposition over optimality of the results, since the independence is our main concern.

Theorem 5.2. *It is independent of ZFC if every 2-colouring of the edges of K^{ω,ω_1} admits a monochromatic copy of K^{ω,ω_1}. It is true in ZFC+$\omega_1 < \mathfrak{p}$ but fails under ZFC+CH.*

Theorem 5.3. *Assuming the consistency of AD, it is independent of ZF+DC if Breaker has a winning strategy in the Maker-Breaker $(K^{\omega_n}, K^{\omega_n})$-game for $n \in \{1,2\}$. He has such winning strategies under ZFC+GCH, while Maker has winning strategies in these games under ZF+DC+AD.*

Theorem 5.4. *Assuming the consistency of AD, it is independent of ZF+DC if Breaker has a winning strategy in the Maker-Breaker $(K^{\omega_1}, K^{\mathsf{club}})$-game.*

We first present the winning strategies for Breaker in Section 5.1 and then the winning strategies for Maker in Section 5.2.

5.1 The Winning Strategies of Breaker under GCH

Proposition 5.5 (ZFC+GCH)**.** *For every infinite cardinal κ, Breaker has a winning strategy in the Maker-Breaker $(K^{\kappa^+}, K^{\kappa,\kappa^+})$-game.*

Proof. Let us assume that K^{κ^+} is represented as the complete graph on the vertex set κ^+. Working under GCH, we fix an enumeration $\{A_\alpha \colon \alpha < \kappa^+\}$ of $[\kappa^+]^\kappa$ and for each $\alpha < \kappa^+$, we pick a surjective function $f_\alpha \colon \kappa \to \{A_\beta \colon \beta \leq \alpha\}$). Whenever Maker plays an edge $\{\beta, \alpha\}$ with $\beta < \alpha$ and there is a $\gamma < \kappa$ such that this is the $(\gamma+1)$-st downwards edge from α she claims, Breaker chooses the smallest $\delta \in f_\alpha(\gamma)$ for which $\{\delta, \alpha\}$ is available, and plays $\{\delta, \alpha\}$ if such a δ exists - otherwise he plays arbitrarily.

Suppose for a contradiction that Maker manages to build a K^{κ,κ^+} (despite Breaker playing as above) and let A be its smaller and B its bigger vertex class. Then there is an $\alpha < \kappa^+$ with $A_\alpha = A$. Fix a $\beta \in B$ with $\beta > \max\{\alpha, \sup A\}$ and let $\gamma < \kappa$ with $f_\beta(\gamma) = A$. At the turn when Maker claims a downwards edge from β for the $(\gamma+1)$-st time, there are still κ many $\delta \in A$ for which $\{\delta, \beta\}$ is available, thus Breaker's next play is $\{\delta, \beta\}$ for the smallest such δ. This contradicts $\{\delta, \beta\} \in E(G_M)$. □

The corresponding negative Ramsey-result can be proved in a similar manner:

Corollary 5.6 (ZFC+GCH)**.** *For every infinite cardinal κ, there exists a 2-colouring of the edge set of K^{κ,κ^+} without a monochromatic copy of K^{κ,κ^+}.*

Proof. Let $\{v_\alpha : \alpha < \kappa^+\}$ be an enumeration of the larger vertex class and let $\{A_\alpha : \alpha < \kappa^+\}$ be an enumeration of $[\kappa^+]^\kappa$. For each $\alpha < \kappa^+$, we colour the edges incident with v_α in such a way that for every $\beta \leq \alpha$ both colours appear among the edges between v_α and A_β. This clearly ensures that no set A can be the smaller vertex class of a monochromatic copy of K^{κ,κ^+} and therefore no such a monochromatic copy exists. □

Observation 5.7. If Breaker has a winning strategy in the Maker-Breaker (G, H)-game, then he also has one in every Maker-Breaker (G', H')-game where G' is a subgraph of G and H' is a supergraph of H.

Since K^{κ,κ^+} is a subgraph of K^{κ^+}, Observation 5.7 guarantees that Proposition 5.5 has the following consequences:

Corollary 5.8 (ZFC+GCH)**.** *For every infinite cardinal κ, Breaker has a winning strategy in the following games:*

(1) *the Maker-Breaker $(K^{\kappa,\kappa^+}, K^{\kappa,\kappa^+})$-game,*

(2) *the Maker-Breaker $(K^{\kappa^+}, K^{\kappa^+})$-game, and*

(3) *the Maker-Breaker $(K^{\kappa^+}, K^{\mathsf{club}})$-game.*

5.2 Winning Strategies for Maker

This section is divided into three subsections. In Subsection 5.2.1 we investigate the $(K^{\omega,\omega_1}, K^{\omega,\omega_1})$-game, in Subsection 5.2.2 we look into the $(K^{\omega_1}, K^{\omega_1})$-game and the $(K^{\omega_2}, K^{\omega_2})$-game and lastly the $(K^{\omega_1}, K^{\mathsf{club}})$-game in Subsection 5.2.3.

5.2.1 The Maker-Breaker $(K^{\omega,\omega_1}, K^{\omega,\omega_1})$-Game

A set $\mathcal{F}$ of sets has the *strong finite intersection property* if the intersection of any finitely many elements of $\mathcal{F}$ is infinite. Given two sets X and Y, we write $X \subseteq^* Y$ if $X \setminus Y$ is finite. A *pseudo-intersection* for a set $\mathcal{F}$ of sets is a set P with $P \subseteq^* F$ for all $F \in \mathcal{F}$. The cardinal $\mathfrak{p}$ is the minimum cardinality of a set $\mathcal{F}$ of subsets of ω that has the strong finite intersection property but does not admit an infinite pseudo-intersection. Clearly $\aleph_0 < \mathfrak{p} \leq 2^{\aleph_0}$ and it is known that $\omega_1 < \mathfrak{p}$ is consistent relative to ZFC (see [Kun11, Lemma III.3.22 on p. 176]).

Proposition 5.9. *Maker has a winning strategy in the Maker-Breaker $(K^{\omega,\omega_1}, K^{\omega,\omega_1})$-game if $\omega_1 < \mathfrak{p}$.*

Proof. Let U and V be the two sides of the bipartite graph K^{ω,ω_1}, where $|U| = \omega$ and $|V| = \omega_1$. We denote the subgraph of G induced by the edges Maker claimed before turn α by G_M^α and we write $N_{G_M^\alpha}(v)$ for the set of the neighbours of v in this graph.

During the game Maker will choose a sequence $\langle v_\alpha \colon \alpha < \omega_1 \rangle$ of distinct vertices from V and a sequence $\langle N_\alpha \colon \alpha < \omega_1 \rangle$ of subsets of U in such a way as to ensure that for any $\alpha < \omega_1$

(1) $N_\alpha \subseteq N_{G_M}^{\omega\cdot(\alpha+1)}(v_\alpha)$, and

(2) the set $\{N_\beta \colon \beta \leq \alpha\}$ has the strong finite intersection property.

Assume that turn $\alpha\cdot\omega$ has just begun for some $\alpha < \omega_1$ and that Maker has constructed suitable v_β and N_β for all $\beta < \alpha$. She picks v_α to be any fresh vertex in V. Using (2) for all $\beta < \alpha$, we know that the set $\{N_\beta \colon \beta < \alpha\}$ has the strong finite intersection property. Let P_α be an infinite pseudo-intersection of this family. In each of the next ω turns, Maker claims an edge $\{u, v_\alpha\}$ with $u \in P_\alpha$. Let N_α be the set of all the endpoints $u \in U$ of these edges. This construction satisfies (1) and (2) for α.

At the end of the game $\{N_\alpha \colon \alpha < \omega_1\}$ has the strong finite intersection property and hence (by the assumption $\omega_1 < \mathfrak{p}$) admits an infinite pseudo-intersection P. By the definition of P, for each $\alpha < \omega_1$, the set $P \setminus N_\alpha$ is finite. Then there exists an uncountable $O \subseteq \omega_1$ and a finite $F \subseteq P$ such that $P \setminus N_\alpha = F$ for every $\alpha \in O$. Finally, $(P \setminus F) \cup \{v_\alpha \colon \alpha \in O\}$ induces a copy of K^{ω,ω_1}, all of whose edges have been claimed by Maker. $\square$

Remark 5.10. The same proof shows that Maker has a winning strategy in the Maker-Breaker $(K^{\omega,\kappa}, K^{\omega,\kappa})$-game for every $\kappa < \mathfrak{p}$ with $\mathsf{cf}(\kappa) > \aleph_0$.

The proof of Proposition 5.9 leads to the following positive Ramsey result:

Corollary 5.11. *If $\omega_1 < \mathfrak{p}$, then any 2-colouring of the edges of K^{ω,ω_1} admits a monochromatic copy of K^{ω,ω_1}.*

Proof. Call the colours red and blue, and call the countable and uncountable sides of the original graph U and V respectively. We pick a free ultrafilter $\mathcal{U}$ on U. Then for each $v \in V$ either the set $N_r(v)$ of the red neighbours of v is in $\mathcal{U}$ or the set $N_b(v)$ of the blue neighbours. We may assume that there is an uncountable $V' \subseteq V$ such that $N_r(v) \in \mathcal{U}$ for each $v \in V'$. Since $\mathcal{U}$ is a free ultrafilter, the family $\{N_r(v) \colon v \in V'\}$ has the strong finite intersection property and therefore (by $\omega_1 < \mathfrak{p}$) admits an infinite pseudo-intersection P. This means that for every $v \in V'$ the set $P \setminus N_r(v)$ is finite. Then there exists an uncountable $V'' \subseteq V'$ and finite $F \subseteq P$ such that $P \setminus N_r(v) = F$ for each $v \in V''$ and hence $(P \setminus F) \cup V''$ induces a red copy of K^{ω,ω_1}. $\square$

Question 5.12. Is it consistent with ZFC+$\aleph_\omega < 2^{\aleph_0}$ that Maker has a winning strategy in the Maker-Breaker $(K^{\omega,\omega_\omega}, K^{\omega,\omega_\omega})$-game?

Theorem 5.1 is implied by the case $\kappa = \omega$ of Corollary 5.8/(1) together with Proposition 5.9. Similarly, Theorem 5.2 follows from Corollary 5.6 and Corollary 5.11.

5.2.2 The Maker-Breaker $(K^{\omega_1}, K^{\omega_1})$-Game and the Maker-Breaker $(K^{\omega_2}, K^{\omega_2})$-Game

Proposition 5.13 (ZF)**.** *If either κ is measurable or $\kappa = \omega$, then Maker has a winning strategy in the Maker-Breaker (K^κ, K^κ)-game.*

Proof. A sub-binary Hausdorff tree is a set theoretic tree T in which each vertex has at most two children and no two vertices at any limit level have the same set of predecessors.

During the game Maker builds a sequence $\langle T_\alpha \colon \alpha \leq \kappa \rangle$ of sub-binary Hausdorff trees with root 0 and $T_\alpha \subseteq \kappa$ of height at most $1+\alpha$ such that

(a) (i) $T_0 = \{0\}$,

(ii) $T_{\alpha+1}$ is obtained from T_α by inserting a new maximal element,

(iii) $T_\alpha = \bigcup_{\beta<\alpha} T_\beta$ if α is a limit ordinal, and

(b) for every distinct $<_{T_\alpha}$-comparable $u, v \in T_\alpha$, the edge $\{u, v\}$ is claimed by Maker in the game.

Suppose that $\alpha = \beta + 1$ and T_β is already defined. Maker picks the smallest ordinal v such that no edge incident with v is claimed and claims edge $\{0, v\}$. Then, for as long as she can, on each following turn she connects v to vertices in T_β in such a way that:

(1) she maintains that the current neighbourhood of v in her graph is a downward closed chain in T_β, and

(2) whenever she claims some $\{u, v\}$, then Breaker has no edge between v and the subtree $T_{\beta,u}$ of T_β rooted at u.

Note that, at any step at which v has a largest Maker-neighbour in T_β and this neighbour has two children in T_β, she can proceed. Moreover, she can also proceed even if there is no such largest Maker-neighbour as long as there is some element of T_β whose predecessors are precisely the Maker-neighbours of v in T_β. Thus, if Maker is unable to continue this process with v, then either v has a largest Maker-neighbour in T_β which has at most one child or else there is no vertex in T_β with precisely the Maker-neighbours of u as its predecessors. In either case we can define $T_{\beta+1}$ by adding v to T_β with its current set of Maker-neighbours as its predecessors, and Maker starts a new phase with a new fresh vertex.

It is enough to show that there is a κ-branch B in T_κ, because then $G_M[B]$ is a copy of K^κ by (b). Since $|T_\kappa| = \kappa$ by (a), we can fix a κ-complete free ultrafilter $\mathcal{U}$ on T_κ.

By transfinite recursion we build a κ-branch. Let $v_0 := 0$. Suppose that there is an $\alpha < \kappa$ such that the $<_{T_\kappa}$-increasing sequence $\langle v_\beta \colon \beta < \alpha \rangle$ is already defined and for each $\beta < \alpha$, $T_{\kappa, v_\beta} \in \mathcal{U}$. If α is a limit ordinal, then since $\bigcap_{\beta<\alpha} T_{v_\beta} \in \mathcal{U}$ by the κ-completeness of $\mathcal{U}$, there is at least one vertex of T with all v_β as predecessors. We define v_α to be the unique minimal such vertex such that $T_{v_\alpha} = \bigcap_{\beta<\alpha} T_{v_\beta} \in \mathcal{U}$. If $\alpha = \beta + 1$, then $T_{\kappa, v_\beta} \in \mathcal{U}$ by assumption. Since T_κ is sub-binary, v_β has a unique child v satisfying $T_{\kappa, v} \in \mathcal{U}$ and we let $v_{\beta+1} := v$. The recursion is done and $\{v_\alpha \colon \alpha < \kappa\}$ is clearly a κ-branch. $\square$

We remark that this strategy is quite flexible and deals also with a number of variants of the Maker-Breaker game. For example, if Breaker is allowed $k < \omega$ moves for every move that Maker picks, simply take a sub-$(k+1)$-regular Hausdorff tree, in which every node has at most $k+1$ children. Furthermore, if in addition Breaker is allowed to go first in every turn, simply weaken the Hausdorff assumption to the requirement that at most $k+1$ vertices at a limit level have the same set of predecessors.

Since ω_1 and ω_2 are measurable cardinals under ZF+DC+AD ([Kan08, Theorems 28.2 and 28.6]), the cases $\kappa \in \{\omega, \omega_1\}$ of Corollary 5.8/(2) and the cases $\kappa \in \{\omega_1, \omega_2\}$ of Proposition 5.13 together imply Theorem 5.3.

5.2.3 The Maker-Breaker $(K^{\omega_1}, K^{\mathsf{club}})$-Game

Proposition 5.14. *Under ZF+DC+AD, Breaker does not have a winning strategy in the Maker-Breaker* $(K^{\omega_1}, K^{\mathsf{club}})$*-game.*

Proof. First of all, the club filter on ω_1 is a countably complete free ultrafilter under ZF+DC+AD (this is explicit in the proof of [Kan08, Theorem 28.2]). Furthermore, it is normal [DH21, Proposition 4.1]. Thus for any 2-colouring of $[\omega_1]^2$ there exists a colour with a monochromatic K^{club} (the standard proof of this for arbitrary normal ultrafilters uses only ZF, see [Jec03, Theorem 10.22]). It follows that if Breaker successfully prevents Maker from building a K^{club}, then he necessarily builds a K^{club} himself.

Suppose for a contradiction that Breaker has a winning strategy. We shall show that Maker can 'steal' this winning strategy. Indeed, Maker picks an arbitrary edge in turn 0 as well as in each limit turn while in successor turns she pretends to be Breaker and claims edges according to his winning strategy. This is a winning strategy for Maker, a contradiction. $\square$

Theorem 5.4 follows from the case $\kappa = \omega$ of Corollary 5.8/(3) and Proposition 5.14.

Remark 5.15. The same strategy stealing argument shows that if κ is a weakly compact cardinal, then Breaker does not have a winning strategy in the Maker-Breaker (K^κ, K^κ)-game.

Remark 5.16. We did not really use the full power of AD, just some consequences that are weaker in the sense of consistency strength than AD itself. The axiom-system ZF+DC+'ω_1 is measurable' is equiconsistent with ZFC+'there exists a measurable cardinal' (see [Jec68]). The club filter being an ultrafilter is a strictly stronger assumption, for more details see p. 3 in [DH21].

5.3 Open Problems

Our results raise the following natural questions:

Question 5.17. Is it consistent with ZFC that neither Maker nor Breaker has a winning strategy in the Maker-Breaker $(K^{\omega,\omega_1}, K^{\omega,\omega_1})$-game?

Question 5.18. Does Breaker have a winning strategy in the Maker-Breaker $(K^{\omega_1}, K^{\omega_1})$-game under ZFC?

Question 5.19. Does Maker have a winning strategy in the Maker-Breaker $(K^{\omega_1}, K^{\mathsf{club}})$-game under ZF+DC+AD?

Part II

Directed and Bidirected Graphs

Chapter 6

Ubiquity of Directed Graphs

In this chapter we start the investigation of ubiquity in directed graphs. Recall from the definition in Subsection 2.3.2 that rays and double rays in this thesis are oriented versions of the undirected rays and undirected double rays in literature, together with linear orders. We begin this Chapter 6 by characterising which rays are ubiquitous regarding the subdigraph relation in Section 6.1. Whenever we write ubiquitous without specifying the relation, we refer to the subdigraph relation. We prove the following theorem:

Theorem 6.1. *A ray is ubiquitous if and only if it has a finite number of turns.*

In Section 6.2 we extend the investigation of ubiquity in digraphs, akin to Halins result [Hal70], to the class of double rays. It turns out that – in contrast to rays – it is not only relevant whether the number of turns is finite or infinite, but also the parity of this number plays a role. Our theorem about double rays reads as follows:

Theorem 6.2. *A double ray with at least one turn is ubiquitous if and only if it has a (finite) odd number of turns.*

Before we start the investigation let us present some open problems concerning the ubiquity of different digraphs. The statement of Theorem 6.2 notably omits one special case that turns out to be quite challenging:

Problem 6.3. Is the consistently oriented double ray, i.e. every vertex has in-degree and out-degree 1, ubiquitous?

More generally, one can investigate ubiquity of digraphs whose underlying undirected graphs are trees. However, even the question of which undirected trees are subgraph-ubiquitous is unsolved and only known for ubiquity with respect to weaker relations such as the topological minor relation. Therefore it might be sensible also to discuss ubiquity of digraphs with respect to weaker relations such as butterfly minors. Moreover, since proving or disproving the ubiquity of consistently oriented double rays is not easy, we propose to initially consider *out-trees*, i.e. trees in which all edges are oriented away from the root.

Problem 6.4. Which out-trees are ubiquitous concerning a fitting notion of ubiquity?

6.1 Ubiquity of Oriented Rays

In this section we develop novel methods that enhance the techniques in the field of ubiquity theory in order to prove Theorem 6.1. For the forward implication of Theorem 6.1 (see Subsection 6.1.1) we follow the proof for Halin's ray ubiquity result in undirected graphs which can be found in [Die17, Theorem 8.2.5 (i)]. To make this possible we require sets of arbitrarily many disjoint copies of an oriented ray that have an additional property, they must be forked. The existence of such forked sets, which we prove in Lemma 6.7, is our key contribution in the proof of the forward implication.

For the backwards implication of Theorem 6.1 (see Subsection 6.1.2) we construct a counterexample from infinitely many disjoint copies of an oriented ray by identifying vertices. In the proof of Theorem 6.10 we extend the technique of identifying vertices according to a recursive process.

6.1.1 Positive Results

In this subsection we prove

Theorem 6.5. *A ray is ubiquitous if it has finitely many turns.*

The proof of Theorem 6.5 will be an easy consequence of Theorem 6.6 together with Lemma 6.7 below. The following Theorem is a variant of Halin's ray ubiquity result for digraphs with an additional restriction on the start vertices of the rays. The proof given here is derived from the proof of Halin's result in [Die17, Theorem 8.2.5 (i)].

Theorem 6.6. *Let D be a digraph, R an out-oriented ray, $\mathcal{F}$ a thick R-tribe in D and $X \subseteq V(D)$ such that each member of $\mathcal{F}$ has its first vertex in X. Then there are infinitely many disjoint out-rays in D whose first vertices are contained in X.*

Proof. We will recursively fix for every $n \in \mathbb{N}^+$ a set $\mathcal{R}^n = \{R_1^n, \dots, R_n^n\}$ of pairwise disjoint out-rays in D and a set of vertices $\{u_1^n, \dots, u_n^n\}$ such that for every $k \in [n]$:

- R_k^n has its first vertex in X,
- $u_k^n \in R_k^n$, and
- $R_k^n u_k^n \subsetneq R_k^{n+1} u_k^{n+1}$.

Then $\{\bigcup_{n \geq k} R_k^n u_k^n \colon k \in \mathbb{N}\}$ is an infinite set of pairwise disjoint out-rays where each of the rays has its first vertex in X.

For $n = 1$ we pick one ray $R_1^1 \in \bigcup \mathcal{F}$, set $\mathcal{R}^1 := \{R_1^1\}$ and pick $u_1^1 \in R_1^1$ arbitrarily. This satisfies the required properties.

Now let $\ell > 1$ and suppose that for all $i \in [\ell]$ there are sets $\mathcal{R}^i$ and $\{u_1^i, \dots, u_i^i\}$ subject to the conditions above. Consider a layer F of $\mathcal{F}$ of size at least $|\bigcup\{R_i^\ell u_i^\ell \colon i \in [\ell]\}| + \ell^2 + 1$. First, we delete from F every ray that meets an oriented path $R_i^\ell u_i^\ell$ for some $i \in [\ell]$. Then there are still at least $\ell^2 + 1$ rays left in F. Next, we repeatedly check whether there is a ray $R_i^\ell \in \mathcal{R}^\ell$ that meets at most ℓ elements of F. If that is the case, we set $R_i^{\ell+1} := R_i^\ell$, choose a vertex $u_i^{\ell+1}$ beyond u_i^ℓ on R_i^ℓ arbitrarily, and delete the at most ℓ many rays from F that have non-empty intersection with $R_i^{\ell+1}$. Suppose that after $m \le \ell$ many steps, every ray in $\mathcal{R}^\ell$ meets either none or more than ℓ many rays from the reduced F, which we will refer to as F'.

Consider the $(\ell - m)$-sized subset $J \subseteq [\ell]$ containing all $j \in [\ell]$ for which $R_j^{\ell+1}$ has not yet been defined. Then any ray R_j^ℓ with $j \in J$ meets more than ℓ rays from F'. We deleted at most $|\bigcup\{R_i^\ell u_i^\ell \colon i \in [\ell]\}|$ rays from F in the first step and at most $m\ell$ in the second step, thus F' has size at least

$$|\{R_i^\ell u_i^\ell \colon i \in [\ell]\}| + \ell^2 + 1 - |\{R_i^\ell u_i^\ell \colon i \in [\ell]\}| - m\ell = (\ell - m)\ell + 1.$$

For any ray R_j^ℓ with $j \in J$ we fix the vertex $c_j \in R_j^\ell$ which is the first intersection of R_j^ℓ with the ℓ-th ray from F' that it meets. Note that c_j lies beyond u_j^ℓ on R_j^ℓ. Then $\bigcup_{j \in J} R_j^\ell c_j$ meets at most $|J|\ell = (\ell - m)\ell$ rays from F'. Therefore, there is at least one ray left in F' that is disjoint from $\bigcup_{j \in J} R_j^\ell c_j$ and we pick this ray as $R_{\ell+1}^{\ell+1}$. We choose an arbitrary vertex $u_{\ell+1}^{\ell+1} \in R_{\ell+1}^{\ell+1}$, define $F^* := F' \setminus \{R_{\ell+1}^{\ell+1}\}$, and write $F^* = \{S_i \colon i \in I\}$ for a suitable index set I.

Now for any ray $S_i \in F^*$ we choose a vertex w_i that lies beyond all vertices of $\bigcup_{j \in J} u_j^\ell R_j^\ell c_j$ on S_i. Consider the finite subdigraph

$$H := \bigcup_{j \in J} u_j^\ell R_j^\ell c_j \cup \bigcup_{i \in I} S_i w_i$$

of D. Additionally, we define $U := \{u_j^\ell \colon j \in J\}$ and $W := \{w_i \colon i \in I\}$. We show that for any set $Z \subseteq V(H)$ of fewer than $\ell - m$ vertices there is a U–W path in $H - Z$. Indeed, Z misses at least one path of the form $u_j^\ell R_j^\ell c_j$ and since there are $\ell \ge \ell - m$ many paths $S_i w_i$ with $u_j^\ell R_j^\ell c_j \cap S_i w_i \neq \emptyset$, at least one such path $S_i w_i$ avoids Z. Let v_j be the first vertex on $u_j^\ell R_j^\ell c_j$ which lies on $S_i w_i$; then $u_j^\ell R_j^\ell v_j S_i w_i$ is a path from u_j^ℓ to w_i: firstly, by the choice of v_j the graph of $u_j^\ell R_j^\ell v_j S_i w_i$ clearly is an oriented path. Secondly, the path $u_j^\ell R_j^\ell v_j$ is directed from u_j^ℓ to v_j since $v_j \in u_j^\ell R_j^\ell$. Thirdly, the path $v_j S_i w_i$ is directed from v_j to w_i since v_j lies in $S_i w_i$. Thus by Menger's Theorem [BJG08, Theorem 7.3.1], there is a set $\mathcal{P}$ of $\ell - m = |J|$ pairwise disjoint U–W paths in H.

For all $j \in J$, we write P_j for the path in $\mathcal{P}$ starting at u_j^ℓ. Let $h \colon J \to I$ such that $w_{h(j)}$ is the endvertex of P_j in W. Now we define

$$R_j^{\ell+1} := R_j^\ell u_j^\ell P_j w_{h(j)} S_{h(j)}$$

and $u_j^{\ell+1} := w_{h(j)}$, which clearly fulfils the required properties. $\square$

Lemma 6.7. *Let D and H be digraphs, $\hat{H} \subseteq H$ a finite subdigraph and let $\mathcal{E}$ be a thick H-tribe in D. Then there is a thick H-tribe $\mathcal{F}$ in D that is forked at $\hat{H}$.*

Proof. For all $n \in \mathbb{N}$, we recursively define a set F_n containing at least n disjoint copies of H in D and a thick subtribe $\mathcal{E}_n$ of $\mathcal{E}$ such that

(i) the H-tribe $\mathcal{F}_n := \{F_0, \dots, F_n\}$ is forked at $\hat{H}$,

(ii) for each $H_1 \in \bigcup \mathcal{E}_n$ and each $H_2 \in \bigcup \mathcal{F}_n$ the digraph $\hat{H}_1$ is disjoint from H_2 and the digraph $\hat{H}_2$ is disjoint from H_1.

In the end, $\{F_n \colon n \in \mathbb{N}\}$ will be a thick H-tribe satisfying the lemma. For the first step we set $F_0 := \emptyset$ and $\mathcal{E}_0 := \mathcal{E}$. Now suppose that $\mathcal{F}_{n-1}$ and $\mathcal{E}_{n-1}$ are already defined. Set $h := |\hat{H}|$ and choose a layer L from $\mathcal{E}_{n-1}$ of size at least $h+n$. We will choose F_n as an n-element subset of L. Then $\mathcal{F}_n$ will be forked at $\hat{H}$ since (i) and (ii) hold for $\mathcal{E}_{n-1}$ and $\mathcal{F}_{n-1}$. Our task is to find a suitable subset F_n of L and a thick subtribe $\mathcal{E}_n$ of $\mathcal{E}$ such that $\bigcup \mathcal{F}_n$ and $\bigcup \mathcal{E}_n$ satisfy (ii). We begin by deleting from each layer $M \neq L$ of $\mathcal{E}_{n-1}$ any element that has non-empty intersection with some $H' \in L$ in its subdigraph $\hat{H}'$. Note that for every digraph $H' \in L$ there are at most $|\hat{H}'| = h$ many digraphs from M which meet $\hat{H}'$. Therefore we delete from every layer of $\mathcal{E}_{n-1}$ at most $h \cdot |L|$ elements and the resulting subtribe $\mathcal{C}$ of $\mathcal{E}_{n-1}$ is still a thick tribe in D. Note that in particular, L is not a layer of $\mathcal{C}$.

Claim. For every $j \in \mathbb{N}$ there is a subset $L_j \subseteq L$ with $|L_j| = n$ and a subset C_j with $|C_j| \geq j$ of a layer of $\mathcal{C}$ such that for any $H_1 \in L_j$ and any $H_2 \in C_j$ the digraph H_1 is disjoint from $\hat{H}_2$ and H_2 is disjoint from $\hat{H}_1$.

Proof. Let $j \in \mathbb{N}$ and C a layer of $\mathcal{C}$ of size at least $j\binom{|L|}{n}$. By the construction of $\mathcal{C}$, we only need to find sets $L_j \subseteq L$ and $C_j \subseteq C$ such that no $H_1 \in L_j$ meets any $H_2 \in C_j$ in its subdigraph $\hat{H}_2$. For every $H' \in C$, at most $|\hat{H}'| = h$ elements of L meet $\hat{H}'$. Since $|L| \geq h+n$, we can choose for every $H' \in C$ a subset of n elements of L such that each of these does not meet $\hat{H}'$. This defines a map $\alpha \colon C \to \mathfrak{L} := \{L' \subseteq L \colon |L'| = n\}$. Since $|C| \geq j\binom{|L|}{n} = j|\mathfrak{L}|$, there is a set $L_j \in \mathfrak{L}$ with $|\alpha^{-1}(L_j)| \geq j$ by pigeonhole principle. Then L_j and $C_j := \alpha^{-1}(L_j)$ are as desired. ■

Since L has only finitely many subsets, there is an infinite strictly increasing sequence $(j_k)_{k \in \mathbb{N}}$ in $\mathbb{N}$ such that the sets L_{j_k} coincide for all $k \in \mathbb{N}$. We choose this as the set F_n. By the claim, $\mathcal{E}_n := \{C_{j_k} \colon k \in \mathbb{N}\}$ is a thick subtribe of $\mathcal{E}_{n-1}$ satisfying (ii). This concludes the proof. □

Proof of Theorem 6.5. Let R be a ray with finitely many turns. This implies that all but finitely many edges are oriented the same way. Furthermore, let D be a digraph and $\mathcal{E}$ a thick R-tribe in D. We show that D contains infinitely many copies of R. By Proposition 2.6 we may assume that all but finitely many edges of R are out-oriented.

Let $\hat{R}$ be the (connected) subdigraph of R that consists precisely of all finite phases of R. By Lemma 6.7, there is a thick subtribe $\mathcal{F}$ of $\mathcal{E}$ that is forked at $\hat{R}$. Consider the set X which contains for any $R' \in \bigcup \mathcal{F}$ the first vertex of the out-ray $R' - \hat{R}'$, and the set $Y := \bigcup_{R' \in \bigcup \mathcal{F}} V(\hat{R}')$. By deleting $\hat{R}'$ from each member R' of $\mathcal{F}$, we obtain a thick $(R - \hat{R})$-tribe in $D - Y$. Hence, by Theorem 6.6 there exists an infinite family $(R_i)_{i \in \mathbb{N}}$ of disjoint out-rays in $D - Y$ such that each R_i starts in a vertex $r_i \in X$. By definition of X, for all $i \in \mathbb{N}$ there is $S_i \in \mathcal{F}$ such that r_i is the first vertex of $S_i - \hat{S}_i$. Note that $\hat{S}_i$ and $\hat{S}_j$ are disjoint for $i \neq j$ since $\mathcal{F}$ is forked at $\hat{R}$. Finally, by combining each initial segment $\hat{S}_i$ with the out-ray R_i, we obtain infinitely many disjoint copies of R in D. □

6.1.2 Negative Results

For the proof that all rays with infinitely many turns are non-ubiquitous, we distinguish between two different cases, partitioned into two subsections. In Theorem 6.8 we show that any ray with bounded representing sequence is non-ubiquitous and in Theorem 6.10 we show that any ray with unbounded representing sequence is non-ubiquitous.

Rays with Bounded Representing Sequence

We prove the following theorem.

Theorem 6.8. *All rays with bounded representing sequence are non-ubiquitous.*

Proof. Let R be an arbitrary ray with bounded representing sequence. Let c be the largest natural number that occurs infinitely often in its representing sequence. We construct a digraph D that contains arbitrarily many but not infinitely many disjoint copies of R. By Proposition 2.6, we may assume that infinitely many phases of length c in R are out-oriented.

For the construction of D, we set

$$I := \{(n, m) \in \mathbb{N}^2 \colon n \leq m\}$$

and let $(R(n, m))_{(n,m) \in I}$ be a family of pairwise disjoint copies of R. Furthermore, set

$$J := \{((n^0, m^0), (n^1, m^1)) \in I^2 \colon m^0 < m^1\}$$

and let $((n_i^0, m_i^0), (n_i^1, m_i^1))_{i \in \mathbb{N}}$ be a sequence in J which contains every element of J infinitely often. We begin by setting

$$D_0 := \bigcup_{(n,m) \in I} R(n, m),$$

which is the disjoint union of infinitely many copies of R. Now we define a sequence $(g_i^0, g_i^1)_{i \in \mathbb{N}}$ of pairwise disjoint pairs of vertices of D_0 recursively with $g_i^0 \in R(n_i^0, m_i^0)$ and $g_i^1 \in R(n_i^1, m_i^1)$ for all $i \in \mathbb{N}$.

If (g_j^0, g_j^1) has been defined for all $j < i$, we pick for $\varepsilon \in \{0,1\}$ the vertex g_i^ε beyond all vertices $g_0^0, g_0^1, \ldots, g_{i-1}^0, g_{i-1}^1$ on $R(n_i^\varepsilon, m_i^\varepsilon)$ with the following properties (see Figure 6.1):

(i) g_i^1 is a turn in $R(n_i^1, m_i^1)$ at the start of an out-oriented phase of length c, and

(ii) g_i^0 is a turn in $R(n_i^0, m_i^0)$ at the end of an out-oriented phase of length c with the property that $|R(n_i^0, m_i^0)g_i^0| > |R(n_i^1, m_i^1)g_i^1|$.

This is possible since $R(n_i^0, m_i^0)$ and $R(n_i^1, m_i^1)$ contain infinitely many out-oriented phases of length c.

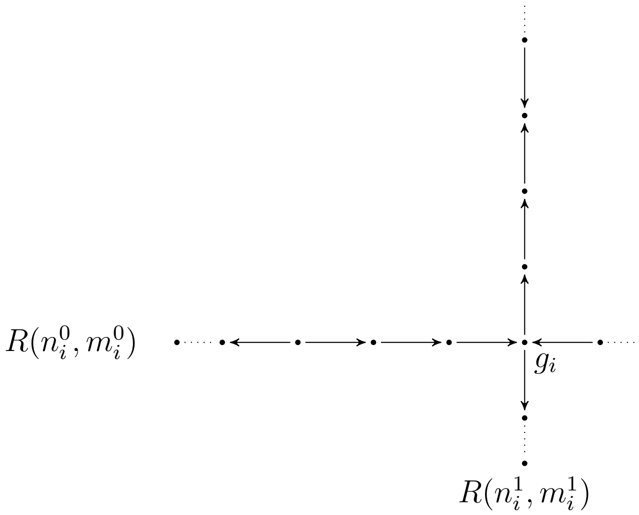

Figure 6.1: Example of an identified vertex $g_i^0 = g_i^1 = g_i$ in D for $c = 3$.

Let D be the digraph obtained from D_0 by identifying for each $i \in \mathbb{N}$ the vertices $g_i^0 \in R(n_i^0, m_i^0)$ and $g_i^1 \in R(n_i^1, m_i^1)$. We write g_i for the vertex $g_i^0 = g_i^1$ in D. By the choice of I and J, we have identified infinitely many vertices of $R(n, m)$ and $R(n', m')$ if $m \neq m'$ and none otherwise in the construction of D. Hence D contains arbitrarily many disjoint copies of R as the rays $R(0, m), R(1, m), \ldots, R(m, m)$ are disjoint for all $m \in \mathbb{N}$.

To prove that D does not contain infinitely many disjoint copies of R, we will show that for any copy R' of R in D there is $(n, m) \in I$ such that a tail of R' coincides with a tail of $R(n, m)$. Assuming this, suppose for a contradiction that there is a family $\mathcal{R}$ of infinitely many pairwise disjoint copies of R in D. Then there are $R(n, m)$ and $R(n', m')$ with $m \neq m'$ and two copies $R_0, R_1 \in \mathcal{R}$ of R in D whose tails coincide with tails of $R(n, m)$ and $R(n', m')$, respectively. But by construction, two rays $R(n, m)$ and $R(n', m')$ in D have disjoint tails if and only if $m = m'$. This contradicts the assumption that R_0 and R_1 are disjoint.

It remains to prove that for any copy R' of R in D there is $(n, m) \in I$ such that a tail of R' coincides with a tail of $R(n, m)$. Begin by fixing an arbitrary copy R' of R in D. Since c is the largest number which occurs infinitely often in the representing sequence of R, there is a tail R'' of R' whose representing sequence contains only numbers up to c and whose initial vertex is a turn of some ray $R(n^*, m^*)$ for $(n^*, m^*) \in I$. In the construction of D we only identified turns, thus each phase of any $R(n', m')$ for $(n', m') \in I$ is either completely traversed by R'' or all edges of this phase are avoided by R''.

Let $i \in \mathbb{N}$ be arbitrary. By properties (i) and (ii), any ray in D traversing both, the phase of $R(n_i^0, m_i^0)g_i$ incident with g_i and an edge of $R(n_i^1, m_i^1)$ incident with g_i, contains a phase of length $> c$ (see Figure 6.1). Clearly, the same holds for a ray traversing both, the phase of $g_i R(n_i^1, m_i^1)$ incident with g_i and an edge of $R(n_i^0, m_i^0)$ incident with g_i. Thus if R'' contains g_i as an inner vertex, exactly one of the following properties holds:

(1) both edges of R'' incident with g_i are contained in $R(n_i^0, m_i^0)$, or

(2) both edges of R'' incident with g_i are contained in $R(n_i^1, m_i^1)$, or

(3) one edge of R'' incident with g_i is contained in $R(n_i^1, m_i^1)g_i$ and one is contained in $g_i R(n_i^0, m_i^0)$.

This feature restricts in which way R'' is embedded into D. First we observe:

Claim. For any $(n, m) \in I$ and any edge $e \in E(R(n, m)) \cap E(R'')$, the rays $R(n, m)$ and R'' traverse e in the same direction.

Proof. Suppose this does not hold. Pick $(n, m) \in I$ and an edge $e \in E(R(n, m)) \cap E(R'')$ with $|R(n, m)a|$ minimal such that e contradicts this property. Let u, v be the endvertices of e such that $|R(n, m)u| < |R(n, m)v|$ applies. By minimality of e, the other edge of R'' incident with u is not contained in $R(n, m)u$. Therefore $u = g_i$ for some $i \in \mathbb{N}$, i.e. there is $(n', m') \in I$ with $m \neq m'$ such that $u \in R(n, m) \cap V(R(n', m'))$. Then (3) holds for R'' at the vertex u. As $e \in E(uR(n, m))$, the inequality $m < m'$ holds by definition of J and the other edge incident with u in R'' is contained in $R(n', m')u$. The rays R'' and $R(n', m')$ traverse this edge in opposite directions, but $|R(n', m')u| < |R(n, m)u|$ holds by property (ii) of the construction. This contradicts the minimality of $|R(n, m)e|$. ■

Now let $m \in \mathbb{N}$ be the smallest number such that there is $(n, m) \in I$ with $E(R'') \cap E(R(n, m)) \neq \emptyset$ and let e be an element of this intersection. We prove that eR'' coincides with $eR(n, m)$. Suppose not and let $(n', m') \neq (n, m) \in I$ and $u \in R(n, m) \cap V(R(n', m'))$ such that u is incident with the first edge f of eR'' not contained in $E(R(n, m))$. Property (3) applies to u. By minimality of m we have $m < m'$ and thus $f \in E(R(n', m')u)$ by (3). Therefore eR'' and $R(n', m')$ traverse the edge f in opposite directions, which contradicts the claim above. Thus R' has a tail that coincides with a tail of $R(n, m)$. This completes the proof. □

Rays with Unbounded Representing Sequence

We begin with the following lemma which we need in the proof of Theorem 6.10.

Lemma 6.9. *Let R be a ray with unbounded representing sequence. Then the tails vR and wR are non-isomorphic for all $v \neq w \in R$.*

Proof. Let R be a ray with unbounded representing sequence. Suppose that there are $v \neq w \in R$ such that w lies beyond v on R and there is an isomorphism $\varphi\colon vR \to wR$. Clearly, we have $\varphi(v) = w$ and the oriented paths $\varphi^n(v)R\varphi^{n+1}(v)$ are isomorphic for all $n \in \mathbb{N}$. Therefore, R has a tail whose representing sequence is periodic. Thus the representing sequence of R is bounded, a contradiction. □

Theorem 6.10. *All rays with unbounded representing sequence are non-ubiquitous.*

The idea of the proof of Theorem 6.10 is akin to that of Theorem 6.8: we build a counterexample recursively by identifying vertices, starting with infinitely many disjoint copies of a ray R. We ensure that the resulting digraph still contains arbitrarily but not infinitely many disjoint copies of R, utilising the structure of the representing sequence. Hence, the setup is similar but the construction differs.

Proof. Let R be a ray with unbounded representing sequence. Define

$$I := \{(n, m) \in \mathbb{N}^2 \colon n \leq m\}$$

and let $(R(n, m))_{(n,m)\in I}$ be a family of pairwise disjoint copies of R and

$$D_0 := \bigcup_{(n,m)\in I} R(n, m).$$

We will recursively define a sequence $(g_i^0, g_i^1)_{i\in\mathbb{N}^+}$ of pairwise disjoint pairs of vertices of D_0. Let D be the digraph obtained from D_0 by identifying g_i^0 with g_i^1 for all $i \in \mathbb{N}$. We will choose $(g_i^0, g_i^1)_{i\in\mathbb{N}^+}$ in such a way that D contains arbitrarily many disjoint copies of R but not infinitely many. Then D witnesses that R is non-ubiquitous.

Let

$$J := \{((n^0, m^0), (n^1, m^1)) \in I^2 \colon m^0 < m^1\}$$

and let $((n_i^0, m_i^0), (n_i^1, m_i^1))_{i\in\mathbb{N}^+}$ be a sequence which contains every element of J infinitely often. We will choose each pair (g_i^0, g_i^1) in such a way that $g_i^\varepsilon \in R(n_i^\varepsilon, m_i^\varepsilon)$ for $\varepsilon \in \{0, 1\}$. Thus we will glue together two rays $R(n, m)$ and $R(n', m')$ infinitely often if $m \neq m'$ and keep them disjoint otherwise. Hence D contains arbitrarily many disjoint copies of R as the rays $R(0, m), R(1, m), \ldots, R(m, m)$ are disjoint for all $m \in \mathbb{N}$. We still need to define $(g_i^0, g_i^1)_{i\in\mathbb{N}^+}$ so that D does not contain infinitely many disjoint copies of R.

We fix an enumeration $\{v_0, v_1, \ldots\}$ of $V(D_0)$. Denote by D_i the digraph obtained from D_0 by identifying the vertices g_ℓ^0 and g_ℓ^1 for all $\ell \leq i$ and write g_ℓ for the vertex $g_\ell^0 = g_\ell^1$. We will make sure that the following holds for all $i \in \mathbb{N}$:

(i) Let $k \leq i$ and let S be an oriented v_k–g_i path in D_i which is isomorphic to an initial segment of R. Then $|S| = |R(n_i^0, m_i^0)g_i|$ or $|S| = |R(n_i^1, m_i^1)g_i|$.

(ii) g_i is not a turn of $R(n_i^0, m_i^0)$ or $R(n_i^1, m_i^1)$ and thus g_i lies in the interior of phases M^0 of $R(n_i^0, m_i^0)$ and M^1 of $R(n_i^1, m_i^1)$. Let $t_0^\varepsilon, t_1^\varepsilon$ be the two turns which are endvertices of M^ε. Then the four numbers $d_{R(n_i^\varepsilon, m_i^\varepsilon)}(t_\delta^\varepsilon, g_i)$ for $\delta, \varepsilon \in \{0, 1\}$ are pairwise distinct. Furthermore, M^0 and M^1 do not contain any g_j with $j \neq i$.

Let us first derive from the existence of a sequence $(g_i^0, g_i^1)_{i \in \mathbb{N}}$ as above that D does not contain infinitely many disjoint copies of R. We show that for every copy R' of R in D there is a pair $(n^*, m^*) \in I$ such that a tail of R' is contained in $R(n^*, m^*)$: we know that R' traverses infinitely many vertices g_j with $j \in \mathbb{N}$ because every ray $R(m, n)$ is glued together with other rays at infinitely many vertices in D. Suppose that R' starts in v_k and let g_i be the first vertex of R' which lies in $\{g_j j \geq k\}$. Then the oriented path $R'g_i$ is a subdigraph of D_i as it contains no vertex g_j with $j > i$. Hence by (i), there is a ray $R(m^*, n^*)$ containing g_i with $|R'g_i| = |R(m^*, n^*)g_i|$. Therefore, the tails g_iR' and $g_iR(m^*, n^*)$ are isomorphic. Then by (ii), iteratively applied to g_i and all vertices of the form $g_{i'}$ on $R(n^*, m^*)$ beyond g_i, the digraphs $g_iR' = g_iR(n^*, m^*)$ coincide in D. Hence indeed, a tail of R' is contained in $R(n^*, m^*)$.

As a consequence, there is no infinite family $\mathcal{R}$ of pairwise disjoint copies of R in D. Otherwise we could find two rays $R(n, m)$, $R(n', m')$ with $m \neq m'$ which both contain a tail of a ray from $\mathcal{R}$. However, $R(n, m)$ and $R(n', m')$ were glued together at infinitely many different vertices in D and thus do not have disjoint tails.

All that remains is to define (g_i^0, g_i^1) for all $i \in \mathbb{N}$. Suppose that (g_j^0, g_j^1) is already defined for all $j < i$ so that (i) and (ii) hold for $i - 1$. We write $R^\varepsilon := R(n_i^\varepsilon, m_i^\varepsilon)$ for $\varepsilon \in \{0, 1\}$. Our task is to specify suitable vertices $g_i^\varepsilon \in R^\varepsilon$. Let x be a vertex of R^0 that lies beyond all phases of R^0 that contain any vertex of $\{g_0, \ldots, g_{i-1}, v_0, \ldots, v_i\}$. Let $\mathcal{P}$ be the set of all oriented $\{v_0, \ldots, v_i\}$–x paths in D_{i-1} which are isomorphic to initial segments of R. The following claim implies that $\mathcal{P}$ is finite:

Claim. For all $j \in \mathbb{N}$ and for all vertices $v, w \in D_j$, the set of oriented v–w paths in D_j is finite.

Proof. We use induction on j. The claim holds for $j = 0$ as D_0 is a disjoint union of rays. Now consider an arbitrary oriented v–w path P in D_j. If P does not use the vertex g_j, then P is also an oriented v–w path in D_{j-1}, and there are only finitely many such oriented paths by induction. Otherwise g_j lies on P, and P consists of an oriented v–g_j path Q concatenated with an oriented g_j–w path Q' in D_j. Since Q and Q' are also oriented paths in D_{j-1}, there are only finitely many possibilities for Q and Q' by induction. ■

We write $\mathcal{Q}$ for the subset of $\mathcal{P}$ consisting of all oriented paths Q with $|Q| \neq |R^0x|$. By Lemma 6.9, for every oriented path $Q \in \mathcal{Q}$ there is a vertex $y \in R^0$ beyond x such that QxR^0y is not isomorphic to an initial segment of R. Since $\mathcal{Q}$ is finite, we can find a vertex z^0 on R^0 such that no oriented path from $\mathcal{Q}$ can be extended to an oriented

$\{v_0, \dots, v_i\}$–z^0 path in D_{i-1} which is isomorphic to an initial segment of R. Similarly, define a vertex $z^1 \in R^1$.

Next, we show that D_i will satisfy (i) for every choice of vertices g_i^ε on R^ε that lie beyond z^ε on R^ε for $\varepsilon \in \{0, 1\}$. So suppose we have already fixed vertices g_i^ε as above and glued them together. Now consider any $k \leq i$ and any oriented v_k–g_i path S in D_i which is isomorphic to an initial segment of R. Since S is also an oriented v_k–g_i path in D_{i-1} and S must contain either the vertex z^0 or z^1, it follows from the choice of z^0 and z^1 that $|S| = |R^0 g_i|$ or $|S| = |R^1 g_i|$, which proves (i).

Finally, we further specify the choice of g_i^0 and g_i^1 so that (ii) holds. Recall that the representing sequence of R is unbounded. Therefore, we can find a phase M^0 of R^0 which is contained in $z^0 R^0$ and has length at least 3. We choose g_i^0 as an interior vertex of M^0 so that g_i^0 has different distances to both endvertices of M^0. Next, find a phase M^1 of R^1 which is contained in $z^1 R^1$ such that $|M^1| \geq 2|M^0| + 1$. Then (ii) is fulfilled for a vertex g_i^1 in M^1 which has distance $|M^0|$ to one endvertex of M^1 and hence distance $> |M^0|$ to its other endvertex. □

6.1.3 Conclusion

We show how to deduce the theorem about ubiquity of oriented rays from the above results.

Theorem 6.1. *A ray is ubiquitous if and only if it has a finite number of turns.*

Proof. The 'if' follows directly from Theorem 6.5. For the 'only if' direction let R be a ray with infinitely many turns. Then R has a representing sequence which can either be bounded or unbounded. In the former case R is non-ubiquitous by Theorem 6.8. In the latter case R is non-ubiquitous by Theorem 6.10. □

6.2 Ubiquity of Oriented Double Rays

In this section we extend the investigation of ubiquity in digraphs to the class of oriented double rays. First we prove that any oriented double ray with an odd number of turns is ubiquitous in Subsection 6.2.1. Second we show that R is non-ubiquitous if it has an even but non-zero or infinite number of turns in Subsection 6.2.2: we deal with the case that R has an even but non-zero number of turns in Subsection 6.2.2. If R has infinitely many turns, we distinguish whether R is non-periodic (Subsection 6.2.2) or periodic (Subsection 6.2.2), which is defined in Subsection 2.3.4.

6.2.1 Positive Results

In this subsection we prove Theorem 6.12, for which we utilise results from Subsection 6.1.1. In particular we need a lemma that is very similar to Theorem 6.6.

Let U be the disjoint union of two out-rays and X a set of two element sets of vertices. We say that U *is rooted in* X if there is $\{x, y\} \in X$, such that x and y are the roots of the two rays which are the components of U. If $U = \{\{x, y\}\}$, we also say that U *is rooted in* $\{x, y\}$.

Lemma 6.11. *Let U be the disjoint union of two out-rays, let D be a digraph and let X be a set of pairs of vertices of D such that there is a thick U-tribe $\mathcal{F}$ in D where all members of $\mathcal{F}$ are rooted in X. Then D contains infinitely many disjoint copies of U that are rooted in X.*

We omit the proof of Lemma 6.11 since it is very similar to the proof of Theorem 6.6.

Theorem 6.12. *Any double ray with an odd number of turns is ubiquitous.*

Proof. Let R be a double ray with an odd number of turns, D a digraph and $\mathcal{E}$ a thick R-tribe in D. We show that D contains infinitely many disjoint copies of R. Let $\hat{R}$ be a finite connected subdigraph of R that contains all turns of R. Since R has an odd number of turns, by deleting $\hat{R}$ from R the digraph R falls apart into a disjoint union U of either two in-rays or two out-rays. By Proposition 2.6, we may assume that the latter is the case.

Next, we apply Lemma 6.7 to D, R, $\hat{R}$ and $\mathcal{E}$, which yields a thick R-tribe $\mathcal{F}$ in D that is forked at $\hat{R}$. Let $\mathcal{F}'$ be the U-tribe resulting from $\mathcal{F}$ by deleting the subgraph corresponding to $\hat{R}$ from each member of $\mathcal{F}$. Let D' be the union of all members of $\mathcal{F}'$. Further, let X be the set of pairs $\{x, y\} \in [V(D)]^2$ for which there is a member of $\mathcal{F}'$ which is rooted in $\{x, y\}$. Thus any member of $\mathcal{F}'$ is rooted in X. Now we apply Lemma 6.11 to D', $\mathcal{F}'$ and X, which yields an infinite set $\mathcal{U}$ of disjoint copies of U in D' that are rooted in X.

For any $U \in \mathcal{U}$ rooted in $\{x, y\}$, there is a member R' of $\mathcal{F}$ such that x and y are the neighbours of $\hat{R}'$ in R'. We join U with $xR'y$, which again is a copy of R in D. Since $\mathcal{F}$ is forked at $\hat{R}$, doing this for any $U \in \mathcal{F}'$ gives an infinite family of disjoint copies of R in D. $\square$

6.2.2 Negative Results

In this subsection we prove the backward implication of Theorem 6.2, which is divided into three different parts. First we show that double rays with a (finite) even, non-zero number of turns are non-ubiquitous in Theorem 6.13. Then we show that non-periodic double rays with infinitely many turns are non-ubiquitous in Theorem 6.17, and lastly that periodic double rays with infinitely many turns are non-ubiquitous in Theorem 6.18.

Double Rays with an Even, Non-Zero Number of Turns

Any double ray R with an even, non-zero number of turns contains an in-ray and an out-ray. By glueing together in- and out-rays of the members of a thick R-tribe in a specific way we can show:

Theorem 6.13. *Any double ray with an even, non-zero number of turns is non-ubiquitous.*

Proof. Let R be a double ray with an even, non-zero number of turns. Let s be the first and t be the last turn of R. Since the number of turns is even, exactly one of Rs and tR is an in-ray and exactly one an out-ray. By possibly reversing the order $\leq_R$, we may assume that the former is an in-ray and the latter an out-ray. Let $p \in \mathbb{N}$ be the length of the longest finite phase of R. We will construct a digraph D containing arbitrarily many but not infinitely many disjoint copies of R.

For the construction of D, we use the auxiliary set

$$I := \{(n, m) \in \mathbb{N}^2 \colon n \leq m\}$$

ordered by the colexicographic order $\leq_{\text{col}}$. For $(n, m) \in I$, we write $(n, m)^+$ for the successor of (n, m) and $(n, m)^-$ for the predecessor (if one exists) of (n, m) under $\leq_{\text{col}}$.

Let $(R(n, m))_{(n,m)\in I}$ be a family of pairwise disjoint copies of R. For $(n, m) \in I$, write $s(n, m)$ for the first turn of $R(n, m)$ and $t(n, m)$ for the last turn of $R(n, m)$.

Next, we define two families of vertices of $R(n, m)$ for every $(n, m) \in I$: let $(v^{(i,j)}_{(n,m)})_{(i,j)\in I}$ be a family of vertices of $R(n, m)s(n, m)$ such that

- the order $\leq_{R(n,m)}$ and the order induced by $\leq_{\text{col}}$ on the superindices are reversed on $\{v^{(i,j)}_{(n,m)} \colon (i, j) \in I\}$,
- the vertices of $\{v^{(i,j)}_{(n,m)} \colon (i, j) \in I\}$ have distance at least $p + 1$ to each other and to $s(n, m)$ in $R(n, m)$.

Let $(w^{(i,j)}_{(n,m)})_{(i,j)\in I}$ be a family of vertices of $t(n, m)R(n, m)$ such that

- the order $\leq_{R(n,m)}$ and the order induced by $\leq_{\text{col}}$ on the superindices coincide on $\{w^{(i,j)}_{(n,m)} \colon (i, j) \in I\}$,
- the vertices of $\{w^{(i,j)}_{(n,m)} \colon (i, j) \in I\}$ have distance at least $p + 1$ to each other and to $t(n, m)$ in $R(n, m)$.

Let D be the digraph constructed from the disjoint union $\bigsqcup_{(n,m)\in I} R(n, m)$ by identifying the two vertices $v^{(k,\ell)}_{(i,j)}$ and $w^{(i,j)}_{(k,\ell)}$ for any $(i, j), (k, \ell) \in I$ with $j \neq \ell$ (see Figure 6.2). The digraph D contains m disjoint copies of R for every $m \in \mathbb{N}$, as the double rays $R(1, m), \dots, R(m, m)$ are disjoint.

It remains to prove that D does not contain infinitely many disjoint copies of R. In Claim 1, we investigate how an out-ray can lie in D, and in Claim 2, we investigate how an in-ray with a fixed root can lie in D. After that, we deduce from the two claims and from the fact that R contains both an out-ray and an in-ray, that D does not contain infinitely many copies of R.

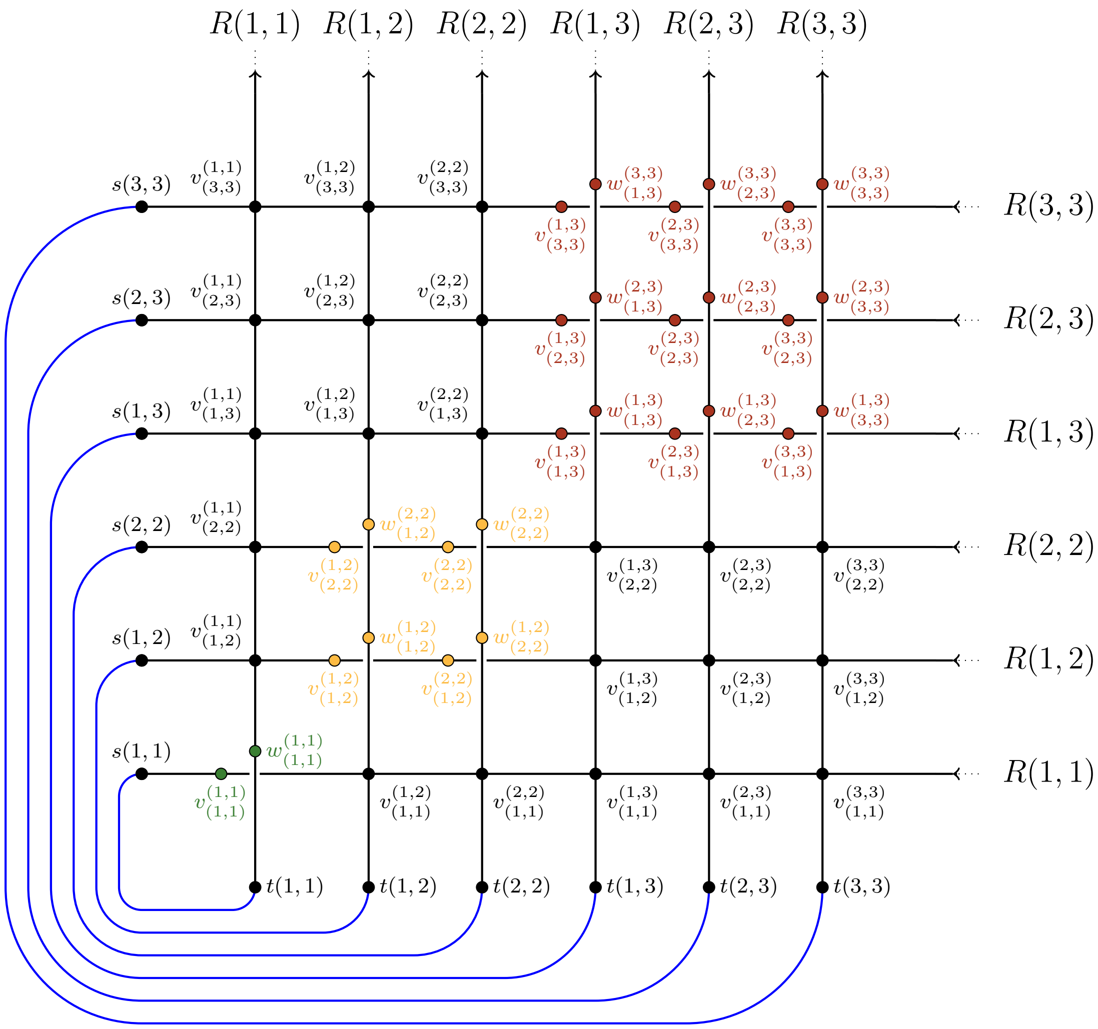

Figure 6.2: The digraph D constructed in the proof of Theorem 6.13. The horizontal lines represent parts of the in-rays $R(n,m)s(n,m)$ and the vertical lines represent parts of the out-rays $t(n,m)R(n,m)$. The blue lines connecting $s(n,m)$ and $t(n,m)$ represent the finite oriented paths $s(n,m)R(n,m)t(n,m)$ that consist of the union of all finite phases of $R(n,m)$.

Claim 1. Any out-ray in D has a tail that coincides with $xR(n,m)$ for some $(n,m) \in I$ and some $x \in V(R(n,m))$.

Proof. Let S be an arbitrary out-ray in D. If S is completely contained in some $R(n,m)$, then we are done.

Otherwise, S must contain a vertex that was identified with another in the construction. Thus there are $(i,j),(k,\ell) \in I$ such that $v^{(k,\ell)}_{(i,j)}$ is contained in S. Let $(n,m) \in I$ be $\leq_{\text{col}}$-minimal with the property that S contains a vertex $v^{(n,m)}_{(i,j)}$ for some $(i,j) \in I$. We show that $v^{(n,m)}_{(i,j)}S$ is a tail of $t(n,m)R(n,m)$.

Suppose not. Let $v := v^{(n,m)}_{(i,j)}$ if $v^{(n,m)}_{(i,j)} \notin V(t(n,m)R(n,m))$ (i.e. if $j = m$), and otherwise let v be the last vertex of $v^{(n,m)}_{(i,j)}S$ such that $v^{(n,m)}_{(i,j)}Sv$ is contained in the out-ray $t(n,m)R(n,m)$. In either case, there is $(k,\ell) \in I$ such that $v = v^{(n,m)}_{(k,\ell)}$ since in the latter case v had to be identified with another vertex in the construction. Then the first edge of vS is contained in $R(k,\ell)$. Since S is an out-ray, it also contains $v_{(k,\ell)}^{(n,m)^-}$ if $(n,m) \neq (1,1)$, or $s(k,\ell)$ if $(n,m) = (1,1)$ (see Figure 6.2). This contradicts either the minimality of (n,m) or the fact that $s(k,\ell)$ is a turn, respectively. ■

Claim 2. For any $(n',m'),(n,m) \in I$, any in-ray in D with root $s(n',m')$ contains a vertex of $\{v_{(i,j)}^{(n,m)} \colon (i,j) \in I\}$.

Proof. Let $(n,m),(n',m') \in I$ and let S be an arbitrary in-ray in D with root $s(n',m')$. As no vertex of $s(n',m')S(n',m')t(n',m')$ has been identified with other vertices and $t(n',m')$ is a turn, S contains the vertex $v_{(n',m')}^{(1,1)}$.

We consider the set $\mathcal{X} := \{v_{(i,j)}^{(k,\ell)} \in V(S) \colon (i,j),(k,\ell) \in I, (k,\ell) \leq_{\mathrm{col}} (n,m)\}$ and let $(k,\ell) \in I$ be $\leq_{\mathrm{col}}$-maximal with the property that an element of $\mathcal{X}$ has superindex (k,ℓ). We show $(k,\ell) = (n,m)$, which implies the statement of Claim 2.

Suppose for a contradiction that $(k,\ell) <_{\mathrm{col}} (n,m)$. Let $(i,j) \in I$ be $\leq_{\mathrm{col}}$-minimal with the property that $v_{(i,j)}^{(k,\ell)}$ is an element of $\mathcal{X}$. The first edge of $v_{(i,j)}^{(k,\ell)}S$ lies either in $R(i,j)s(i,j)$ or in $t(k,\ell)R(k,\ell)$. In the first case, the in-ray S must also contain $v_{(i,j)}^{(k,\ell)^+}$ (see Figure 6.2), contradicting the maximality of (k,ℓ). In the second case, S also contains $t(k,\ell)$ if $v_{(i,j)}^{(k,\ell)}$ is the first vertex of $t(k,\ell)R(k,\ell)$ that has been identified, or S contains a vertex $v_{(i',j')}^{(k,\ell)}$ for $(i',j') <_{\mathrm{col}} (i,j)$ otherwise (see Figure 6.2). This contradicts either the fact that $t(k,\ell)$ is a turn or the minimality of (i,j), respectively. ■

Let $\hat{R}$ be an arbitrary copy of R in D. By Claim 1, there is $(n,m) \in I$ and $x \in V(R(n,m))$ such that an out-ray in $\hat{R}$ coincides with $xR(n,m)$. To prove that D cannot contain infinitely many disjoint copies of R, it suffices to show that any copy of R in D has an in-ray that starts in some $s(n',m')$. Then by Claim 2, every copy of R in D contains a vertex of $\{v_{(i,j)}^{(n,m)} \colon (i,j) \in I\}$, contradicting that the set $\{v_{(i,j)}^{(n,m)} \colon (i,j) \in I\} \setminus V(xR(n,m))$ is finite.

Let $\tilde{R}$ be any copy of R in D and let $\tilde{R}y$ be the unique phase of $\tilde{R}$ that forms an in-ray. By construction of D, any phase of length at most p of $\tilde{R}$ is contained in $s(i,j)R(i,j)t(i,j)$ for some $(i,j) \in I$, which immediately yields $y \in \{s(n',m'), t(n',m')\}$ for some $(n',m') \in I$. Since $\tilde{R}y$ is an in-ray, we must have $y = s(n',m')$ as desired. This completes the proof. □

Non-Periodic Double Rays with Infinitely Many Turns

In Lemma 6.14, Lemma 6.15 and Theorem 6.16 we investigate symmetry properties of non-periodic double rays and show that any such double ray R has a tail which is isomorphic to only very specific other tails of R. In Theorem 6.17, we use this result to reduce the non-ubiquity of non-periodic double rays with infinitely many turns to the non-ubiquity of rays with infinitely many turns (Theorem 6.1).

Lemma 6.14. *For every non-periodic double ray R there is $v^* \in V(R)$ such that Rv^* is non-periodic.*

Proof. Suppose for a contradiction that Rw is periodic for all $w \in V(R)$. Let $p \in \mathbb{N}$ be minimal among the periodicities of Rw for all $w \in V(R)$ and let $v \in V(R)$ such that Rv has periodicity p. We show that Rw has periodicity p for any $w \in V(R)$. This then implies that R is periodic with periodicity p, a contradiction.

For $w >_R v$, let f be any non-trivial endomorphism of Rw, which exists since Rw is periodic. By concatenating f with itself multiple times if necessary, we obtain an endomorphism of Rw whose image is contained in Rv. Thus it remains to prove that Ru has periodicity p for any $u <_R v$. Since Ru is a tail of Rv, Ru has periodicity at most p, and by minimality of p, Ru has periodicity exactly p. □

Lemma 6.15. *For every non-periodic double ray R there is $v^* \in V(R)$ such that for all $v \geq_R v^*$ and all $w \in V(R)$:*

$$Rv \cong Rw \Rightarrow v = w\,.$$

Proof. We choose v^* such that Rv^* is non-periodic according to Lemma 6.14. Let $v \geq_R v^*$ and $w \in V(R)$ such that $Rv \cong Rw$. If $w <_R v$, we can restrict the isomorphism $Rv \to Rw$ to a non-trivial endomorphism of Rv^*, contradicting that Rv^* is non-periodic. If $w >_R v$, we can restrict the isomorphism $Rw \to Rv$ to a non-trivial endomorphism of Rv^*, which again is a contradiction. Thus $v = w$ holds. □

Theorem 6.16. *For every non-periodic double ray R there is $\hat{v} \in V(R)$ such that for all $w \in V(R)$:*

(1) $R\hat{v} \not\leq \hat{v}R$,

(2) $R\hat{v} \cong Rw \Rightarrow \hat{v}R \cong wR$, *and*

(3) $R\hat{v} \cong wR \Rightarrow \hat{v}R \cong Rw$.

Proof. Let v^* be as in Lemma 6.15.

Case 1: There is $v \geq_R v^*$ such that $Rv \ncong wR$ for all $w \in V(R)$. In this case we set $\hat{v} := v$, which directly implies (1) and (3). For (2), since $\hat{v} \geq v^*$ and v^* was picked with the property of Lemma 6.15, $R\hat{v} \cong Rw$ implies $\hat{v} = w$ and thus $\hat{v}R = wR$.

Case 2: For all $v \geq_R v^*$ there is $\alpha(v) \in V(R)$ such that $Rv \cong \alpha(v)R$.

Claim 1. For every $v \geq_R v^*$, the vertex $\alpha(v)$ is unique.

Proof. Suppose for a contradiction that there are $\alpha(v) <_R \alpha(v)' \in V(R)$ such that $\alpha(v)R \cong Rv \cong \alpha(v)'R$. Since $\alpha(v)'R$ is a proper tail of $\alpha(v)R$, it follows that $Rv \cong \alpha(v)'R$ is isomorphic to a proper tail of $Rv \cong \alpha(v)R$. Thus there exists a non-trivial endomorphism of Rv, which contradicts that v^* has the property of Lemma 6.15. ■

Claim 2. There is $\hat{v} \geq_R v^*$ such that $\alpha(\hat{v}) <_R \hat{v}$.

Proof. Let v and v' be any vertices of R with $v^* \leq_R v <_R v'$. Since Rv is a proper tail of Rv', it follows that $\alpha(v)R \cong Rv$ is isomorphic to a proper tail of $\alpha(v')R \cong Rv'$. Thus there is a vertex $w >_R \alpha(v')$ such that $wR \cong \alpha(v)R$. Then $w = \alpha(v)$ by Claim 1. In conclusion, we have established that $\alpha(v) >_R \alpha(v')$ whenever $v <_R v'$. Therefore, it is possible to pick a vertex $\hat{v}$ which is sufficiently large with respect to $\leq_R$, so that $\alpha(\hat{v}) <_R \hat{v}$. ■

We show that any vertex $\hat{v}$ as in Claim 2 satisfies properties (1) to (3). For (1), assume that $R\hat{v} \leq \hat{v}R$. This means that there is $v' \geq_R \hat{v}$ with $R\hat{v} \cong v'R$. A contradiction since by Claim 1 $\alpha(\hat{v})$ is unique, but $\alpha(\hat{v}) <_R \hat{v} \leq_R v'$ by choice of $\hat{v}$. For (2), recall that $\hat{v} \geq_R v^*$. Thus, Lemma 6.15 implies that the only vertex $w \in V(R)$ with $R\hat{v} \cong Rw$ is $\hat{v}$. This proves (2) since clearly $\hat{v}R \cong \hat{v}R$. For (3), as $\alpha(\hat{v})$ is unique by Claim 1, it is enough to prove $\hat{v}R \cong R\alpha(\hat{v})$. Let $\psi\colon R\hat{v} \to \alpha(\hat{v})R$ be an isomorphism, which exists by choice of $\alpha(\hat{v})$. Then ψ maps $\alpha(\hat{v})$ to $\hat{v}$, since the distance of $\alpha(\hat{v}) \in V(R\hat{v})$ to the root of $R\hat{v}$ and the distance of $\hat{v} \in V(\alpha(\hat{v})R)$ to the root of $\alpha(\hat{v})R$ are the same. Thus ψ can be restricted to an isomorphism $R\alpha(\hat{v}) \to \hat{v}R$.

This completes the proof. □

Now we combine Theorem 6.1 and Theorem 6.16 to prove:

Theorem 6.17. *Any non-periodic double ray R with infinitely many turns is non-ubiquitous.*

Proof. Let R be any such double ray. We construct a digraph D that contains arbitrarily many but not infinitely many copies of R. Without loss of generality, for every $v \in V(R)$ the ray vR contains infinitely many turns (otherwise reverse the order $\leq_R$). Let $\hat{v} \in V(R)$ be as in Theorem 6.16. As $\hat{v}R$ contains infinitely many turns, there is a digraph D' containing arbitrarily many but not infinitely many disjoint copies of $\hat{v}R$ by Theorem 6.1. We construct D from D' and a family $(S_x)_{x \in V(D')}$ of disjoint copies of $R\hat{v}$ by identifying the root of S_x with x for each $x \in V(D')$.

By construction, D contains arbitrarily many disjoint copies of R. We have to show that D does not contain infinitely many disjoint copies of R, which implies the theorem. It suffices to prove that each copy of R in D has a tail isomorphic to $\hat{v}R$ that is contained in the subdigraph D' of D. Then D cannot contain infinitely many disjoint copies of R since D' does not contain infinitely many disjoint copies of $\hat{v}R$.

Let $\tilde{R}$ be any copy of R in D. If $\tilde{R}$ is completely contained in D', we are done. Thus we can suppose that there is $x \in V(D')$ and $w \in V(\tilde{R})$ such that either $S_x = \tilde{R}w$ or $S_x = w\tilde{R}$.

In the former case, we have $R\hat{v} \cong S_x = \tilde{R}w$ and thus $\hat{v}R \cong w\tilde{R}$ by Theorem 6.16 (2). It follows from (1) that $R\hat{v} \not\leq \hat{v}R \cong w\tilde{R}$. Hence $w\tilde{R}$ cannot have a tail in any S_y for

$y \in V(D')$. Thus $w\tilde{R}$ is the desired tail of $\tilde{R}$ which is isomorphic to $\hat{v}R$ and contained in D'.

Similarly, in the latter case, we have $R\hat{v} \cong S_x = w\tilde{R}$ and thus $\hat{v}R \cong \tilde{R}w$ by (3). It follows from (1) that $R\hat{v} \not\leq \hat{v}R \cong \tilde{R}w$. Hence $\tilde{R}w$ cannot have a tail in any S_y for $y \in V(D')$ and $\tilde{R}w$ is the desired tail of $\tilde{R}$. $\square$

Periodic Double Rays with Infinitely Many Turns

Let R be a periodic double ray with infinitely many turns and let $\hat{R}, \tilde{R}$ be disjoint copies of R. By periodicity of R, one can show that identifying a turn of $\hat{R}$ of out-degree 2 and a turn of $\tilde{R}$ of in-degree 2 results in a digraph in which a copy of R has to be completely contained in either $\hat{R}$ or $\tilde{R}$. We use this fact to prove:

Theorem 6.18. *Any periodic double ray with infinitely many turns is non-ubiquitous.*

Proof. Let R be any periodic double ray with infinitely many turns and denote the periodicity of R by $p \in \mathbb{N}$. We will construct a digraph D containing arbitrarily many but not infinitely many copies of R.

We set
$$I := \{(n, m) \in \mathbb{N}^2 \colon n \leq m\}$$
and let $(R(n, m))_{(n,m)\in I}$ be a family of pairwise disjoint copies of R. Let D be the digraph constructed from the disjoint union $\bigsqcup_{(n,m)\in I} R(n, m)$ by identifying pairwise disjoint pairs of vertices such that for any $(n, m), (n', m') \in I$:

- **(i)** no vertices of $R(n, m)$ and $R(n', m')$ have been identified with each other if $m = m'$,
- **(ii)** exactly one vertex of $R(n, m)$ and exactly one vertex of $R(n', m')$ have been identified if $m \neq m'$,
- **(iii)** if $v \in R(n, m)$ and $w \in R(n', m')$ have been identified with each other, then either the out-degree of v in $R(n, m)$ is 2 and the out-degree of v in $R(n', m')$ is 0 or vice versa, and
- **(iv)** two vertices $v \neq w \in R(n, m)$ that have been identified with other vertices have distance at least p in $R(n, m)$.

A graph D satisfying (i) to (iv) can be constructed by enumerating all unordered pairs $\{R(n, m), R(n', m')\}$ of double rays with $m \neq m'$ and recursively identifying suitable turns of the two rays in each pair.

The digraph D contains arbitrarily many disjoint copies of R, as the double rays $R(1, m), \ldots, R(m, m)$ are disjoint for any $m \in \mathbb{N}$ by (i). To prove that D does not contain infinitely many disjoint copies of R, it suffices to show that any copy of R in D is of the form $R(n, m)$ for some $(n, m) \in I$: Then any infinite family of disjoint copies of R in D would contain two rays $R(n,m)$, $R(n',m')$ with $m \neq m'$ by definition of I. However, $R(n,m)$ and $R(n',m')$ are not disjoint in D by (ii).

Suppose for a contradiction that there is a copy $\hat{R}$ of R in D that is not contained in some $R(n,m)$ for $(n,m) \in I$. Then there are $(n,m) \neq (n',m') \in I$ and $v \in V(\hat{R})$ such that one edge of $\hat{R}$ incident with v is contained in $R(n,m)$ and the other edge of $\hat{R}$ incident with v is contained in $R(n',m')$. We assume without loss of generality that the first edge of $v\hat{R}$ is contained in $vR(n,m)$ (and not in $R(n,m)v$, $R(n',m')v$ or $vR(n',m')$).

Since $R(n,m)$ has periodicity p, the edge e of $R(n,m)$ preceding v has the same orientation as the p-th edge e' of $R(n,m)$ succeeding e. Similarly, the edge f of $\hat{R}$ preceding v has the same orientation as the p-th edge f' of $\hat{R}$ succeeding f. As no vertices of distance at most $p-1$ to v in D other than v were identified by (iv), the first p edges of $v\hat{R}$ coincide with the first p edges of $vR(n,m)$ and in particular we have $e' = f'$. Hence the edges e, f either both point towards v or both point away from v, contradicting (iii) since $e \in E(R(n',m'))$ and $f \in E(R(n,m))$. □

6.2.3 Conclusion

Finally, we show how to deduce the theorem about ubiquity of oriented double rays from the results of this section.

Theorem 6.2. *A double ray with at least one turn is ubiquitous if and only if it has a (finite) odd number of turns.*

Proof. The 'if' direction is immediate from Theorem 6.12. For the 'only if' direction let R be a double ray with at least one turn, but not an odd number of turns. If R has an even number of turns, then we are done by Theorem 6.13. Otherwise, R has infinitely many turns and R is either non-periodic or periodic. In the former case R is non-ubiquitous by Theorem 6.17. In the latter case R is non-ubiquitous by Theorem 6.18. □

Chapter 7

Flames

The starting point of our investigation is the following theorem of Lovász.

Theorem 7.1 (Lovász, consequence of [Lov73, Theorem 2])**.** *Let D be a digraph with $r \in V(D)$. Then there is a spanning subdigraph L of D in which for every $v \in V(D) - r$ the following three quantities are equal: the local connectivities $\kappa_D(r, v)$ and $\kappa_L(r, v)$, and the in-degree of v in L.*

The optimality of an L satisfying the min-max criteria from Lovász' Theorem may instead also be captured by the following structural characterisation: for every $v \in V - r$ there is a system $\mathcal{P}_v$ of internally disjoint r–v paths in L covering all incoming edges of v in L such that one can choose from each $P \in \mathcal{P}_v$ either an edge or an internal vertex in such a way that the resulting set meets every r–v path of D. The positive result for countably infinite digraphs based on this structural infinite generalisation was proved by Joó in [Joó19b, Theorem 1.2]. He later proved in [EGJ21, Theorem 1.3] together with Erde and Gollin a strengthening of this result stating that every vertex-flame of a countable rooted digraph can be extended into a large one.

The main result of this chapter leaves countable digraphs behind and handles the smallest uncountable case:

Theorem 7.2. *Every rooted digraph of size at most $\aleph_1$ admits a large vertex-flame.*

As in the case of the infinite version of Menger's Theorem (i.e. Theorem 2.3), the construction and the necessary arguments get significantly more complex when the digraph in question is uncountable. Although several of our tools can be used to approach the problem for arbitrarily large digraphs, our proof relies strongly on the fact that there is an enumeration of the vertex set in which the proper initial segments are countable. We expect that Theorem 7.2 remains true without any size restriction on the digraph but we feel that, despite solving the smallest uncountable case, a complete understanding of the problem is still far away.

Conjecture 7.3. Every rooted digraph admits a large vertex-flame.

The following edge-variant of the problem is wide open even in the countable case, but known for finite digraphs even in a fractional variant with edge-capacities and 'flow-connectivity' (see [Joó21, Theorem 4.1]).

Question 7.4. Let D be a countable digraph with $r \in V(D)$. Is it always possible to find a spanning subdigraph L of D such that for every $v \in V(D) - r$ there is a system $\mathcal{P}_v$ of *edge-disjoint* r–v paths in L covering all incoming edges of v in L such that one can choose exactly one *edge* from each $P \in \mathcal{P}_v$ in such a way the resulting edge set is an rv-cut in D?

We first prove a number of lemmas in Section 7.1 that we will use in Section 7.2, which is the proof of Theorem 7.2.

7.1 Preparations

All *digraphs* D in this chapter are simple and have no incoming edges to their 'root vertex' r whenever they have such a root.

7.1.1 Elementary Submodels

Elementary submodels are defined for first order structures in logic but for simplicity let us talk only about the special case we use. An elementary submodel of a set A is an $M \subseteq A$ such that for every first order formula $\varphi(x_1, \dots, x_n)$ in the language of set theory (with free variables $x_1, \dots, x_n$) and for every $a_1, \dots, a_n \in M$, the statement $\varphi(a_1, \dots, a_n)$ is true in the first order structure $(A, \in|_{A \times A})$ if and only if it is true in $(M, \in|_{M \times M})$. Elementary submodels provide a powerful method in topology, infinite combinatorics and in other fields to cut up uncountable structures into smaller 'well-behaved' pieces. In these applications A usually consists of the sets whose transitive closure is of cardinality less than λ (denoted by $H(\lambda)$), where λ is chosen in such a way that $H(\lambda)$ contains all sets that are relevant in the proof. By *elementary submodel* we always mean an elementary submodel of $H(\lambda)$ for a large enough λ. For a detailed introduction for elementary submodel techniques and their applications in infinite combinatorics we refer to [Sou11]. For an elementary submodel M and digraph $D = (V, E)$ we let $D \cap M := (V \cap M, E \cap M)$.

7.1.2 A Reduction to Quasi-Vertex-Flames

Lemma 7.5 ([Joó19b, Lemma 2.1]). *For every rooted digraph D, there is a quasi-vertex-flame $F \subseteq D$ such that whenever an $L \subseteq F$ is F-vertex-large it is D-vertex-large as well.*

Corollary 7.6. *One may assume without loss of generality in the proof of Theorem 7.2 that D is a quasi-vertex-flame.*

7.1.3 Linkability of Finite Sets from r

Lemma 7.7 ([Joó19b, Claim 3.14]). *If a finite $U \subseteq V - r$ is linked from r in D and L is large, then U is linked from r in L as well.*

Corollary 7.8 ([Joó19b, Lemma 2.3]). *If D is a quasi-vertex-flame and L is large, then L is also a quasi-vertex-flame.*

7.1.4 Variants of Pym's Theorem

We need some derivatives of the following theorem due to Pym [Pym69].

Theorem 7.9 (Pym's Theorem [Pym69, The Linkage Theorem]). *Let $X, Y \subseteq V$, furthermore, let $\mathcal{P}$ and $\mathcal{Q}$ be disjoint systems of X–Y paths. Then there is a system $\mathcal{R}$ of disjoint X–Y paths such that $V^-(\mathcal{R}) \supseteq V^-(\mathcal{P})$ and $V^+(\mathcal{R}) \supseteq V^+(\mathcal{Q})$, moreover, each $R \in \mathcal{R}$ is either in $\mathcal{P} \cup \mathcal{Q}$ or there are $P \in \mathcal{P}$, $Q \in \mathcal{Q}$ and $v_R \in V(P) \cap V(Q)$ such that $R = Pv_RQ$.*

Corollary 7.10. *Suppose that $\mathcal{P}$ links S to v and $\mathcal{Q}$ is a v-infan with $V(\mathcal{Q}) \cap S = V^-(\mathcal{Q})$. Then there is a v-infan $\mathcal{R}$ with $V^-(\mathcal{R}) = S$ covering $E^+(\mathcal{Q})$, furthermore, each $R \in \mathcal{R}$ is either in $\mathcal{P} \cup \mathcal{Q}$ or there are $P \in \mathcal{P}$, $Q \in \mathcal{Q}$ and $v_R \in V(P) \cap V(Q)$ such that $R = Pv_RQ$.*

We need one version of the theorem in which $r \in S$ and more than one path in $\mathcal{P}$ and in $\mathcal{R}$ may start in r. This variant can be reduced to Corollary 7.10 by splitting r into a vertex set $V_r := \{r_e \colon e \in \mathsf{out}_D(r)\}$ where r_e inherits the single outgoing edge e of r.

Corollary 7.11. *Suppose that $\mathcal{P}$ is a system of S–v paths with $v \notin S$ such that $V(P_0) \cap V(P_1) - v \subseteq \{r\}$ for every $P_0 \neq P_1$ from $\mathcal{P}$ and suppose $\mathcal{Q}$ is a v-infan with $V(\mathcal{Q}) \cap S = V^-(\mathcal{Q})$. Then there is a system $\mathcal{R}$ of S–v paths with $V(R_0) \cap V(R_1) - v \subseteq \{r\}$ for every $R_0 \neq R_1$ from $\mathcal{R}$ covering $V^-(\mathcal{P}) \cup E^+(\mathcal{Q})$, furthermore, each $R \in \mathcal{R}$ is either in $\mathcal{P} \cup \mathcal{Q}$ or there are $P \in \mathcal{P}$, $Q \in \mathcal{Q}$ and $v_R \in V(P) \cap V(Q)$ such that $R = Pv_RQ$.*

We also deduce that for every $I \in \mathcal{G}_D(v)$ we find a path-system witnessing this in $\mathfrak{P}_D(v, S)$.

Corollary 7.12. *Let $v \in V - r$ be given and let $S \in \mathfrak{S}_D(v)$ and $I \in \mathcal{G}_D(v)$. There is $\mathcal{R} \in \mathfrak{P}_D(v, S)$ with $I - rv \subseteq E^+(\mathcal{R})$.*

Proof. Let $\mathcal{P} \in \mathfrak{P}_D(v, S)$ and let $\mathcal{Q}$ be a witness for $I \in \mathcal{G}_D(v)$. We define $\mathcal{Q}'$ and $\mathcal{P}'$ to be the set of terminal segments of the paths in $\mathcal{Q}$ and $\mathcal{P}$ from the last intersection with S respectively. We then apply Corollary 7.10 with $\mathcal{P}'$ and $\mathcal{Q}'$ and yield an (S, v)-path-system $\mathcal{R}'$ with $V^-()\mathcal{R}' = S$ and $I \subseteq E^+(\mathcal{R}')$. Lastly, we extend the paths in $\mathcal{R}'$ backwards to r via the initial segments of the paths in $\mathcal{P}$ up to S to obtain a path-system $\mathcal{R}$ as desired. □

7.1.5 Preservation of the Vertex-Flame Property

Lemma 7.13 ([EGJ21, Lemma 4.10]). *Suppose that there is $I \in \mathcal{G}_D(w)$ such that $(I + f) \in \mathcal{G}_D(w)$ for every $f \in \mathsf{in}_D(w) \setminus I$. Assume that there is $uv \in E(D)$ with $u \neq r$, $v \neq w$ in such a way that $I \notin \mathcal{G}_{D-uv}(w)$. Then there exists a set $S \subseteq V - r$ with $v \in S$ which is linked from r by a path-system $\mathcal{P}$, such that S separates $N^-_D(v) - u$ from r. In particular, uv is the last edge of some $P_{uv} \in \mathcal{P}$.*

We are interested only in the special case where $I = \mathsf{in}_D(w)$:

Corollary 7.14. *Suppose that $\mathsf{in}_D(w) \in \mathcal{G}_D(w)$ and there is $uv \in E(D)$ with $u \neq r$, $v \neq w$ for which $\mathsf{in}_D(w) \notin \mathcal{G}_{D-uv}(w)$. Then there exists a set $S \subseteq V - r$ with $v \in S$ which is linked from r by a path-system $\mathcal{P}$, such that S separates $N^-_D(v) - u$ from r. In particular, uv is the last edge of some $P_{uv} \in \mathcal{P}$.*

A digraph D has the *G-quasi-vertex-flame property* for some $G \subseteq D$ at $v \in V - r$ if $\mathcal{G}_D(v)$ contains every $I \subseteq \mathsf{in}_D(v)$ for which $I \setminus \mathsf{in}_G(v)$ is finite. (For a single v only the edges $\mathsf{in}_G(v)$ are relevant for the G-quasi-vertex-flame property at v.) We need a statement that closely resembles [EGJ21, Lemma 2.10]. In fact, even though the statement of Lemma 7.15 appears stronger on first glance since we drop the condition that D should have the G-quasi-flame property at every vertex, Lemma 7.15 is obtained with the proof in [EGJ21, Lemma 2.10].

Lemma 7.15 ([EGJ21, Lemma 2.10]). *Assume that $D = (V, E)$ is a countable r-rooted digraph, $v \in V - r$ and $G \subseteq D$. Then there is an $I^* \in \mathcal{G}_D(v)$ such that $D \restriction_v I^*$ has the G-quasi-vertex-flame property for every $u \in V - r$ for which D has this property.*

On the one hand, we are interested only in cases where G has a certain special form. On the other hand we want to weaken the assumption that D is countable. The following variant will be suitable for our purpose:

Corollary 7.16. *Assume that $D = (V, E)$ is an r-rooted digraph, $v \in V - r$ and $W \subseteq V - v - r$ is a countable set such that D has the vertex-flame property at every $w \in W$. Then there is an $I^* \in \mathcal{G}_D(v)$ such that $D \restriction_v I^*$ also has the vertex-flame property for every $w \in W$.*

Proof. Let $G := (V, \bigcup_{w \in W} \mathsf{in}_D(w))$. We pick a countable elementary submodel M with $v, r, D, G, W \in M$ and apply Lemma 7.15 with $D \cap M$, v and $G \cap M$. This yields an $I^* \in \mathcal{G}_{D \cap M}(v)$ such that $D \cap M \restriction_v I^*$ has the flame property at every $w \in W$.

We shall show that $D \restriction_v I^*$ has the flame property at every $w \in W$. Let $w \in W$ be fixed. Let us pick a $\mathcal{P}$ witnessing the flame property of $D \cap M \restriction_v I^*$ at w and a $\mathcal{Q} \in M$ showing the flame property of D at w. We claim that $\mathcal{R} := \mathcal{P} \cup (\mathcal{Q} \setminus M)$ witnesses the flame property of $D \restriction_v I^*$ at w. Indeed, for any path $Q \in \mathcal{Q}$ using a vertex $u \in (V - r - w) \cap M$, we must have $Q \in M$ because Q is definable from u and $\mathcal{Q}$. Since Q is finite, $Q \subseteq D \cap M$ follows. This shows that $\mathcal{R}$ is an internally disjoint path-system. Furthermore, if $e \in \mathsf{in}_D(w)$, then e is the last edge of a path in $\mathcal{P}$ or in $\mathcal{Q} \setminus M$ depending on if $e \in M$. This completes the proof of the corollary. $\square$

7.1.6 Preservation of Largeness

We introduce some terminology that we are going to use only locally to prove a number of lemmas applying previous results. For a vertex set $X \subseteq V - r$, the *entrance* of X with respect to D is

$$\mathsf{ent}_D(X) := \{v \in X \colon \text{there is } uv \in E(D) \text{ with } u \notin X\},$$

and $\mathsf{int}_D(X)$ stands for its *interior* $X \setminus \mathsf{ent}_D(X)$. A set $B \subseteq V - r$ is a *v-bubble* with respect to D if there exists a v-infan $\mathcal{P} = \{P_u : u \in \mathsf{ent}_D(B) - v\}$ in $D[B]$ where P_u starts at u. Let us denote the set of v-bubbles in D by $\mathsf{bubb}_D(v)$. Clearly $\{v\} \in \mathsf{bubb}_D(v)$ since either the trivial path consisting of the single vertex v or the empty set is a witness for it depending on if $v \in \mathsf{ent}_D(\{v\})$.

Lemma 7.17 (Bubble uniting lemma, [Joó19b, Lemma 3.5]). *Let α be an ordinal number. Suppose that $\langle B_\beta : \beta < \alpha \rangle$ is a sequence where $B_\beta \in \mathsf{bubb}_D(v_\beta)$ for some $v_\beta \in V - r$. Let us denote $\bigcup_{\gamma<\beta} B_\gamma$ by $B_{<\beta}$. If for each $\beta < \alpha$ either $v_\beta = v_0$ or $v_\beta \in \mathsf{int}_D(B_{<\beta})$, then $B_{<\alpha} \in \mathsf{bubb}_D(v_0)$.*

Note that for $S \in \mathfrak{S}_D(v)$, the set $B_{D,S,v}$ of vertices that are separated from r by S in $D - rv$ form a v-bubble with $\mathsf{ent}_{D-rv}(B_{D,S,v}) = S$ such that $N^-_{D-rv}(v) \subseteq B_{D,S,v}$.

Corollary 7.18. *There is a $\subseteq$-largest v-bubble $B_{D,v}$ in D for every $v \in V - r$ and it contains $N^-_{D-rv}(v)$.*

Lemma 7.19 ([Joó19b, Lemma 3.10]). *A spanning subdigraph L of D is large if and only if $u \in B_{L,v}$ for every $uv \in E(D) \setminus E(L)$. Furthermore, if L is large and $v \in V - r$, then $\mathsf{ent}_{L-rv}(B_{L,v}) = \mathsf{ent}_{D-rv}(B_{L,v}) \in \mathfrak{S}_D(v)$.*

Note that this also shows $S_{L,v} = \mathsf{ent}_{D-rv}(B_{L,v})$.

Corollary 7.20. [Joó19b, Lemma 2.2] *Assume that $\mathsf{out}_D(r) \subseteq L \subseteq D$ such that for every $v \in V - r$ with $\mathsf{in}_L(v) \subsetneq \mathsf{in}_D(v)$ there is a $\mathcal{P} \in \mathfrak{P}_D(v)$ that lies in L. Then L is large.*

We call a vertex set $A \subseteq V - r$ an *anti-bubble* in D if $\mathsf{ent}_D(A)$ is linked from r in D. Note that the family of anti-bubbles are closed under arbitrarily large intersection. For $S \in \mathfrak{S}_D(v)$ the set $B_{D,S,v}$ is not just a v-bubble but also an anti-bubble that contains $N^-_{D-rv}(v)$. Moreover, if X is a v-bubble and also an anti-bubble and contains $N^-_{D-rv}(v)$, then $\mathsf{ent}_{D-rv}(X) \in \mathfrak{S}_D(v)$. Let $A_{D,v}$ be the intersection of all anti-bubbles in D containing $\{v\} \cup N^-_{D-rv}(v)$.

Proposition 7.21. *For every $v \in V - r$, $A_{D,v}$ is a v-bubble in D, furthermore, we have $T_{D,v} = \mathsf{ent}_{D-rv}(A_{D,v}) \in \mathfrak{S}_D(v)$.*

Proof. We apply Theorem 2.3 (Aharoni-Berger) in $D[A_{D,v}]$ with $\mathsf{ent}_{D-rv}(A_{D,v})$ and $N_D^-(v)$. If the resulting separation S is not $\mathsf{ent}_{D-rv}(A_{D,v})$ itself, then $B_{D,S,v} \subsetneq A_{D,v}$ is an anti-bubble containing $\{v\} \cup N_{D-rv}^-(v)$, which contradicts the minimality of $A_{D,v}$. Thus $S = \mathsf{ent}_{D-rv}(A_{D,v})$, and hence the path-system given by Theorem 2.3 witnesses that $A_{D,v}$ is a v-bubble. Since $A_{D,v}$ is an anti-bubble as well and contains $N_{D-rv}^-(v)$, we have $\mathsf{ent}_{D-rv}(A_{D,v}) \in \mathfrak{S}_D(v)$. It now follows from the definition that $\mathsf{ent}_{D-rv}(A_{D,v}) = T_{D,v}$ □

Before we proceed, we need another lemma. One of the standard proofs of Menger's Theorem is based on the so called Augmenting Walk Lemma. For a given disjoint system $\mathcal{P}$ of X–Y paths it either provides a bigger such system or an X–Y-separation consisting of exactly one vertex from each of the paths in $\mathcal{P}$. The infinite generalisation of this lemma (see [Die17, Lemmas 3.3.2 and 3.3.3]) was an important tool in the proof of the infinite version of Menger's Theorem (i.e. Theorem 2.3). There are several variants of the lemma depending on whether the paths are edge-disjoint or vertex-disjoint, whether we consider graphs or digraphs etc. but the proofs of these variants are essentially the same. We make use of the following variant:

Lemma 7.22 (Augmenting walk)**.** *Assume that $D = (V, E)$ is a digraph, $X \subseteq V$ and $v \in V \setminus X$. Let $\mathcal{P}$ be a v-infan with $V(\mathcal{P}) \cap X = V^-(\mathcal{P})$. Then there is either an S that separates v from X consisting of a unique $v_P \in V(P) - v$ for every $P \in \mathcal{P}$ or there is a v-infan $\mathcal{Q}$ with $V(\mathcal{Q}) \cap X = V^-(\mathcal{Q})$ such that $|\mathcal{P} \setminus \mathcal{Q}| + 1 = |\mathcal{Q} \setminus \mathcal{P}| < \aleph_0$ and $V^-(\mathcal{Q}) \supseteq V^-(\mathcal{P})$.*

We say that the augmentation is *successful* if the second case occurs and we say that it is *unsuccessful* otherwise.

Lemma 7.23. *Assume that $I \subseteq \mathsf{in}_D(v)$ such that $T_{D,v}$ remains linked to v in $D' := D \restriction_v I$. Then for every $u \in V - r$ every $S \in \mathfrak{S}_D(u)$ remains linked to u in D'.*

Proof. Let $u \in V - r - v$ and $S \in \mathfrak{S}(u)$ be given and let $\mathcal{P}$ be a path-system that links S to u in D. We may assume that there is some $e \in E(D) \setminus E(D')$ such that there is a $P_e \in \mathcal{P}$ through e since otherwise $\mathcal{P}$ is a path-system in D' as well and we are done.

We apply Lemma 7.22 in D' with S, u and $\mathcal{P} - P_e$. If the augmentation is successful, the resulting path system witnesses that S is linked to u in D' and we are done. Therefore suppose that the augmentation is unsuccessful. Then we can choose a unique $v_P \in V(P) - u$ from each $P \in \mathcal{P} - P_e$ such that the resulting S' separates u from S in D'. Then $B_{D',S',u}$ is a u-bubble in D' by definition.

Next, we want to show that $B := A_{D,v} \cup B_{D',S',u} \in \mathsf{bubb}_{D'}(u)$ via Lemma 7.17. In order to do this we need to show $A_{D,v} \in \mathsf{bubb}_{D'}(v)$ and $v \in \mathsf{int}_{D'}(B_{D',S',u})$. For $A_{D,v} \in \mathsf{bubb}_{D'}(v)$, recall that $T_{D,v}$ remains linked to v in D' by assumption. Since we only deleted incoming edges of v and $A_{D,v} \supseteq N^-_{D-rv}(v)$ by definition, $A_{D,v}$ is a v-bubble in D'. For $v \in \mathsf{int}_{D'}(B_{D',S',u})$ note that the terminal segment vP_eu of P_e must lie in $B_{D',S',u}$ since otherwise S' would not separate u from S in D'. This implies $v \in B_{D',S',u}$ in particular. Since $v \notin S'$ by construction but $\mathsf{ent}_{D'}(B_{D',S',u}) = S'$ by definition, we can conclude $v \in \mathsf{int}_{D'}(B_{D',S',u})$. Thus we really may apply Lemma 7.22 to $A_{D,v}$ and $B_{D',S',u}$ and obtain $B \in \mathsf{bubb}_{D'}(u)$.

For the construction of a path system witnessing that S is linked to u in D', we show next that $B \subseteq B_{D,S,u}$ by showing this for $A_{D,v}$ and $B_{D',S',u}$ separately. For $A_{D,v} \subseteq B_{D,S,u}$, note that v is the head of e, as we only deleted incoming edges at v. Since $u \neq v$, $V^-(\mathcal{P}) = \mathsf{ent}_D(B_{D,S,u}) - u$ and the paths in $\mathcal{P}$ are pairwise disjoint, this implies $v \in \mathsf{int}_D(B_{D,S,u})$, which in turn implies $\{v\} \cup N^-_{D-rv}(v) \subseteq B_{D,S,u}$. Furthermore, $S \in \mathfrak{S}_D(u)$ implies that $\mathsf{ent}_D(B_{D,S,u})$ is linked from r in D. Thus $B_{D,S,u}$ is an anti-bubble in D that contains $\{v\} \cup N^-_{D-rv}(v)$ and as $A_{D,v}$ is the smallest such anti-bubble by definition, we have $A_{D,v} \subseteq B_{D,S,u}$ as desired. To see $B_{D',S',u} \subseteq B_{D,S,u}$, first note that $S' \subseteq B_{D,S,u}$: as $B_{D,S,u}$ is the set of vertices that are separated from r by S in D, no path from $\mathcal{P}$ can contain a vertex outside of $B_{D,S,u}$, thus as $S' \subseteq V(\bigcup \mathcal{P})$ by construction, S' in fact is a subset of $B_{D,S,u}$. Now suppose for a contradiction that there is a vertex $w \in B_{D',S',u} \setminus B_{D,S,u}$. As $w \notin B_{D,S,u}$, there is an r–w path Q in D that avoids S by definition of $B_{D,S,u}$. Since $w \in B_{D',S',u}$, there is no such path in $D' - S'$, which means that Q either meets a vertex of S' or an edge $e \in E(D) \setminus E(D')$. In the former case the fact that $S' \subseteq B_{D,S,u}$ together with $S = \mathsf{ent}_D(B_{D,S,u})$ which is true by definition, it must meet a vertex of S, a contradiction. In the latter case v is an internal vertex of Q and as $v \in \mathsf{int}_D(B_{D,S,u})$, it cannot be disjoint from S, again a contradiction.

Since $\{v\} \cup N^-_{D-rv}(v) \subseteq A_{D,v}$ by definition, every edge in $E(D) \setminus E(D')$ is spanned by $A_{D,v}$. Thus we can build a u-infan in D' that starts in S as follows: take the initial segments of the paths in $\mathcal{P}$ until the first vertex in B and extend them forward from B to u by using the fact that $B \in \mathsf{bubb}_{D'}(u)$. The resulting path system links S to u in D', as desired. $\square$

Lemma 7.24. *Assume that L is large and $\mathcal{Q}_v \in \mathfrak{P}_L(v, S_{L,v})$ for some $v \in V - r$. Then $L' := L \restriction_v E^+(\mathcal{Q}_v)$ is large, moreover, $S_{L',u} = S_{L,u}$ for every $u \in V - r$.*

Proof. Since L is large, Lemma 7.19 and Corollary 7.18 ensure that $N^-_{D-ru}(u) \subseteq B_{L,u}$ for every $u \in V - r$. In particular, all edges in $E(L) \setminus E(L')$ are spanned by $B_{L,v}$. We are going to prove $B_{L,u} \in \mathsf{bubb}_{L'}(u)$ for every $u \in V - r$ later and first show that this is sufficient. Note that this implies $B_{L',u} \supseteq B_{L,u}$ for $u \in V - r$ because $B_{L',u}$ is the $\subseteq$-largest element of $\mathsf{bubb}_{L'}(u)$, which implies largeness of L' by Lemma 7.19. To prove the moreover part, let $u \in V - r$ be given. First suppose $\mathsf{ent}_{L-ru}(B_{L',u}) = \mathsf{ent}_{L'-ru}(B_{L',u})$. Then $\mathsf{ent}_L(B_{L',u}) - u = \mathsf{ent}_{L'}(B_{L',u}) - u$, which implies $B_{L',u} \in \mathsf{bubb}_L(u)$ and hence $B_{L,u} \supseteq B_{L',u}$, therefore $B_{L,u} = B_{L',u}$. But then

$$S_{L,u} = \mathsf{ent}_{L-ru}(B_{L,u}) = \mathsf{ent}_{L-ru}(B_{L',u}) = \mathsf{ent}_{L'-ru}(B_{L',u}) = S_{L',u}\,,$$

which we wanted to prove. Now suppose that $\mathsf{ent}_{L-ru}(B_{L',u}) \neq \mathsf{ent}_{L'-ru}(B_{L',u})$ for a contradiction. This means that we must have $\mathsf{ent}_{L-ru}(B_{L',u}) \supsetneq \mathsf{ent}_{L'-ru}(B_{L',u})$ with $\mathsf{ent}_{L-ru}(B_{L',u}) \setminus \mathsf{ent}_{L'-ru}(B_{L',u}) = \{v\}$. This implies $v \in \mathsf{int}_{L'-ru}(B_{L',u})$ and that there is some $wv \in E(L) \setminus E(L')$ with $w \notin B_{L',u}$. Then either $v \in \mathsf{int}_{L'}(B_{L',u})$ or $u = v$. Thus by applying Lemma 7.17 with $B_{L',u}$ and $B_{L,v} \in \mathrm{bubb}_{L'}(v)$ we conclude that $B_{L',u} \cup B_{L,v} \in \mathrm{bubb}_{L'}(u)$. This implies that actually $B_{L',u} \supseteq B_{L,v}$ because $B_{L',u}$ is the $\subseteq$-largest element of $\mathrm{bubb}_{L'}(u)$ by definition. But then $w \in B_{L,v} \subseteq B_{L',u} \not\ni w$, a contradiction.

Now we prove that $B_{L,u} \in \mathsf{bubb}_{L'}(u)$ for every $u \in V - r$. For $u = v$ this is witnessed by the terminal segments of the paths in $\mathcal{Q}_v$ from $S_{L,v}$. Thus let $u \in V - r - v$ and let $\mathcal{P}$ be a path-system that witnesses $B_{L,u} \in \mathsf{bubb}_L(u)$. We can assume that $\mathcal{P}$ uses some $e := wv \in E(L) \setminus E(L')$, since otherwise $\mathcal{P}$ ensures that $B_{L,u} \in \mathsf{bubb}_{L'}(u)$ and we are done. Let P_e be the unique path in $\mathcal{P}$ through e. The head v of e must be in $\mathsf{int}_L(B_{L,u})$ because $V^-(\mathcal{P}) = \mathsf{ent}_L(B_{L,u}) - u$ with $u \neq v$ and the paths in $\mathcal{P}$ are pairwise disjoint. Then $B_{L,u} \supseteq B_{L,v}$ since otherwise Lemma 7.17 would give $B_{L,u} \subsetneq (B_{L,u} \cup B_{L,v}) \in \mathsf{bubb}_L(u)$ which is a contradiction. We apply Lemma 7.22 in L' with $\mathsf{ent}_L(B_{L,u}) - u, u$ and $\mathcal{P} - P_e$. If the augmentation is successful, the resulting path-system must lie in $L'[B_{L,u}]$ and witnesses $B_{L,u} \in \mathsf{bubb}_{L'}(u)$ thus we are done.

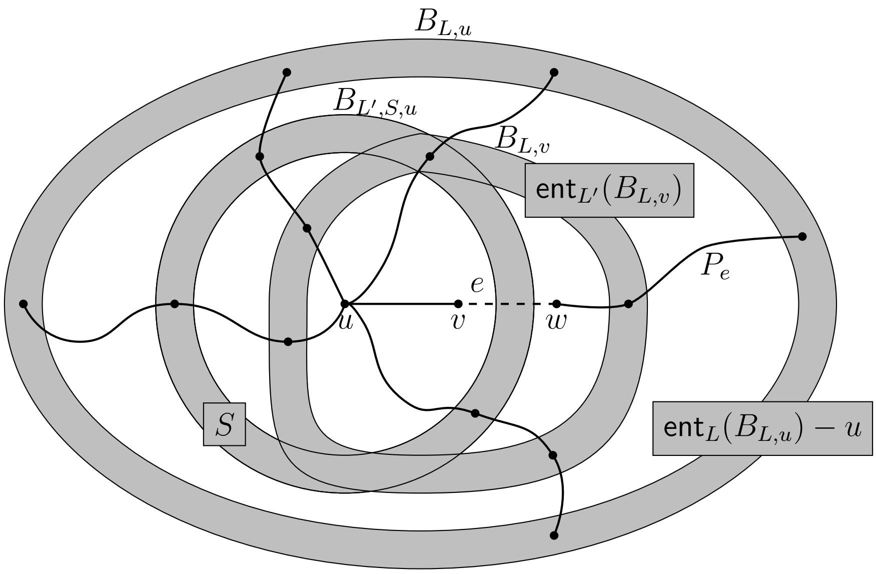

Figure 7.1: The situation when the augmentation in the proof of Lemma 7.24 is unsuccessful.

Suppose that the augmentation is unsuccessful, we depict this situation in Figure 7.1. Then we can choose a unique $v_P \in V(P) - u$ from each $P \in \mathcal{P} - P_e$ such that the resulting S separates u from $\mathsf{ent}_L(B_{L,u}) - u$ in L'. We know that the terminal segment vP_eu of P_e must lie in $B_{L',S,u}$ since otherwise S would not separate u from $\mathsf{ent}_L(B_{L,u}) - u$ in L'. Thus in particular $v \in B_{L',S,u}$ with $v \notin S$, i.e. $v \in \mathsf{int}_{L'}(B_{L',S,u})$. But then Lemma 7.17 ensures $B := B_{L',S,u} \cup B_{L,v} \in \mathsf{bubb}_{L'}(u)$. Let $\mathcal{P}'$ consist of the initial segments of the paths in $\mathcal{P}$ until the first vertex in B. All edges in $E(L) \setminus E(L')$ are spanned by B because they are spanned by $B_{L,v}$ and $B \supseteq B_{L,v}$ by construction, therefore $\mathcal{P}'$ is a path-system in L' (and not just in L). But then by using the fact $B \in \mathsf{bubb}_{L'}(u)$, each path in $\mathcal{P}'$ can be continued forward in L' to reach u in such a way that the resulting paths meet only at u. This path-system witnesses $B_{L,u} \in \mathsf{bubb}_{L'}(u)$ which completes the proof. $\square$

7.2 Proof of the Main Result

7.2.1 Definitions and a Sketch of the Construction

Let an r-rooted digraph $D = (V, E)$ of size $\aleph_1$ be fixed. First of all, we can assume by Corollary 7.6 that D is a quasi-vertex-flame. Let $\langle M_\alpha : \alpha \leq \omega_1 \rangle$ be a sequence such that

- $M_0 = \emptyset$,
- M_α is an elementary submodel for each $\alpha > 0$,
- $M_\alpha = \bigcup_{\beta<\alpha} M_\beta$ if α is a limit ordinal,
- M_α is countable for $\alpha < \omega_1$, and
- $D, r, \langle M_\beta : \beta \leq \alpha \rangle, \alpha \in M_{\alpha+1}$ for $\alpha < \omega_1$.

Observation 7.25. $M_\beta \cup \{M_\beta\} \subseteq M_\alpha$ for $\beta < \alpha \leq \omega_1$.

Let $M^\alpha := M_{\alpha+1} \setminus M_\alpha$ for $\alpha < \omega_1$ and we define $V_\alpha := V \cap M_\alpha$, $D_\alpha := D \cap M_\alpha = D[V_\alpha]$ for $\alpha \leq \omega_1$ as well as

$$\begin{aligned} V^\alpha &:= V \cap M^\alpha, \\ D^\alpha &:= (D \cap M_{\alpha+1}) \setminus [E(D_\alpha) \cup \mathsf{out}_D(V_\alpha - r)], \text{ and} \\ D^{\alpha\leq} &:= D \setminus [E(D_\alpha) \cup \mathsf{out}_D(V_\alpha - r)] \end{aligned}$$

for $\alpha < \omega_1$. We choose enumerations $V_\alpha - r = \{v_{\alpha,n} \colon n < \omega\}$ and $V^\alpha - r = \{v^\alpha_n \colon n < \omega\}$ for $\alpha < \omega_1$ (technically we fix a choice function $c \in M_1$ and we choose the enumerations accordingly). Recall that every countable ordinal number can be written uniquely in the form $\omega\alpha + n$ where $n < \omega$. We obtain an enumeration $V - r = \{v_\xi \colon \xi < \omega_1\}$ by letting $v_{\omega\alpha+n} := v^\alpha_n$. Observe that $V_\alpha - r = \{v_\xi \colon \xi < \omega\alpha\}$ for $\alpha \leq \omega_1$. We shall construct a sequence $\langle L_\xi \colon \xi \leq \omega_1 \rangle$ of large subdigraphs of D with

- $L_0 = D$,
- we obtain $L_{\xi+1}$ by the deletion of some of the incoming edges of v_ξ from L_ξ, and
- $L_\nu = \bigcap_{\xi<\nu} L_\xi$ if ν is a limit ordinal.

Before giving the complete list of properties of the recursive construction, we need some furhter definitions. We also explain roughly the ideas behind what we are going to do in order to make it easier to follow the formal proof afterwards. First of all, L_{ω_1} will be a large vertex-flame which completes the proof of Theorem 7.2. For $v = v_{\omega\beta+n}$, let $S_v := S_{L_{\omega\beta},v}$. After $L_{\omega\beta}$ is defined, the separations S_v will be defined for the following countably many vertices, namely for the vertices in V^β. By guaranteeing that L_ξ is large for every ξ, we will automatically ensure $S_v \in \mathfrak{S}_{D,v}$ by Lemma 7.19. We strive to end up with path-systems $\mathcal{P}_v \in \mathfrak{P}_{L_{\omega_1}}(v, S_v)$ with $E^+(\mathcal{P}_v) = \mathsf{in}_{L_{\omega_1}}(v) - rv$ in L_{ω_1} for every $v \in V - r$. The path-systems $\mathcal{P}_v$ witness that L_{ω_1} is indeed a large vertex-flame. For each $v \in V - r$ we build the path-system $\mathcal{P}_v$ 'layer by layer' according to our chain of elementary submodels (see Figure 7.2) in the following sense. If $v = v_{\omega\beta+n}$, then we construct first a segment $\mathcal{P}_{v,\beta+1} \in \mathfrak{P}_{L_{\omega_1}\cap D_{\beta+1}}(v, S_v \cap V_{\beta+1})$ with $E^+(\mathcal{P}_{v,\beta+1}) = \mathsf{in}_{L_{\omega_1}\cap D_{\beta+1}}(v) - rv$. In every new layer we extend this by a new segment: for every γ with $\beta + 1 \leq \gamma < \omega_1$ we construct a path-system $\mathcal{P}_v^\gamma \in \mathfrak{P}_{L_{\omega_1}\cap D^\gamma}(v, S_v \cap V^\gamma)$ with $E^+(\mathcal{P}_v^\gamma) = \mathsf{in}_{L_{\omega_1}\cap D^\gamma}(v)$. Since by the definition of D^γ the paths in $\mathcal{P}_v^\gamma$ are internally disjoint from V_γ, the new segments never share any internal vertex with the previously constructed segments. By letting

$$\mathcal{P}_{v,\alpha} := \mathcal{P}_{v,\beta+1} \cup \bigcup_{\beta+1\leq\gamma<\alpha} \mathcal{P}_v^\gamma$$

for $\beta + 1 \leq \alpha \leq \omega_1$, the path-system $\mathcal{P}_v := \mathcal{P}_{v,\omega_1}$ will be as desired.

Instead of constructing $\mathcal{P}_{v,\beta+1}$ and $\mathcal{P}_v^\gamma$ 'directly' we are going to build some supersets of them and throw away the surplus paths. For this we let $L^{\omega\beta+n} := L_{\omega\beta+n} \cap D^{\beta\leq}$.

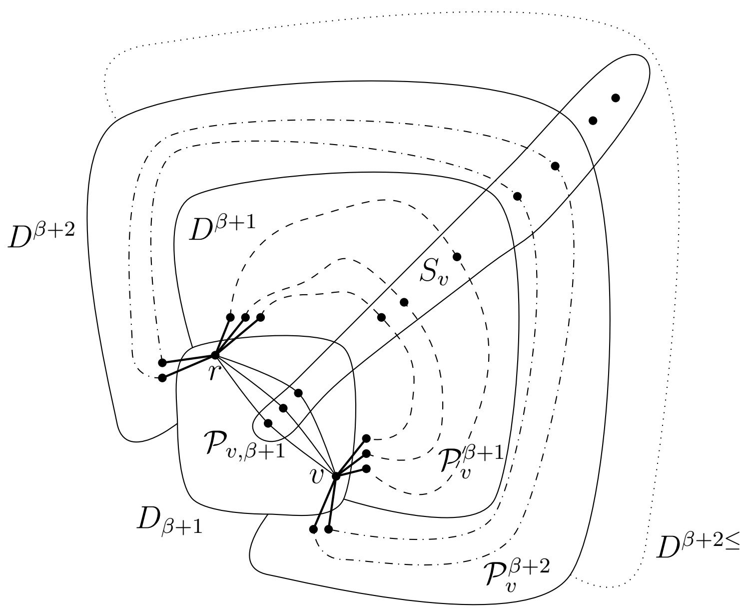

Figure 7.2: A sketch of the strategy to build $\mathcal{P}_v$.

7.2.2 The Conditions of the Recursion

Let us now make the construction precise. We shall define by transfinite recursion sequences

$$\langle \mathcal{Q}_\xi : \xi < \omega_1 \rangle, \ \langle \mathcal{Q}^\xi : \omega \le \xi < \omega_1 \rangle \ \text{ and } \ \langle L_\xi : \xi \le \omega_1 \rangle$$

satisfying the following properties:

(1) $L_0 = D$,

(2) $\mathcal{Q}_\xi \in \mathfrak{P}_{L_\xi}(v_\xi, S_{v_\xi})$,

(3) $L_{\xi+1} := L_\xi \restriction_{v_\xi} E^+(\mathcal{Q}_\xi)$,

(4) $L_\nu = \bigcap_{\xi<\nu} L_\xi$ if ν is a limit ordinal,

(5) L_ξ is D-large,

(6) $S_{L_{\omega\alpha+n}, v} = S_v$ for every $\omega\alpha + n < \omega_1$ and $v \in V^\alpha$,

(7) for every $\omega\alpha + n < \omega_1$ and $v \in V_\alpha - r$:

- **(a)** $S_v \setminus V_\alpha \in \mathfrak{S}_{L^{\omega\alpha+n}}(v)$,
- **(b)** $L^{\omega\alpha+n}$ has the vertex-flame property at v,
- **(c)** if $v = v_{\alpha,n}$, then $\mathcal{Q}^{\omega\alpha+n} \in \mathfrak{P}_{L^{\omega\alpha+n}}(v, S_v \setminus V_\alpha)$ with $E^+(\mathcal{Q}^{\omega\alpha+n}) = \mathsf{in}_{L^{\omega\alpha+n}}(v)$,

(8) $\langle \mathcal{Q}_\xi : \xi < \nu \rangle, \ \langle \mathcal{Q}^\xi : \omega \le \xi < \nu \rangle, \ \langle L_\xi : \xi < \nu \rangle \in M_{\alpha+1}$ for $\nu = \omega\alpha + n < \omega_1$, and

(9) for $v = v_{\omega\beta+m}$:

$$\bigcup_{n<m} \mathsf{in}_{\mathcal{Q}_{\omega\beta+n}}(v) \subseteq E^+(\mathcal{Q}_{\omega\beta+m}) + rv \text{ if } \beta = 0 \text{ and}$$

$$\bigcup_{n\le m} \mathsf{in}_{\mathcal{Q}^{\omega\beta+n}}(v) \cup \bigcup_{n<m} \mathsf{in}_{\mathcal{Q}_{\omega\beta+n}}(v) \subseteq E^+(\mathcal{Q}_{\omega\beta+m}) + rv \text{ if } \beta > 0\,.$$

Note that L_ν is uniquely determined by $\langle \mathcal{Q}_\xi : \xi < \nu \rangle$ (see properties (1) to (4)), thus we are going to always have at most one suitable choice for L_ξ but we still need to check if it respects the conditions. The preservation of property (8) will follow immediately from the fact that the definitions of $\langle \mathcal{Q}_\xi : \xi < \omega\alpha + n \rangle$ and $\langle \mathcal{Q}^\xi : \omega \le \xi < \omega\alpha + n \rangle$ rely only on parameters that are in $M_{\alpha+1}$, namely $D, r, \langle M_\beta, \beta \le \alpha \rangle$, c and vertices $v_{\omega\alpha+k}, v_{\alpha,k}$ for $k < n$ (where c is some fixed choice function).

7.2.3 The Path-Systems $\mathcal{P}_{v,\beta+1}$ and $\mathcal{P}_v^\gamma$

Let $\nu \leq \omega_1$ and suppose that $\mathcal{Q}_\xi$ and L_ξ are defined for $\xi < \nu$ and $\mathcal{Q}^\xi$ is defined for $\omega \leq \xi < \nu$ and none of the conditions (1) to (9) is violated so far. Let $v = v_{\omega\beta+n}$ for some $\omega\beta + n < \nu$ be fixed. We define $\mathcal{P}_{v,\beta+1} := \mathcal{Q}_{\omega\beta+n} \cap M_{\beta+1}$. By the definition of the enumeration $\{v_{\gamma,n} \colon n < \omega\}$, for every γ with $\beta + 1 \leq \gamma < \omega_1$ there is a unique $m_\gamma < \omega$ with $v_{\gamma,m_\gamma} = v$. We let $\mathcal{P}_v^\gamma := \mathcal{Q}^{\omega\gamma+m_\gamma} \cap M_{\gamma+1}$ whenever $\beta+1 \leq \gamma$ and $\omega\gamma + m_\gamma < \nu$.

Property (9) is designed to prevent the deletion of edges of the path-systems $\mathcal{P}_{v,\beta+1}$ and $\mathcal{P}_v^\gamma$:

Lemma 7.26. *Let $v = v_{\omega\beta+n}$ for some $\omega\beta + n < \nu$. The path-systems $\mathcal{P}_{v,\beta+1}$ and $\mathcal{P}_v^\gamma$ lie in L_ξ for every $\xi < \nu$.*

Proof. Since $\mathcal{Q}_{\omega\beta+n}$ is a path-system in $L_{\omega\beta+n}$ (see property (2)), so is $\mathcal{P}_{v,\beta+1}$. It follows from properties (1) to (4) that L_ξ is a $\subseteq$-decreasing function of ξ which implies that $\mathcal{P}_{v,\beta+1}$ is a path-system in L_ξ for $\xi \leq \omega\beta + n$. By (3) we do not delete any edges of $\mathcal{Q}_{\omega\beta+n} \supseteq \mathcal{P}_{v,\beta+1}$ when we obtain $L_{\omega\beta+n+1}$ from $L_{\omega\beta+n}$. Property (9) applied to the vertices $v_{\omega\beta+m}$ with $n < m < \omega$ together with (3) guarantees that none of the edges of $\mathcal{P}_{v,\beta+1}$ is deleted when we construct $L_{\omega\beta+m}$ for $m > n$. Thus $\mathcal{P}_{v,\beta+1}$ is a path-system in $L_{\omega(\beta+1)}$ as well by (4). Whenever $P \in M_{\beta+1}$ is a path, we have $V(P) \subseteq M_{\beta+1}$, therefore $\mathcal{P}_{v,\beta+1}$ lies completely in $D_{\beta+1}$ and after step $\omega(\beta+1)$ we delete only edges e whose head is in $V \setminus V_{\beta+1}$. Thus the path-system $\mathcal{P}_{v,\beta+1}$ lies in L_ξ for every ξ with $\omega(\beta+1) < \xi < \nu$ as well. The proof for $\mathcal{P}_v^\gamma$ goes similarly. □

Proposition 7.27. *If $v = v_{\omega\beta+n}$ for some $\omega\beta + n < \nu$, then $\mathcal{P}_{v,\beta+1} \in \mathfrak{P}_{D_{\beta+1}}(v, S_v \cap V_{\beta+1})$. Furthermore, $\mathsf{in}_{L_\xi}(v) \cap E(D_{\beta+1}) - rv = E^+(\mathcal{P}_{v,\beta+1})$ for every $\xi \in (\omega\beta + n, \nu)$.*

Proof. Note that $\omega\beta + n \in M_{\beta+1}$ because $\beta \in M_{\beta+1}$ by assumption. Hence by (8), $\mathcal{Q}_{\omega\beta+n} \in M_{\beta+1}$. Each $P \in \mathcal{Q}_{\omega\beta+n}$ which has an internal vertex u in $V_{\beta+1}$ is definable from $\mathcal{Q}_{\omega\beta+n}$ and u and therefore must be in $M_{\beta+1}$. This means that for each $P \in \mathcal{Q}_{\omega\beta+n}$ either $V(P) \subseteq V_{\beta+1}$ or P is internally disjoint from $V_{\beta+1}$. Thus by (2) it follows that $\mathcal{P}_{v,\beta+1} \in \mathfrak{P}_{D_{\beta+1}}(v, S_v \cap V_{\beta+1})$. Moreover, by (3), $\mathcal{P}_{v,\beta+1}$ covers $\mathsf{in}_{L_{\omega\beta+n+1}}(v) \cap E(D_{\beta+1}) - rv$ which is the same as $\mathsf{in}_{L_\xi}(v) \cap E(D_{\beta+1}) - rv$ whenever $\omega\beta + n + 1 \leq \xi < \nu$ by properties (1) to (4). □

Proposition 7.28. *If $v = v_{\gamma,m_\gamma} = v_{\omega\beta+n}$ with $\omega\gamma + m_\gamma < \nu$, then $\mathcal{P}_v^\gamma \in \mathfrak{P}_{D^\gamma}(v, S_v \cap V^\gamma)$. Furthermore, $\mathsf{in}_{L_\xi \cap D^\gamma}(v) = E^+(\mathcal{P}_v^\gamma)$ for every $\xi \in (\omega\beta + n, \nu)$.*

Proof. Note that $\beta < \gamma$ because $v \in V_\gamma$ by $v = v_{\gamma,m_\gamma}$ and $\beta + 1$ is the smallest ordinal with $v \in V_{\beta+1}$ by $v = v_{\omega\beta+n}$ (see the definition of the enumerations after Observation 7.25). Since $\mathcal{Q}^{\omega\gamma+m_\gamma} \in M_{\gamma+1}$, we obtain via property (7c) that $\mathcal{P}_v^\gamma \in \mathfrak{P}_{D^\gamma}(v, (S_v \setminus V_\gamma) \cap V_{\gamma+1})$ and $E^+(\mathcal{P}_v^\gamma) = \mathsf{in}_{L^{\omega\gamma+m_\gamma} \cap D_{\gamma+1}}(v)$. The first part of the proposition follows by observing that $(S_v \setminus V_\gamma) \cap V_{\gamma+1} = S_v \cap V^\gamma$ by definition. The second part follows from the fact that: $L^{\omega\gamma+m_\gamma} \cap D_{\gamma+1} = L^{\omega\gamma+m_\gamma} \cap D^\gamma$ (which is also a direct consequence of the corresponding definitions). Finally $\mathsf{in}_{L^\xi}(v)$ remains the same for every $\xi \in (\omega\beta + n, \nu)$ by properties (1) to (4). □

Lemma 7.29. *If $v = v_{\omega\beta+n}$ for some $\omega\beta + n < \nu$ and $\beta < \alpha < \omega_1$, then we have $\mathcal{Q}_{\omega\beta+n} \setminus M_\alpha \in \mathfrak{P}_{L^{\omega\alpha}}(v, S_v \setminus V_\alpha)$. Furthermore, $\mathsf{in}_{L_\xi}(v) \cap E(D^{\alpha\leq}) = E^+(\mathcal{Q}_{\omega\beta+n} \setminus M_\alpha)$ for every $\xi \in (\omega\beta + n, \nu)$.*

Proof. First of all, $\mathcal{Q}_{\omega\beta+n}$ lies in $L_{\omega\beta+n}$ according to (2) and we have already seen that $\mathcal{Q}_{\omega\beta+n} \in M_{\beta+1} \subseteq M_\alpha$. Thus each path in $\mathcal{Q}_{\omega\beta+n} \setminus M_\alpha$ is internally disjoint from V_α. This gives $\mathcal{Q}_{\omega\beta+n} \setminus M_\alpha \in \mathfrak{P}_{L^{\omega\beta+n}}(v, S_v \setminus V_\alpha)$ via (2). In order to obtain $L^{\omega\alpha}$ from $L^{\omega\beta+n}$, we delete only edges whose heads are in $V_\alpha - r$. The only such edges in $E(\mathcal{Q}_{\omega\beta+n} \setminus M_\alpha)$ are also in $E^+(\mathcal{Q}_{\omega\beta+n})$, but we do not delete any of those by (3), therefore $\mathcal{Q}_{\omega\beta+n} \setminus M_\alpha \in \mathfrak{P}_{L^{\omega\alpha}}(v, S_v \setminus V_\alpha)$. The second part follows directly from properties (1) to (4). □

7.2.4 Limit Step

Suppose now that $\nu = \omega\alpha$ for some $\alpha \leq \omega_1$ and as earlier assume that $\mathcal{Q}_\xi$ and L_ξ are defined for $\xi < \nu$ and $\mathcal{Q}^\xi$ is defined for $\omega \leq \xi < \nu$ and none of the conditions (1) to (9) is violated so far. If $\alpha = 0$, then let $L_0 := D$ which is our only possible choice by (1) and this choice does not violate any of the conditions for trivial reasons. If $\alpha > 0$, then by (4) our only option is to let $L_{\omega\alpha} := \bigcap_{\xi<\omega\alpha} L_\xi$. We need to check that none of (1) to (9) is violated. Note that $\mathcal{Q}^{\omega\alpha}$ and $\mathcal{Q}_{\omega\alpha}$ will be defined from $L_{\omega\alpha}$ in step $\omega\alpha + 1$, thus we need not check any clause that refers to either of them. The preservation of (1) to (4) is clear. We intend to apply Corollary 7.20 to demonstrate the largeness of $L_{\omega\alpha}$ i.e. that (5) holds. Clearly $\mathsf{out}_D(r) \subseteq L_{\omega\alpha}$ because we never delete any outgoing edge of r (this was built into the definition of '$\restriction_v$'). Note that $\mathsf{in}_{L_{\omega\alpha}}(v) \subsetneq \mathsf{in}_D(v)$ may happen only for $v \in V_\alpha - r$. For every $v \in V_\alpha - r$ there is some $\beta < \alpha$ and $n < \omega$ such that $v = v_{\omega\beta+n}$. We know by (5) that $L_{\omega\beta}$ is large but then $S_{L_{\omega\beta},v}$ (which is S_v by definition) is in $\mathfrak{S}_D(v)$ (see Lemma 7.19). We define

$$\mathcal{P}_{v,\alpha} := \mathcal{P}_{v,\beta+1} \cup \bigcup_{\beta+1\leq\gamma<\alpha} \mathcal{P}_v^\gamma \,.$$

Note that $\mathsf{in}_{L_{\omega\alpha}}(v) \cap E(D_{\beta+1}) - rv = E^+(\mathcal{P}_{v,\beta+1})$ and $\mathsf{in}_{L_{\omega\alpha} \cap D^\gamma}(v) = E^+(\mathcal{P}_v^\gamma)$ follow from Proposition 7.27 and Proposition 7.28 respectively via $L_{\omega\alpha} = \bigcap_{\xi<\nu} L_\xi$. Thus these propositions have the following consequence:

Corollary 7.30. *For every $v \in V_\alpha - r$, $\mathcal{P}_{v,\alpha} \in \mathfrak{P}_{D_\alpha}(v, S_v \cap V_\alpha)$ with $E^+(\mathcal{P}_{v,\alpha}) = \mathsf{in}_{L_{\omega\alpha}}(v) \cap E(D_\alpha) - rv$.*

It follows from Lemma 7.29 and Corollary 7.30 that if $v = v_{\omega\beta+n} \in V_\alpha - r$, then

$$\mathcal{P}_{v,\alpha} \cup (\mathcal{Q}_{\omega\beta+n} \setminus M_\alpha) \in \mathfrak{P}_{L_{\omega\alpha}}(v, S_v) \,.$$

Thus $L_{\omega\alpha}$ is large by Corollary 7.20. Property (6) for $n = 0$ is true by the definition of S_v. Property (7a) and (7b) for $v = v_{\omega\beta+n}$ is witnessed by $\mathcal{Q}_{\omega\beta+n} \setminus M_\alpha$ according to Lemma 7.29. Property (8) is maintained, because as we already argued the transfinite recursion so far can be carried out in $M_{\alpha+1}$ since it relies only on the parameters $D, r, \langle \mathcal{Q}_\xi : \xi < \omega\alpha \rangle, c \in M_{\alpha+1}$. Finally, we do not check (7c) and (9), since they refer to $\mathcal{Q}_{\omega\alpha}$ and $\mathcal{Q}^{\omega\alpha}$.

Lemma 7.31. *$L_{\omega\alpha} \cap D_\alpha$ is a D_α-large vertex-flame. In particular, if $\alpha = \omega_1$, then L_{ω_1} is a large vertex-flame.*

Proof. We know by Corollary 7.30 that $\mathcal{P}_{v,\alpha} \in \mathfrak{P}_{D_\alpha}(v, S_v \cap V_\alpha)$ with $E^+(\mathcal{P}_{v,\alpha}) = \mathsf{in}_{L_{\omega\alpha}}(v) \cap E(D_\alpha) - rv$. Note that $S_v \cap V_\alpha$ separates v from r in $D_\alpha - rv$ because so does S_v in $D - rv$. Thus $S_v \cap V_\alpha \in \mathfrak{S}_{D_\alpha}(v)$ is witnessed by $\mathcal{P}_{v,\alpha}$. Since the path-systems $\mathcal{P}_{v,\alpha}$ for $v \in V_\alpha - r$ lie in $L_{\omega\alpha}$ (see Lemma 7.26 and (4)) they show that $L_{\omega\alpha} \cap D_\alpha$ is a D_α-large vertex-flame. □

We will make use of the following consequence.

Corollary 7.32. *For every finite $U \subseteq V_\alpha - r$ which is linked from r in D, it is also linked from r in $L_{\omega\alpha} \cap D_\alpha$.*

Proof. This follows directly from Lemma 7.31 via Lemma 7.7. □

7.2.5 Successor Step

Suppose that there is some $\omega\alpha + n < \omega_1$ such that the following are already defined without violating the conditions:

- L_ξ for $\xi \leq \omega\alpha + n$,
- $\mathcal{Q}_\xi$ for $\xi < \omega\alpha + n$, and
- $\mathcal{Q}^\xi$ for $\omega \leq \xi < \omega\alpha + n$.

Note that in the first of these three properties we have 'less than or equal' while there is a strict inequality for the other two. The reason for that being that $L_{\omega\alpha+n}$ is always defined from the choices for $\mathcal{Q}_\xi$ for $\xi < \omega\alpha + n$, see (3) and (4), while $\mathcal{Q}_{\omega\alpha+n}$ and $\mathcal{Q}^{\omega\alpha+n}$ are defined in step $\omega\alpha + n + 1$. In particular, if $n = 0$, then $L_{\omega\alpha}$ is defined, see Subsection 7.2.4, while $\mathcal{Q}_{\omega\alpha}$ and $\mathcal{Q}^{\omega\alpha}$ will be defined from it in this subsection.

Let $v = v_{\omega\alpha+n}$. Note that $L_{\omega\alpha+n}$ is a quasi-vertex-flame by Corollary 7.8 because D is a quasi-vertex-flame by assumption and $L_{\omega\alpha+n}$ is large by (5).

Suppose first that $\alpha = 0$. Let $I := \bigcup_{k<n} \mathsf{in}_{\mathcal{Q}_k}(v)$. Since $|I| \leq n$, we have $I \in \mathcal{G}_{L_n}(v)$ by the quasi-vertex-flame property. We have $S_n \in \mathfrak{S}_D(v)$ by (6). By applying Corollary 7.12 we pick a $\mathcal{Q}_n \in \mathfrak{P}_{L_n}(v, S_v)$ that covers $I - rv$. Recall that $\mathcal{Q}^\xi$ will only be defined for $\xi \geq \omega$, so we do not define $\mathcal{Q}^n$. We define $L_{n+1} := L_n \restriction_v E^+(\mathcal{Q}_n)$. Preservation of properties (2), (3) and (9) follows directly from the construction. Conditions (1)

and (4) do not demand anything new for this step. Properties (5) and (6) are preserved by Lemma 7.24. Since $V_0 = \emptyset$, (7) says nothing so far. The definition of $\mathcal{Q}_n$ and L_{n+1} used only L_n, v_n and the choice function c as parameters all of which are in M_1, thus (8) is preserved.

Assume now that $\alpha > 0$. In this case (7) also has demands. We choose $\mathcal{Q}^{\omega\alpha+n}$ in accordance with (7c) which is possible by combining (7a) and (7b) via Corollary 7.12. To fulfil (7a) and (7b) for $\omega\alpha+n+1$ we are going to pick an I according to the following claim.

Claim 7.33. There exists an $I \in \mathcal{G}_{L_{\omega\alpha+n}}(v)$ such that

(I) $S_u \setminus V_\alpha \in \mathfrak{S}_{L^{\omega\alpha+n}\restriction_v I}(u)$ for every $u \in V_\alpha - r$,

(II) $L^{\omega\alpha+n} \restriction_v I$ has the vertex-flame property for every $u \in V_\alpha - r$, and

(III) $I \supseteq \bigcup_{k\leq n} \mathsf{in}_{Q^{\omega\alpha+k}}(v) \cup \bigcup_{k<n} \mathsf{in}_{Q_{\omega\alpha+k}}(v) =: F$.

Suppose that we already know Claim 7.33. By Corollary 7.12 we can pick a path-system $\mathcal{Q}_{\omega\alpha+n} \in \mathfrak{P}_{L_{\omega\alpha+n}}(v, S_v)$ that covers $I - rv$. We define $L_{\omega\alpha+n+1} := L_{\omega\alpha+n} \restriction_v E^+(\mathcal{Q}_{\omega\alpha+n})$. Conditions (1) to (6) are preserved for the same reason as in the case $\alpha = 0$. Note that properties (I) and (II) of the desired set I are increasing in the sense that if they hold for some I, then they remain true for every $I' \supseteq I$. Indeed, the path-systems witnessing these properties for I also witness them with respect to I'. Conditions (I) and (II) guarantee (7a) and (7b) respectively for $\omega\alpha+n+1$. Preservation of (9) is ensured by (III). The definition of $\mathcal{Q}_{\omega\alpha+n}$, $\mathcal{Q}^{\omega\alpha+n}$ and $L_{\omega\alpha+n+1}$ rely only on the parameters $L_{\omega\alpha+n}, v_{\omega\alpha+n}, c$ and M_α all of which are in $M_{\alpha+1}$, thus we keep (8) as well.

Proof of Claim 7.33. Since F is finite ($|F| \leq 2n+1$ follows directly from its definition) and $L_{\omega\alpha+n}$ is a quasi-vertex-flame, we have $F \in \mathcal{G}_{L_{\omega\alpha+n}}(v)$. We claim that it is possible to choose a witness $\mathcal{Q}_F = \{Q_e : e \in F\}$ for $F \in \mathcal{G}_{L_{\omega\alpha+n}}(v)$ where e is the last edge of Q_e in such a way that whenever a path in $\mathcal{Q}_F$ leaves V_α it never returns, in other words no path in $\mathcal{Q}_F$ has an edge in $\mathsf{in}_D(V_\alpha)$. Indeed, suppose that $\mathcal{Q}'_F$ is an arbitrary witness for $F \in \mathcal{G}_{L_{\omega\alpha+n}}(v)$ and let U_F be the set of the last common vertices of the paths in $\mathcal{Q}'_F$ with V_α. Then $U_F - r$ is a finite subset of $V_\alpha - r$ which is linked from r in D. But then, since M_α is an elementary submodel, $U_F - r$ is linked from r in D_α as well. It follows from Lemma 7.7 via the D_α-largeness of $L_{\omega\alpha} \cap D_\alpha$ (see Lemma 7.31) and $L_{\omega\alpha} \cap D_\alpha = L_{\omega\alpha+n} \cap D_\alpha$ that $U_F - r$ remains linked from r in $L_{\omega\alpha+n} \cap D_\alpha$. This means that we can replace the initial segments of the paths in $\mathcal{Q}'_F$ up to U_F in $L_{\omega\alpha+n} \cap D_\alpha$ in such a way that these new initial segments have vertices only in V_α. This modification of $\mathcal{Q}'_F$ provides the desired $\mathcal{Q}_F$.

We build an auxiliary digraph A by adding a 'dummy' vertex w_u for every $u \in V_\alpha - r$ to $L^{\omega\alpha+n}$ whose in-neighbours are $S_u \setminus V_\alpha$ and which has no out-neighbours. Let $W := (V_\alpha - r) \cup \{w_u : u \in V_\alpha - r\}$. Then A has the vertex-flame property at every $w \in W$ by properties (7a) and (7b), moreover, W is countable. We are going to choose

I in such a way that $A \restriction_v I$ also has the vertex flame property for every $w \in W$. For the original digraph $L^{\omega\alpha+n}$ this means that $L^{\omega\alpha+n} \restriction_v I$ has the vertex-flame property for every $u \in V_\alpha - r$ (demanded by (II)), furthermore, the preservation of the vertex-flame property for the dummy vertex w_u ensures that $S_u \setminus V_\alpha$ remains linked from r in $L^{\omega\alpha+n} \restriction_v I$ which can be thought of as 'half' of condition (I).

By applying Corollary 7.16 with $A, v = v_{\omega\alpha+n}$ and W, we obtain $I^* \in \mathcal{G}_A(v) = \mathcal{G}_{L^{\omega\alpha+n}}(v)$ such that $A \restriction_v I^*$ has the vertex-flame property for every $w \in W$. Let $\mathcal{Q}$ be a system of internally disjoint r–v paths in $L_{\omega\alpha+n}$ such that

- **(i)** $F \subseteq E^+(\mathcal{Q}) \subseteq F \cup I^*$,
- **(ii)** $I^* \setminus E^+(\mathcal{Q})$ is finite,
- **(iii)** Whenever some $Q \in \mathcal{Q}$ is not a path in $L^{\omega\alpha+n}$, then $Q \in \mathcal{Q}_F$, and
- **(iv)** $I^* \setminus E^+(\mathcal{Q})$ is minimal among path-systems satisfying (i) to (iii).

We first show that $\mathcal{Q}$ is well-defined: if we take a path-system $\mathcal{Q}_{I^*}$ witnessing $I^* \in \mathcal{G}_A(v)$, then, since $\mathcal{Q}_F$ is finite, there is a co-finite subset $\mathcal{Q}'_{I^*}$ of $\mathcal{Q}_{I^*}$ for which the path-system $\mathcal{Q}'_{I^*} \cup \mathcal{Q}_F$ is internally disjoint and hence satisfies (i) to (iii).

We claim that $E^+(\mathcal{Q})$ still has the property that I^* had, namely that $A \restriction_v E^+(\mathcal{Q})$ has the vertex-flame property for every $w \in W$. Suppose for a contradiction that $A \restriction_v E^+(\mathcal{Q})$ does not have the vertex-flame property at some $w \in W$. Note that we necessarily must have $w \neq v$, because $\mathcal{Q}$ witnesses $E^+(\mathcal{Q}) \in \mathcal{G}_{\restriction_v} E^+(\mathcal{Q})(v)$. Let $\mathcal{P}_w$ be a witness for $\mathsf{in}_A(w) \in \mathcal{G}_{A\restriction_v I^*}(w)$. Then there is precisely one path $P \in \mathcal{P}_w$ that uses precisely one edge $uv \in I^* \setminus E^+(\mathcal{Q})$. Thus $\mathcal{P}_w$ witnesses $\mathsf{in}_{A\restriction_v E^+(\mathcal{Q})+uv}(w) \in \mathcal{G}_{A\restriction_v E^+(\mathcal{Q})+uv}(w)$. Note also that $u \neq r$, moreover, $rv \notin E(D)$ since otherwise $rv \in E(A)$ and hence the initial segment Pv of P can be replaced by the single edge rv and this shows that $\mathsf{in}_A(w) \in \mathcal{G}_{A\restriction_v E^+(\mathcal{Q})}(w)$, which contradicts the choice of w. Thus we may apply Corollary 7.14 with $A \restriction_v E^+(\mathcal{Q}) + uv$, w and uv and we obtain a vertex set $S \ni v$ which is linked from r in $A \restriction_v E^+(\mathcal{Q}) + uv$ by a path-system $\mathcal{P}_S$, such that S separates $N^-_{A\restriction_v E^+(\mathcal{Q})+uv}(v) - u$ from r. In particular, uv is the last edge of some $P_{uv} \in \mathcal{P}_S$. We can assume $S \cap V_\alpha = \emptyset$ by taking $S := S \setminus V_\alpha$ instead, because the vertices in $V_\alpha - r$ do not have outgoing edges in A and hence it is still a separator. We modify $\mathcal{Q}$ in the following way. Whenever $Q \in \mathcal{Q}$ does not meet $S - v$, then we let $Q' := Q$. Note that since S separates $N^-_{A\restriction_v E^+(\mathcal{Q})+uv}(v) - u$ from r and no path of $\mathcal{Q}$ uses one of the edges in $\mathsf{in}_A(v) \setminus E^+(\mathcal{Q})$ by definition, such a path is not a path in $L^{\omega\alpha+n}$, thus by (iii) it is a path of $\mathcal{Q}_F$. Any other path $Q \in \mathcal{Q} \setminus \mathcal{Q}_F$ meets $S - v$. In this case we take the last common vertex v_Q of Q with $S - v$ and replace the initial segment Qv_Q by the unique path in $\mathcal{P}_S$ that terminates at v_Q to obtain Q'. Since the paths in $\mathcal{P}_S$ are paths in $A \restriction_v E^+(\mathcal{Q})$, the paths are disjoint from $Q \in \mathcal{Q}_F$ by the construction of A. Likewise, no path of $\mathcal{P}_S$ can share a vertex with one of the segments $v_Q Q$ other than v_Q, since their union then would contain an r–v-path in $A \restriction_v E^+(\mathcal{Q}) + uv$ avoiding $S - v$. Thus the constructed paths are internally disjoint and $\mathcal{Q}' := \{Q' \colon Q \in \mathcal{Q}\} \cup \{P_{uv}\}$ satisfies (i) to (iii) and witnesses via uv that $\mathcal{Q}$ does not satisfy (iv), a contradiction.

Choosing I to be $E^+(\mathcal{Q})$ is 'almost' suitable. Indeed, (II) and (III) would be satisfied as well as 'half' of (I). We shall define I as a superset of $E^+(\mathcal{Q})$ guaranteeing the 'other half' of (I), namely that $S_u \setminus V_\alpha$ remains linked to u in $L^{\omega\alpha+n} \upharpoonright_v I$ for every $u \in V_\alpha - r$. Note that for $T := T_{L^{\omega\alpha+n},v}$ (see the definition in Subsection 2.3.5) we have $T \cap V_\alpha = \emptyset$ because in $L^{\omega\alpha+n}$ the vertices in $V_\alpha - r$ have no outgoing edges and T is a minimal separation. We are going to choose I in such a way that T remains linked to v in $L^{\omega\alpha+n} \upharpoonright_v I$. By Lemma 7.23 this ensures that $S_u \setminus V_\alpha$ remains linked to u in $L^{\omega\alpha+n} \upharpoonright_v I$ for every $u \in V_\alpha - r$. Let $\mathcal{P} \in \mathfrak{P}_{L^{\omega\alpha+n}}(v, T)$ and let $\mathcal{P}'$ consists of the terminal segments of the paths in $\mathcal{P}$ from T. For $Q \in \mathcal{Q}$, let Q' be the terminal segment of Q from the last common vertex with $T \cup V_\alpha$. Corollary 7.11 applied in digraph $L_{\omega\alpha+n}$ with vertex set $T \cup V_\alpha$ and vertex v together with path-systems $\mathcal{P}'$ and $\mathcal{Q}'$ provides a system $\mathcal{R}'$ (see Figure 7.3) of $T \cup V_\alpha$–v paths such that $V(R_0) \cap V(R_1) \subseteq \{r, v\}$ for every distinct $R_0, R_1 \in \mathcal{R}'$,

$$I := E^+(\mathcal{R}') \supseteq E^+(\mathcal{Q}') = E^+(\mathcal{Q})$$

and for every $t \in T$ there is a unique $R_t \in \mathcal{R}'$ with first vertex t. We claim that paths R_t lie completely in the subdigraph $L^{\omega\alpha+n}$ of $L_{\omega\alpha+n}$. This is true by definition for $R_t \in \mathcal{P}'$. If this is not the case, then R_t consists of the initial segment of some $P \in \mathcal{P}'$ up to some v_R and the terminal segment of some $Q \in \mathcal{Q}'$ from v_R. If Q itself lies in $L^{\omega\alpha+n}$ then we are done again. If this is not the case, then Q is a terminal segment of a path in $\mathcal{Q}_F$ (see (iii)). Clearly $v_R \notin V_\alpha$ because no path in $\mathcal{P}'$ meets V_α. But then the terminal segment of Q from v_R lies entirely in $L^{\omega\alpha+n}$ because the paths in $\mathcal{Q}_F$ never return to V_α once they left it. Therefore R_t lies in $L^{\omega\alpha+n}$ in all possible cases. Thus the paths $\{R_t \colon t \in T\}$ link T to v in $L^{\omega\alpha+n}$.

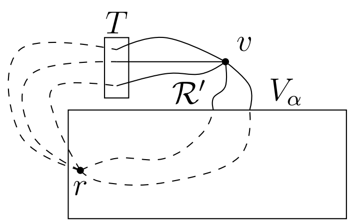

Figure 7.3: The path-system $\mathcal{R}$.

We extend the paths in $\mathcal{R}'$ backwards to obtain a path-system $\mathcal{R}$ witnessing $I \in \mathcal{G}_{L_{\omega\alpha+n}}(v)$. For $t \in T$, we take the initial segment of the unique $P_t \in \mathcal{P}$ through t until t. These extended paths meet V_α only at r because in $L^{\omega\alpha+n}$ the vertices in $V_\alpha - r$ have no outgoing edges. There are only finitely many paths R in $\mathcal{R}'$ whose first vertex is in V_α, each of which is a terminal segment of a path $Q_R \in \mathcal{Q}$. By property (iii) these Q_R are in $\mathcal{Q}_F$. As a backward extension of such an R we choose simply Q_R itself. The new initial segments in this case have vertices only in V_α and therefore meet the initial segments added to paths R_t only at r. Thus the resulting $\mathcal{R}$ really is internally disjoint which completes the proof of Claim 7.33. ■

Chapter 8

Directed Separations

In this chapter we introduce *torsoids*, a canonical structure in matching covered graphs, corresponding to the bricks and braces of the graph. This allows a more fine-grained understanding of the structure of finite and infinite directed graphs with respect to their 1-separations. We begin by proving a number of foundational results about matching covered graphs that will be used throughout the chapter, sometimes even implicitly.

First we establish a number of elementary results, structured into Section 8.1 and Section 8.2. Then we introduce torsoids in Section 8.3. We investigate how they interact with tight sets in Section 8.4 and with particular tight cut contractions in Section 8.5.

8.1 Basic Facts about Matching covered Graphs

Proposition 8.1. *The tight sets in a matching covered bipartite graph G with respect to a perfect matching M correspond one-to-one to the 1-separations in $D(G, M)$.*

Proof. Let X be a tight set in G. Then M has exactly one edge e with one endpoint in X and the other in $V(G) \setminus X$. Let $X' \subseteq D := D(G, M)$ be the set corresponding to the edges of M lying in $X \setminus e$ and let v_e be the vertex obtained by contracting the edge e. We show that $(X' \cup \{v_e\}, V(D) \setminus X')$ is a 1-separation.

Without loss of generality assume the unique vertex in $e \cap X$ is in V_1. Then there is no edge between a vertex in $X \cap V_0$ and $V(G) \setminus X$, for the following reason. Suppose there was such an edge, then, as G is matching covered, there is a matching $M' \in \mathcal{M}(G, M)$ containing it. We construct a path in X alternating between the two matchings M and M' starting with this edge. Then every edge from M' ends in a vertex of V_0 and thus every edge of M we add starts in a vertex of V_0. Therefore we never add e, and never close a cycle. But this implies that the symmetric difference between M and M' is infinite, a contradiction. Thus all edges between X and $V(G) \setminus X$ have their endpoint in X within the partition class V_1. Therefore there are no edges from $(V(D) \setminus X') \setminus (X' \cup \{v_e\})$ to $(X' \cup \{v_e\}) \setminus (V(D) \setminus X')$.

Let (A, B) be a 1-separation in D with separation vertex v such that there are no edges from $B \setminus A$ to $A \setminus B$. Let $v_0 \in V_0$ and $v_1 \in V_1$ be the two vertices such that the edge $(v_0, v_1) \in M$ gets contracted to v.

Consider the set X obtained by taking all vertices in G that are contracted to a vertex in $A \setminus B$ together with the vertex v_0. We claim that X is tight. Suppose towards a contradiction that it is not. First, we consider the case that there is a matching $M' \in \mathcal{M}(G, M)$ having no edge with one endpoint in X and one in $V(G) \setminus X$. Then consider the component C of the symmetric difference between M and M'. As v_0 has no neighbour in X it is matched to by M and M' has no edge leaving X, C cannot be a cycle, this contradicts, by Observation 2.4, that $M' \in \mathcal{M}(G, M)$.

Second, we consider the case that there is a matching $M' \in \mathcal{M}(G, M)$ with more than one edge having exactly one endpoint in X. For every edge of $M' \setminus \{(v_0, v_1)\}$ having exactly one endpoint in X consider the component of the symmetric difference between M and M' containing it. As at most one of the edges can be contained in a cycle with (v_0, v_1), there is at least one that is contained in an infinite path, contradicting that $M' \in \mathcal{M}(G, M)$. □

Lemma 8.2. *Let $\mathcal{P}$ be a tight set partition of a matching covered graph G, then $\mathsf{coll}(\mathcal{P})$ is matching covered.*

Proof. Let M be the perfect matching of G with respect to which G is matching covered. For every perfect matching $M' \in \mathcal{M}(G, M)$ we define

$$M'(\mathcal{P}) := \{\{P_i, P_j\} \colon \text{there exist } x_i \in P_i \text{ and } x_j \in P_j \text{ with } \{x_i, x_j\} \in M'\}.$$

We first show that this yields a perfect matching in the collapse.

Claim 1. For every $M' \in \mathcal{M}(G, M)$ the obtained set of edges $M'(\mathcal{P})$ is a perfect matching of $\mathsf{coll}(\mathcal{P})$.

Proof. As $P \in \mathcal{P}$ is tight $|M' \cap \partial_G(P)| = 1$. Thus, every $P \in V(\mathsf{coll}(\mathcal{P}))$ lies in exactly one edge of $M'(\mathcal{P})$. ■

Next we show that the collapse is matching covered.

Claim 2. The collapse $\mathsf{coll}(\mathcal{P})$ is matching covered with respect to $M(\mathcal{P})$.

Proof. Let $\{P_i, P_j\} \in E(\mathsf{coll}(\mathcal{P}))$. By definition of collapses, there are vertices $x_i \in P_i$ and $x_j \in P_j$ with $\{x_i, x_j\} \in E(G)$. So there is a matching $M' \in \mathcal{M}(G, M)$ with $\{x_i, x_j\} \in M'$. Because the symmetric difference between M and M' is finite, so is the symmetric difference between $M'(\mathcal{P})$ and $M(\mathcal{P})$. Thus we obtain $M'(\mathcal{P}) \in \mathcal{M}(\mathsf{coll}(\mathcal{P}), M(\mathcal{P}))$. ■

Together Claim 1 and Claim 2 imply the statement. □

Lemma 8.3. *Let $\mathcal{P}$ be a tight set partition of a matching covered graph G and let X be a subset of $\mathcal{P}$. Then X is tight in $\mathsf{coll}(\mathcal{P})$ if and only if $\bigcup X$ is tight in G.*

Proof. Let M be the perfect matching of G with respect to which G is matching covered. By Lemma 8.2, this implies that $\mathsf{coll}(\mathcal{P})$ is matching covered with respect to the matching $M(\mathcal{P})$.

Suppose first of all that X is tight in $\mathsf{coll}(\mathcal{P})$. Then for any $M' \in \mathcal{M}(G, M)$ the matching $M'(\mathcal{P})$ has exactly one edge in $\partial_{\mathsf{coll}(\mathcal{P})}(X)$. Thus, there is exactly one set $P \in X$ and exactly one $Q \in V(\mathsf{coll}(\mathcal{P})) \setminus X$ with $PQ \in E(\mathsf{coll}(\mathcal{P})) \cap M'(\mathcal{P})$. Thus there are vertices $x \in P$ and $y \in Q$ with $xy \in E(G) \cap M'$. As P and Q are tight sets, there is exactly one such edge, thus $|M' \cap \partial_G(\bigcup X)| = 1$.

Conversely, suppose that $\bigcup X$ is tight in G, and consider any $M' \in \mathcal{M}(\mathsf{coll}(\mathcal{P}), M(\mathcal{P}))$. Let the finitely many edges in $M' \setminus M(\mathcal{P})$ be P_1Q_1, P_2Q_2, $\ldots P_nQ_n$. For $i \leq n$ let e_i be any edge of G with one endpoint in P_i and the other in Q_i, let M_i be any element of $\mathcal{M}(G, M)$ containing e_i and let N_i be the set of edges in M_i with both endpoints in $P_i \cup Q_i$. Let $N := \bigcup_{i=1}^n N_i$. Finally, let M'' be obtained from M by removing all edges with an endpoint in any P_i or Q_i and adding the edges in N.

By construction M'' is a perfect matching of G whose symmetric difference with M is a subset of $\bigcup_{i=1}^n (M \triangle M_i)$ and so is finite. So $M'' \in \mathcal{M}(G, M)$, and so there is a unique edge in $M'' \cap \partial(\bigcup X)$. Since by construction also $M''(\mathcal{P}) = M'$, this implies that there is a unique edge in $M' \cap \partial(X)$. Since M' was arbitrary, this implies that X is tight in $\mathsf{coll}(\mathcal{P})$, as required. $\square$

8.1.1 Even and Odd Sets in Infinite Graphs

Let G be a matching covered graph with respect to a perfect matching M. A set $X \subseteq V(G)$ has a *parity* if $|\partial_G(X) \cap M|$ is finite. For a set X with a parity such that $|\partial_G(X) \cap M|$ is even, we say that X is *even* and if $|\partial_G(X) \cap M|$ is odd, we call X *odd*.

Lemma 8.4. *Let G be a matching covered graph with respect to a perfect matching M. If $X \subseteq V(G)$ is odd, then every perfect matching $M' \in \mathcal{M}(G, M)$ contains a finite, odd number of edges with exactly one endpoint in X. Similarly, if $X \subseteq V(G)$ is even, then every perfect matching $M' \in \mathcal{M}(G, M)$ contains a finite, even number of edges with exactly one endpoint in X.*

Proof. Let $M' \in \mathcal{M}(G, M)$ and $X \subseteq V(G)$. As the symmetric difference between M and M' is finite, M' has finitely many edges with exactly one endpoint in X if and only if M does.

Consider the case that both have finitely many edges with exactly one endpoint in X. We split the set of edges $\partial(X) \cap (M \cup M')$ in two disjoint parts: the ones lying in M-M'-alternating cycles and the ones lying in $M \cap M'$. As every cycle has an even number of edges in $\partial(X)$, it either contains an even number of edges from both matchings or an odd number of edges from both matchings. Thus, M and M' have the same parity of edges lying in M-M'-alternating cycles. Clearly, the parity of edges from both matchings lying in $M \cap M'$ is also the same. Thus, the intersections $\partial(X) \cap M$ and $\partial(X) \cap M'$ have the same parity. $\square$

In the following lemma we make some simple observations about even and odd sets and show that these definitions behave intuitively. We use these properties throughout the chapter, often without explicit reference to this lemma.

Lemma 8.5. *Let G be a matching covered graph with respect to a perfect matching M and X and Y two subsets of $V(G)$.*

- **(P1)** *If X and Y are disjoint and both odd, then $X \cup Y$ is even.*
- **(P2)** *If X and Y are disjoint and X is even while Y is odd, then $X \cup Y$ is odd.*
- **(P3)** *If X and Y are disjoint and both even, then $X \cup Y$ is even.*
- **(P4)** *If X and Y both have a parity, then $X \cap Y$ has a parity.*
- **(P5)** *If X and Y both have a parity, then $X \setminus Y$ has a parity.*

Proof. We prove the points separately.

(P1) Assume that X and Y are disjoint and both odd. We can partition the edges in $\partial(X \cup Y) \cap M$ into three parts: edges with one endpoint in X and the other in Y, remaining edges with one endpoint in X and remaining edges with one endpoint in Y.

If there is an even number of edges in M with one endpoint in X and the other in Y, then the remaining two sets are both odd and thus $\partial(X \cup Y) \cap M$ is even. If there is an odd number of edges in M with one endpoint in X and the other in Y, then the remaining two sets are both even and thus $\partial(X \cup Y) \cap M$ is even.

(P2), (P3) Can be proved with arguments akin to (P1).

(P4) Consider any $e \in M$ with exactly one endpoint in $X \cap Y$. If $e \subset X$, then e has exactly one endpoint in Y, if $e \subset Y$, then e has exactly one endpoint in X and if $e \not\subset X$ and $e \not\subset Y$, then its other endpoint is in $V(G) \setminus (X \cup Y)$. Thus, in every case e lies in at least one of $M \cap \partial(X)$ or $M \cap \partial(Y)$ and therefore, there are only finitely many such edges.

(P5) Consider any $e \in M$ with exactly one endpoint in $X \setminus Y$. If $e \subset X$, then e has exactly one endpoint in Y, thus e lies in $M \cap \partial(Y)$. If $e \not\subset X$, then e lies in $M \cap \partial(X)$. Therefore there are only finitely many such edges. $\square$

8.2 Crossing Tight Sets and Tight Set Partitions

For the notion of torsoids, which we formally introduce in Section 8.3, we have to understand how tight sets interact with each other. We begin with the investigation of crossing tight sets in Subsection 8.2.1 and based on this, we define and study passable sets in Subsection 8.2.3. Thereafter, we explore how tight set partitions interact with tight sets in Subsection 8.2.4 and introduce a relation of tight set partitions that we call *correspondence* in Subsection 8.2.5.

8.2.1 Basic Facts about Crossing Tight Sets

For this subsection we fix a matching covered graph G. The first few results in this subsection were already known for finite matching covered graphs [ELP82, LP09, Lov87], and their proofs for infinite matching covered graphs do not contain any new ideas. However, we still give these proofs here.

Lemma 8.6. *Let X and X' be tight sets with $X \cap X'$ odd. Then both $X \cap X'$ and $X \cup X'$ are tight sets and there is no edge with endpoints in both $X \setminus X'$ and $X' \setminus X$. If $X \setminus X' \neq \emptyset$ then there is an edge with endpoints in $X \cap X'$ and $X \setminus X'$.*

Proof. First we show that $X \cap X'$ is tight. For any matching M,

$$|M \cap \partial(X \cap X')| \leq |M \cap \partial(X)| + |M \cap \partial(X')| = 2\,,$$

but since $X \cap X'$ is odd, $|M \cap \partial(X \cap X')|$ must also be odd and so must be equal to one. An identical argument shows that $X \cup X'$ is tight.

Next, suppose for a contradiction that there is an edge e with endpoints in $X \setminus X'$ and $X' \setminus X$. Let M be any matching containing e, and let f be the unique edge in $M \cap \partial(X \cap X')$. As $f \in \partial(X)$ and $\partial(X)$ contains both e and f, contradicting tightness of X. Thus there can be no such an edge e.

Finally, suppose that $X \setminus X'$ is non-empty. By connectivity of G there must be some edge $e \in \partial(X \setminus X')$. Let M be a matching containing e. Then $|M \cap \partial(X \setminus X')|$ is a positive even number, so it is at least 2. Every edge in $M \cap \partial(X \setminus X')$ must be in one of $M \cap \partial(X)$ and $M \cap \partial(X')$, and each of those sets contains only one edge, so one of the edges of $M \cap \partial(X \setminus X')$ must be in $M \cap \partial(X')$. Then that edge has one endpoint in $X \setminus X'$ and the other in X'. We have just seen that the endpoint in X' cannot be in $X' \setminus X$, and so it must be in $X \cap X'$. □

Lemma 8.7. *Let X, X' and X'' be tight sets such that $X \cap X'$ and $X \cap X''$ are even and disjoint. Then there is no edge between these two sets.*

Proof. Applying Lemma 8.6 to X and the complement of X', we see that $X \setminus X'$ is tight. But then applying the same lemma to $X \setminus X'$ and the complement of X'' we get the desired result. □

Lemma 8.8. *Let X, X' and X'' be tight sets such that $X \cap X' = X \cap X'' = X' \cap X''$ is even. Then $X \cap X' = X \cap X'' = X' \cap X''$ is empty.*

Proof. Suppose not for a contradiction. By Lemma 8.6 applied to X and the complement of X' there is an edge e with one end in $X \setminus X'$ and the other in $X \cap X' = X' \cap X''$. But this contradicts Lemma 8.6 applied to X' and the complement of X''. □

Lemma 8.9. *Let $(X_i)_{i \in I}$ be a family of tight sets such that the union of any two of them is also tight. Then $X := \bigcup_{i \in I} X_i$ is either also tight or is the whole vertex set.*

Proof. Suppose for a contradiction that there is a matching M containing two edges e and f in the boundary of X. Choose $i, j \in I$ such that the endpoints of e and f in X are in X_i and X_j respectively. Then both e and f are also in the boundary of $X_i \cup X_j$, contradicting the tightness of this set.

Thus, any matching M contains at most one edge in $\partial(X)$. If X is not the whole vertex set then there is a matching containing some edge, and so exactly one edge, in the boundary of X. Thus X is odd, and any matching contains exactly one edge in the boundary of X as required. □

Applying this to the complements of a family of tight sets also gives the following statement.

Corollary 8.10. *Let $(X_i)_{i \in I}$ be a family of tight sets such that the intersection of any two of them is odd. Then $X := \bigcap_{i \in I} X_i$ is either also tight or else is empty.* □

Lemma 8.11. *Let X be a tight set and $(X_i)_{i \in I}$ a family of tight sets whose intersections with X are even and disjoint. Then $X \setminus \bigcup_{i \in I} X_i$ is tight.*

Proof. By repeated applications of Lemma 8.6, $(X \setminus X_i)_{i \in I}$ is a family of tight sets and any intersection of two of its elements is tight. So by applying Corollary 8.10 to this family we see that $X \setminus \bigcup_{i \in I} X_i$ is either tight or empty.

Suppose for a contradiction that it is empty. Let M be a perfect matching, let e be the edge of M in the boundary of X and let $i \in I$ be chosen such that the endpoint of e in X is in X_i. Since $X \cap X_i$ is even there must be some other edge $f \in M$ in the boundary of X_i, and the endpoint of f outside X_i must lie in some X_j with $j \neq i$. But this contradicts Lemma 8.7. □

8.2.2 Characterising Cycles

Lemma 8.12. *Let G be a matching covered graph with at least six vertices and with a cyclic order on its vertex set such that any set of three consecutive vertices is tight. Then G is a cycle.*

Proof. Suppose for a contradiction that there is an edge e joining vertices v and w which are not adjacent in the cyclic order. Let M be a matching containing e. Let x

and y be the neighbours of v in the cyclic order and let x' and y' be the neighbours of x and y in the cyclic order other than v. Since G has at least 6 vertices, $x' \neq y'$ and so without loss of generality $w \neq x'$. Since e is in the boundary of $\{x, v, y\}$, no other edge of M can be, so the edge in M incident to x can only be xy. But then both e and xy are in the boundary of the tight set $\{x', x, v\}$, giving the desired contradiction. $\square$

For a cycle C we call a non-empty proper subset I of $V(C)$ an *interval* if $C[I]$ is connected.

Remark 8.13. Let C be a matching covered cycle. A set $X \subseteq V(C)$ is tight in C if and only if it is an interval in C with odd cardinality.

Lemma 8.14. *Let G be a matching covered graph and let $X = \{p_1, p_2, p_3\}$ be a tight set such that collapsing X in G*[1] *gives rise to a graph H whose vertices can be ordered cyclically such that any set of three consecutive vertices is tight. Let the neighbours of X in H be q and q'. Suppose further that both $\{q, p_1, p_2\}$ and $\{p_2, p_3, q'\}$ are tight. Then G is a cycle.*

Proof. Let r and r' be the neighbours of q and q' other than X after collapsing X. Note that since collapsing X in G gives rise to a cycle, G has at least six vertices and $r, r' \notin \{p_1, p_2, p_3, q, q'\}$. By Lemma 8.12 together with Remark 8.13 it suffices to show that $\{r, q, p_1\}$ and $\{p_3, q', r'\}$ are tight, and by symmetry it suffices to show that the first of these is. But this follows by applying Lemma 8.6 to $\{r, q, p_1, p_2, p_3\}$ and the complement of $\{p_2, p_3, q'\}$. $\square$

8.2.3 Passable Sets

We now need a notion capturing when we can make small modifications to tight set partitions by moving some elements from one partition class to another. As in Subsection 8.2.1 we fix a matching covered graph G for this subsection.

Definition 8.15. A set S is *passable* for P if $P \setminus S$ and $P \cup S$ are tight. It is *passable between P and Q* if $S \subseteq P \cup Q$ and S is passable for both.

Remark 8.16. Note that any passable set is even, since both $P \setminus S$ and $P \cup S$ are tight and therefore odd.

Lemma 8.17. *Let S and S' be sets which are passable between disjoint tight sets P and Q whose union is not the whole vertex set. Then $P \cap S \cap S'$ is even.*

Proof. Suppose not for a contradiction. We may replace S with $S \cap P$ if necessary, since it is passable between P and Q as well. Thus we may assume without loss of generality that S is a subset of P. Similarly we may assume that S' is also a subset of P. By

[1]this means that we take the collapse of the tight set partition consisting of X and the singletons of all vertices not in X.

assumption $P \cap S \cap S' = S \cap S'$ is odd, which implies that $S \setminus S'$ and $S' \setminus S$ are also odd, since S and S' are passable.

The set $P \setminus S$ is tight since S is passable for P. Let A be $V(G) \setminus (P \cup Q)$. Now applying Lemma 8.6 again to $P \setminus S$ and $Q \cup S'$ we see that their union $(P \cup Q) \setminus (S \setminus S')$ is tight, hence so is its complement $A \cup (S \setminus S')$. Similarly $A \cup (S' \setminus S)$ is tight. Finally, $A \cup Q$ is tight since it is the complement of P. Applying Lemma 8.8 to these three tight sets yields the desired contradiction, since by assumption A is non-empty. □

Lemma 8.18. *Let S and S' be sets which are passable between disjoint tight sets P and Q whose union is not the whole vertex set. Then $S \cup S'$ is also passable between P and Q.*

Proof. By symmetry it suffices to show that $S \cup S'$ is passable for P. We know that $P \cap S$ and $P \cap S'$ are even since $P \setminus S$ and $P \setminus S'$ are odd, and $P \cap S \cap S'$ is even by Lemma 8.17. $P \cap (S \cup S')$ is even and therefore $P \setminus (S \cup S')$ is odd. So applying Lemma 8.6 to $P \setminus S$ and $P \setminus S'$ shows that $P \setminus (S \cup S')$ is tight.

A symmetric argument shows that $Q \cap (S \cup S')$ is also even and so $P \cup (S \cap S')$ is odd. Thus, applying Lemma 8.6 to $P \cup S$ and $P \cup S'$ shows that $P \cup S \cup S'$ is tight. Thus, $S \cup S'$ is passable for P, as required. □

Lemma 8.19. *Let P and Q be disjoint tight sets whose union is not the whole vertex set. Let $\mathcal{S}$ be any set of sets which are passable between P and Q. Then $S := \bigcup \mathcal{S}$ is itself passable between P and Q. In particular, if $\mathcal{S}$ is the set of all sets which are passable between P and Q then S is the largest such set.*

Proof. By symmetry it suffices to show that S is passable for P. $(P \cup T)_{T \in \mathcal{S}}$ is a family of tight sets and by Lemma 8.18 the union of any two of them is also tight, so by Lemma 8.9 the union of the whole family is tight. A similar argument shows that $Q \cup S$ is also tight, and so $P \setminus S = P \setminus (Q \cup S)$ is tight by Lemma 8.6 applied to the complement of P and to $Q \cup S$. □

Lemma 8.20. *Let P be a tight set, let S be a passable set for P and let $(S_i)_{i \in I}$ be a family of passable sets for P such that S and all of the S_i are disjoint and such that $P \cup S \cup \bigcup_{i \in I} S_i$ is not the whole vertex set. Suppose further that all sets of the form $(P \cup S) \setminus S_i$ or $(P \cup S_i) \setminus S$ are tight. Then S is passable for both $P \cup \bigcup_{i \in I} S_i$ and $P \setminus \bigcup_{i \in I} S_i$.*

Proof. The set $(P \cup \bigcup_{i \in I} S_i) \cup S$ is tight by Lemma 8.9 applied to the sets $X_i := (P \cup S_i) \cup S$, which are all tight by Lemma 8.6. Similarly, the set $(P \cup \bigcup_{i \in I} S_i) \setminus S$ is tight by Lemma 8.9 applied to the sets $X_i := (P \cup S_i) \setminus S$. This shows that S is passable for $P \cup \bigcup_{i \in I} S_i$.

The set $(P \setminus \bigcup_{i \in I} S_i) \cup S$ is tight by Lemma 8.11 applied to $P \cup S$ and to the X_i given by the complements of the tight sets $(P \cup S) \setminus S_i$. Similarly $(P \setminus \bigcup_{i \in I} S_i) \setminus S$ is tight by Lemma 8.11 applied to $P \setminus S$ with the same choice of X_i. This shows that S is passable for $P \setminus \bigcup_{i \in I} S_i$. □

8.2.4 The Interaction between Tight Sets and Tight Set Partitions

Definition 8.21. Let G be a matching covered graph. Let $\mathcal{P}$ be a tight set partition of G and let X be a tight set in G. We denote by $\mathsf{odd}_{\mathcal{P}}(X)$ the set of all elements of $\mathcal{P}$ whose intersection with X is odd.

For the rest of this subsection, fix a matching covered graph G.

Lemma 8.22. *Let $\mathcal{P}$ be a tight set partition of G and let X be a tight set. Then $\bigcup \mathsf{odd}_{\mathcal{P}}(X)$ is a tight set. Furthermore, $\mathsf{odd}_{\mathcal{P}}(X) \notin \{\emptyset, \mathcal{P}\}$.*

Proof. Let Y be the complement of X. Then by Lemma 8.11 the set $Y \setminus \bigcup \mathsf{odd}_{\mathcal{P}}(X)$ is tight, and therefore so is its complement $X \cup \bigcup \mathsf{odd}_{\mathcal{P}}(X)$. Applying Lemma 8.11 again, we get that

$$(X \cup \bigcup \mathsf{odd}_{\mathcal{P}}(X)) \setminus \bigcup (\mathcal{P} \setminus \mathsf{odd}_{\mathcal{P}}(X))$$

is tight, but this set is just $\bigcup \mathsf{odd}_{\mathcal{P}}(X)$, giving the desired result.

Since tight sets are neither empty nor the whole vertex set, and $\bigcup \mathsf{odd}_{\mathcal{P}}(X)$ is one, $\mathsf{odd}_{\mathcal{P}}(X)$ is neither empty nor $\mathcal{P}$. $\square$

Lemma 8.23. *Let $\mathcal{P}$ be a tight set partition of G such that $\mathsf{coll}(\mathcal{P})$ is a BoB, and let X be a tight set. Then $\mathsf{odd}_{\mathcal{P}}(X)$ is either a singleton or the complement of $\mathsf{odd}_{\mathcal{P}}(X)$ in $\mathcal{P}$ is a singleton.*

Proof. By Lemma 8.22 the set $\bigcup \mathsf{odd}_{\mathcal{P}}(X)$ is tight, meaning that $\mathsf{odd}_{\mathcal{P}}(X)$ is a tight set in $\mathsf{coll}(\mathcal{P})$ by Lemma 8.3. Since $\mathsf{coll}(\mathcal{P})$ is a BoB, the result follows. $\square$

Lemma 8.24. *Let $\mathcal{P}$ be a tight set partition of G such that $\mathsf{coll}(\mathcal{P})$ is a cycle, and let X be a tight set in G. Let $n := |\mathcal{P}|$ and $m := |\mathsf{odd}_{\mathcal{P}}(X)|$.*

Then there is a cyclic enumeration $P_1, \ldots, P_n$ of the vertices in $\mathsf{coll}(\mathcal{P})$ such that $P_i \in \mathsf{odd}_{\mathcal{P}}(X)$ if and only if $i \in [m]$. Furthermore,

$$\bigcup_{1<i<m} P_i \subset X \subset V(G) \setminus \bigcup_{m+1<i<n} P_i$$

holds. If $3 \leq m \leq n-3$ holds, then $P_1 \setminus X, P_n \cap X$ are passable between P_1 and P_n, and $P_m \setminus X, P_{m+1} \cap X$ are passable between P_m and P_{m+1}.

Proof. The set $\bigcup \mathsf{odd}_{\mathcal{P}}(X)$ is tight by Lemma 8.22 and thus $\mathsf{odd}_{\mathcal{P}}(X)$ is tight in $\mathsf{coll}(\mathcal{P})$ by Lemma 8.3. By Remark 8.13, $\mathsf{odd}_{\mathcal{P}}(X)$ is an interval in $\mathsf{coll}(\mathcal{P})$. Thus we can choose the desired enumeration of the vertices in $\mathsf{coll}(\mathcal{P})$.

Let $1 < i < m$ be arbitrary. We prove that $P_i \subset X$. Suppose for a contradiction that $P_i \setminus X \neq \emptyset$. As $P_i \in \mathsf{odd}_{\mathcal{P}}(X)$, $P_i \cap X$ is odd, which implies that $(V(G) \setminus X) \cap (V(G) \setminus P_i)$ is odd. Applying Lemma 8.6 to $V(G) \setminus X$ and $V(G) \setminus P_i$ shows that there is an edge e with endpoints in $(V(G) \setminus X) \setminus (V(G) \setminus P_i) = P_i \setminus X$ and $(V(G) \setminus X) \cap (V(G) \setminus P_i) =$

$V(G) \setminus (X \cup P_i)$. As $\mathsf{coll}(\mathcal{P})$ is cyclic, there is $j \in \{i-1, i+1\}$ such that e has an endpoint in P_j. The set $X' := X \cup P_j$ is a tight set by Lemma 8.6, as $P_j \in \mathsf{odd}_{\mathcal{P}}(X)$. The edge e is an edge with endpoints in $P_i \setminus X'$ and $X' \setminus P_i$. This contradicts Lemma 8.6, as $P_i \cap X' = P_i \cap X$ is odd. Therefore, $P_i \subset X$. Applying the same argument to the complement of X shows that $P_i \cap X = \emptyset$ for all $m+1 < i < n$. Thus the desired subset relation is true.

From now on we suppose that $3 \leq m \leq n-3$ holds. We show that $P_n \cap X$ is passable between P_1 and P_n. By considering the complement of X and/or reversing the enumeration one can show that the other sets are passable as well. $P_n \cup (X \cap P_n) = P_n$ and $P_1 \setminus (X \cap P_n) = P_1$ are tight by definition. $P_n \cap X$ is even since $P_n \notin \mathsf{odd}_{\mathcal{P}}(X)$, thus $P_n \setminus X = (V(G) \setminus X) \cap P_n$ is odd and thus tight by Lemma 8.6. To see that $P_1 \cup (X \cap P_n)$ is tight we first note that by Remark 8.13, the set $P_1 \cup P_n \cup P_{n-1}$ is tight. Consider the set $X' := X \cup P_1$, which is tight by Lemma 8.6 since $P_1 \in \mathsf{odd}_{\mathcal{P}}(X)$. By assumption, $m < n-1$ and thus $X' \cap (P_1 \cup P_n \cup P_{n-1}) = P_1 \cup (X \cap P_n)$ is odd and therefore tight. This concludes the proof that $X \cap P_n$ is passable between P_1 and P_n. $\square$

8.2.5 Correspondences between Tight Set Partitions

Throughout this subsection we work with a fixed matching covered graph G.

Definition 8.25. Let $\mathcal{P}$ and $\mathcal{Q}$ be tight set partitions of G. A *correspondence between* $\mathcal{P}$ *and* $\mathcal{Q}$ is a bijection $\rho\colon \mathcal{P} \to \mathcal{Q}$ such that for any $P \in \mathcal{P}$ and $Q \in \mathcal{Q}$ we have that $P \cap Q$ is odd if and only if $Q = \rho(P)$. If there is such a ρ then we say that $\mathcal{P}$ *is in correspondence with* $\mathcal{Q}$.

This relation is fundamental for our later constructions - roughly, two tight set partitions arising from different tight set decompositions encode 'the same' BoB precisely when they are in correspondence.

To begin to make this more precise, we first note some basic properties of correspondences of tight set partitions.

Lemma 8.26. *Let* $\rho\colon \mathcal{P} \to \mathcal{Q}$ *be a correspondence of tight set partitions of* G *and let* P *and* P' *be distinct elements of* $\mathcal{P}$. *Then* $P \cap \rho(P')$ *is passable from* P *to* P'. *If it is non-empty then* P *and* P' *are neighbours in* $\mathsf{coll}(\mathcal{P})$.

Proof. To show passability, we have to show that the following four sets are tight:

$$P \setminus (P \cap \rho(P')),\ P \cup (P \cap \rho(P')),\ P' \setminus (P \cap \rho(P')) \text{ and } P' \cup (P \cap \rho(P'))$$

The second and third of these are just P and P'. The first is $P \setminus \rho(P')$, which is tight, since $\rho(P') \cap P$ is even by definition of ρ, thus $P \cap (V(G) \setminus \rho(P'))$ is odd and then $P \cap (V(G) \setminus \rho(P')) = P \setminus \rho(P')$ is tight by Lemma 8.6. To show that the fourth is tight, we first note that it is equal to

$$(P' \cup \rho(P')) \setminus \bigcup_{P'' \in \mathcal{P} \setminus \{P, P'\}} P'',$$

which is tight by Lemma 8.11 since $P' \cup \rho(P')$ is tight by Lemma 8.6, since $P' \cap \rho(P')$ is odd by the definition of ρ.

Now suppose that $P \cap \rho(P')$ is non-empty. Applying Lemma 8.6 to $P' \cup (P \cap \rho(P'))$ and the complement of P we see that there is an edge from $P \cap \rho(P')$ to P', which witnesses that P and P' are neighbours in $\mathsf{coll}(\mathcal{P})$. $\square$

Lemma 8.27. *Let $\rho\colon \mathcal{P} \to \mathcal{Q}$ be a correspondence of tight set partitions of G. Then ρ is also a graph isomorphism from $\mathsf{coll}(\mathcal{P})$ to $\mathsf{coll}(\mathcal{Q})$.*

Proof. First we show that ρ^{-1} is a graph homomorphism. Let e be an edge joining $\rho(P)$ to $\rho(P')$ in $\mathsf{coll}(\mathcal{Q})$. Let v be the endpoint of e in $\rho(P)$ and w the endpoint in $\rho(P')$.

Suppose for a contradiction that v is in an element P'' of $\mathcal{P}$ other than P and P'. Applying Lemma 8.7 to P'', $\rho(P)$ and $\rho(P')$, we see that w cannot also be in P''. But then we get the desired contradiction by applying Lemma 8.6 to P'' and the complement of $\rho(P)$.

Thus v must lie in $P \cup P'$, and a symmetric argument shows that w is also in $P \cup P'$. If $w \in P$ then $P \cap \rho(P')$ is non-empty and so by Lemma 8.26 P and P' are neighbours in $\mathsf{coll}(\mathcal{P})$. The same result follows if v is in P', since then $P' \cap \rho(P)$ is non-empty. If neither of these are applicable then v is in P and w is in P', so e itself witnesses that P and P' are neighbours in $\mathsf{coll}(\mathcal{P})$. So in any case they are neighbours, completing the proof that ρ^{-1} is a graph homomorphism.

Since ρ^{-1} is also a correspondence from $\mathcal{Q}$ to $\mathcal{P}$, the same argument applied to ρ^{-1} shows that ρ is also a graph homomorphism, and thus a graph isomorphism. $\square$

8.3 Torsoids

The relation of being in correspondence is not an equivalence relation, as illustrated in Figure 8.1. However, our next aim is to show that its restriction to some important classes of tight set partitions is an equivalence relation, and to introduce some canonical objects, called torsoids, displaying the equivalence classes.

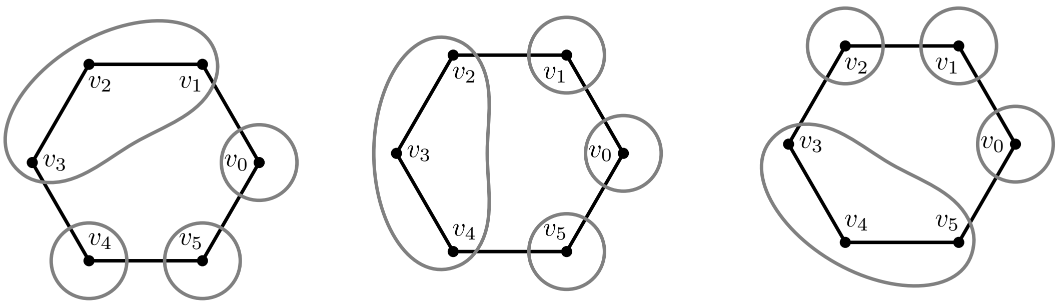

Figure 8.1: The tight set partitions on the left and the right are both in correspondence with the one in the middle. But they are not in correspondence with each other. Thus, being in correspondence is not transitive.

The example in Figure 8.1 shows that being in correspondence is not even an equivalence relation when restricted to tight set partitions whose collapses are BoBs, since the collapses of the tight set partitions in the figure are all cycles of length 4. However, the cycle of length 4 is the only BoB which is problematic in this sense; we will show that being in correspondence is an equivalence relation on all other tight set partitions that collapse to BoBs.

Given two partitions $\mathcal{P}, \mathcal{P}'$ of some set, we say $\mathcal{P}$ *refines* $\mathcal{P}'$ if every partition class of $\mathcal{P}$ is subset of a partition class of $\mathcal{P}'$. The problem with the tight set partitions in Figure 8.1 is that they can be refined to larger cycles. More precisely, we say that a tight set partition $\mathcal{P}$ is *cyclic* if $\mathsf{coll}(\mathcal{P})$ is a cycle, and *maximal cyclic* if it is cyclic but cannot be refined to a finer cyclic tight set partition. In the first part of this section we show that being in correspondence is an equivalence relation on maximal cyclic tight set partitions. We call a tight set partition that either is maximal cyclic or whose collapse is a BoB other than C_4 *torsoid-inducing*. The torsoid-inducing tight set partitions, as their name suggests, play a central role in the theory that is built throughout this chapter.

For the rest of the chapter let G be a fixed matching covered graph. As a preliminary step, we show that this class of tight set partitions is closed under correspondence. First we establish a special case.

Lemma 8.28. *Let $\mathcal{P}$ be a maximal cyclic tight set partition of G and let P and Q be consecutive vertices on the cycle $\mathsf{coll}(\mathcal{P})$. Let $S \subseteq P$ be a passable set between P and Q. Then $\mathcal{P}' := (\mathcal{P} \setminus \{P, Q\}) \cup \{P \setminus S, Q \cup S\}$ is also maximal cyclic.*

Proof. Suppose not for a contradiction. It is certainly cyclic by Lemma 8.27, so it cannot be maximal. Choose a proper cyclic refinement $\mathcal{Q}'$ with as few vertices as possible. We will construct a tight set partition $\mathcal{Q}$ refining $\mathcal{P}$ which is in correspondence with $\mathcal{Q}'$ and therefore also cyclic by Lemma 8.27. This contradicts the fact that $\mathcal{P}$ is maximal cyclic, giving the desired contradiction.

Since all elements of $\mathcal{Q}$ and $\mathcal{Q}'$ are odd sets, the minimality of $\mathcal{Q}'$ implies $|\mathcal{Q}'| = |\mathcal{Q}| + 2$. Thus, $\mathcal{Q}'$ must be of the form $(\mathcal{P}' \setminus \{X\}) \cup \{X_1, X_2, X_3\}$ for some $X \in \mathcal{P}'$. If X is not equal to either of $P \setminus S$ or $Q \cup S$ then we can simply set

$$\mathcal{Q} := (\mathcal{Q}' \setminus \{P \setminus S, Q \cup S\}) \cup \{P, Q\}.$$

Thus X is one of $P \setminus S$ or $Q \cup S$. Suppose first that it is $P \setminus S$. Let R be the neighbour of Q other than P in $\mathsf{coll}(\mathcal{P})$, and suppose without loss of generality that X_3 is a neighbour of $Q \cup S$ in $\mathcal{Q}'$. Then $X_3 \cup S$ is tight by Lemma 8.6 applied to $X_3 \cup (Q \cup S) \cup R$ and P, so we can set

$$\mathcal{Q} := (\mathcal{Q}' \setminus \{X_3, Q \cup S\}) \cup \{X_3 \cup S, Q\}.$$

Now assume instead that X is $Q \cup S$. Once again, let R be the neighbour of Q other than P in $\mathsf{coll}(\mathcal{P})$. Suppose without loss of generality that X_1 is a neighbour of $P \setminus S$

and X_3 is a neighbour of R in $\mathcal{Q}'$. Let e be an edge with one endpoint v in X_3 and the other endpoint in R. By Lemma 8.6 applied to $V(G) \setminus P$ and $Q \cup S$, v cannot lie in $X_3 \cap S$ and so must lie in $X_3 \cap Q$. If $X_3 \setminus Q$ is odd then applying Lemma 8.6 to X_3 and the complement of Q contradicts the existence of e. So $X_3 \setminus Q$ must be even, meaning that $X_3 \cap Q$ is odd.

Since $(X_1 \cap Q) \cup (X_2 \cap Q) \cup (X_3 \cap Q) = Q$ is odd, it follows that $(X_1 \cap Q) \cup (X_2 \cap Q)$ is even and so $X_1 \cap Q$ and $X_2 \cap Q$ have the same parity. There are now two cases, depending on whether both are odd or both are even.

If both are odd we can note that all sets $X_i \cap Q$ are tight by Lemma 8.6 and so we can set

$$\mathcal{Q} := (\mathcal{Q}' \setminus \{P \setminus S, X_1, X_2, X_3\}) \cup \{P, X_1 \cap Q, X_2 \cap Q, X_3 \cap Q\}\,.$$

If both are even then $X_1 \cap S$ is tight by Lemma 8.6 applied to X_1 and P, and $S \setminus X_1$ is tight by Lemma 8.6 applied to P and $X_2 \cup X_3 \cup R$. So we can set

$$\mathcal{Q} := (\mathcal{Q}' \setminus \{X_1, X_2, X_3\}) \cup \{X_1 \cap S, S \setminus X_1, Q\}\,. \qquad \square$$

Lemma 8.29. *Let $\mathcal{P}$ and $\mathcal{Q}$ be tight set partitions of G such that there is a correspondence $\rho\colon \mathcal{P} \to \mathcal{Q}$. If $\mathcal{P}$ is maximal cyclic then so is $\mathcal{Q}$.*

Proof. Since $\mathcal{P}$ is cyclic, we can enumerate its elements as P_1, P_2, $\dots P_n$ for some n. For any i, let $Q_i := \rho(P_i)$. We will argue by induction on $|\{(i,j) \in [n]^2 \colon i \neq j \text{ but } P_i \cap Q_j \neq \emptyset\}|$. The base case is that this set is empty. In that case for any $i \leq n$ we must have $P_i \cap Q_j = \emptyset$ for all $j \neq i$ and so $P_i \subseteq Q_i$. Similarly $Q_i \cap P_j = \emptyset$ for all $j \neq i$ and so $Q_i \subseteq P_i$. Thus in this case $P_i = Q_i$ for all $i \leq n$ and so $\mathcal{Q} = \mathcal{P}$, giving the desired result.

For the induction step, let $(i,j) \in [n]^2$ be such that $i \neq j$ but $P_i \cap Q_j \neq \emptyset$. By Lemma 8.26 the set $S := P_i \cap Q_j$ is passable from P_i to P_j, and P_i and P_j are consecutive on $\mathsf{coll}(\mathcal{P})$. We set $P_i' := P_i \setminus S$, $P_j' := P_j \cup S$ and $P_k' := P_k$ for all other k. Let $\mathcal{P}' := \{P_i' \colon i \leq n\}$. By Lemma 8.28, $\mathcal{P}'$ is also maximal cyclic. For any $(i', j') \in [n]^2$ other than (i,j) we have $P_{i'}' \cap Q_{j'} = P_{i'} \cap Q_{j'}$, whereas $P_i' \cap Q_j = \emptyset$. So we are done by applying the induction hypothesis to $\mathcal{P}'$, $\mathcal{Q}$ and ρ'. $\square$

Now we return to the construction of the objects displaying the correspondence equivalence classes. The construction relies on the following key lemma.

Lemma 8.30. *Let $\mathcal{P}$ be a torsoid-inducing tight set partition of G. Let P, Q and Q' be distinct elements of $\mathcal{P}$. Let S be passable between P and Q and let S' be passable between P and Q'. Then $S \cap S' = \emptyset$.*

Proof. First we suppose for a contradiction that $S \cap S'$ is odd. We begin by establishing some notation. Let $P_1 := P \setminus S'$, $P_2 := S \cap S'$ and $P_3 := (P \cap S') \setminus S$. Thus P_1, P_2 and P_3 partition P. All three of these are tight sets. P_1 is tight by passability of S'. P_2 is tight by Lemma 8.6 applied to $Q \cup S$ and $Q' \cup S'$. Finally, $P \setminus S$ is tight because S is passable for P, and P_3 is tight by Lemma 8.6 applied to $P \setminus S$ and the complement of P_1.

We know that $P = P_1 \cup P_2 \cup P_3$ is tight. We can show that also $Q \cup P_1 \cup P_2$ is tight by applying Lemma 8.6 to $Q \cup S$ and P_1. Finally, we know that $P_2 \cup P_3 \cup Q' = Q' \cup S'$ is tight.

There are now three cases:

Case 1: $\mathsf{coll}(\mathcal{P})$ is a BoB with more than four vertices. Applying Lemma 8.6 to $Q \cup S$ and $Q' \cup S'$ shows that the set $Q \cup S \cup Q' \cup S'$ is tight. But $\mathsf{odd}_{\mathcal{P}}(Q \cup S \cup Q' \cup S') = \{P, Q, Q'\}$ is neither a singleton nor the complement of a singleton in $\mathcal{P}$, contradicting Lemma 8.23.

Case 2: $\mathsf{coll}(\mathcal{P})$ is a BoB with four vertices but is not a cycle. In this case, applying Lemma 8.14 to $\mathsf{coll}((\mathcal{P} \setminus \{P\}) \cup \{P_1, P_2, P_3\})$ shows that this graph is a cycle of length 6, and hence $\mathsf{coll}(\mathcal{P})$ is a cycle of length 4, contradicting our assumptions.

Case 3: $\mathcal{P}$ is maximal cyclic. Applying Lemma 8.14 to $\mathsf{coll}((\mathcal{P} \setminus \{P\}) \cup \{P_1, P_2, P_3\})$ shows that this graph is a cycle, contradicting maximality of $\mathcal{P}$.

Since we reached a contradiction in all three cases, our assumption that $S \cap S'$ is odd must have been false. So $S \cap S'$ is even.

By Lemma 8.6 applied to $Q \cup S$ and the complement of $Q' \cup S'$, we know that $Q \cup (S \setminus S')$ is tight. Similarly $Q' \cup (S' \setminus S)$ is tight. Applying Lemma 8.11 to P, $Q \cup (S \setminus S')$ and $Q' \cup (S' \setminus S)$ yields that $P' := (P \setminus (S \cup S')) \cup (S \cap S')$ is tight. Applying Lemma 8.8 to $Q \cup S$, $Q' \cup S'$ and P', we see that their intersection $S \cap S'$ is empty as required. $\square$

This allows us to define canonical objects which determine the correspondence equivalence classes for such tight set partitions.

Definition 8.31. A *torsoid* (H, ε) in a matching covered graph G is a pair (H, ε) with:

- **(T1)** H is a matching covered graph on at least 4 vertices that is a BoB or a cycle,
- **(T2)** the elements of $V(H)$ are tight sets of G,
- **(T3)** $\varepsilon : E(H) \to 2^{V(G)}$,
- **(T4)** $V(H) \cup \mathsf{im}(\varepsilon)$ is a near partition of $V(G)$, where $\mathsf{im}(\varepsilon)$ is the image of ε,
- **(T5)** for $vw \in E(H)$, there is an edge from $v \cup \varepsilon(vw)$ to w and $\varepsilon(vw)$ is largest among all subsets of $v \cup \varepsilon(vw) \cup w$ that are passable for both v and w,
- **(T6)** for $vw \notin E(H)$, there is no edge from v to w in G, and

(T7) if H is a cycle and v is a vertex of H with neighbours u and w then there is no partition of $\varepsilon(uv) \cup v \cup \varepsilon(vw)$ into tight sets P_1, P_2, P_3 such that both $u \cup P_1 \cup P_2$ and $P_2 \cup P_3 \cup w$ are tight.

We call $\mathcal{T}$ *cyclic* if H is a cycle and *non-cyclic* otherwise.

A *correspondence* between $\mathcal{T}$ and a tight set partition $\mathcal{P}$ is a bijection $\sigma \colon V(H) \to \mathcal{P}$ such that for every $v \in V(H)$ we have

$$v \subseteq \sigma(v) \subseteq v \cup \bigcup_{w \in N_H(v)} \varepsilon(vw) \, .$$

We call σ *strong* if for every vertex v of H the set $\sigma(v)$ is a union of v and some of the sets $\varepsilon(vw)$ with w a neighbour of v in H.

We say $\mathcal{P}$ is in (strong) correspondence with $\mathcal{T}$ if there is such a (strong) correspondence σ.

There is another slightly different perspective on the set of tight set partitions in strong correspondence with $\mathcal{T}$. For any graph H, we say a function $\kappa \colon E(H) \to V(H)$ is a *choice function* for H if $\kappa(e)$ is an endvertex of e for all edges e of H. For any pair $\mathcal{T} = (H, \varepsilon)$ of a graph and a function defined on its edge set and any choice function κ for H, we take $\mathcal{P}(\mathcal{T}, \kappa)$ to be the set of all sets of the form $v \cup \bigcup_{\kappa(e)=v} \varepsilon(e)$.

Lemma 8.32. *For any torsoid $\mathcal{T} = (H, \varepsilon)$ in G and any choice function κ for H, the set $\mathcal{P}(\mathcal{T}, \kappa)$ is a tight set partition which is in strong correspondence with $\mathcal{T}$. Furthermore, all tight set partitions in strong correspondence with $\mathcal{T}$ arise in this way.*

Proof. All elements of $\mathcal{P}(\mathcal{T}, \kappa)$ are tight by (T5) and Lemma 8.9, and they partition $V(G)$ by (T4). The function σ sending each element of $\mathcal{P}(\mathcal{T}, \kappa)$ to the unique vertex of H which it includes is a strong correspondence.

For the last part, given a strong correspondence σ between $\mathcal{T}$ and a tight set partition $\mathcal{P}$, let κ be the function sending each edge e of H to the unique vertex v with $\varepsilon(e) \subseteq \sigma(v)$. Then $\mathcal{P} = \mathcal{P}(\mathcal{T}, \kappa)$. □

Proposition 8.33. *Let $\mathcal{T} = (H, \varepsilon)$ be a torsoid in G and $uv \in E(H)$. Any edge in G with an endvertex in $\varepsilon(uv)$ has both endvertices in $u \cup v \cup \varepsilon(uv)$.*

Proof. This follows by applying Lemma 8.6 to $u \cup \varepsilon(uv)$ and the complement of $\varepsilon(uv) \cup v$. □

Lemma 8.34. *If σ is a strong correspondence between a torsoid $\mathcal{T} = (H, \varepsilon)$ in G and a tight set partition $\mathcal{P}$ of G then σ is a graph isomorphism from H to $\mathsf{coll}(\mathcal{P})$.*

Proof. This is immediate from (T5), (T6) and Proposition 8.33. □

We can now show that any tight set partition of the kinds discussed above induces a torsoid.

Definition 8.35. Let $\mathcal{P}$ be a torsoid-inducing tight set partition of G. For any edge PQ of $\mathsf{coll}(\mathcal{P})$ we define $\delta_{\mathcal{P}}(PQ)$ to be the largest passable set between P and Q (this exists by Lemma 8.19). For $P \in \mathcal{P}$, we define $\tau_{\mathcal{P}}(P)$ to be $P \setminus \bigcup_{Q \in N_{\mathsf{coll}(\mathcal{P})}(P)} \delta(PQ)$.

Let $H_{\mathcal{P}}$ be the unique graph on the image of $\tau_{\mathcal{P}}$ making $\tau_{\mathcal{P}}$ a graph isomorphism from $\mathsf{coll}(\mathcal{P})$ to $H_{\mathcal{P}}$. Let $\sigma_{\mathcal{P}}$ be the inverse of $\tau_{\mathcal{P}}$. Let $\varepsilon_{\mathcal{P}} \colon E(H) \to 2^{V(G)}; vw \mapsto \delta_{\mathcal{P}}(\sigma_{\mathcal{P}}(v)\sigma_{\mathcal{P}}(w))$. Then we call $\mathcal{T}_{\mathcal{P}} := (H_{\mathcal{P}}, \varepsilon_{\mathcal{P}})$ the *induced torsoid* of $\mathcal{P}$.

Theorem 8.36. *In the context of Definition 8.35, $\mathcal{T}_{\mathcal{P}}$ is a torsoid and $\sigma_{\mathcal{P}}$ is a correspondence.*

Proof. Conditions (T1), (T3) and (T6) are clear from the construction. (T2) follows from Lemma 8.30 and Lemma 8.11. (T4) follows from Lemma 8.30.

For (T5), first we note that passability of $\varepsilon_{\mathcal{P}}(vw)$ follows from Lemma 8.6 together with Lemma 8.20. Similarly any set $S \subseteq v \cup \varepsilon_{\mathcal{P}}(vw) \cup w$ which is passable for both v and w is also passable between $\sigma_{\mathcal{P}}(v)$ and $\sigma_{\mathcal{P}}(w)$ by Lemma 8.20, hence $\varepsilon_{\mathcal{P}}(vw)$ is maximal among such sets. To see that there is an edge from $v \cup \varepsilon_{\mathcal{P}}(vw)$ to w, let e be any edge from $\sigma_{\mathcal{P}}(v)$ to $\sigma_{\mathcal{P}}(w)$ in G. For any $u \in N_{H_{\mathcal{P}}}(v) \setminus \{w\}$, the endpoint of e in $\sigma_{\mathcal{P}}(v)$ cannot be in $\varepsilon_{\mathcal{P}}(uv)$ by Lemma 8.6 applied to $u \cup \varepsilon_{\mathcal{P}}(uv)$ and the complement of $\varepsilon_{\mathcal{P}}(uv) \cup v$. So it must be in $v \cup \varepsilon_{\mathcal{P}}(vw)$. Similarly the other endpoint of e must be in $\varepsilon_{\mathcal{P}}(vw) \cup w$. If either of these endpoints is in $\varepsilon_{\mathcal{P}}(vw)$ then we get the desired edge by Lemma 8.6 applied to $v \cup \varepsilon_{\mathcal{P}}(vw)$ and the complement of $\varepsilon_{\mathcal{P}}(vw) \cup w$. Otherwise we can take e itself as the desired edge.

For (T7), suppose $H_{\mathcal{P}}$ is a cycle and let κ be a choice function for $H_{\mathcal{P}}$ such that $\kappa^{-1}(u)$ and $\kappa^{-1}(w)$ are empty. Then $\mathcal{P}(\mathcal{T}, \kappa)$ is a tight set partition by (T5) and Lemma 8.6 applied to the sets of the form $\kappa(e) \cup \varepsilon(e)$, and is in correspondence with $\mathcal{P}$ by construction. By (T5) and (T6), $\mathsf{coll}(\mathcal{P}(\mathcal{T}, \kappa))$ is a cycle, and hence so is $\mathcal{P}$ by Lemma 8.27. So $\mathcal{P}$ is cyclic, and by assumption it must be maximal cyclic. Thus $\mathcal{P}(\mathcal{T}, \kappa)$ is also maximal cyclic by Lemma 8.29. But sets P_1, P_2 and P_3 as in (T7) would contradict maximality due to Lemma 8.14. So they cannot exist.

Finally, it follows directly from the construction that $\sigma_{\mathcal{P}}$ is a correspondence from $\mathcal{T}_{\mathcal{P}}$ to $\mathcal{P}$. □

Theorem 8.37. *Let $\mathcal{P}$ and $\mathcal{Q}$ be torsoid-inducing tight set partitions of G. Then $\mathcal{P}$ and $\mathcal{Q}$ are in correspondence if and only if $\mathcal{T}_{\mathcal{P}} = \mathcal{T}_{\mathcal{Q}}$. In particular, being in correspondence is an equivalence relation for such tight set partitions.*

Proof. Suppose first that there is a correspondence $\rho \colon \mathcal{P} \to \mathcal{Q}$. Then, by Lemma 8.26, any set of the form $P \cap \rho(Q)$ with $P \neq Q$ is passable between P and Q. Using Lemma 8.20 twice, we may therefore show first that $\delta_{\mathcal{P}}(PQ)$ is passable for $P \cap \rho(P) = P \setminus \bigcup_{R \in N_{\mathsf{coll}(\mathcal{P})}(P) \setminus \{Q\}} (P \cap \rho(R))$ and then that it is passable for $\rho(P)$. A similar argument shows that it is passable for $\rho(Q)$, so we have $\delta_{\mathcal{P}}(PQ) \subseteq \delta_{\mathcal{Q}}(\rho(P)\rho(Q))$. A similar argument applied to ρ^{-1} implies the reverse inclusion, so we have $\delta_{\mathcal{P}}(PQ) = \delta_{\mathcal{Q}}(\rho(P)\rho(Q))$.

Now let x be any element of $\tau_{\mathcal{P}}(P)$ for any $P \in \mathcal{P}$. Let Q be the element of $\mathcal{P}$ with $x \in \rho(Q)$. Then Q cannot be different from P, since then it would have to be a neighbour of P in $\mathsf{coll}(\mathcal{P})$ by Lemma 8.26, so we would have $x \in P \cap \rho(Q) \subseteq \delta_{\mathcal{P}}(PQ)$, contradicting the definition of $\tau_{\mathcal{P}}(P)$. Thus $x \in \rho(P)$. By construction, for any neighbour $\rho(Q)$ of $\rho(P)$ in $\mathsf{coll}(\mathcal{Q})$ we have $x \notin \delta_{\mathcal{P}}(PQ) = \delta_{\mathcal{Q}}(\rho(P)\rho(Q))$. Thus $x \in \tau_{\mathcal{Q}}(\rho(P))$. This argument shows that $\tau_{\mathcal{P}}(P) \subseteq \tau_{\mathcal{Q}}(\rho(P))$. A similar argument applied to ρ^{-1} implies the reverse inclusion, so we have $\tau_{\mathcal{P}}(P) = \tau_{\mathcal{Q}}(\rho(P))$ for any P.

This immediately implies that $H_{\mathcal{P}}$ and $H_{\mathcal{Q}}$ are equal. Finally, for any edge vw of $H_{\mathcal{P}}$ with $v = \tau_{\mathcal{P}}(P)$ and $w = \tau_{\mathcal{P}}(Q)$, we have

$$\varepsilon_{\mathcal{P}}(vw) = \delta_{\mathcal{P}}(PQ) = \delta_{\mathcal{Q}}(\rho(P)\rho(Q)) = \varepsilon_{\mathcal{Q}}(vw)\,.$$

This completes the proof that $\mathcal{T}_{\mathcal{P}} = \mathcal{T}_{\mathcal{Q}}$.

For the reverse direction, suppose that $\mathcal{T}_{\mathcal{P}} = \mathcal{T}_{\mathcal{Q}}$, and let this torsoid be given by (H, ε). Let $\rho := \sigma_{\mathcal{Q}} \cdot \tau_{\mathcal{P}}$. Let PQ be any edge of $\mathsf{coll}(\mathcal{P})$. Set $v := \tau_{\mathcal{P}}(P)$ and $w := \tau_{\mathcal{P}}(Q)$. Then by Lemma 8.11 applied to P and the sets $\delta_{\mathcal{P}}(P, Q')$ for all other neighbours Q' of P, the set $v \cup (P \cap \varepsilon(vw))$ is tight. Similarly the set $w \cup (Q \cap \varepsilon(vw))$ is tight. Thus $P \cap \varepsilon(vw)$ is passable between v and $\varepsilon(vw) \cup w$. A similar argument shows that $\rho(P) \cap \varepsilon(vw)$ is also passable between these sets. So by Lemma 8.17 their intersection $P \cap \rho(P) \cap \varepsilon(vw)$ is even. Since $P \cap \varepsilon(vw)$ is even, it follows that $(P \cap \varepsilon(vw)) \setminus \rho(P) = P \cap \rho(Q) \cap \varepsilon(vw)$ is also even. Since by construction any element of $P \cap \rho(Q)$ must lie in $\varepsilon(vw)$, this implies that $P \cap \rho(Q)$ is even.

Taking stock of the argument so far, if $Q \neq P$ is a neighbour of P then $P \cap \rho(Q)$ is even. But if Q is not a neighbour of P then by construction $P \cap \rho(Q)$ is empty and so also even. Thus the set $\mathcal{Q}_P$ of all elements of $\mathcal{Q}$ whose intersection with P is odd is either empty or equal to $\{\rho(P)\}$. By Lemma 8.22 its union is tight, so it cannot be empty. Thus it is $\{\rho(P)\}$. That is, $\rho(P)$ is the unique element of $\mathcal{Q}$ whose intersection with P is odd. Since this is true for all $P \in \mathcal{P}$, ρ is a correspondence from $\mathcal{P}$ to $\mathcal{Q}$. $\square$

8.4 Relation of Tight Sets to Torsoids

In this section we investigate how torsoids interact with tight cuts and tight sets. We introduce three classes of tight cuts with respect to a given torsoid and prove that these three classes partition the set of tight cuts. This partition induces a partition of the tight sets.

For a torsoid $\mathcal{T} = (H, \varepsilon)$ in G and a tight set $X \subseteq V(G)$ we define $\mathsf{odd}_H(X)$ to be the set of vertices of H that have odd intersection with X. For a tight cut $C = \partial(X)$ we define $\theta_{\mathcal{T}}(C) := \min\{|\mathsf{odd}_H(X)|, |\mathsf{odd}_H(V(G) \setminus X)|\}$.

The constant $\theta_{\mathcal{T}}(C)$ determines to which of the three classes a tight cut C belongs: C resides at an edge ($\theta_{\mathcal{T}}(C) = 0$), resides at a vertex ($\theta_{\mathcal{T}}(C) = 1$) or resides at an interval ($\theta_{\mathcal{T}}(C) > 1$). Note that as $\mathsf{odd}_H(V(G) \setminus X) = V(H) \setminus \mathsf{odd}_H(X)$ holds, the constant $\theta_{\mathcal{T}}(C)$ is bounded by $\frac{|V(H)|}{2}$.

Lemma 8.38. *Let $\mathcal{T} = (H, \varepsilon)$ be a torsoid in G, $a \in V(H)$ and X a tight set in G. If $X \cap a$ is even then $X \setminus a$ is tight and if $X \cap a$ is odd then $X \cup a$ is tight.*

Proof. First suppose that $X \cap a$ is even. Then $V(G) \setminus X$ and a have odd intersection thus $(V(G) \setminus X) \cup a$ is tight by Lemma 8.6 and so is $V(G) \setminus ((V(G) \setminus X) \cup a) = X \setminus a$. Now suppose that $X \cap a$ is odd. Then $X \cup a$ is tight by Lemma 8.6. □

Lemma 8.39. *Let $\mathcal{T} = (H, \varepsilon)$ be a torsoid in G, $ab \in E(H)$ and X a tight set in G. If $X \cap (a \cup \varepsilon(ab))$ is odd, $X \cup a \cup \varepsilon(ab)$ is tight. If $X \cap (a \cup \varepsilon(ab))$ is even, $X \setminus (a \cup \varepsilon(ab))$ is tight.*

Proof. If $X \cap (a \cup \varepsilon(ab))$ is odd, the tight sets X and $a \cup \varepsilon(ab)$ intersect oddly. Thus $X \cup a \cup \varepsilon(ab)$ is tight by Lemma 8.6. If $X \cap (a \cup \varepsilon(ab))$ is even, apply the prior result to the complement of X. □

Definition 8.40. Let $\mathcal{T} = (H, \varepsilon)$ be a torsoid in G and let e be an edge of H. Then we say that a tight cut C *resides at an edge e in $\mathcal{T}$* if there is a tight set $X \subseteq \varepsilon(e)$ such that $C = \partial(X)$ and we call X a *$\mathcal{T}$-edge-resident at e*.

See Figure 8.3 for an example of a BoB with an edge-resident.

Lemma 8.41. *Let $\mathcal{T} = (H, \varepsilon)$ be a torsoid in G. Let $X \subseteq V(G)$ be a tight set and $uv \in E(H)$ such that $X \cap \varepsilon(uv)$ is odd. Then $\partial(X)$ resides at the edge uv in $\mathcal{T}$.*

Proof. Let κ_1 be any choice function such that $\kappa_1(uv) = u$, and $\kappa_1(e) \notin \{u, v\}$ for any $uv \neq e \in E(H)$. Let κ_2 be the choice function which agrees with κ_1 except at uv, where $\kappa_2(uv) = v$. Consider $\mathcal{P}^i := \mathcal{P}(\mathcal{T}, \kappa_i)$ for $i \in [2]$. We set $u_1 := u \cup \varepsilon(uv), v_1 := v$ and $u_2 := u, v_2 := v \cup \varepsilon(uv)$. Note that $u_i, v_i \in \mathcal{P}^i$ for $i \in [2]$. Furthermore, $\mathcal{P}^1 \setminus \{u_1, v_1\} = \mathcal{P}^2 \setminus \{u_2, v_2\}$.

Let us first show that X intersects either both u, v evenly or both oddly. Without loss of generality we suppose for a contradiction that X intersects u oddly and v evenly. There are now three cases.

Case 1: H is a BoB with more than four vertices. In this case $|\mathsf{odd}_{\mathcal{P}^1}(X)|$ and $|\mathsf{odd}_{\mathcal{P}^2}(X)|$ are elements of $\{1, |V(H)| - 1\}$ by Lemma 8.23. By construction, $|\mathsf{odd}_{\mathcal{P}^1}(X)| + 2 = |\mathsf{odd}_{\mathcal{P}^2}(X)|$ holds and this gives a contradiction since $|V(H)| \geq 6$.

Case 2: H is a BoB with four vertices but is not a cycle. Note that K^4 is the only BoB on four vertices other than a cycle. At first, we assume that there exists $s \in V(H) \setminus \{u, v\}$ such that $Y_s := s \cup \varepsilon(us) \cup u \cup (X \cap \varepsilon(uv))$ is tight and show that this gives a contradiction. Afterwards, we show that such a vertex s exists. Under these assumptions let t be the vertex of H distinct from u, v, s and consider the sets $Z_1 := v \cup \varepsilon(vs)$ and $Z_2 := t \cup \bigcup_{w \in E(H) \setminus \{t\}} \varepsilon(wt)$, $Z_3 := s \cup \varepsilon(us)$, $Z_4 := u$, $Z_5 := X \cap \varepsilon(uv)$, $Z_6 := \varepsilon(uv) \setminus X$ in this cyclic order. The sets Z_1, Z_2, Z_3, Z_4 are tight by definition. The set Z_5 is tight by applying Lemma 8.6 to $X \setminus v$ and

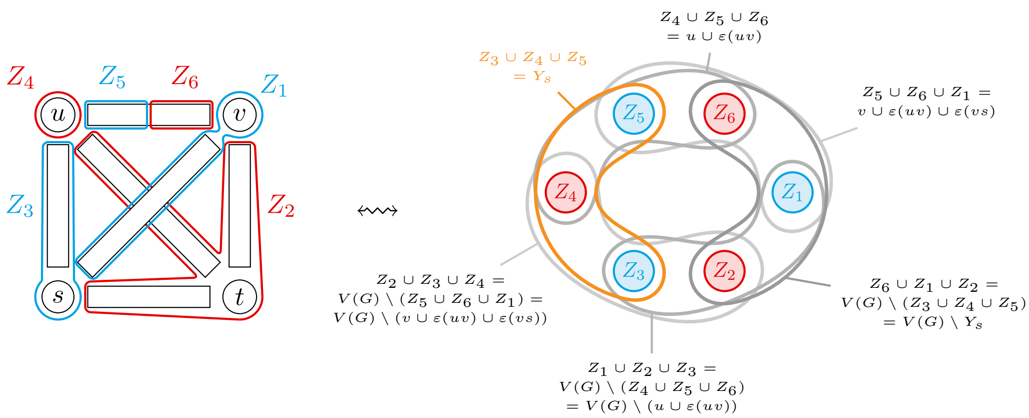

Figure 8.2: In the context of Lemma 8.41 if H is a K^4 and Y_s is a tight set as depicted, we deduce that the edge ut of H cannot exist for a contradiction.

$v \cup \varepsilon(uv)$. The set Z_6 is tight by applying Lemma 8.6 to $V(G) \setminus (X \cup u)$ and $u \cup \varepsilon(uv)$.

Furthermore, $Z_3 \cup Z_4 \cup Z_5 = Y_s$ is tight by assumption. Also $Z_4 \cup Z_5 \cup Z_6 = u \cup \varepsilon(uv)$ and $Z_5 \cup Z_6 \cup Z_1 = v \cup \varepsilon(uv) \cup \varepsilon(vs)$ are tight. As $Z_6 \cup Z_1 \cup Z_2$, $Z_1 \cup Z_2 \cup Z_3$ and $Z_2 \cup Z_3 \cup Z_4$ are complements of the first three sets respectively, they are also tight. Thus these sets satisfy the conditions of Lemma 8.12 and therefore $\mathsf{coll}(\{Z_i \colon i \in [6]\})$ is a cycle. This contradicts that ut is an edge of H.

It remains to prove that there exists $s \in V(H) \setminus \{u, v\}$ such that Y_s is indeed tight. The set $X' := X \cup \bigcup \mathsf{odd}_H(X) \setminus \bigcup (V(H) \setminus \mathsf{odd}_H(X))$ is tight by repeated application of Lemma 8.38 to X. By construction, X' has the property that any vertex $v \in V(H)$ is either contained in X' or disjoint from X'. Next we apply Lemma 8.39 repeatedly to X' and any edge $ab \in E(H)$ with $a \in \{s, t\}$ to obtain a tight set X'' such that $X' \cap (u \cup v \cup \varepsilon(uv)) = X'' \cap (u \cup v \cup \varepsilon(uv))$ and for any edge $e \in E(H) \setminus \{uv\}$ we have $X'' \cap \varepsilon(e) \in \{\emptyset, \varepsilon(e)\}$. Furthermore, any vertex $v \in V(H)$ is also either contained in X'' or disjoint from X'' by construction.

Note that X'' is odd, X'' intersects precisely one edge of $\mathcal{T}$ oddly and X'' intersects u oddly and v evenly. Thus there exists exactly one vertex $s \in V(H) \setminus \{u, v\}$ such that $s \subset X''$. Therefore the tight set X'' has the following properties: $X'' \cap \varepsilon(uv) = X \cap \varepsilon(uv)$, $\mathsf{odd}_H(X'') = \{u, s\}$ and for any $e \in E(H) \setminus \{uv\}$ we have $X'' \cap \varepsilon(e) \in \{\emptyset, \varepsilon(e)\}$.

We apply Lemma 8.39 to X'', the vertex s and the edge us to obtain a tight set X''' with $X''' = X'' \cup \varepsilon(us)$. Similarly we apply Lemma 8.39 repeatedly to X''', the vertex v with the edges vs, vt and the vertex t with the edge st to obtain a tight set that coincides with Y_s, by construction. This completes this case.

Case 3: H is a cycle of length four. We find a contradiction to (T7). Let w be the neighbour of v in H distinct from u. Set $P_1 := \varepsilon(uv) \cap X$, $P_2 := \varepsilon(uv) \setminus X$ and

$P_3 := v \cup \varepsilon(vw)$. By construction, P_1, P_2, P_3 partition $\varepsilon(uv) \cup v \cup \varepsilon(vw)$. The set P_3 is a tight set as $\varepsilon(vw)$ is passable. We consider the tight set $X' := (X \cup u) \setminus v$. Now $P_1 = X' \cap (v \cup \varepsilon(uv))$ and $P_2 = (V(G) \setminus X') \cap (u \cup \varepsilon(uv))$ are tight sets. Furthermore by construction, $u \cup P_1 \cup P_2$ is a tight set. It remains to prove that $P_2 \cup P_3 \cup w$ is a tight set.

We suppose without loss of generality that $\varepsilon(vw) \cup w$ is a partition class of $\mathcal{P}_1$ (otherwise modify $\mathcal{P}_1$ slightly). By choice, $V(G) \setminus X'$ intersects u_1 oddly, but does not contain u_1. Furthermore, $V(G) \setminus X'$ intersects v_1 oddly. Applying Lemma 8.24 to $V(G) \setminus X'$ and $\mathcal{P}^1$, we see that $\mathsf{odd}_{\mathcal{P}^1}(V(G) \setminus X')$ is an interval with $u_1 = u \cup \varepsilon(uv)$ as an endpoint and containing v. Since this interval has odd length, it must also contain the element $\varepsilon(vw) \cup w$ of $\mathcal{P}^1$ containing w. By construction, $\mathsf{odd}_{\mathcal{P}^1}(V(G) \setminus X') \cap (u \cup \varepsilon(uv) \cup v \cup \varepsilon(vw) \cup w)$ is odd, and by Lemma 8.6 tight. Note that $\mathsf{odd}_{\mathcal{P}^1}(V(G) \setminus X') \cap (u \cup \varepsilon(uv) \cup v \cup \varepsilon(vw) \cup w) = P_2 \cup P_3 \cup w$.

Thus we can suppose that X intersects either both u, v oddly or both evenly. This implies that X intersects exactly one of u_i, v_i oddly for $i \in [2]$. We prove that X intersects either all elements of $\mathcal{P}^i \setminus \{u_i, v_i\}$ evenly or all these elements oddly for $i \in [2]$. If H is a BoB this follows directly from Lemma 8.23. If H is a cycle, we apply Lemma 8.24. Then $\mathsf{odd}_{\mathcal{P}^1}(X)$ and $\mathsf{odd}_{\mathcal{P}^2}(X)$ are intervals in $\mathsf{coll}(\mathcal{P}^1)$ and $\mathsf{coll}(\mathcal{P}^2)$. As X intersects u_1 oddly if and only if X intersects v_2 oddly and similarly for v_1 and u_2, X has to intersect the elements of $\mathcal{P}^i \setminus \{u_i, v_i\}$ either all oddly or all evenly.

From now on we suppose without loss of generality that X intersects all elements of $\mathcal{P}^i \setminus \{u_i, v_i\}$ evenly (otherwise consider the complement of X). Furthermore, we suppose without loss of generality that X intersects u_1 and v_2 oddly (otherwise exchange κ_1 and κ_2). Then $\mathsf{odd}_{\mathcal{P}}(X)^1 = \{u_1\}$ and $\mathsf{odd}_{\mathcal{P}}(X)^2 = \{v_2\}$ holds. We define the following two tight set partitions:

$$\tilde{\mathcal{P}}^1 := \{P \setminus X : P \in \mathcal{P}^1 \setminus \{u_1\}\} \cup \{u_1 \cup X\}, \text{ and}$$
$$\tilde{\mathcal{P}}^2 := \{P \setminus X : P \in \mathcal{P}^2 \setminus \{v_2\}\} \cup \{v_2 \cup X\}.$$

By Lemma 8.6 these are indeed tight set partitions. By construction, $\tilde{\mathcal{P}}^1$ and $\mathcal{P}^1$ correspond. Likewise $\tilde{\mathcal{P}}^2$ and $\mathcal{P}^2$ correspond. By Lemma 8.27 and Lemma 8.29, $\mathsf{coll}(\tilde{\mathcal{P}}^1)$ and $\mathsf{coll}(\tilde{\mathcal{P}}^1)$ are torsoid-inducing. Thus we can apply Theorem 8.37, which shows that, $\mathcal{T}_{\tilde{\mathcal{P}}_1} = \mathcal{T} = \mathcal{T}_{\tilde{\mathcal{P}}_2}$ holds. This implies that $u_1 \cup X \subseteq u \cup \bigcup_{w \in N_H(u)} \varepsilon(uw)$ and $v_2 \cup X \subseteq v \cup \bigcup_{w \in N_H(v)} \varepsilon(vw)$. Then

$$X \subset (u \cup \bigcup_{w \in N_H(u)} \varepsilon(uw)) \cap (v \cup \bigcup_{w \in N_H(v)} \varepsilon(vw)) = \varepsilon(uv)$$

holds. This proves that X is an edge-resident. Thus $\partial(X)$ resides at an edge of $\mathcal{T}$. □

Lemma 8.42. *Let $\mathcal{T}$ be a torsoid in G and $C \subseteq E(G)$ any tight cut. The tight cut C resides at an edge in $\mathcal{T}$ if and only if $\theta_{\mathcal{T}}(C) = 0$ holds.*

Proof. By definition, any tight cut C of G that resides at an edge of H in $\mathcal{T}$ satisfies $\theta_{\mathcal{T}}(C) = 0$. For the only if direction let C be a tight cut in G with $\theta_{\mathcal{T}}(C) = 0$. Let $X \subseteq V(G)$ such that $C = \partial(X)$. By Lemma 8.38 we can assume that $X \cap v = \emptyset$ for any $v \in V(H)$. We show that X has odd intersection with $\varepsilon(e)$ for some $e \in E(H)$. If H is a cycle, this holds true by parity of X. If H is a BoB, let $\mathcal{P}$ be any tight set partition in strong correspondence with $\mathcal{T}$. By Lemma 8.23, there is $P \in \mathcal{P}$ such that $X \cap P$ is odd. Suppose for a contradiction that X has even intersection with $\varepsilon(e)$ for all $e \in E(H)$ with $\varepsilon(e) \subset P$. By Lemma 8.11, $X \setminus P = X \setminus \bigcup_{e \in E(H)\colon \varepsilon(e) \subset P} \varepsilon(e) \cap X$ is a tight set. This contradicts the fact that $X \setminus P$ is even by choice of P. Thus there exists $e \in E(H)$ such that $X \cap \varepsilon(e)$ is odd. This implies via Lemma 8.41 that $\partial(X) = C$ resides at the edge e in $\mathcal{T}$. $\square$

Corollary 8.43. *Let $\mathcal{T} = (H, \varepsilon)$ be a torsoid in G and X a tight set with $0 \neq |\mathsf{odd}_H(X)| \neq |V(H)|$. Then X intersects $\varepsilon(e)$ evenly for any $e \in E(H)$.*

Definition 8.44. Let $\mathcal{T} = (H, \varepsilon)$ be a torsoid in G and let v be a vertex of H. Then we say that a tight cut C *resides at a vertex v in $\mathcal{T}$* if there is a tight set $X \subseteq v \cup \bigcup_{v \in e \in E(H)} \varepsilon(e)$ with $X \cap \varepsilon(e)$ even for all $v \in e \in E(H)$ and such that $C = \partial(X)$, and we call X a *$\mathcal{T}$-vertex-resident at v*. We call X a *proper $\mathcal{T}$-vertex-resident* if $v \subseteq X$. Then C *resides properly at v*.

See Figure 8.3 for an example of a BoB with a vertex-resident.

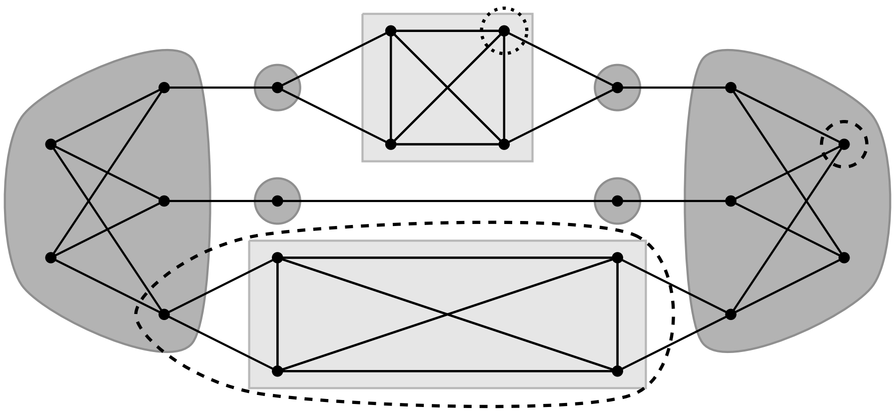

Figure 8.3: A graph with a possible 6-vertex torsoid $\mathcal{T}$ that is a BoB and some $\mathcal{T}$-residents. The vertex sets enclosed by smooth outlines shaded in darker grey represent vertices of $\mathcal{T}$ and the two sets enclosed by rectangular outlines filled in lighter grey represent edges of $\mathcal{T}$. The vertex sets enclosed by dashed lines represent $\mathcal{T}$-vertex-residents and the vertex set enclosed by a dotted line represents a $\mathcal{T}$-edge-resident.

Lemma 8.45. *Let $\mathcal{T}$ be a torsoid in G. A tight cut $C \subseteq E(G)$ resides at a vertex of $\mathcal{T}$ if and only if $\theta_{\mathcal{T}}(C) = 1$.*

Proof. By definition, any $\mathcal{T}$-vertex-resident C satisfies $\theta_{\mathcal{T}}(C) = 1$. For the only if direction let C be a tight cut in G with $\theta_{\mathcal{T}}(C) = 1$, let $X \subseteq V(G)$ be the tight set with $\partial(X) = C$ and $|V(H)_X| = 1$, and let $v \in V(H)$ be the unique vertex such that $X \cap v$ is odd. By Corollary 8.43, $X \cap \varepsilon(e)$ is even for any $e \in E(H)$. It remains to prove that X avoids $w \neq v$ and $\varepsilon(f)$ for $f \in E(H)$ with $v \notin f$. Consider an arbitrary tight set partition $\mathcal{P}$ in strong correspondence with $\mathcal{T}$ and let $Q \in \mathcal{P}$ with $v \subseteq Q$. As $\mathsf{odd}_H(X) = \{v\}$ and X intersects $\varepsilon(f)$ evenly for $f \in E(H)$, one verifies that Q is the unique element of $\mathcal{P}$ that intersects X oddly. Construct a partition $\mathcal{P}' := \{P \setminus X : v \nsubseteq P \in \mathcal{P}\} \cup \{Q \cup X\}$. The partition $\mathcal{P}'$ is indeed a tight set partition by Lemma 8.6. Furthermore, $\mathcal{P}$ and $\mathcal{P}'$ correspond. By Lemma 8.27 and Lemma 8.29, $\mathsf{coll}(\mathcal{P}')$ is also torsoid-inducing. Thus we can apply Theorem 8.37, which shows that, $\mathcal{T}_{\mathcal{P}'} = \mathcal{T}$ holds. Thus $Q \cup X \subseteq v \cup \bigcup_{v \in e} \varepsilon(e)$. This shows that X is a $\mathcal{T}$-vertex-resident, thus C resides at the vertex v in $\mathcal{T}$. □

Lemma 8.46. *Let $\mathcal{T} = (H, \varepsilon)$ be a non-cyclic torsoid in G. Any tight cut $C \subseteq E(G)$ satisfies $\theta_{\mathcal{T}}(C) \leq 1$.*

Proof. Suppose for a contradiction that C is a tight cut in G with $1 < \theta_{\mathcal{T}}(C)$. Let $X \subseteq V(G)$ be a vertex set with $\partial(X) = C$ and $\mathcal{P}$ be any tight cut partition in strong correspondence with $\mathcal{T}$. By Corollary 8.43, the intersection $X \cap \varepsilon(e)$ is even for any $e \in E(H)$. One verifies that X intersects an element $P \in \mathcal{P}$ oddly if, and only if there is $v \in \mathsf{odd}_H(X)$ with $v \subseteq P$. Thus $|\mathsf{odd}_{\mathcal{P}}(X)| = \theta_{\mathcal{T}}(C)$. This contradicts Lemma 8.23, as H is a BoB. □

The results from this section so far yield the following corollary:

Corollary 8.47. *Let $\mathcal{T}$ be a non-cyclic torsoid in G. Any tight cut of G either resides at a vertex or an edge in $\mathcal{T}$. Furthermore, for any set X that is tight in G either X or its complement is a $\mathcal{T}$-vertex-resident or a $\mathcal{T}$-edge-resident.* □

Definition 8.48. Let $\mathcal{T} = (H, \varepsilon)$ be a cyclic torsoid in G and let $I \subseteq V(H)$ be an interval in H with $3 \leq |I| \leq |V(H)| - 3$. Then we say that a tight cut C *resides at the interval I in $\mathcal{T}$* if there is a tight set $X \subseteq V(G)$ with

$$\bigcup I \cup \bigcup_{\substack{vw \in E(H):\\ v,w \in I}} \varepsilon(vw) \quad \subseteq \quad X \quad \subseteq \quad \bigcup I \cup \bigcup_{\substack{vw \in E(H):\\ v \in I}} \varepsilon(vw)$$

such that $C = \partial(X)$ and we call X a *$\mathcal{T}$-interval-resident at I*.

See Figure 8.4 for an example of a torsoid with an interval-resident. Note that by Lemma 8.24, any $\mathcal{T}$-interval-resident at an interval I intersects any set $\varepsilon(uv)$ evenly for $u \in I$ and $v \notin I$.

Lemma 8.49. *Let $\mathcal{T} = (H, \varepsilon)$ be a cyclic torsoid in G and $C \subseteq E(G)$ any tight set. The set C resides at an interval in $\mathcal{T}$ if and only if $3 \leq \theta_{\mathcal{T}}(C)$ holds.*

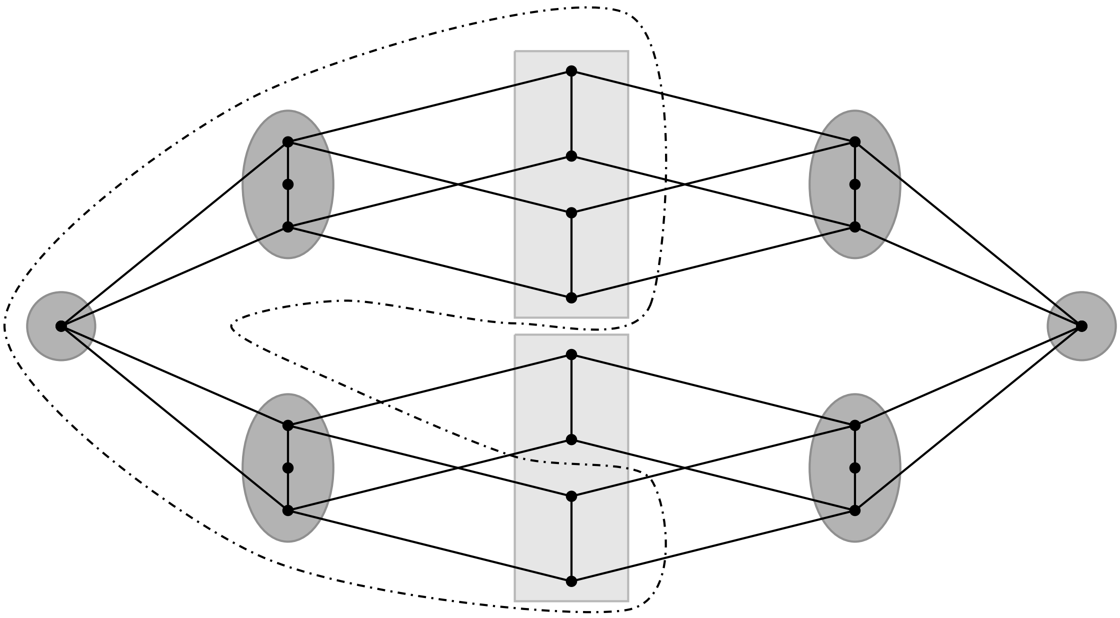

Figure 8.4: A graph with a possible torsoid $\mathcal{T}$ that is a C_6 and a $\mathcal{T}$-interval-resident. The vertex sets enclosed by smooth outlines shaded in darker grey represent vertices of the torsoid $\mathcal{T}$ and the two sets enclosed by rectangular outlines filled in lighter grey represent edges of $\mathcal{T}$. The vertex set enclosed by the dash dotted line represents a $\mathcal{T}$-interval-resident.

Proof. By definition, any tight cut C that resides at an interval in $\mathcal{T}$ satisfies $3 \leq \theta_{\mathcal{T}}(C)$. For the only if direction let C be a tight cut with $3 \leq \theta_{\mathcal{T}}(C)$, let $X \subseteq V(G)$ be the tight set with $\partial(X) = C$, $\mathsf{odd}_H(X) = \theta_{\mathcal{T}}(C)$ and let $\mathcal{P}$ be any tight set partition in strong correspondence with $\mathcal{T}$. Set $n := |\mathcal{P}| = |V(H)|$ and $m := |\mathsf{odd}_{\mathcal{P}}(X)|$. We apply Lemma 8.24 to obtain a cyclic enumeration $P_1, \ldots, P_n$ of $\mathcal{P}$ such that $\mathsf{odd}_{\mathcal{P}}(X) = \{P_1, \ldots, P_m\}$. Let $v_i \in V(H)$ with $v_i \subseteq P_i$ for $i \in [n]$. We assumed $\theta_{\mathcal{T}}(C) \neq 0$, this implies that $|\mathsf{odd}_H(X)| \notin \{0, |V(H)|\}$, thus X has even intersection with $\varepsilon(f)$ for any $f \in E(H)$ by Corollary 8.43. Thus for $i \in [n]$ we have $v_i \in V(H)_X$ if and only if $P_i \in \mathsf{odd}_{\mathcal{P}}(X)$. This implies $m = \theta_{\mathcal{T}}(C)$ and we can deduce by Lemma 8.24 that $X \setminus \bigcup \mathsf{odd}_{\mathcal{P}}(X)$ can be partitioned into S', S'' such that S' is passable between P_n, P_1 and S'' is passable between P_m, P_{m+1}. By Lemma 8.20, S' is passable between v_n and v_1. Thus it is contained in $\varepsilon(v_n v_1)$. Similarly S'' is contained in $\varepsilon(v_m, v_{m+1})$. By the same argument $\bigcup \mathsf{odd}_{\mathcal{P}}(X) \setminus X$ can be partitioned into passable sets between v_n, v_1 and v_m, v_{m+1}. This completes the proof. □

Note that for any tight cut C the number $\theta_{\mathcal{T}}(C)$ is either 0 or odd.

Corollary 8.50. *Let $\mathcal{T}$ be a cyclic torsoid in G. Any tight cut of G either resides at a vertex, an edge or an interval in $\mathcal{T}$. Furthermore, for any tight set X of G either X or its complement is a $\mathcal{T}$-vertex-resident, a $\mathcal{T}$-edge-resident, or a $\mathcal{T}$-interval-resident.* □

8.4.1 Correspondences between Torsoids and Tight Set Partitions

Using the results of the last subsection, we can show that not only strong correspondences, but arbitrary correspondences between torsoids and tight set partitions have some good properties.

Lemma 8.51. *Let $\mathcal{P}$ be a tight set partition of G such that either $\mathsf{coll}(\mathcal{P})$ is a BoB or $\mathcal{P}$ is a maximal cycle. Let σ be a correspondence from a torsoid $\mathcal{T} = (H, \varepsilon)$ in G to $\mathcal{P}$ and let vw be an edge of H. Then $\varepsilon(vw)$ is the largest set passable between $\sigma(v)$ and $\sigma(w)$.*

Proof. Let $P = \sigma(v)$ and $Q = \sigma(w)$. We begin by showing that $\varepsilon(vw)$ is passable between P and Q. By symmetry it is enough to show that it is passable for P. By Corollary 8.47 and Corollary 8.50, P must be a vertex cut with respect to (H, ε). Thus $P \setminus \varepsilon(vw)$ is tight by Lemma 8.6 applied to P and the complement of $w \cup \varepsilon(vw)$, and $P \cup \varepsilon(vw)$ is tight by Lemma 8.6 applied to P and $v \cup \varepsilon(vw)$.

Now let S be the largest set passable between P and Q. From the arguments in the last paragraph we already know that for any neighbour x of v other than w in H the set $\varepsilon(vx)$ is passable for $\mathcal{P}$. So by Lemma 8.30 and Lemma 8.20 the set S is passable for $P \setminus \bigcup_{x \in N_H(v) \setminus \{w\}} \varepsilon(vx)$. Since $\varepsilon(vw) \subseteq S$ this implies that S is also passable for $P \setminus \bigcup_{x \in N_H(v)} \varepsilon(vx)$. This latter set is just equal to v because $v \subseteq P \subseteq v \cup \bigcup_{x \in N_H(v)} \varepsilon(vx)$.

A similar argument shows that S is passable for w, and by Lemma 8.30 it is a subset of $v \cup \varepsilon(vw) \cup w$, so we have $S \subseteq \varepsilon(vw)$. Since, as we saw in the previous paragraph, $\varepsilon(vw)$ is passable between P and Q, we also have $\varepsilon(vw) \subseteq S$. Thus $S = \varepsilon(vw)$, as required. □

Theorem 8.52. *Let $\mathcal{P}$ be a torsoid-inducing tight set partition of G. Then the only torsoid $\mathcal{T}$ in correspondence with $\mathcal{P}$ is $\mathcal{T}_{\mathcal{P}}$ and the only such correspondence is $\sigma_{\mathcal{P}}$.*

Proof. Let σ be any correspondence from any torsoid $\mathcal{T} = (H, \varepsilon)$ to $\mathcal{P}$. By Lemma 8.51, for any edge vw of H we have $\varepsilon(vw) = \delta_{\mathcal{P}}(\sigma(v)\sigma(w))$. For any vertex v of H, since $v \subseteq \sigma(v) \subseteq v \cup \bigcup_{w \in N_H(v)} \varepsilon(vw)$ we have

$$v = \sigma(v) \setminus \bigcup_{w \in N_H(v)} \varepsilon(vw) = \sigma(v) \setminus \bigcup_{Q \in N_{\mathsf{coll}(\mathcal{P})}(\sigma(v))} \delta_{\mathcal{P}}(\sigma(v)Q) = \tau_{\mathcal{P}}(\sigma(v))\,.$$

Thus $\tau_{\mathcal{P}} = \sigma^{-1}$ and so $\sigma_{\mathcal{P}} = \sigma$. It follows that $H_{\mathcal{P}} = H$ and $\varepsilon_{\mathcal{P}} = \varepsilon$, giving the desired result. □

Corollary 8.53. *If σ is a correspondence from a torsoid $\mathcal{T} = (H, \varepsilon)$ in G to a tight set partition $\mathcal{P}$ of G then it is a graph isomorphism between H and $\mathsf{coll}(\mathcal{P})$.* □

8.5 Relation of Torsos to Torsoids

In this section we consider specific tight cut contractions of a graph with respect to a fixed maximal family of nested tight cuts and investigate how they relate to our concept of torsoids. The results of this section emphasise that torsoids are a global tool capturing all of these tight cut contractions independently of the precise choice of a tight cut family.

Definition 8.54. Let $\mathcal{C}$ be a maximal family of nested tight cuts in G. A *maximal star* of $\mathcal{C}$ is a tight set partition $\mathcal{P}$ of size at least 4 such that $\mathsf{coll}(\mathcal{P})$ is a BoB and $\partial(P) \in \mathcal{C}$ for every $P \in \mathcal{P}$. Then we call $\mathsf{coll}(\mathcal{P})$ a *torso of* $\mathcal{C}$ *at the maximal star* $\mathcal{P}$, or a *torso* for short. Furthermore, we simply call a torso at some maximal star of some maximal family of nested tight cuts in G a *torso in* G. If a torso is a C_4, we call it a C_4*-torso*. Otherwise, the torso is a BoB other than C_4 and we call it a *non-*C_4*-torso*.

We want to emphasise the significant difference between torsoids and torsos. Torsos are completely dependent on a fixed family of tight cuts $\mathcal{C}$, they capture how this specific family behaves and can only be used to describe the properties of $\mathcal{C}$. This is different for torsoids, a torsoid is a global object that captures multiple ways to choose maximal families of nested tight cuts by keeping passable sets on its edges rather than assigning them rigidly and by bundling any maximal cyclic structure in one torsoid instead of considering it as different torsos.

The relation of torsos and torsoids that we investigate in this section is defined as follows:

Definition 8.55. Let $\mathcal{S}$ be a torso in G and $\mathcal{T}$ a torsoid in G. We say $\mathcal{S}$ *cleaves* $\mathcal{T}$ if every vertex of $\mathcal{S}$ contains a vertex of $\mathcal{T}$.

In Subsection 8.5.1 we prove that each torso cleaves exactly one torsoid. However, in the other direction there might be several torsos that cleave a given torsoid, as cyclic torsoids always induce maximal cyclic tight set partitions in contrast to cyclic torsos whose induced tight set partition can be non-maximal cyclic. So, in Subsection 8.5.2 we prove the main result of this section that describes how many torsos cleave a given torsoid.

Observation 8.56. Note that all vertices of a torso $\mathcal{S}$ cleaving a torsoid $\mathcal{T}$ are either proper $\mathcal{T}$-vertex-residents or $\mathcal{T}$-interval-residents. Thus every vertex of $\mathcal{T}$ is contained in a vertex of $\mathcal{S}$. Therefore there exists a partition in correspondence to $\mathcal{T}$ that refines the partition $V(\mathcal{S})$.

8.5.1 Every Torso Cleaves a Unique Torsoid

Lemma 8.57. *Let* $\mathcal{S}$ *be a torso in* G *that cleaves a torsoid* $\mathcal{T}$ *in* G*. Then* $\mathcal{T}$ *is cyclic if and only if* $\mathcal{S}$ *is a* C_4*.*

Proof. If $\mathcal{T}$ is a cyclic torsoid, then every vertex of $\mathcal{S}$ is either a proper $\mathcal{T}$-vertex-resident or a $\mathcal{T}$-interval-resident. Then $\mathcal{S}$ is a cycle. As $\mathcal{S}$ is a BoB, it is a C_4.

If $\mathcal{T}$ is a non-cyclic torsoid, then every vertex of $\mathcal{S}'$ is a proper $\mathcal{T}'$-vertex-resident. This implies that the partition $V(\mathcal{S}')$ corresponds to $\mathcal{T}$. Therefore $\mathcal{S}'$ is not a cycle and thus a non-C_4-torso. □

Lemma 8.58. *Let $\mathcal{S}$ be a torso in G. There exists exactly one torsoid $\mathcal{T}$ in G such that $\mathcal{S}$ cleaves $\mathcal{T}$.*

Proof. If $\mathcal{S}$ is non-C_4, it is a BoB other than C_4 and thus by Definition 8.35 and Theorem 8.36, $\mathcal{T}_{V(\mathcal{S})}$ is a torsoid. By construction, $\mathcal{S}$ cleaves $\mathcal{T}_{V(\mathcal{S})}$. Furthermore, any torsoid $\mathcal{T}$ cleaved by $\mathcal{S}$ is non-cyclic by Lemma 8.57. Thus any vertex of $\mathcal{S}$ is a proper $\mathcal{T}$-vertex-resident and thus $\mathcal{T} = \mathcal{T}_{V(\mathcal{S})}$, since $\mathcal{T}$ and $V(\mathcal{S})$ correspond. Thus we can assume that $\mathcal{S}$ is a C_4-torso.

We use refinement to show that there exists a torsoid that $\mathcal{S}$ cleaves: given any C_4-torso and its underlying tight set partition $\mathcal{P}_1 := \mathcal{P}$, we begin by, if possible, refining a partition class $P \in \mathcal{P}_1$ in such a way that we split it up into three smaller sets to obtain a new partition $\mathcal{P}_2$ such that $\mathsf{coll}(\mathcal{P}_2)$ is again a cycle. Note that if P can be split up into finitely many smaller sets it can always be split up into exactly three, since in a cyclic partition any odd union of consecutive partition classes is a tight cut. We repeat this as long as possible.

Claim 1. The process stops after finitely many refinement steps.

Proof. Suppose for a contradiction that infinitely many refining steps are possible. Then there are partitions $(\mathcal{P}_i)_{i\in\mathbb{N}}$ and tight sets $P_1 \supset P_2 \supset P_3 \supset \dots$ where $P_i \in \mathcal{P}_i$, since any $\mathcal{P}_i$ is finite. Without loss of generality we may assume that $\mathcal{P}_i$ is a refinement of $\mathcal{P}_{i-1}$ in such a way that only P_{i-1} is refined (otherwise obtain such a partition by unifying some consecutive tight sets of $\mathcal{P}_i$). Furthermore we may assume that for each of the $\mathcal{P}_i$ we fixed a circular ordering of the partition sets in such a way that the order on $\mathcal{P}_i$ induces the one on $\mathcal{P}_{i-1}$.

Let Q be a partition set of $\mathcal{P}_1$ other than P_1 or one of its neighbours in $\mathsf{coll}(\mathcal{P}_1)$. By assumption, for each $i \in \mathbb{N}$ we have $Q \in \mathcal{P}_i$ and furthermore the circular ordering on $\mathcal{P}_i$ induces a linear ordering on $\mathcal{P}_i \setminus \{Q\}$. For $\sigma \in \{-, +\}$ we define P_i^σ to be the union of all sets in $\mathcal{P}_i \setminus \{Q\}$ that are $\{\text{smaller}, \text{larger}\}$ than P_i in this linear order. By definition P_i^σ is a subset of P_{i+1}^σ and we define $P_\infty^\sigma := \bigcup_{i\in\mathbb{N}} P_i^\sigma$ for $\sigma \in \{+, -\}$. Note that this definition implies that P_∞^+ and P_∞^- are disjoint. Next, it is clear that $|\mathcal{P}_i| \geq 4$, thus we may pick a path R from P_1^- to P_1^+ in G that is disjoint from Q. For one $\sigma \in \{+, -\}$ we have $P_i^\sigma \subsetneq P_{i+1}^\sigma$ for infinitely many distinct $i \in \mathbb{N}$, without loss of generality let $\sigma = -$. Then there is an edge vw of R with $v \in P_\infty^-$ and $w \notin P_\infty^-$. Furthermore, there are natural numbers $k < \ell$ such that k is the smallest number with $v \in P_k^-$ and $P_k^- \subsetneq P_\ell^-$. Therefore vw witnesses that there is a chord in $\mathsf{coll}(\mathcal{P}_\ell)$, a contradiction to the assumption that $\mathsf{coll}(\mathcal{P}_\ell)$ is a cycle. Thus the refinement stops at some $n \in \mathbb{N}$, as desired. ∎

By Claim 1 there is $n \in \mathbb{N}$ and a sequence $(\mathcal{P}_i)_{i\in[n]}$ as described above such that $\mathcal{P}_n$ is maximal cyclic. Thus $\mathcal{P}_n$ induces a torsoid $\mathcal{T}_{\mathcal{P}_n}$ by Theorem 8.36, which is cleaved by $\mathcal{S}$.

Let $\mathcal{T}$ be any torsoid that is cleaved by $\mathcal{S}$. By Lemma 8.57, the torsoid $\mathcal{T}$ is cyclic. We show:

Claim 2. For each $i \in [n]$ every $P \in \mathcal{P}_i$ contains a vertex of $\mathcal{T}$.

Claim 2 implies that every $P \in \mathcal{P}_i$ is either a $\mathcal{T}$-interval-resident or a proper $\mathcal{T}$-vertex-resident and there exists a partition in correspondence with $\mathcal{T}$ that refines $\mathcal{P}_i$. Let $\mathcal{Q}$ be the partition in correspondence with $\mathcal{T}$ that refines $\mathcal{P}_n$. Since $\mathcal{P}_n$ is maximal cyclic and $\mathcal{Q}$ is cyclic, $\mathcal{Q}$ coincides with $\mathcal{P}_n$ and thus $\mathcal{T} = \mathcal{T}_{\mathcal{P}_n}$ holds. Therefore $\mathcal{T}_{\mathcal{P}_n}$ is the unique torsoid cleaved by $\mathcal{S}$.

Proof of Claim 2. We prove by induction on $i \in [n]$. By construction, $\mathcal{P} = \mathcal{P}_1$ has the desired property. Suppose that any partition class of $\mathcal{P}_{i-1}$ contains a vertex of $\mathcal{T}$. We prove that also $\mathcal{P}_i$ has this property.

Let $P \in \mathcal{P}_{i-1}$ and $P_1, P_2, P_3 \in \mathcal{P}_i$ in cyclic order such that $P = P_1 \sqcup P_2 \sqcup P_3$ is a disjoint union with $\mathcal{P}_i = \{P_1, P_2, P_3\} \cup (\mathcal{P}_{i-1} \setminus \{P\})$. By Observation 8.56, there exists a tight set partition $\mathcal{R}$ in correspondence to $\mathcal{T}$ that refines $\mathcal{P}_{i-1}$. We define $\mathcal{P}_i' := \{P\} \cup \{R \in \mathcal{R} \colon R \subset V(G) \setminus P\}$. As $\mathcal{R}$ is a refinement of $\mathcal{P}_i'$ and $\mathsf{coll}(\mathcal{R})$ is cyclic, also $\mathsf{coll}(\mathcal{P}_i')$ is cyclic. We consider $\mathcal{P}_i'' := \{P_1, P_2, P_3\} \cup \{R \in \mathcal{R} \colon R \subset V(G) \setminus P\}$. By construction, $\mathcal{P}_i''$ is a refinement of $\mathcal{P}_i'$ and $\mathcal{P}_i$.

We prove that $\mathcal{P}_i''$ is also cyclic. Let $\hat{P}, \tilde{P}$ be the neighbours of P in $\mathcal{P}_{i-1}$. Let $\hat{R} \subseteq \hat{P}$ and $\tilde{R} \subseteq \tilde{P}$ be the neighbours of P in $\mathcal{P}_i'$. Without loss of generality, $\partial(P) = E(P_1, \hat{P}) \sqcup E(P_3, \tilde{P})$, since $\mathcal{P}_i$ is cyclic. Then $\partial(P) = E(P, \hat{R}) \sqcup E(P, \tilde{R})$, since $\mathcal{P}_i'$ is cyclic. This implies

$$\partial(P) = (E(P, \hat{R}) \cap E(P_1, \hat{P})) \sqcup (E(P, \tilde{R}) \cap E(P_3, \tilde{P})) = E(P_1, \hat{R}) \sqcup E(P_3, \hat{P})\,,$$

which proves that $\mathcal{P}_i''$ is cyclic.

If P is a proper $\mathcal{T}$-vertex-resident, then $\mathcal{P}_i'$ corresponds to $\mathcal{T}$. Therefore $\mathcal{P}_i'$ is maximal cyclic, which contradicts that $\mathcal{P}_i''$ is cyclic and refines $\mathcal{P}_i'$. Thus P is a $\mathcal{T}$-interval-resident. In particular, P contains at least 3 vertices of $\mathcal{T}$. Suppose towards a contradiction that one of P_1, P_2, P_3 does not contain a vertex of $\mathcal{T}$, i.e. it is either a non-proper $\mathcal{T}$-vertex-resident or a $\mathcal{T}$-edge-resident. If it is a non-proper $\mathcal{T}$-vertex-resident, then this non-proper $\mathcal{T}$-vertex-resident contains a proper, odd subset of a vertex of $\mathcal{T}$. Since $\mathcal{T}$-edge-residents and $\mathcal{T}$-interval-residents either are supersets of this vertex or avoid this vertex, all three tight sets P_1, P_2, P_3 are $\mathcal{T}$-vertex-residents at that same vertex. Thus P_1, P_2, P_3 do not contain any other vertex of $\mathcal{T}$, a contradiction as $P = P_1 \cup P_2 \cup P_3$.

Thus, one of P_1, P_2, P_3 is a $\mathcal{T}$-edge-resident. By parity, two of P_1, P_2, P_3 are $\mathcal{T}$-edge-residents at the same edge. The other tight set is a $\mathcal{T}$-interval-resident, since P contains at least 3 vertices of $\mathcal{T}$. The boundary of this interval-resident contains an edge of $\partial(P)$, as the two other tight sets are edge-resident at the same edge. Thus, either P_1 or P_3 is the interval-resident and we suppose without loss of generality the former.

We prove that $\mathcal{R}$ is not maximal cyclic, a contradiction. Note that $P_2, P_3 \subset P$ are $\mathcal{T}$-edge-residents at an edge $uw \in E(H)$ with $u \subset P$ and $w \cap P = \emptyset$. Therefore, both are contained in the element $R \in \mathcal{R}$ that contains u. We show that $\mathcal{R}^* = \{R \cap P_1, P_2, P_3\} \cup \mathcal{R} \setminus \{R\}$ is a tight set partition such that $\mathsf{coll}(\mathcal{R}^*)$ is cyclic. By construction, $E(P_1, P_2) = E(P_1 \cap R, P_2)$ holds. Then since $\mathcal{P}_i''$ and $\mathcal{R}$ are cyclic, also $\mathcal{R}^*$ is cyclic, which provides the desired contradiction. ■

This completes the proof. □

8.5.2 Torsos Cleaving a Fixed Torsoid

As we have established that every torso cleaves exactly one torsoid we can introduce the following: let $\mathcal{C}$ be a maximal family of nested tight cuts in G, $\mathfrak{S}$ be the set of all torsos of $\mathcal{C}$ and $\mathfrak{T}$ the set of all torsoids in G. We define the map $\kappa_{\mathcal{C}} \colon \mathfrak{S} \longrightarrow \mathfrak{T}$ sending $\mathcal{S} \in \mathfrak{S}$ to the unique torsoid $\mathcal{T}$ such that $\mathcal{S}$ cleaves $\mathcal{T}$.

Now we want to know what we can say about the torsos that cleave a given torsoid. This question leads to the main result of this section, which reads as follows:

Theorem 8.59. *Let $\mathcal{C}$ be a maximal family of nested tight cuts in G. Then the $\kappa_{\mathcal{C}}$-pre-image of a non-cyclic torsoid consists of one non-C_4 torso and the $\kappa_{\mathcal{C}}$-pre-image of a cyclic torsoid with n vertices consists of $\frac{n}{2} - 1$ C_4-torsos.*

First we consider the special case of Theorem 8.59 for BoBs and cycles. Note that every BoB and every cycle contains a unique torsoid. Furthermore, any torso in a BoB or in a cycle cleaves this unique torsoid.

Proposition 8.60. *Let H be a matching covered graph with a maximal family $\mathcal{C}$ of nested tight cuts in H. If H is a BoB, there is exactly one torso of $\mathcal{C}$. If H is a cycle of even length n, there are exactly $\frac{n}{2} - 1$ torsos of $\mathcal{C}$.*

Proof. If H is a BoB, the set $\mathcal{C}$ contains only trivial tight sets and thus H itself is the unique torso of $\mathcal{C}$. If H is a cycle, we prove the statement by strong induction on n. For $n = 4$ the statement is true, as C_4 is a BoB. Let H be a cycle of length $n \geq 6$ and suppose the statement is true for all cycles of length at most $n - 2$. Let $\partial(X) \in \mathcal{C}$ be any non-trivial tight cut. Set $\mathcal{C}_1 := \{C' \in \mathcal{C} \colon \exists Y \subseteq V(G) \setminus X \text{ s.t. } C' = \partial(Y)\}$ and $\mathcal{C}_2 := \{C' \in \mathcal{C} \colon \exists Y \subseteq X \text{ s.t. } C' = \partial(Y)\}$. Note that $\mathcal{C} = \mathcal{C}_1 \cup \mathcal{C}_2$ and $\{\partial(X)\} = \mathcal{C}_1 \cap \mathcal{C}_2$.

Let H_1 be the cycle obtained by contracting X and consider the set $\hat{\mathcal{C}_1}$ of nested tight cuts induced by $\mathcal{C}_1$. It is maximal by construction. Let H_2 be the cycle obtained by contracting $V(H) \setminus X$ and consider the the set $\hat{\mathcal{C}}_2$ of nested tight cuts induced by $\mathcal{C}_2$, which is also maximal. There is a canonical bijection between the set of torsos of $\mathcal{C}$ and the set of torsos of $\hat{\mathcal{C}}_1$ and $\hat{\mathcal{C}}_2$. Applying the induction hypothesis to $\mathcal{C}_1, H_1$ and $\mathcal{C}_2, H_2$ gives the desired statement for n. □

In preparation of the proof for the general case of Theorem 8.59 we prove a couple of propositions that enable us to define a bijection between the $\kappa_{\mathcal{C}}$-pre-image of a torsoid $\mathcal{T} = (H, \varepsilon)$ and the set of torsos of a specific maximal family of nested tight cuts in H. Note that Proposition 8.60 determines the size of the latter set, as H is either a BoB or a cycle.

Firstly, we show that the vertices of a torso cleaving $\mathcal{T}$ are $(\mathcal{T}, \vec{\mathcal{C}})$-maximal which we define below. Secondly, we introduce the specific family of nested tight cuts in H and show that it is indeed maximal.

Now we return to the graph G itself and consider the residents on vertices and intervals of a torsoid $\mathcal{T} = (H, \varepsilon)$ in G. Recall that $\mathsf{odd}_H(X)$ is the set of vertices in H that intersect X oddly. We call a proper $\mathcal{T}$-vertex-resident or a $\mathcal{T}$-interval-resident $X \in \vec{\mathcal{C}}$ $(\mathcal{T}, \vec{\mathcal{C}})$-*maximal*, if $Y \subseteq X$ for every tight set $Y \in \vec{\mathcal{C}}$ with $\mathsf{odd}_H(Y) = \mathsf{odd}_H(X)$. Note that $\bigcup \mathsf{odd}_H(X) \subseteq X$ holds for every proper $\mathcal{T}$-vertex-resident or $\mathcal{T}$-interval-resident X.

The next two statements prove that for every vertex and every interval that has a resident such maximal residents exist.

Proposition 8.61. *Let $\mathcal{C}$ be a maximal family of nested tight cuts in G and $\mathcal{T} = (H, \varepsilon)$ a torsoid in G. Then for every $v \in V(H)$ there is a $(\mathcal{T}, \vec{\mathcal{C}})$-maximal proper $\mathcal{T}$-vertex-resident at v.*

Proof. Consider the set $A := \{Y \in \vec{\mathcal{C}} \colon \mathsf{odd}_H(Y) = \{v\}\}$. We show that $X := \bigcup A$ is the desired set. By construction, the set A contains only $\mathcal{T}$-vertex-residents. Thus X avoids all vertices of $\mathcal{T}$ apart from v. Furthermore, the set X contains the set v as every single vertex in v forms a $\mathcal{T}$-vertex-resident, whose cut is clearly nested with $\mathcal{C}$. Therefore $\mathsf{odd}_H(X) = \{v\}$ holds. Thus X is a proper $\mathcal{T}$-vertex-resident. We have to prove that X is tight and an element of $\vec{\mathcal{C}}$.

As $v \subseteq X$, $X = \bigcup_{Y \in A}(Y \cup v)$ holds. For all $Y', Y'' \in A$ the sets $Y' \cup v, Y'' \cup v$ are tight. Since Y', Y'' are nested, either the sets Y', Y'' are disjoint or one contains the other. In the former case, $(Y' \cup v) \cap (Y'' \cup v) = v$ is odd. In the latter case, $(Y' \cup v) \cap (Y'' \cup v) \in \{Y' \cup v, Y'' \cup v\}$, which is also odd. In either case Lemma 8.6 implies that the union of the two is a tight set. Then, by Lemma 8.9, X is tight.

It remains to show that $\partial(X)$ is nested with every tight cut in $\mathcal{C}$. Let C be an arbitrary tight cut in $\mathcal{C}$. Note that C is nested with every tight cut induced by an element of A. If C resides at v, there is $Y \in A$ such that $C = \partial(Y)$ and thus $\partial(X)$ and C are nested by construction. If C resides at another vertex, its resident avoids all elements of A and therefore $\partial(X)$ and C are nested. If C resides at an edge, let Z

be the edge-resident of C. Either Z is contained in an element of A, or Z is disjoint to all elements of A. In the former case $Z \subseteq X$ holds, in the latter case $Z \cap X = \emptyset$ holds. In both cases $\partial(X)$ and C are nested. If C resides at an interval, let Z be the interval-resident of C containing v. For any $Y \in A$, the tight cuts $\partial(Y)$ and $C = \partial(Z)$ are nested and Y, Z have non-empty intersection. Therefore Y is contained in Z. Thus X is contained in Z and $\partial(X)$ and C are nested. $\square$

Proposition 8.62. *Let $\mathcal{C}$ be a maximal family of nested tight cuts in G and $\mathcal{T} = (H, \varepsilon)$ be a cyclic torsoid in G. Furthermore, let I be an interval in H such that there exists $Y \in \vec{\mathcal{C}}$ with $\mathsf{odd}_H(Y) = I$. Then there is a $(\mathcal{T}, \vec{\mathcal{C}})$-maximal $\mathcal{T}$-interval-resident X with $\mathsf{odd}_H(X) = I$.*

Proof. Consider the set $A := \{Y \in \vec{\mathcal{C}} \colon \mathsf{odd}_H(Y) = I\}$. We prove that $X := \bigcup A$ is the desired set. As A contains only interval-residents, the set X contains the vertices in I and avoids all other vertices of $\mathcal{T}$, which implies $\mathsf{odd}_H(X) = I$. For any two elements $Y', Y'' \in A$ one is contained in the other, since $\bigcup I \subseteq Y', Y''$. Therefore Y', Y'' intersect oddly and X is tight by Lemma 8.9. To complete the proof we have to show that the set X is contained in $\vec{\mathcal{C}}$.

Let $\partial(Z) \in \mathcal{C}$ be an arbitrary tight cut. It suffices to prove that $\partial(Z)$ and $\partial(X)$ are nested. If Z is contained in an element of A, it is contained in X and thus $\partial(X)$ and $\partial(Z)$ are nested. We suppose that Z is not contained in any element of A. If Z is disjoint to all elements of A, the sets Z and X are disjoint and thus $\partial(X)$ and $\partial(Z)$ are nested. Therefore we suppose that there is $Y \in A$ such that Y and Z have non-empty intersection. As $\partial(Z)$ and $\partial(Y)$ are nested, $Y \subset Z$ holds by assumption. This implies $\bigcup I \subset Z$. Therefore any $Y' \in A$ intersects Z and as $\partial(Y'), \partial(Z)$ are nested, Y' is contained in Z by assumption. Thus $X \subseteq Z$, which shows that $\partial(X)$ and $\partial(Z)$ are nested. $\square$

Next we show that the vertices of a torso cleaving $\mathcal{T}$ are indeed such maximal residents.

Proposition 8.63. *Let $\mathcal{C}$ be a maximal family of nested tight cuts in G and $\mathcal{T}$ any torsoid in G. Then every vertex of a torso of $\mathcal{C}$ cleaving $\mathcal{T}$ is $(\mathcal{T}, \vec{\mathcal{C}})$-maximal.*

Proof. Suppose towards a contradiction that there is a torso $\mathcal{S}$ cleaving $\mathcal{T}$ with a vertex X such that X is not $(\mathcal{T}, \vec{\mathcal{C}})$-maximal. By Proposition 8.61 and Proposition 8.62, there is a $(\mathcal{T}, \vec{\mathcal{C}})$-maximal tight set Y in $\vec{\mathcal{C}}$ with $\mathsf{odd}_H(X) = \mathsf{odd}_H(Y)$. Then $X \subset Y$ holds and thus there exists a vertex X' of $\mathcal{S}$ that contains an element of $Y \setminus X$. Since $\partial(Y)$ and $\partial(X')$ are nested, X' is contained in Y.

By construction, $\bigcup \mathsf{odd}_H(X) = \bigcup \mathsf{odd}_H(Y)$ is contained in X, and Y avoids any vertex of $V(H) \setminus \mathsf{odd}_H(Y)$. Therefore X' does not contain any vertex of H, which contradicts the definition of cleaving. $\square$

In the following we turn our attention to H and define the specific family of nested tight cuts in H for the proof of Theorem 8.59:

Definition 8.64. Let $\mathcal{C}$ be a maximal family of nested tight cuts in G and $\mathcal{T} = (H, \varepsilon)$ a torsoid in G. We define the set

$$\mathcal{C}^{\mathcal{T}} := \{\partial(\mathsf{odd}_H(X)) \colon \partial(X) \in \mathcal{C} \text{ and } \partial(X) \text{ resides at an interval or properly at a vertex of } \mathcal{T}\}\,.$$

Proposition 8.65. *Let $\mathcal{C}$ be a maximal family of nested tight cuts in G and $\mathcal{T} = (H, \varepsilon)$ a torsoid in G. Then $\mathcal{C}^{\mathcal{T}}$ is a maximal family of nested tight cuts in H.*

Note that $\partial(X) = \partial(V(G) \setminus X)$ and in particular $\partial(\mathsf{odd}_H(X)) = \partial(\mathsf{odd}_H(V(G) \setminus X))$ for every $X \subset V(G)$.

Proof. The elements of $\mathcal{C}^{\mathcal{T}}$ are nested, since $\bigcup \mathsf{odd}_H(X) \subseteq X$ and $\bigcup \mathsf{odd}_H(V(G) \setminus X) \subseteq V(G) \setminus X$ for any proper $\mathcal{T}$-vertex-resident or $\mathcal{T}$-interval-resident X. The cuts in $\mathcal{C}^{\mathcal{T}}$ are indeed tight: for any cut $C \in \mathcal{C}^{\mathcal{T}}$ there is either a $\mathcal{T}$-vertex-resident or an $\mathcal{T}$-interval-resident X such that $C = \partial(\mathsf{odd}_H(X))$. Then $\mathsf{odd}_H(X)$ is a single vertex if H is a BoB. If H is a cycle, $\mathsf{odd}_H(X)$ is either a single vertex or an odd interval in H. This implies that $\partial(\mathsf{odd}_H(X))$ is a tight set.

It remains to prove that $\mathcal{C}^{\mathcal{T}}$ is maximal. By Proposition 8.61, for every vertex $v \in V(H)$ there is a proper vertex-resident in $\mathcal{C}$ containing v and therefore all trivial tight cuts of H are contained in $\mathcal{C}^{\mathcal{T}}$. If H is a BoB, we are done. Thus we can suppose that H is a cycle.

We suppose for a contradiction that there exists a tight cut in H nested with $\mathcal{C}^{\mathcal{T}}$ but not contained in $\mathcal{C}^{\mathcal{T}}$. We already observed that $\mathcal{C}^{\mathcal{T}}$ contains all trivial tight cuts. Thus this tight cut is non-trivial and therefore of the form $\partial(I)$ for an interval I with $3 \le |I| \le |V(H)| - 3$. Let i_1 be the first and i_2 be last element of I (regarding the cyclic order of $V(H)$). We construct a $\mathcal{T}$-interval-resident Y with $\partial(Y) \in \mathcal{C}$ that fulfils $\mathsf{odd}_H(Y) = I$. This contradicts that $\partial(I) \notin \mathcal{C}^{\mathcal{T}}$.

For $k \in \{1, 2\}$ consider all $(\mathcal{T}, \vec{\mathcal{C}})$-maximal proper $\mathcal{T}$-vertex-residents and $\mathcal{T}$-interval-residents containing i_k and let Y_k be the subset-maximal one. By Proposition 8.61 and Proposition 8.62, such Y_k exists. It contains all tight sets containing i_k. As $\partial(I)$ is nested with $\partial(\mathsf{odd}_H(Y_k)) \in \mathcal{C}^{\mathcal{T}}$, $\mathsf{odd}_H(Y_k) \subseteq I$ holds for $k \in \{1, 2\}$. Furthermore, since $\partial(I)$ is not contained in $\mathcal{C}^{\mathcal{T}}$, $\mathsf{odd}_H(Y_k) \subset I$ holds for $k \in \{1, 2\}$. Then $i_1 \cap Y_2 = \emptyset$ and $i_2 \cap Y_1 = \emptyset$, since Y_1, Y_2 are $\mathcal{T}$-interval-residents or $\mathcal{T}$-vertex-residents. As $\partial(Y_1), \partial(Y_2)$ are nested and $i_1 \subseteq Y_1, i_2 \subseteq Y_2$ but $i_1 \cap Y_2 = \emptyset, i_2 \cap Y_1 = \emptyset$, the sets Y_1, Y_2 are disjoint.

We consider the set $A := \bigcup_{vw \in E(H) \colon v,w \in I} v \cup w \cup \varepsilon(vw)$, which is tight by construction. Since Y_k is a $\mathcal{T}$-interval-resident or a proper $\mathcal{T}$-vertex-resident that holds $\mathsf{odd}_H(Y_k) \subset I$, Y_k intersects A oddly. We set $Y := A \cup Y_1 \cup Y_2$, which is tight by Lemma 8.6. By construction, $\mathsf{odd}_H(Y) = I$ holds. We show that $\partial(Y)$ is nested with $\mathcal{C}$. Then it is contained in $\mathcal{C}$, which gives the desired contradiction.

Suppose that there exists $\partial(Z) \in \mathcal{C}$ which crosses $\partial(Y)$. Then Z crosses Y. Let $k \in \{1, 2\}$. As Z crosses Y, Z is not a subset of Y_k. Since Y_k is the subset-maximal proper $\mathcal{T}$-vertex-resident or $\mathcal{T}$-interval-resident containing i_k, Z is not a superset of Y_k. Thus Z avoids Y_k, as Z, Y_k are nested. Therefore it contains elements of both

$$V(G) \setminus Y \subseteq \bigcup (V(H) \setminus I) \cup \bigcup_{\substack{vw \in E(H): \\ v \notin I}} \varepsilon(vw)\,, \text{ and}$$

$$Y \setminus \bigcup \{Y_1, Y_2\} \subseteq \bigcup (I \setminus \{i_1, i_2\}) \cup \bigcup_{\substack{vw \in E(H): \\ v \in I \setminus \{i_1, i_2\}}} \varepsilon(vw)\,.$$

The tight set Z clearly can neither be a vertex-resident nor an edge-resident. Since X avoids $i_1 \cup i_2 \subseteq Y_1 \cup Y_2$ it neither can be an interval-resident. $\square$

To obtain the main theorem of this section about the number of torsos cleaving a torsoid, it now remains to prove that there is a bijection between the $\kappa_{\mathcal{C}}$-pre-image of a torsoid $\mathcal{T}$, i.e. the set of torsos of $\mathcal{C}$ that cleave $\mathcal{T}$, and the set of torsos of $\mathcal{C}^{\mathcal{T}}$.

Theorem 8.59. *Let $\mathcal{C}$ be a maximal family of nested tight cuts in G. Then the $\kappa_{\mathcal{C}}$-pre-image of a non-cyclic torsoid consists of one non-C_4 torso and the $\kappa_{\mathcal{C}}$-pre-image of a cyclic torsoid with n vertices consists of $\frac{n}{2} - 1$ C_4-torsos.*

Before we get to the proof of the theorem, we show one further result: that every torso of $\mathcal{C}$ cleaving $\mathcal{T}$ maps canonically to a torso of $\mathcal{C}^{\mathcal{T}}$.

Proposition 8.66. *Let $\mathcal{C}$ be a maximal family of nested tight cuts in G. Then every torso $\mathcal{S}$ of $\mathcal{C}$ cleaving a torsoid $\mathcal{T} = (H, \varepsilon)$ in G is isomorphic to $\mathsf{coll}((\mathsf{odd}_H(X))_{X \in V(\mathcal{S})})$.*

Proof. By Observation 8.56 any vertex of $\mathcal{S}$ is either a proper $\mathcal{T}$-vertex-resident or a $\mathcal{T}$-interval-resident. Thus $(\mathsf{odd}_H(X))_{X \in V(\mathcal{S})}$ is indeed a partition of $V(H)$. It remains to show that for every $X, Y \in V(\mathcal{S})$ there is an edge in $\partial_G(X) \cap \partial_G(Y)$ if and only if there is $xy \in E(H)$ with $x \in X$ and $y \in Y$.

Let $X, Y \in V(\mathcal{S})$ be chosen arbitrarily. If there is an edge e in $\partial_G(X) \cap \partial_G(Y)$, then there is $uv \in E(H)$ such that both endvertices of e are contained in $u \cup \varepsilon(uv) \cup v$ by (T6) and Proposition 8.33. Since any vertex of $\mathcal{S}$ is either a proper vertex-resident or an interval-resident, X and Y each contain exactly one of u, v, which proves the forward implication.

For the backwards implication, note that there is a tight set partition $\mathcal{P}$ corresponding to $\mathcal{T}$ that refines $V(\mathcal{S})$, as any vertex of $\mathcal{S}$ is either a proper vertex-resident or an interval-resident. If there is $xy \in E(H)$ with $x \in X$ and $y \in Y$, let $P_x, P_y \in \mathcal{P}$ such that $x \in P_x$ and $y \in P_y$. By Theorem 8.36, there is an edge between P_x and P_y in $\mathsf{coll}(\mathcal{P})$. This implies that there is an edge between P_x and P_y in G and therefore an edge in $\partial_G(X) \cap \partial_G(Y)$. $\square$

Proof of Theorem 8.59. Let $\mathcal{C}^{\mathcal{T}}$ be as in Proposition 8.65. By Proposition 8.60, the number of torsos of $\mathcal{C}^{\mathcal{T}}$ is 1 if H is a BoB and $\frac{n}{2} - 1$ if H is a cycle of length n. It remains to prove that there is a bijection between the set of torsos of $\mathcal{C}$ cleaving $\mathcal{T}$ and the set of torsos of $\mathcal{C}^{\mathcal{T}}$.

Given a torso $\mathcal{S}$ of $\mathcal{C}$ cleaving $\mathcal{T}$ the graph $\mathsf{coll}((\mathsf{odd}_H(X))_{X \in V(\mathcal{S})})$ is a torso of $\mathcal{C}^{\mathcal{T}}$ by Proposition 8.66. Thus the map β sending a torso $\mathcal{S}$ of $\mathcal{C}$ cleaving $\mathcal{T}$ to the torso of $\mathcal{C}^{\mathcal{T}}$ at $(\mathsf{odd}_H(X))_{X \in V(\mathcal{S})}$ is well-defined.

We show that the map β is a bijection. Let $\mathcal{S}'$ be an arbitrary torso of $\mathcal{C}^{\mathcal{T}}$. For $J \in V(\mathcal{S}')$ let X_J be the unique $(\mathcal{T}, \vec{\mathcal{C}})$-maximal tight set with $\mathsf{odd}_H(X_J) = J$. By Proposition 8.63, the maximal star $(X_J)_{J \in V(\mathcal{S}')}$ is the only one whose torso is mapped to $\mathcal{S}'$ by β. Therefore β is injective.

For surjectivity we have to show that $(X_J)_{J \in V(\mathcal{S}')}$ is indeed a maximal star of $\mathcal{C}$. It remains to prove that $(X_J)_{J \in V(\mathcal{S}')}$ is a partition of $V(G)$. For $J, K \in V(\mathcal{S}')$ the sets X_J, X_K are not contained in each other and $\partial(X_J), \partial(X_K)$ are nested, therefore X_J, X_K are disjoint. Furthermore, $\bigcup V(H) \subseteq \bigcup_{J \in V(\mathcal{S}')} X_J$ holds. We have to prove that $\bigcup_{e \in E(H)} \varepsilon(e) \subseteq \bigcup_{J \in V(\mathcal{S}')} X_J$ holds.

Let $vw \in E(H)$ be arbitrary. If v, w are contained in X_J for some $J \in V(\mathcal{S}')$, then also $\varepsilon(vw)$ is contained in X_J, as X_J is a $\mathcal{T}$-interval-resident. Therefore we can assume that there are $K \neq L \in V(\mathcal{S}')$ with $v \in K, w \in L$. Suppose towards a contradiction that $Y := \varepsilon(vw) \setminus (X_K \cup X_L)$ is non-empty. Then $X_K \cup Y = X_K \cup ((v \cup \varepsilon(vw)) \cap (V(G) \setminus X_L))$ is a tight set. We show that $X_K \cup Y$ is contained in $\vec{\mathcal{C}}$, which contradicts the fact that X_K is $(\mathcal{T}, \vec{\mathcal{C}})$-maximal.

For every vertex-resident or interval-resident $Z \in \vec{\mathcal{C}}$ either $V(G)_Z$ or $V(G)_{V(G) \setminus Z}$ is contained in some $J \in V(\mathcal{S}')$, since $\mathcal{S}'$ is a torso of $\mathcal{C}^{\mathcal{T}}$. As X_J is $(\mathcal{T}, \vec{\mathcal{C}})$-maximal, Z or its complement is contained in X_J. Therefore no vertex-resident and no interval-resident of $\vec{\mathcal{C}}$ crosses $X_K \cup Y$. As no $\mathcal{T}$-edge-resident Z of $\vec{\mathcal{C}}$ at vw crosses X_K or X_L, Z is contained in either X_K, X_L or Y and thus nested with $X_K \cup Y$. Every $\mathcal{T}$-edge-resident at a distinct edge is nested with $X_K \cup Y$, as it is nested with X_K and $Y \subset \varepsilon(vw)$. Thus $X_K \cup Y$ is nested with $\vec{\mathcal{C}}$ and therefore contained in $\vec{\mathcal{C}}$. This gives the desired contradiction.

Indeed, $(X_J)_{J \in V(\mathcal{S}')}$ is a partition of $V(G)$. There are no non-trivial tight cuts in $(X_J)_{J \in V(\mathcal{S}')}$ since there are no non-trivial tight cuts in $\mathsf{coll}(V(\mathcal{S}'))$ Therefore $(X_J)_{J \in V(\mathcal{S}')}$ is a maximal star of $\mathcal{C}$, which completes the proof that the map β is a bijection. $\square$

Chapter 9

Menger's Theorem in Bidirected Graphs

The main result of this chapter reads as follows:

Theorem 9.1. *Let $\mathcal{X}$ and $\mathcal{Y}$ be sets of signed vertices of a bidirected graph B, and suppose that $\mathcal{X}$ is clean. Let $P_1, \ldots, P_k$ be vertex-disjoint $\mathcal{X}$–$\mathcal{Y}$ paths in B where P_i starts in $v_i \in V(\mathcal{X})$ for $i \in [k]$. Then precisely one of the following is true:*

(1) *There is a set S of k vertices of B such that $B - S$ contains no $\mathcal{X}$–$\mathcal{Y}$ path.*

(2) *There are $k+1$ vertex-disjoint $\mathcal{X}$–$\mathcal{Y}$ paths $P'_1, \ldots, P'_{k+1}$ in B where P'_i starts in v_i for $i \in [k]$.*

The general strategy of our proof of Theorem 9.1 follows the idea of [BGH01], but needs several extensions in order to circumnavigate the complication in bidirected graphs. In particular, we introduce 'appendages' as our main tool to overcome the walk-path problem for clean sets of vertices in Section 9.3. This allows us to first prove an edge-version of Menger's Theorem, Theorem 9.17, in Section 9.4 from which we then deduce the vertex-version, Theorem 9.1, in Section 9.5. Our proof of Theorem 9.1 can also be turned into a polynomial-time algorithm which finds a maximal set of disjoint paths, we prove this in Section 9.6. We begin by highlighting some connectivity properties of bidirected graphs in Section 9.1 and proving that Menger's original theorem cannot be translated to bidirected graphs verbatim in Section 9.2.

In this chapter in every figure involving bidirected graphs, we depict the signs of halfedges in a particular manner, by drawing the signs onto the edge, which results in a bar perpendicular to the edge at incident vertices with sign $+$ and none at incident vertices with sign $-$, see e.g. Figure 9.2. Note that in this chapter undirected and directed graphs are allowed to have parallel edges, but they do not have loops.

9.1 Connectivity in Bidirected Graphs

A key difference between bidirected and both undirected and directed graphs which complicates understanding their connectivity properties considerably lies in the relation of walks, trails and paths: unlike for (un)directed graphs, the existence of a v–w walk between two vertices v and w in a bidirected graph B does not imply that there exists a v–w trail in B, and similarly a v–w trail does not guarantee the existence of a v–w path (see [Wie22, Figure 9.1] for examples).

We use an analogous notion for the concatenation of trails at edges: for example if the union $Q_1\vec{e} \cup \vec{e}Q_2\vec{f} \cup \vec{f}Q_3$ of three trails is again a trail, we denote it as $Q_1\vec{e}Q_2\vec{f}Q_3$. We write similarly $P_1xP_2yP_3$ for paths where x, y may be both vertices or oriented edges.

The complicated connectivity structure of bidirected graphs also manifests in various different notions of strong connectivity (see [Wie22, Section 9.2] for an overview). We discuss two of them, namely 'strongly connected' and 'circularly connected', in this section. Let us first define 'strongly connected':

Definition 9.2. We say a bidirected graph B is *strongly connected* if for any two vertices v and w of B there are signs α and β such that B contains both a (v, α)–(w, β) path and a $(v, -\alpha)$–$(w, -\beta)$ path.

It turns out that a condition that appears weaker at first glance is equivalent to the one in this definition:

Lemma 9.3. *Let B be a bidirected graph such that for any two vertices v and w of B there are a $(v, +)$–w path and a $(v, -)$–w path in B. Then B is strongly connected.*

Proof. Let v and w be vertices of B. We must find paths joining them as in the definition of strong connectivity, Definition 9.2. By assumption there are signs β_1 and β_2 such that B contains a $(v, +)$–(w, β_1) path and a $(v, -)$–(w, β_2) path. Similarly, applying the assumption to w and v then reversing the paths shows that there are signs α_1 and α_2 such that B contains a (v, α_1)–$(w, +)$ path and a (v, α_2)–$(w, -)$ path.

If $\beta_2 = -\beta_1$ then we are done, so we may suppose $\beta_1 = \beta_2$. Similarly if $\alpha_2 = -\alpha_1$ then we are done, so we may suppose $\alpha_1 = \alpha_2$. Setting $\alpha := -\alpha_1$ and $\beta := \beta_1$ we are done. □

We now turn to 'circularly connected'. For this, a *cycle* in a bidirected graph $B = (G, \sigma)$ is a trail $C = v_0\vec{e}_1v_1\vec{e}_2v_2 \dots v_{\ell-1}\vec{e}_\ell v_\ell$ in B whose vertices are all distinct except $v_0 = v_\ell$ where we also have $\sigma(v_0, e_1) \neq \sigma(v_\ell, e_\ell)$.

Definition 9.4. Let $B = (G, \sigma)$ be a bidirected graph and consider the undirected graph $H := (V(G), F)$ where F is the set of edges of G that lie on some cycle of B. We refer to the connected components of H as the *circular components* of B. If H is connected we call B *circularly connected*.

We now aim to show that 'strongly connected' and 'circularly connected' in fact describe the same type of connectivity. This result was already known (see for example [Wie22, Section 9.2]), but we include a proof here for the convenience of the reader.

Theorem 9.5. *A bidirected graph $B = (G, \sigma)$ is strongly connected if and only if it is circularly connected.*

For the proof of Theorem 9.5, we make use of a lemma that describes which types of connectivity are forced by the interaction of paths between two vertices that differ in the signs at their endpoints:

Lemma 9.6. *Let v and w be vertices of a bidirected graph B. Let α and β be signs such that B contains both a (v, α)–(w, β) path P and a $(v, -\alpha)$–$(w, -\beta)$ path Q. Then at least one of the following statements holds:*

- *v and w lie in the same circular component of B.*
- *There is a (v, α)–$(w, -\beta)$ path in B.*

Proof. The proof proceeds by induction on the sum of the lengths of the paths P and Q. If this sum is 0, then $v = w$ and therefore the first statement holds. Otherwise, let x be the first vertex of P other than v which also lies on Q. Suppose that xP is an (x, γ)–(w, β) path and that xQ is an (x, δ)–$(w, -\beta)$ path.

Case 1: $\gamma = \delta$: In this case, PxQ is a (v, α)–$(w, -\beta)$ path in B.

Case 2: $\gamma = -\delta$: In this case we can apply the induction hypothesis to x, w, γ, β and the paths xP and xQ in the bidirected graph B' given by the union of xP and xQ. There are two possibilities.

Case 2.1: x and w lie in the same circular component of B: As the path Px is a (v, α)–$(x, -\gamma)$ path, Qx is a $(v, -\alpha)$–(x, γ) path, and these paths are disjoint except at their endvertices, their union is a cycle. This cycle witnesses that v is in the same circular component as x, hence also as w.

Case 2.2: There is an (x, γ)–$(w, -\beta)$ path R in B': In this case, the concatenation PxR is a (v, α)–$(w, -\beta)$ path in B.

□

Proof of Theorem 9.5. Suppose first that B is circularly connected. By Lemma 9.3 it is enough to show that for any vertices v and w of B there are both a $(v, +)$–w path and a $(v, -)$–w path. Let $v \in V(B)$ and $\alpha \in \{+, -\}$ be arbitrary. We set $X \subseteq V(B)$ to be the set of all vertices x for which there is a (v, α)–x path. Note that the set X is non-empty since it contains v. We show that X is the whole vertex set of B. Suppose not, then by circular connectivity there is some edge e of B contained in some cycle C of B which joins some $x \in X$ and some $y \notin X$.

There is a (v,α)–x path P by definition of X. Let z be the first vertex of P on C and let $\beta \in \{+,-\}$ such that the (possibly trivial) path Pz ends in (z,β). Then there is a $(z,-\beta)$–y path Q contained in C. Thus, PzQ is a (v,α)–y path, contradicting $y \notin X$. Since this is true for any $v \in V(B)$ and any $\alpha \in \{+,-\}$, this completes this direction of the proof.

Now suppose instead that B is strongly connected, and suppose for a contradiction that it is not circularly connected. Then it must contain an edge e which joins two vertices x and y in different circular components. Let $\sigma(x,e) = -\alpha$ and $\sigma(y,e) = \beta$. Suppose there is no (x,α)–$(y,-\beta)$ path. Since B is strongly connected, there must be an (x,α)–(y,β) path and an $(x,-\alpha)$–$(y,-\beta)$ path. Applying Lemma 9.6, we yield an (x,α)–$(y,-\beta)$ path in B. That path together with e is a cycle in B, contradicting our assumption that x and y lie in different circular components. □

9.2 Counterexample for General Menger Theorem

In this section we show that Menger's Theorem, Theorem 2.1, does not hold true in bidirected graphs if we transfer the statement verbatim. Even if we allow the separating set S to have any fixed size k for some $k \in \mathbb{N}$, the statement is false:

Theorem 9.7. *For each number $k \in \mathbb{N}$, there exists a bidirected graph B and disjoint sets $\mathcal{X}$ and $\mathcal{Y}$ of signed vertices of B, each of size at least k, such that there are no two disjoint $\mathcal{X}$–$\mathcal{Y}$ paths and for each subset $S \subseteq V(B)$ of size k there exists an $\mathcal{X}$–$\mathcal{Y}$ path in $B - S$.*

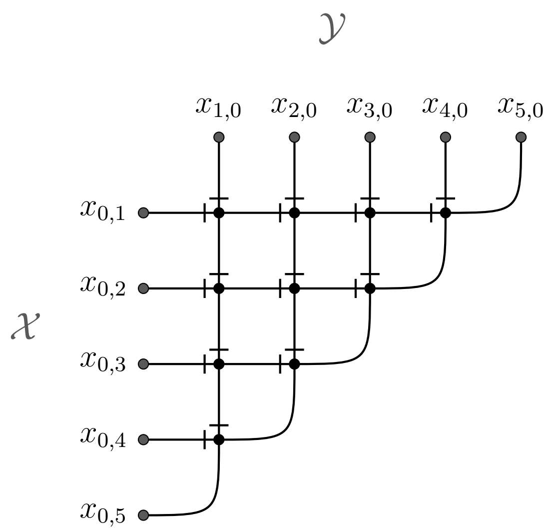

Figure 9.1: A bidirected graph containing neither two vertex-disjoint $\mathcal{X}$–$\mathcal{Y}$ paths nor a vertex set W of size ≤ 2 such that in $B - W$ there are no $\mathcal{X}$–$\mathcal{Y}$ paths.

For the proof of Theorem 9.7 we will rely on the following well-known topological lemma, which follows directly from the fact that a complete graph on five vertices is not planar:

Lemma 9.8. *Let x_1, x_2, y_1 and y_2 be points appearing in the clockwise order on the boundary of the closed disc. Then any arc from x_1 to y_1 meets any arc from x_2 to y_2.* □

Proof of Theorem 9.7. The vertices of B will be given by vertices $x_{i,j}$ indexed by pairs (i, j) of natural numbers with $1 \leq i+j \leq 2k+1$. We add edges of three kinds. For any natural numbers i and j with $1 \leq i$ and $i+j \leq 2k$ we add an edge that is incident to $x_{i,j}$ with sign $-$ and incident to $x_{i,j+1}$ with sign $+$. Similarly for any natural numbers i and j with $1 \leq j$ and $i+j \leq 2k$ we add an edge that is incident to $x_{i,j}$ with sign $-$ and incident to $x_{i+1,j}$ with sign $+$. Finally, for any $i \leq 2k$ we add an edge incident to $x_{i,2k+1-i}$ and $x_{i+1,2k-i}$, with sign $-$ at both vertices. We call the edges of these three kinds *vertical*, *horizontal* and *diagonal edges* respectively – see Figure 9.1 for an illustration where $k = 2$. We set $\mathcal{X} := \{x_{0,j} \colon 1 \leq j \leq 2k+1\}$ and $\mathcal{Y} := \{x_{i,0} \colon 1 \leq i \leq 2k+1\}$.

We show that for any set $S \subseteq V(B)$ of size k there is an $\mathcal{X}$–$\mathcal{Y}$ path in $B - S$. First we construct a sequence of $\mathcal{X}$–$\mathcal{Y}$ paths $P_1, P_2, \dots, P_{2k+1}$ by taking the path P_i to be the concatenation of the path from $x_{0,i}$ to $x_{2k+1-i,i}$ consisting of horizontal edges, followed by the diagonal edge from $x_{2k+1-i,i}$ to $x_{2k+2-i,i-1}$ and then the path consisting of vertical edges to $x_{2k+2-i,0}$. Note that no vertex appears in more than two of these paths, so the number of these paths meeting S is at most $2k$. Since there are $2k+1$ of these paths, one of them avoids S.

Now suppose for a contradiction that there are two disjoint $\mathcal{X}$–$\mathcal{Y}$ paths P and Q. Let the initial vertices of P and Q be x_P and x_Q and let their final vertices be y_P and y_Q. Since deleting all the diagonal edges from B leaves a digraph in which no directed edge points into $\mathcal{X}$ or $\mathcal{Y}$, both P and Q must contain diagonal edges. Let z_P and z_Q be endvertices of those diagonal edges on P and Q respectively which are distinct from x_P, x_Q, y_P and y_Q. We may embed B in the disc in such a way that on the boundary we have, in clockwise order, first the elements of $\mathcal{X}$, then those of $\mathcal{Y}$, then all other $x_{i,j}$ with $i + j = 2k + 1$. Without loss of generality we may assume that in the clockwise order after $\mathcal{Y}$ on the boundary of the disc we first reach z_Q then z_P. Then disjointness of the arcs induced by the paths $z_P P y_P$ and $x_Q Q z_Q$ contradicts Lemma 9.8, completing the proof. □

We will show in Section 9.5 that the vertex-version of Menger's Theorem in the context of bidirected graphs follows from the edge-version, meaning that the counterexample above also implies the existence of a counterexample to the edge-version of Menger's Theorem in our context. It is also not difficult to construct counterexamples to the vertex-version directly with a slight modification of the construction above, as illustrated in Figure 9.2.

Theorem 9.9. *For each number $k \in \mathbb{N}$, there exists a bidirected graph B and distinct vertices x and y, such that there are no two edge-disjoint x–y paths and for each subset $S \subseteq E(B)$ of size k there exists an x–y path in $B - S$.* $\square$

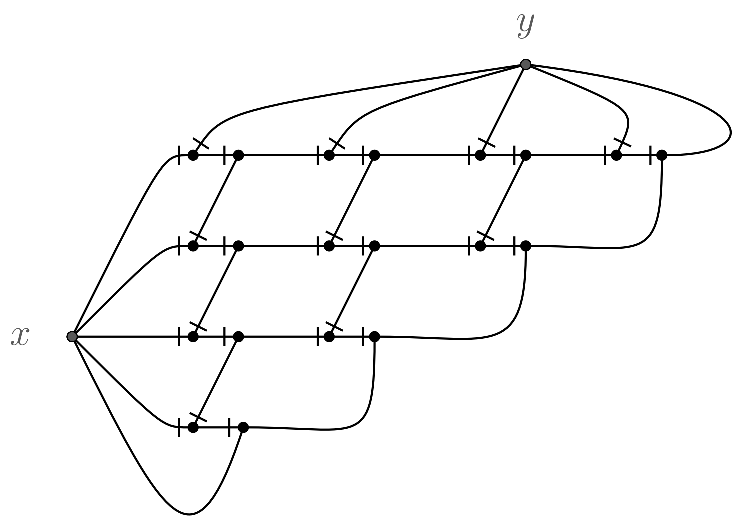

Figure 9.2: A bidirected graph containing neither two edge-disjoint x–y paths nor an edge set W of size ≤ 2 such that in $B - W$ there is no x–y path.

9.3 Appendages of Paths

In this section, we introduce the main tool for the proof of the edge-version of Menger's Theorem, Theorem 9.17. Given an arbitrary path P of a bidirected graph B starting in a vertex x, we define a set $A(P, x)$ of edges that includes much of the structure of B in the vicinity of P:

Definition 9.10. An edge set $A \subseteq E(B)$ is (P, x)-*admissible* if for any $\vec{e} \in \overleftrightarrow{A}$ there is an $\vec{e}$–x trail in $A \cup E(P)$. The *appendage* $A(P, x)$ of P and x is the union of all (P, x)-admissible sets.

Remark 9.11. Any union of (P, x)-admissible sets is again (P, x)-admissible. Thus, $A(P, x)$ is the maximal (P, x)-admissible set.

In the vertex-version of Menger's Theorem, Theorem 9.1, we restrict to bidirected graphs B and sets $\mathcal{X}$ of signed vertices such that B contains no non-trivial path that starts and ends in $\mathcal{X}$. However, for an edge-version of Menger's Theorem, forbidding such paths is not sufficient: joining the counterexample of Theorem 9.9 with a bidirected graph akin to the one depicted in Figure 9.3 by identifying the graphs at the respective vertices called x gives rise to a bidirected graph with no x'–x' trails whose internal vertices are all distinct[1]. But there are neither two disjoint x'–y paths nor is there a small set S of edges such that $B - S$ contains no x'–y path.

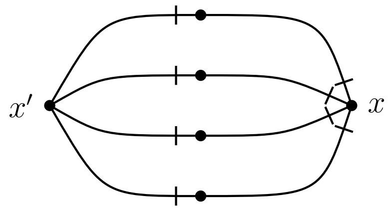

Figure 9.3: A bidirected graph with no x'–x' trails whose internal vertices are all distinct and the property that for any edge set W of size at most 3 there exists an x'-x path in $B - W$.

For the edge-version of Menger's Theorem, we thus forbid not only non-trivial x–x trails whose vertices are all distinct, but even non-trivial x–x trails in general:

Definition 9.12. A vertex x of a bidirected graph B is *edge-clean* if there exists no non-trivial x–x trail in B.

From now on we assume x to be edge-clean. Our aim is to show three essential properties of the appendage $A(P, x)$ that makes it a key tool in our proof of Theorem 9.17, the edge-version of Menger's Theorem in bidirected graphs. First, $A(P, x)$ is edge-disjoint from all paths starting in x that are edge-disjoint from P. Second, for any $v \in V(A(P, v)) \cup V(P)$ and any $\alpha \in \{+, -\}$ if there is an x–(v, α) path in B then there also is one in $A(P, x) \cup E(P)$ that coincides with P in the first edge. Third, any x–y walk contains an x–y path if it can be partitioned into a path in $A(P, x) \cup E(P)$ and a path in $E(B) \setminus (A(P, x) \cup E(P))$.

In the proof of Theorem 9.17, we replace the set $A(P, x) \cup E(P)$ by some auxiliary edges for a fixed path P. In this construction we do not want to remove paths starting in x that are edge-disjoint to P. Therefore it is essential that all paths starting in x which are disjoint to $E(P)$ are also disjoint to $A(P, x)$:

Lemma 9.13. *Let x be edge-clean and let Q be a path in B starting in x that is edge-disjoint from P. Then $E(Q) \cap A(P, x) = \emptyset$.*

Proof. Suppose not for a contradiction, and let $\vec{e}$ be the first edge of Q in $A(P, x)$. Since $A(P, x)$ is (P, x)-admissible, there is an $\vec{e}$–x-trail R in $A(P, x) \cup E(P)$. But then $Q\vec{e}R$ is a non-trivial x–x trail, contradicting the edge-cleanness of x. □

Corollary 9.14. *Let x be edge-clean. Then $A(P, x) \cup E(P)$ contains exactly one edge incident to x.*

Proof. Suppose not and let $e \in A(P, x) \setminus E(P)$ be incident to x. Then the path consisting of the single edge e contradicts Lemma 9.13. □

[1]We use this term rather than "x'–x' path", as the definition of a path requires all vertices to be distinct.

Now we turn our attention to the second essential property of appendages, which provides the existence of specific paths in $A(P, x) \cup E(P)$. We use this property in the proof of the edge-version of Menger's Theorem to ensure that certain paths are contained in $A(P, x) \cup E(P)$ by redirecting them if necessary.

Lemma 9.15. *Let x be edge-clean and let (v, α) be a signed vertex with $v \in V(A(P, x)) \cup V(P)$. Suppose that there is an x–(v, α) path in B. Then there is such a path in $A(P, x) \cup E(P)$.*

Proof. Let Q be an x–(v, α) path in B chosen to minimise its set of edges outside $A(P, x) \cup E(P)$. We show that Q is contained in $A(P, x) \cup E(P)$. Suppose not for a contradiction, and let $\vec{e}$ be the first edge of Q outside $A(P, x) \cup E(P)$.

Case 1: some vertex of Q after $\vec{e}$ lies on P. In this case, let w be the first such vertex along P. Note that $w \neq x$ since Q is a path. Let p be the edge preceding w on P and let q be the edge preceding w on Q.

Case 1.1: $\sigma(p, w) = \sigma(q, w)$. The path PwQ is an x–(v, α) path using fewer edges outside $A(P, x) \cup E(P)$ than Q, since Pw is disjoint from wQ per choice of w and $e \notin E(PwQ)$. This contradicts the minimality of Q.

Case 1.2: $\sigma(p, w) = -\sigma(q, w)$. We will show that the set $A(P, x) \cup E(Qw)$ is (P, x)-admissible. That is, for any oriented edge $\vec{f}$ with $f \in A(P, x) \cup E(Qw)$ there is an $\vec{f}$–x trail in $A(P, x) \cup E(Qw) \cup E(P)$. If $f \in A(P, x)$ this is clear, so suppose not. Thus, f is an edge of $\vec{e}Qw$. If $\vec{f} \in \vec{E}(Q)$ then $\vec{f}QwP^-$ is a suitable trail, by choice of w. If $\overleftarrow{f} \in \vec{E}(Q)$ then $\vec{f}\,Q^-$ is a suitable trail. Thus, $A(P, x) \cup E(Qw)$ is (P, x)-admissible, and so it is a subset of $A(P, x)$. This contradicts the assumption that e is not contained in $A(P, x)$.

Case 2: no vertex of Q after $\vec{e}$ lies on P. In this case, let w be the first vertex along Q after $\vec{e}$ in $V(A(P, x)) \cup V(P)$. There is such a vertex since v is a candidate. Let q be the edge preceding w on Q. Our first aim is to show that there is $\vec{a} \in \overleftrightarrow{A(P, x)}$ with startvertex w and such that $\sigma(a, w) = -\sigma(q, w)$. Indeed, as $w \in V(A(P, v))$ there is $\vec{b} \in \overleftrightarrow{A(P, x)}$ so that its endvertex is w. If $\sigma(b, w) = -\sigma(q, w)$ then we can just set $\vec{a} := \overleftarrow{b}$. Otherwise, we can take $\vec{a}$ as the second edge of any $\vec{b}$–x trail in $A(P, x) \cup E(P)$. Let R be any $\vec{a}$–x trail in $A(P, x) \cup E(P)$.

We will now show that $A(P, x) \cup E(Qw)$ is (P, x)-admissible, that is, for any oriented edge $\vec{f}$ with $f \in A(P, x) \cup E(Qw)$, there is an $\vec{f}$–x trail in $A(P, x) \cup E(Qw) \cup E(P)$. If $f \in A(P, x)$ this is clear, so suppose not. Thus, f is an edge of $\vec{e}Qw$. If $\vec{f} \in \vec{E}(Q)$, then we find a suitable trail by following Q from $\vec{f}$ to w and switching onto $\vec{a}R$. If $\overleftarrow{f} \in \vec{E}(Q)$, then $\vec{f}\,Q^-$ is a suitable trail. Thus, $A(P, x) \cup E(Qw)$ is (P, x)-admissible, and so it is a subset of $A(P, x)$. This contradicts our assumption that e is not contained in $A(P, x)$.

□

In the proof of Theorem 9.17, we will find some of the desired x–y paths in certain x–y walks. As mentioned in the introduction, an x–y walk in a bidirected graph does not necessarily contain an x–y path. We here provide a sufficient condition for the existence of an x–y path in an x–y walk:

Lemma 9.16. *Let y and (z, α) be (signed) vertices of B. Let Q be an x–(z, α) path and let R be a $(z, -\alpha)$–y path. If all edges of Q except possibly the last edge are contained in $E(P) \cup A(P, x)$ and R avoids $E(P) \cup A(P, x)$, then there is an x–y path S in $E(Q) \cup E(R)$. Furthermore, if R avoids x, then the first edge of S and the first edge of Q coincide.*

Proof. Let q be the first vertex of Q in $V(R)$. If $q = x$, then qR is the desired path. If $q = z$, then Q and R intersect only in $q = z$ and thus QzR is the desired path. Thus, suppose that q is an internal vertex of Q. Let e be the edge of Q preceding q and f be the edge of R succeeding q. If $\sigma(q, e) = -\sigma(q, f)$, then QqR forms the desired path.

Suppose for a contradiction that instead we have $\sigma(q, e) = \sigma(q, f)$. Then $qR^- zQ^- x$ and $qQzRqQ^- x$ are trails. We show that $E(Rq) \subseteq A(P, x)$, which contradicts the assumption that R avoids $A(P, x)$. More precisely, we prove that $A(P, x) \cup E(qQ) \cup E(Rq)$ is (P, x)-admissible: let $\vec{e} \in \vec{E}(qQ) \cup \vec{E}(Rq)$. Then $\vec{e} \in \vec{E}(qR^- zQ^- x)$ or $\vec{e} \in \vec{E}(qQzRqQ^- x)$. Note that both trails are contained in $A(P, x) \cup E(qQ) \cup E(Rq)$, as $E(Qq) \subseteq A(P, x)$ holds since $q \neq z$. This completes the proof. □

9.4 Menger's Theorem for Edge-Disjoint Paths

In this section we prove an edge-version of Menger's Theorem in bidirected graphs, from which we will then deduce our main result, the vertex-version given by Theorem 9.1, in the subsequent Section 9.5:

Theorem 9.17. *Let x and y be distinct vertices of a bidirected graph B, and suppose that x is edge-clean. Let $P_1, \dots, P_k$ be edge-disjoint x–y paths in B where $\vec{e}_i$ is the first edge of P_i for $i \in [k]$. Then precisely one of the following is true:*

- **(1)** *there is a set S of k edges of B such that $B - S$ contains no x–y path, or*
- **(2)** *there are $k + 1$ edge-disjoint x–y paths $P'_1, \dots, P'_{k+1}$ such that the first edge of P'_i is $\vec{e}_i$ for $i \in [k]$.*

The proof of Theorem 9.17 is inspired by the Böhme, Göring and Harant's proof [BGH01] of Menger's Theorem for digraphs. Its technical finesse makes it possible to handle the complex structure of bidirected graphs.

Proof. The proof is by strong induction on the sum of the lengths of the paths P_i. We may assume without loss of generality that $\sigma(e, x) = -$ for any edge e incident with x, since changing these signs doesn't affect what counts as an x–y path nor the edge-cleanness of x.

If the set $\{e_1, e_2, \ldots, e_k\}$ is as in (1) then we are done, so suppose not. Thus, there is an x–y path P_{k+1} containing no e_i. If this path is edge-disjoint from the P_i with $i \leq k$ then we are done, so suppose not. Let $\vec{e}$ be the first edge of P_{k+1} which lies on some other P_i. Without loss of generality we may assume that $e \in E(P_k)$. If $\overleftarrow{e}$ were in $\vec{E}(P_k)$ then $P_{k+1}\vec{e}P_k^-$ would be an x–x trail, contradicting the edge-cleanness of x. So we must instead have $\vec{e} \in \vec{E}(P_k)$. Let v be the startvertex of $\vec{e}$ and w its endvertex. We define R_1 as $P_k v$ and R_2 as $P_{k+1}v$. Note that R_1 and R_2 have length at least one.

All we will use about R_2 in the following argument is that it is edge-disjoint from all P_i and that it can be extended as a path by adding the edge e. Note that if we had begun with the paths $P_1, \ldots, P_{k-1}$ and $R_2 v P_k$ then R_1 would have the same properties with respect to this choice of paths. We will exploit this symmetry repeatedly in the following argument.

Let $A_i := A(R_i, x) \cup E(R_i)$. Note that by Lemma 9.13 A_1 is disjoint from $E(R_2 v P_k)$ and A_2 is disjoint from $E(R_1 v P_k)$. Similarly both A_1 and A_2 are disjoint from $E(P_i)$ for $i < k$. To apply the induction hypothesis we construct a suitable graph $\hat{B}$ using the sets A_1 and A_2.

We obtain $\hat{B}$ from B by modifying it in the following ways (see Figure 9.4):

- removing the edge e,
- adding a new edge $\hat{e}$ with endvertices x, w and $\sigma(\hat{e}, x) = -$ and $\sigma(\hat{e}, w) = \sigma(e, w)$,
- removing all edges in $A_1 \cup A_2$,
- adding a new vertex a,
- adding a new edge e_a from x to a with $\sigma(e_a, x) = -$ and $\sigma(e_a, a) = +$,
- adding, for each $i \in \{1, 2\}$ and each signed vertex (z, α) for which there is a non-trivial x–(z, α) path P with $E(P) \subseteq A_i$ (not just a trail) a new edge f with endvertices a, z and $\sigma(a, f) = -$ and with $\sigma(z, f) = \alpha$. We refer to f as $e(P)$ and to z as z_P.

We set $\vec{\hat{e}}$ and $\vec{e}_a$ to be the orientation of $\hat{e}$ and e_a that point away from x. Furthermore, we set $\vec{e}(P)$ to be the orientation of $e(P)$ that points away from a and $\vec{e}$ the orientation of e that points away from v. For $i < k$ let $\hat{P}_i := P_i$, and let $\hat{P}_k := \vec{\hat{e}} w P_k$. Note that $\hat{P}_k$ is shorter than P_k.

We now want to apply the induction hypothesis to $\hat{B}$ and the paths $\hat{P}_i$ for $i \leq k$. Since the sum of the lengths of the paths has decreased, we just need to check that x is still edge-clean in $\hat{B}$. So suppose for a contradiction that there is an x–x trail $\hat{Q}$ in $\hat{B}$.

Since all edges incident to x are incident with the same sign, no internal edge of $\hat{Q}$ is incident with x. Similarly, since e_a is the only edge incident to a with sign $+$, if $E(\hat{Q})$

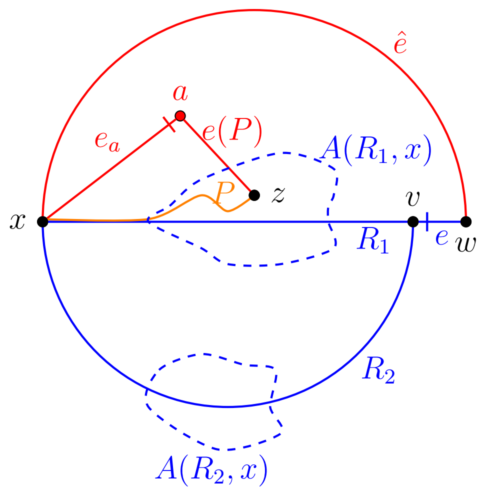

Figure 9.4: Construction of the auxiliary graph $\hat{B}$ in the proof of Theorem 9.17.

contains any $e(P)$, then that $e(P)$ must be adjacent to e_a, which in turn must be the first or last edge of $\hat{Q}$.

Our aim now is to construct an x–x trail Q in B, thus contradicting the edge-cleanness of x in B. Let $\vec{q}_1$ be the first edge and $\overleftarrow{q}_2$ the last edge of $\hat{Q}$. Replacing $\hat{Q}$ with its reversal and relabelling these edges if necessary, we have the following cases:

Case 1: both q_1 and q_2 are edges of B. We can set $Q := \hat{Q}$.

Case 2: $q_1 = \hat{e}$ and q_2 is an edge of B. We can set $Q := P_k w \hat{Q}$.

Case 3: $q_1 = e_a$ and q_2 is an edge of B. The second edge of $\hat{Q}$ must be of the form $\vec{e}(P)$ for some non-trivial path P in some A_i. So we can set $Q := P z_P \hat{Q}$.

Case 4: $q_1 = e_a$ and $q_2 = \hat{e}$. The second edge of $\hat{Q}$ must be of the form $\vec{e}(P)$ for some non-trivial path P. Without loss of generality $E(P)$ is contained in A_1. So we can set $Q := P z_P \hat{Q} w \overleftarrow{e} v R_2^-$.

In any case we reach the desired contradiction, so we can conclude that x is still edge-clean in $\hat{B}$. Thus, we can apply the induction hypothesis. This gives us two cases:

Case 1: there is a set $\hat{S}$ of k edges of $\hat{B}$ such that $\hat{B} - \hat{S}$ does not contain an x–y path. Since the paths $\hat{P}_i$ are edge-disjoint, $\hat{S}$ must consist of one edge from each of these paths and so cannot contain e_a or any $e(P)$. Let $S := \hat{S}$ if $\hat{e} \notin \hat{S}$ and let $S := (S \setminus \{\hat{e}\}) \cup \{e\}$ otherwise. Clearly S has size k. We show that $B - S$ contains no x–y path.

Suppose for a contradiction that there is such a path Q. To obtain our contradiction we construct an x–y path $\hat{Q}$ in $\hat{B} - \hat{S}$.

Case 1.1: $E(Q)$ is disjoint from $A_1 \cup A_2 \cup \{e\}$. We can set $\hat{Q} := Q$.

Case 1.2: $\vec{e} \in \vec{E}(Q)$ and $E(wQ)$ is disjoint from $A_1 \cup A_2$. We have $e \notin S$, which implies $\hat{e} \notin \hat{S}$. Thus, we can set $\hat{Q} := x \vec{\hat{e}} w Q$.

Case 1.3: otherwise. Since $v \in V(A_1)$ the fact that we are not in Case 1.1 implies that Q meets $V(A_1 \cup A_2) \setminus \{x\}$. Let z be the last vertex of Q in this set. Without loss of generality we have $z \in V(A_1) \setminus \{x\}$. The fact that we are not in Case 1.2 implies that even if e appears on Q some vertex in $V(A_1 \cup A_2) \setminus \{x\}$ (possibly v) must come after it. So the set $E(zQ)$ does not contain e. Let α be the sign with which Q arrives at z. Applying Lemma 9.15 to the x–(z, α) path Qz we see that there must be an x–(z, α) path P in A_1. Then we can set $\hat{Q} := x\vec{e}_a a\vec{e}(P)zQ$.

In each case, we find the desired contradiction, completing the proof of this case.

Case 2: there are edge-disjoint x–y paths $\hat{P}'_1, \hat{P}'_2, \ldots, \hat{P}'_{k+1}$ such that the first edge of $\hat{P}'_i$ is $\vec{e}_i$ for $i < k$ and the first edge of $\hat{P}'_k$ is $\vec{\hat{e}}$. In this case, we will construct edge-disjoint x–y paths $P'_1, P'_2, \ldots, P'_{k+1}$ in B such that the first edge of P'_i is $\vec{e}_i$ for $i \leq k$.

For $i < k$ the path $\hat{P}'_i$ avoids the edges $\hat{e}$ and e_a since $\vec{e}_i$ is the first edge of $\hat{P}'_i$, and therefore it avoids also any edge $e(P)$. Thus, $\hat{P}'_i$ is a path of B, and we set $P_i := \hat{P}'_i$ for any $i < k$. It remains to define P'_k and P'_{k+1}.

Case 2.1: the first edge of $\hat{P}'_{k+1}$ is not $\vec{e}_a$. Then $\hat{P}'_{k+1}$ is also a path of B as it avoids $\hat{e}$ since $\hat{e} \in P_k$. We set $P'_{k+1} := \hat{P}'_{k+1}$ and let P'_k be the path obtained by applying Lemma 9.16 to $R_1 v\vec{e}w$ and $w\hat{P}'_k$ with respect to A_1.

Case 2.2: the first edge of $\hat{P}'_{k+1}$ is $\vec{e}_a$ and its second edge is $\vec{e}(P)$ for some path P in A_2. We let P'_{k+1} be the path obtained by applying Lemma 9.16 to $Pz_P, z_P\hat{P}'_{k+1}$ with respect to A_2 and define P'_k as in the prior case.

Case 2.3: the first edge of $\hat{P}'_{k+1}$ is $\vec{e}_a$ and its second edge is $\vec{e}(P)$ for some path P in A_1. Let P'_k be the path that we obtain by applying Lemma 9.16 to $Pz_P, z_P\hat{P}'_{k+1}$ with respect to A_1 and let P'_{k+1} be the path obtained by applying Lemma 9.16 to $R_2 v\vec{e}w$ and $w\hat{P}'_k$ with respect to A_2.

Note that in any of these cases the paths $P'_1, \ldots, P'_{k+1}$ are edge-disjoint. To verify that $\vec{e}_k$ is indeed the first edge of P'_k, it suffices to show that in the construction of P'_k we applied Lemma 9.16 to a path whose first edge is $\vec{e}_k$ and a path that avoids x. The first path is Pz_P, if P is in A_1, and otherwise it is $R_1 v\vec{e}w$. Thus, the first edge of the first path is contained in A_1, and by Corollary 9.14 it is $\vec{e}_k$. The second path avoids the vertex x by choice. Thus, the paths $P'_1, \ldots, P'_{k+1}$ are as desired.

□

Theorem 9.17 implies Menger's Theorem for edge-disjoint paths in both undirected and directed graphs. To show this, we may regard any graph as a digraph by replacing each edge e with the two directed edges $\vec{e}$ and $\overleftarrow{e}$ (see [BJG08] for an in-depth explanation). Similarly, we may regard any digraph as a bidirected graph by viewing a directed edge $\vec{e}$ from x to y as an edge e with endvertices x and y and signs $-$ at x and $+$ at y.

We then obtain the desired version of Menger's Theorem for vertices x and y in a digraph D by considering the digraph D' obtained from D by deleting all edges that point towards x. Note that x is edge-clean in D' since no trail ends in x. If there is no x–y path in $D' - S$ for some $S \subseteq V(D')$, then there is no x–y path in $D - S$. Furthermore, any x–y path in D' is also an x–y path in D.

9.5 Menger's Theorem for Vertex-Disjoint Paths

In this section we deduce the vertex-disjoint version of Menger's Theorem for bidirected graphs, Theorem 9.1, from the above shown edge-version, Theorem 9.17. Theorem 9.1 will be formulated in terms of vertex-disjoint $\mathcal{X}$–$\mathcal{Y}$ paths between two sets $\mathcal{X}$ and $\mathcal{Y}$ of signed vertices of a bidirected graph B. Analogously to the edge-version where the startvertex x of the paths had to be edge-clean, we have to require the set $\mathcal{X}$ of signed vertices to be *clean*, which is defined as follows.

Definition 9.18. A set $\mathcal{X}$ of signed vertices of a bidirected graph B is *clean* if B contains no non-trivial path starting and ending in $\mathcal{X}$.

We remark that a clean set $\mathcal{X}$ of signed vertices may contain both $(x, +)$ and $(x, -)$ for a vertex x.

Let us now recall the vertex-disjoint version of Menger's Theorem for bidirected graphs from the introduction:

Theorem 9.1. *Let $\mathcal{X}$ and $\mathcal{Y}$ be sets of signed vertices of a bidirected graph B, and suppose that $\mathcal{X}$ is clean. Let $P_1, \ldots, P_k$ be vertex-disjoint $\mathcal{X}$–$\mathcal{Y}$ paths in B where P_i starts in $v_i \in V(\mathcal{X})$ for $i \in [k]$. Then precisely one of the following is true:*

(1) *There is a set S of k vertices of B such that $B - S$ contains no $\mathcal{X}$–$\mathcal{Y}$ path.*

(2) *There are $k+1$ vertex-disjoint $\mathcal{X}$–$\mathcal{Y}$ paths $P'_1, \ldots, P'_{k+1}$ in B where P'_i starts in v_i for $i \in [k]$.*

Note that in Theorem 9.1 we only require $\mathcal{X}$ to be clean, but not $\mathcal{Y}$ – just as we only require x (and not y) to be edge-clean in Theorem 9.17; indeed, we do not need any assumptions on $\mathcal{Y}$. We further note that Theorem 9.1 cannot be strengthened by fixing the signed startvertices of the paths P_i rather than just their startvertices, see Figure 9.5 for a counterexample. Also note that the simplification that we do not allow bidirected graphs to have distinct edges e and f that have the same endvertices and the same signs at them does not affect Theorem 9.1, which still holds without this assumption. Indeed, one can easily deduce the statement for bidirected graphs with parallel edges by applying Theorem 9.1 to the bidirected graph obtained by subdividing each edge and giving the two edges incident to a newly arising subdivision vertex distinct signs at it.

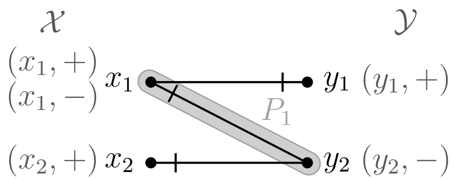

Figure 9.5: This bidirected graph contains two disjoint $\mathcal{X}$–$\mathcal{Y}$ paths, namely $x_2(x_2y_2,x_2,y_2)y_2$ and $x_1(x_1y_1,x_1,y_1)y_1$. However, while $P_1 := x_1(x_1y_2,x_1,y_2)y_2$ is an $\mathcal{X}$–$\mathcal{Y}$ path that starts in $(x_1,+) \in \mathcal{X}$, the graph contains no two disjoint $\mathcal{X}$–$\mathcal{Y}$ paths such that one of them starts in $(x_1,+)$.

Towards a proof of Theorem 9.1, let us first show that we may assume the $\mathcal{X}$–$\mathcal{Y}$ paths in B to be the only paths starting in $\mathcal{X}$ and ending in $\mathcal{Y}$, respectively. More precisely, by removing some edges from B and changing some signs at vertices of $V(\mathcal{X}) \cup V(\mathcal{Y})$, all non-trivial paths starting in $\mathcal{X}$ and ending in $\mathcal{Y}$ are indeed $\mathcal{X}$–$\mathcal{Y}$ paths, i.e. they are internally disjoint from $V(\mathcal{X}) \cup V(\mathcal{Y})$ and contain no trivial $\mathcal{X}$–$\mathcal{Y}$ paths:

Proposition 9.19. *Let B be a bidirected graph and $\mathcal{X}, \mathcal{Y}$ sets of signed vertices in B. Then there exists a bidirected graph B' with $V(B') = V(B)$, $E(B') \subseteq E(B)$ and a set of signed vertices $\mathcal{X}' \subseteq \mathcal{X}$ in B' such that*

- **(1)** *any $\mathcal{X}'$–$\mathcal{Y}$ path in B' is an $\mathcal{X}$–$\mathcal{Y}$ path in B and vice versa,*
- **(2)** *any path in B' starting in $\mathcal{X}'$ and ending in $\mathcal{Y}$ is an $\mathcal{X}'$–$\mathcal{Y}$ path, and*
- **(3)** *there is no trivial $\mathcal{X}'$–$\mathcal{X}'$ path in B'.*

Moreover, if $\mathcal{X}$ is clean in B, then $\mathcal{X}'$ is clean in B'.

Note that, by (1), Theorem 9.1, or more generally any version of Menger's Theorem concerning vertex-disjoint $\mathcal{X}$–$\mathcal{Y}$ paths, holds for B, $\mathcal{X}$ and $\mathcal{Y}$ if and only if it holds for B', $\mathcal{X}'$ and $\mathcal{Y}$.

Proof. We construct B' and $\mathcal{X}'$ from B and $\mathcal{X}$ by performing the following modifications for every $v \in V(\mathcal{X}) \cup V(\mathcal{Y})$ simultaneously:

- **(a)** If there is $\alpha \in \{+,-\}$ with $(v,\alpha) \notin \mathcal{X} \cup \mathcal{Y}$, we delete all edges incident to v that have sign α at v.
- **(b)** If $(v,+)$ and $(v,-)$ are either both contained in $\mathcal{X} \setminus \mathcal{Y}$ or both contained in $\mathcal{Y} \setminus \mathcal{X}$, we change the signs of all edges incident to v at v to $+$. Furthermore, if $(v,-) \in \mathcal{X}$, then we remove $(v,-)$ from $\mathcal{X}$.

(c) If there is $\alpha \in \{+,-\}$ such that $(v,\alpha) \in \mathcal{X}$ and $(v,-\alpha) \in \mathcal{Y}$, we isolate v, i.e. we delete all edges incident with v. If $(v,-\alpha) \in \mathcal{X}$, then we also delete $(v,-\alpha)$ from $\mathcal{X}$.

With this construction, we in particular obtain the following properties of B':

(i) $V(\mathcal{X}) \cup V(\mathcal{Y}) = V(\mathcal{X}') \cup V(\mathcal{Y})$, and all edges incident to $v \in V(\mathcal{X}') \cup V(\mathcal{Y})$ have the same sign at v in B'.

(ii) If an edge of B' is incident to a vertex $w \notin V(\mathcal{X}') \cup V(\mathcal{Y})$, then it has the same sign at w in B' and in B.

(iii) If a vertex v forms a trivial $\mathcal{X}$–$\mathcal{Y}$ path in B, then it is an isolated vertex in B' and forms a trivial $\mathcal{X}'$–$\mathcal{Y}$ path in B'.

We begin by proving (1). For trivial paths, the statement is true by (iii) and since $\mathcal{X}' \subseteq \mathcal{X}$. So let P be a non-trivial $\mathcal{X}$–$\mathcal{Y}$ path in B. Then neither its start- nor its endvertex forms a trivial $\mathcal{X}'$–$\mathcal{Y}$ path in B' since they each do not form a trivial $\mathcal{X}$–$\mathcal{Y}$ path in B. In particular, we did not apply modification (c) to the start- or endvertex of P. As all internal vertices of P are disjoint to $V(\mathcal{X}) \cup V(\mathcal{Y})$, we did not apply any modification to these vertices. This implies that all edges of P are edges in B'. By (ii) P is thus also a path in B'. Clearly, it is internally disjoint to $V(\mathcal{X}') \cup V(\mathcal{Y})$. Finally, since we applied either modification (a) or (b) to its start- and endvertex, the path P in B' starts in $\mathcal{X}'$ and ends in $\mathcal{Y}$.

Let P' be a non-trivial $\mathcal{X}'$–$\mathcal{Y}$ path in B'. Then the path P' is, by (i) internally disjoint to $V(\mathcal{X}) \cup V(\mathcal{Y}) = V(\mathcal{X}') \cup V(\mathcal{Y})$ and thus also a path in B by (ii). As (1) holds for trivial paths, neither its start- nor its endvertex forms a trivial $\mathcal{X}$–$\mathcal{Y}$ path in B. In particular, we did not apply modification (c) to its start- or endvertex. Thus, we applied either modification (a) or (b) to its start- and endvertex. This implies that P' in B starts in $\mathcal{X}$ and ends in $\mathcal{Y}$, which completes the proof of (1).

For (2), consider a non-trivial path P' in B' that starts in $\mathcal{X}'$ and ends in $\mathcal{Y}$. To show that P' is an $\mathcal{X}'$–$\mathcal{Y}$ path, we need to verify two properties: first, no internal vertex of P' is contained in $V(\mathcal{X}') \cup V(\mathcal{Y})$, which is true by (i). Second, P' contains no trivial $\mathcal{X}'$–$\mathcal{Y}$ path. Indeed, if a vertex $v \in P'$ is a trivial $\mathcal{X}'$–$\mathcal{Y}$ path in B', then it is also a trivial $\mathcal{X}$–$\mathcal{Y}$ path in B by (1) and hence isolated by (iii), which contradicts that P' is non-trivial.

For (3), note that the modifications (a) to (c) ensure that $(v,\alpha) \in \mathcal{X}'$ implies $(v,-\alpha) \notin \mathcal{X}'$.

For the 'moreover'-part, assume that $\mathcal{X}$ is clean in B and suppose for a contradiction that x is not edge-clean in B'. Then there exists a non-trivial path P in B' that starts and ends in $\mathcal{X}'$. By (i), no internal vertex of P is contained in $V(\mathcal{X}) \cup V(\mathcal{Y})$. Furthermore, by (ii), P still forms a path in B. Thus, P is a non-trivial path in B starting and ending in $\mathcal{X}$, which contradicts the cleanness of $\mathcal{X}$ in B. $\square$

From now on, let B be a bidirected graph and let $\mathcal{X}$ and $\mathcal{Y}$ be sets of signed vertices of B. By Proposition 9.19 we may assume that B satisfies (2) and (3). To prove Theorem 9.1 for B, $\mathcal{X}$ and $\mathcal{Y}$, we now carefully construct an auxiliary graph $\hat{B}$ with two specified vertices x and y, such that Theorem 9.17, our edge-disjoint version of Menger's Theorem in bidirected graphs, applied to $\hat{B}$, x and y allows us to deduce Theorem 9.1 for B, $\mathcal{X}$ and $\mathcal{Y}$.

The construction of $\hat{B}$, however, turns out to be even more general: it provides a framework to deduce a vertex-version of Menger's Theorem in bidirected graphs from an edge-version. In particular, our construction of $\hat{B}$ is independent of the cleanness assumption on $\mathcal{X}$ as in Theorem 9.1, and the cleanness of $\mathcal{X}$ in B will only be used to make Theorem 9.17 applicable to $x \in \hat{B}$ in that it implies x to be edge-clean.

In order to make the above approach work, we construct $\hat{B}$, x and y such that the vertices x and y of $\hat{B}$ represent the sets $\mathcal{X}$ and $\mathcal{Y}$ of signed vertices of B, respectively. They do so in such a way that any set of vertex-disjoint $\mathcal{X}$–$\mathcal{Y}$ paths in B correspond to edge-disjoint x–y paths in $\hat{B}$ and vice versa.

To show that a vertex-version of Menger's Theorem for B, $\mathcal{X}$ and $\mathcal{Y}$ can be deduced from an edge-version of Menger's Theorem for $\hat{B}$, x and y, we then verify two properties of our construction. We first show that an edge-separator for x and y in $\hat{B}$ corresponds to a vertex-separator for $\mathcal{X}$ and $\mathcal{Y}$ in B of at most the same size. Secondly, we consider vertex-disjoint $\mathcal{X}$–$\mathcal{Y}$ paths $P_1, \dots, P_k$ in B. By our construction, they correspond to edge-disjoint x–y paths $\hat{P}_1, \dots, \hat{P}_k$ in $\hat{B}$. We then show that $k+1$ edge-disjoint x–y paths in $\hat{B}$ such that all but one start in the same edges as $\hat{P}_1, \dots, \hat{P}_k$ correspond to $k+1$ vertex-disjoint $\mathcal{X}$–$\mathcal{Y}$ paths such that all but one start in the same vertices as $P_1, \dots, P_k$.

Finally, we show that our construction indeed allows us to apply Theorem 9.17 to $\hat{B}$, x and y under the assumptions of Theorem 9.1 on B, $\mathcal{X}$ and $\mathcal{Y}$. More precisely, we prove that the cleanness of $\mathcal{X}$ in B implies that x is edge-clean in $\hat{B}$.

Let us now construct $\hat{B} = (\hat{G}, \hat{\sigma})$ from $B = (G, \sigma)$ (see Figure 9.6 for an illustration of this construction). For every vertex $v \in B$, the vertex set of $\hat{B}$ contains two vertices v^+ and v^- which are connected in $\hat{B}$ by an edge with sign $+$ at v^- and sign $-$ at v^+. An edge $e \in E(B)$ which is incident to u and v in B then transfers to an edge $\hat{e} \in E(\hat{B})$ that is incident to $u^{\sigma(u,e)}$ and $v^{\sigma(v,e)}$. The half-edges of $\hat{e}$ have the same signs as the respective half-edges of e, that is $\hat{\sigma}(u^{\sigma(u,e)}, \hat{e}) := \sigma(u, e)$ and $\hat{\sigma}(v^{\sigma(v,e)}, \hat{e}) := \sigma(v, e)$.

We finally add two new vertices x and y to $V(\hat{B})$, which shall represent the sets $\mathcal{X}$ and $\mathcal{Y}$, respectively, as follows. The vertex x is adjacent to those $v^\alpha \in V(\hat{B})$ with $(v, -\alpha) \in \mathcal{X}$, and the respective edge in $\hat{B}$ has sign α at v^α and sign $-$ at x. Analogously, the vertex y is adjacent to every $v^\beta \in V(\hat{B})$ with $(v, -\beta) \in \mathcal{Y}$, and the respective edge has sign β at v^β and sign $-$ at y.

Observe that with this construction of $\hat{B} = (\hat{G}, \hat{\sigma})$, the graph $\hat{G}$ has no parallel edges. For notational simplicity, we will thus identify an (oriented) edge in $\hat{B}$ with the (ordered) pair of its endvertices in what follows.

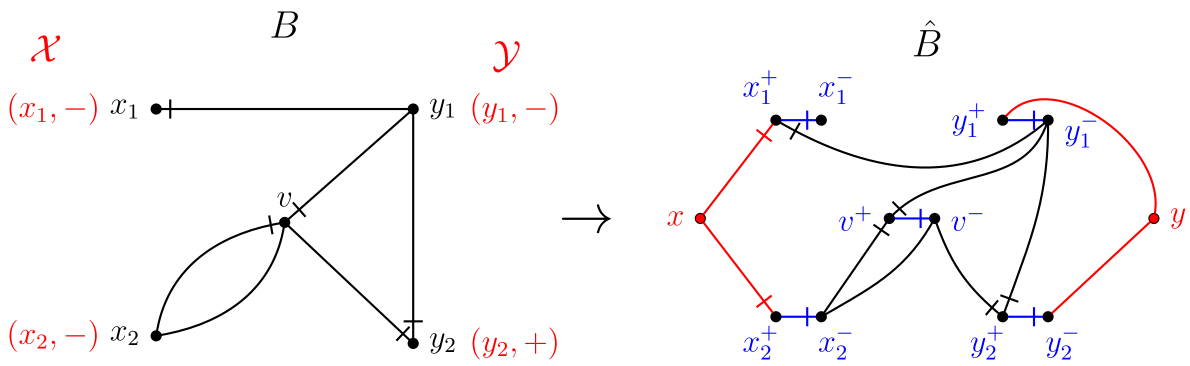

Figure 9.6: The construction of $\hat{B}$ from the bidirected graph B with specified sets $\mathcal{X}$ and $\mathcal{Y}$ of signed vertices as in Section 9.5.

Now we turn our attention to the relation between paths in B and paths in $\hat{B}$. Let $P = v_1\vec{e}_1 v_2 \dots v_{n-1}\vec{e}_{n-1} v_n$ be any $\mathcal{X}$–$\mathcal{Y}$ path in B, and write $\alpha_\ell := \sigma(v_\ell, e_\ell)$ and $\beta_\ell := \sigma(v_{\ell+1}, e_\ell)$ for all $\ell < n$. Since P is an $\mathcal{X}$–$\mathcal{Y}$ path in B, we have $\beta_\ell = -\alpha_{\ell+1}$ for all $\ell < n$ as well as $(v_1, \alpha_1) \in \mathcal{X}$ and $(v_n, \beta_{n-1}) \in \mathcal{Y}$. Thus, the construction of $\hat{B}$ guarantees that

$$\begin{aligned}\hat{P} :=& x(x, v_1^{-\alpha_1})v_1^{-\alpha_1}(v_1^{-\alpha_1}, v_1^{\alpha_1})v_1^{\alpha_1}(v_1^{\alpha_1}, v_2^{\beta_1})v_2^{\beta_1} \dots \\ & \dots v_{n-1}^{\alpha_{n-1}}(v_{n-1}^{\alpha_{n-1}}, v_n^{\beta_{n-1}})v_n^{\beta_{n-1}}(v_n^{\beta_{n-1}}, v_n^{-\beta_{n-1}})v_n^{-\beta_{n-1}}(v_n^{-\beta_{n-1}}, y)y\end{aligned}$$

is an x–y path in $\hat{B}$. Moreover, if $\mathcal{X}$–$\mathcal{Y}$ paths $P_1, \dots, P_k$ are vertex-disjoint, then the corresponding x–y paths $\hat{P}_1, \dots, \hat{P}_k$ in $\hat{B}$ are internally vertex-disjoint and in particular edge-disjoint.

With this transfer of $\mathcal{X}$–$\mathcal{Y}$ paths in B to x–y paths in $\hat{B}$ at hand, we show that an edge-separator of x and y in $\hat{B}$ yields a vertex-separator of $\mathcal{X}$ and $\mathcal{Y}$ in B of at most the same size:

Lemma 9.20. *If there is a set $\hat{S} \subseteq E(\hat{B})$ of size at most k such that there is no x–y path in $\hat{B} - \hat{S}$, then there is a set $S \subseteq V(B)$ of size at most k such that there is no $\mathcal{X}$–$\mathcal{Y}$ path in $B - S$.*

Proof. Let $\hat{S}$ be a set of at most k edges in $\hat{B}$ such that there is no x–y path in $\hat{B} - \hat{S}$. We define a set S of k vertices of B as follows:

- if $\hat{e} \in \hat{S}$ has the form $v^\alpha v^{-\alpha}$, then we add v to S,
- if $\hat{e} \in \hat{S}$ has the form $v^\alpha w^\beta$, then we arbitrarily add one of v and w to S, and
- if $\hat{e} \in \hat{S}$ has the form $v^\alpha x$ or $v^\alpha y$, then we add v to S.

Then there is no $\mathcal{X}$–$\mathcal{Y}$ path P in $B - S$, as the corresponding path $\hat{P}$ in $\hat{B}$ would clearly avoid $\hat{S}$. $\square$

We now turn to the transfer of x–y trails in $\hat{B}$ to $\mathcal{X}$–$\mathcal{Y}$ paths in B. For this, let us first observe that such trails admit a very particular structure due to the construction of $\hat{B}$:

Proposition 9.21. *Let T be either an x–y trail or a non-trivial x–x trail in $\hat{B}$. Then*

(1) *every internal vertex of T has the form v^α for some $v \in V(B)$ and $\alpha \in \{+,-\}$,*

(2) *every internal vertex of T is met by T exactly once, and*

(3) *precisely every second edge of T is of the form $(v^\alpha, v^{-\alpha})$.*

Proof. All edges incident to x or y have sign $-$ at x or y, so every internal vertex of a trail in $\hat{B}$ has the form v^α for some $v \in V(B)$ and $\alpha \in \{+,-\}$, which shows (1). For every vertex of this form v^α (i.e. every vertex of $\hat{B}$ other than x and y), there exists precisely one edge in $\hat{B}$ which has sign $-\alpha$ at v^α, namely the edge $v^\alpha v^{-\alpha}$. This implies that T uses the edge $(v^\alpha, v^{-\alpha})$ either immediately before or after meeting v^α. In particular, (3) holds. Since v^α is neither the start- nor the endvertex of T, it is met by T exactly once, which proves (2). □

Now let $\hat{T} := v_1\vec{e}_1v_2\vec{e}_2 \dots \vec{e}_{n-2}v_{n-1}\vec{e}_{n-1}v_n$ be either an x–y trail or a non-trivial x–x trail in $\hat{B}$. By Proposition 9.21 (1) and (2), the trail $\hat{T}$ meets all vertices of $\hat{B}$ but x at most once. Thus, its subtrail $\hat{T}' := \vec{e}_2\hat{T}\vec{e}_{n-2}$ is a path.

Now replace every subtrail of $\hat{T}'$ that has the form $v^\alpha(v^\alpha, v^{-\alpha})v^{-\alpha}$ by the vertex $v \in B$; note in particular that both $v_2\vec{e}_2v_3$ and $v_{n-2}\vec{e}_{n-2}v_{n-1}$ have this form. Moreover, replace any edge of $\hat{T}'$ that has the form (v^α, w^β) for distinct $v, w \in V(B)$ by the oriented edge in B that has sign α at v and sign β at w (such an edge exists by the construction of $\hat{B}$). Proposition 9.21 (3) then guarantees that the constructed P is indeed a path in B.

We show that P starts in $\mathcal{X}$; the case of P ending in $\mathcal{X}$ or in $\mathcal{Y}$ is symmetrical. The only neighbours of x in $\hat{B}$ are the vertices v^α with $(v, -\alpha) \in \mathcal{X}$. Proposition 9.21 (3) then implies that if $v_2 = v^\alpha$, then $\vec{e}_2 = (v^\alpha, v^{-\alpha})$ and so $\hat{e}_3$ has sign $-\alpha = -\hat{\sigma}(v_3, \hat{e}_2)$ at $v_3 = v^{-\alpha}$. Thus, our construction of P guarantees that P starts in $(v, -\alpha) \in \mathcal{X}$, as desired.

So if $\hat{T}$ is an x–y trail, then P starts in $\mathcal{X}$ and ends in $\mathcal{Y}$ and it is indeed an $\mathcal{X}$–$\mathcal{Y}$ path in B, since we assumed B to satisfy (2) in Proposition 9.19. And if $\hat{T}$ is a non-trivial x–x trail, then P starts and ends in $\mathcal{X}$ and it is non-trivial, since B also satisfies (3) by assumption.

This transfer of x–y trails from $\hat{B}$ to $\mathcal{X}$–$\mathcal{Y}$ paths in B now interacts with the transfer from B to $\hat{B}$ precisely in the desired way:

Lemma 9.22. *Let $P_1, \dots, P_k$ be $\mathcal{X}$–$\mathcal{Y}$ paths in B, and let $\hat{P}_1, \dots, \hat{P}_k$ be the corresponding x–y paths in $\hat{B}$. If there are $k+1$ edge-disjoint x–y paths $\hat{P}'_1, \dots, \hat{P}'_{k+1}$ in $\hat{B} = (\hat{G}, \hat{\sigma})$ where $\hat{P}'_i$ starts in the same edge as $\hat{P}_i$ for $i \in [k]$, then there are $k+1$ vertex-disjoint $\mathcal{X}$–$\mathcal{Y}$ paths $P'_1, \dots, P'_{k+1}$ in B where P'_i starts in the same vertex as P_i for $i \in [k]$.*

Proof. Let $P'_1, \dots, P'_{k+1}$ be the $\mathcal{X}$–$\mathcal{Y}$ paths in B obtained from $\hat{P}'_1, \dots, \hat{P}'_{k+1}$. Let $i \in [k]$ be arbitrary, and let v^α be the endvertex of the first edge of $\hat{P}'_i$. By construction, the path P'_i starts in the vertex v. Now our assumption on $\hat{P}'_i$ implies that the path $\hat{P}'_i$ starts in the same edge as $\hat{P}_i$. By the construction of $\hat{P}_i$ from P_i, the endvertex v^α of the first edge of $\hat{P}_i$ has the property that v is the startvertex of P_i. Thus, P_i and P'_i start in the same vertex, as desired.

It remains to prove that the paths $P'_1, \dots, P'_{k+1}$ are vertex-disjoint: this follows from Proposition 9.21 (3) together with our construction yields that if two distinct P'_i and P'_j would contain the same vertex v, then both $\hat{P}'_i$ and $\hat{P}'_j$ would have to contain the edge $v^\alpha v^{-\alpha}$, which contradicts that they are edge-disjoint. $\square$

With Lemma 9.20 and Lemma 9.22, we have now shown that the construction of $\hat{B}$ indeed provides a general framework to deduce a vertex-version of Menger's Theorem in bidirected graphs form an edge-version. We can thus finally deduce Theorem 9.1 from Theorem 9.17 by checking that x is edge-clean in $\hat{B}$ if $\mathcal{X}$ is clean in B:

Proof of Theorem 9.1. By Proposition 9.19, we may assume that B, $\mathcal{X}$ and $\mathcal{Y}$ satisfy (2) and (3), and this does not affect the cleanness of $\mathcal{X}$ in B. Let $\hat{B}$, x and y be constructed from B, $\mathcal{X}$ and $\mathcal{Y}$ as above, and suppose for a contradiction that x is not edge-clean, i.e. that there exists a non-trivial x–x trail $\hat{T}$ in $\hat{B}$. As shown above, the trail $\hat{T}$ then gives rise to a non-trivial path in B that starts and ends in $\mathcal{X}$, which contradicts the cleanness of $\mathcal{X}$ in B. Thus, we can apply Theorem 9.17 to $\hat{B}$, x, y and the k edge-disjoint x–y paths $\hat{P}_1, \dots, \hat{P}_k$ in $\hat{B}$ corresponding to the k vertex-disjoint $\mathcal{X}$–$\mathcal{Y}$ paths $P_1, \dots, P_k$ in B. Depending on the outcome of Theorem 9.17, Lemma 9.20 or Lemma 9.22 then complete the proof. $\square$

We remark that one can deduce a version of Theorem 9.1 in which an $\mathcal{X}$–$\mathcal{Y}$ path does not have to be internally disjoint from $V(\mathcal{X}) \cup V(\mathcal{Y})$ and may contain trivial $\mathcal{X}$–$\mathcal{Y}$ paths: apply Theorem 9.1 to the bidirected graph that is obtained by adding a new vertex for every $v \in V(\mathcal{X})$ which is joined to v by an edge with sign $-\alpha$ at v and sign α at the newly added vertex for any $(v, \alpha) \in \mathcal{X}$, and by adding vertices and edges for $v \in V(\mathcal{Y})$ in the same way. Then take $\mathcal{X}'$ and $\mathcal{Y}'$ as the respective sets of newly added vertices.

Theorem 9.1 implies Menger's Theorem for vertex-disjoint paths in both undirected graphs and directed graphs. As explained at the end of Section 9.4, we may regard any undirected or directed graph as a bidirected graph. We then obtain the desired version of Menger's Theorem for vertex sets X and Y in a digraph D by considering the sets $\mathcal{X} := \{(x, -) \mid x \in X\}$ and $\mathcal{Y} = \{(y, +) \mid y \in Y\}$ of signed vertices. Note that $\mathcal{X}$ is indeed clean since every path in the digraph D which ends in X does so in $(x, +)$ for some $x \in X$ and hence not in $\mathcal{X}$.

9.6 Polynomial Time Algorithm

In this section we show that the proof of the edge-version of Menger's Theorem for bidirected graphs, Theorem 9.17, can be turned into a polynomial time algorithm. More precisely, we show the following statement:

Theorem 9.23. *There exists a polynomial time algorithm that, given distinct vertices x and y of a bidirected graph B and k edge-disjoint x–y paths $P_1, \dots, P_k$ in B where P_i starts in $\vec{e}_i$ for $i \in [k]$, finds either $k+1$ edge-disjoint x–y paths $P'_1, \dots, P'_{k+1}$ in B where P'_i starts in $\vec{e}_i$ for $i \in [k]$ or a set S of k edges of B such that $B - S$ does not contain any x–y path.*

As a consequence, there exists a polynomial time algorithm for Theorem 9.1 since the transfer process introduced in Section 9.5 can be computed in polynomial time.

For the proof of Theorem 9.23 note that the proof of Theorem 9.17 uses a recursive procedure by induction on the sum of the lengths of the paths P_i; it is thus enough to show that each recursion step runs in polynomial time as the number of recursion steps is bounded by $|E(B)|$.

To show that each recursion step runs in polynomial time, we will in particular need polynomial time algorithms to find paths and trails between two fixed signed vertices as well as the appendage of a given path. One key ingredient to this is Edmonds' celebrated Blossom Algorithm:

Theorem 9.24 ([Edm65]). *There exists a polynomial time algorithm which, given a graph G, either finds a perfect matching in G or determines that G has no perfect matching.*

We first show that a path between two signed vertices can be found in polynomial time.

Lemma 9.25. *There exists a polynomial time algorithm that, given a bidirected graph B and signed vertices (x, α) and (y, β) of B with $x \neq y$, either finds an (x, α)–(y, β) path in B or determines that no such path exists.*

Proof. We first describe an algorithm which runs in polynomial time and then prove its correctness.

The algorithm starts by checking if there exists an edge in B with sign α at x and sign β at y, which then forms our desired path. If no such edge exists, then the algorithm considers the graph $B' := B - \{x, y\}$ and computes the set $\mathcal{X}'$ of all signed vertices (v, γ) of B' for which there exists an (x, α)–(v, γ) edge and the set $\mathcal{Y}'$ which is defined analogously with respect to y. As in Section 9.5, the algorithm then constructs the auxiliary graph $\hat{B}' = (\hat{G}', \hat{\sigma}')$ from B', $\mathcal{X}'$ and $\mathcal{Y}'$; let x' and y' be the two vertices of $\hat{B}'$ representing $\mathcal{X}'$ and $\mathcal{Y}'$, respectively. Note that the construction of $\hat{B}'$ can be done in polynomial time.

The algorithm then uses the polynomial time algorithm of Theorem 9.24 to construct a perfect matching of $\hat{G}'$ if possible. If there is a perfect matching, then the algorithm computes an (x,α)–(y,β) path in B from it in polynomial time (see below). Otherwise, the algorithm returns that there is no (x,α)–(y,β) path in B. This concludes our description of the polynomial time algorithm.

For the proof of its correctness, let $\hat{e}_v$ denote the unique edge of $\hat{G}'$ between v^+ and v^- for every vertex $v \in V(B')$. Then $M := \{\hat{e}_v \mid v \in V(B')\}$ is a matching in $\hat{G}'$, and every vertex of $\hat{G}'$ except from x' and y' is incident to an edge in M.

If there is a perfect matching M' in $\hat{G}'$, then let H' be the subgraph of $\hat{G}'$ with edge set $M \cup M'$. By the definition of M, the component of H' containing x' is an x'–y' path $\hat{P}$. The construction of $\hat{B}'$ together with the definition of M implies that such a path $\hat{P}$ in $\hat{G}'$ induces an x'–y' path $\hat{P}'$ in $\hat{B}'$. As in the proof of Theorem 9.1, this path $\hat{P}'$ then induces an $\mathcal{X}'$–$\mathcal{Y}'$ path P' in B'. The definition of $\mathcal{X}'$ and $\mathcal{Y}'$ then allows us to extend the path P' to an (x,α)–(y,β) path in B.

Conversely, we show that if there exists an (x,α)–(y,β) path in B, then there is a perfect matching in $\hat{G}$. Let $P = x\vec{e}_1 v_1 \dots v_{\ell-1}\vec{e}_\ell y$ be an (x,α)–(y,β) path in B. Then $P' = v_1 \dots v_{\ell-1}$ is an $\mathcal{X}'$–$\mathcal{Y}'$ path in B' by the definition of $\mathcal{X}'$ and $\mathcal{Y}'$. As in the proof of Theorem 9.1, the path P' induces an x'–y' path $\hat{P}'$ in $\hat{B}'$. The respective path $\hat{Q}'$ in $\hat{G}'$ alternates between edges in M and not in M. Note that all the vertices of $\hat{G}'$ but the startvertex and endvertex of $\hat{Q}'$ are incident to edges in M. This means that no vertex of $\hat{Q}'$ is incident to an edge in $M \setminus E(\hat{Q}')$. It follows that the symmetric difference $(M \setminus E(\hat{Q}')) \cup (E(\hat{Q}') \setminus M)$ is a perfect matching in $\hat{G}'$, as desired. □

We next extend Lemma 9.25 from paths to trails. Our main tool for this is the line graph of a bidirected graph which we introduce below. For this we need one more definition: an *orientation of a bidirected graph* B is a map $\nu\colon E(B) \to \overleftrightarrow{E}(B)$ such that $\nu(e) \in \{\vec{e}, \overleftarrow{e}\}$ for any $e \in E(B)$.

Definition 9.26. Given a bidirected graph $B = (G,\sigma)$ and an orientation ν of B, we define the *line graph of B with respect to ν*, denoted as $L := L_\nu = (H,\tau)$, as follows. The graph H has vertex set $E(G)$ and contains an edge $\{e_1,e_2\} \in E(H)$ with label v for any two distinct edges $e_1, e_2 \in E(G)$ and a vertex $v \in e_1 \cap e_2$ satisfying $\sigma(v,e_1) = -\sigma(v,e_2)$. The signing τ of the half-edges $\mathbf{E}(H)$ assigns to an edge $\{e_1,e_2\} \in E(H)$ with label v the sign

$$\tau(e_i,\{e_1,e_2\}) := \begin{cases} + & \text{if } v \text{ is endvertex of } \nu(e_i), \\ - & \text{if } v \text{ is startvertex of } \nu(e_i), \end{cases}$$

for $i = 1,2$.

Fix an arbitrary orientation of a bidirected graph B, and let L be the corresponding line graph of B. Just as for line graphs of (un)directed graphs, the above definition of L implies that walks in B induce walks in L and vice versa. Indeed, a walk $v_0\vec{e}_1 v_1 \dots v_{n-1}\vec{e}_n v_n$ in B induces the walk $e_1\vec{a}_1 e_2 \dots e_{n-1}\vec{a}_{n-1} e_n$ in L where $\vec{a}_i$ is the directed edge from e_i to e_{i+1} with label v_i for any $i \in [n-1]$. Vice versa, consider a walk $W = e_1\vec{a}_1 e_2 \dots e_{n-1}\vec{a}_{n-1} e_n$ in L which has label v_i at the edge a_i for any $i \in [n-1]$, and let $v_0 \in e_1 \setminus \{v_1\}$ and $v_n \in e_n \setminus \{v_{n-1}\}$. Writing $\vec{e}_i = (e_i, v_{i-1}, v_i)$ for $1 \leq i \leq n$, the walk W in L then induces the walk $v_0\vec{e}_1 v_1 \dots \vec{e}_n v_n$ in B.

The above correspondences between walks in B and walks in L in particular imply that trails in B induce paths in L and paths in L induce trails in B. We can thus compute trails in B in polynomial time by applying Lemma 9.25 to a suitable line graph of B.

Corollary 9.27. *There exists a polynomial time algorithm that, given a bidirected graph B and signed vertices (x, α) and (y, β) of B with $x \neq y$, either finds an (x, α)–(y, β) trail in B or determines that no such trail exists.*

Proof. Choose an orientation ν of B such that the startvertex of $\nu(e)$ is x for all $e \in E(B)$ incident to x and such that the endvertex of $\nu(f)$ is y for all $f \in E(B)$ incident to y. Let L be the line graph of B with respect to ν. Then the result follows by applying Lemma 9.25 to every pair $(e, +)$ and $(f, -)$ of signed vertices of L such that in B the edge e has sign α at x and the edge f has sign β at y; the choice of ν here implies that any such $(e, +)$–$(f, -)$ path in L indeed translates to an x–y trail in B starting in x with sign α and ending in y with sign β. $\square$

As our polynomial time algorithm for Theorem 9.23 follows the proof of Theorem 9.17, we will also need to compute the appendage of a path in polynomial time. For this, we first prove a lemma which reduces this problem to finding 'ear trails' in polynomial time which are defined as follows.

Let P be a path in a bidirected graph B, and let v be the startvertex of P. Let A be a (P, v)-admissible set. The set $\mathrm{Bon}(P, v, A)$ is then defined to contain $(v, +)$ and $(v, -)$ as well as all signed vertices (w, α) of B for which there exists an oriented edge $\vec{e} \in \vec{A} \cup \vec{E}(P)$ that points at w with sign $-\alpha$. Note that for each $(w, \alpha) \in \mathrm{Bon}(P, v, A)$, there is a v–$(w, -\alpha)$ trail in $A \cup E(P)$. A non-trivial trail T in B is a (P, v, A)-*quasi-ear trail* if both the signed start- and endvertex of T are contained in $\mathrm{Bon}(P, v, A)$ and the first and last edge of T are not in $\overleftrightarrow{A} \cup \overleftrightarrow{E}(P)$. If this latter condition holds additionally for all edges of T, that is $E(T) \cap (A \cup E(P)) = \emptyset$, then T is an (P, v, A)-*ear trail.*

We observe that for every (P, v, A)-ear trail T in B, the set $A \cup E(T)$ is (P, v)-admissible. Indeed, the signed endvertex (w, α) of T is in $\mathrm{Bon}(P, v, A)$ and thus there exists a v–$(w, -\alpha)$ trail T_w in $A \cup E(P)$. So for every $\vec{e} \in \vec{E}(T)$ the concatenation of $\vec{e}T$ and T_w^- is a $\vec{e}$–v trail in $A \cup E(P) \cup E(T)$. Similarly, there exists an $\vec{e}$–v trail in $A \cup E(P) \cup E(T)$ for every $\vec{e} \in \vec{E}(T^-)$.

Lemma 9.28. *Let P be a path in a bidirected graph B, and let v be the startvertex of P. For every non-maximal (P, v)-admissible set $A \subsetneq \mathcal{A}(P, v)$, there is either an edge $e \in E(P) \setminus A$ such that $A \cup \{e\}$ is (P, v)-admissible or there exists an (P, v, A)-ear trail.*

Proof. Assume that there is no edge $e \in E(P) \setminus A$ such that $A \cup \{e\}$ is (P, v)-admissible. This implies the following observation which we will use repeatedly throughout this proof: let Q be a v–$\vec{e}$ trail in $A(P, v) \cup E(P)$ for some orientation $\vec{e}$ of e. Then $\vec{e} \notin \vec{E}(P^-) \setminus \overleftrightarrow{A}$ as otherwise $A \cup \{e\}$ is (P, v)-admissible witnessed by $Q\vec{e}$ and $P\overleftarrow{e}$, a contradiction to the assumption.

The proof now proceeds in two steps. We first find a (P, v, A)-quasi-ear trail in B. Secondly, we show that every (P, v, A)-quasi-ear trail contains a (P, v, A)-ear trail as a subtrail.

Since A is a non-maximal (P, v)-admissible set, there exists $e \in A(P, v) \setminus A$. The above observation then implies that no orientation $\vec{e}$ of e can be in $\vec{E}(P^-)$, since $e \in A(P, v)$ implies the existence of a v–$\vec{e}$ trail in $A(P, v) \cup E(P)$; so we have $e \in A(P, v) \setminus (A \cup E(P))$.

Fix an arbitrary orientation $\vec{e}$ of e. Since $e \in A(P, v)$, there is a v–$\vec{e}$ trail T in $A(P, v) \cup E(P)$, and we let $\vec{f}$ be the first edge on T not contained in $\overleftrightarrow{A} \cup \vec{E}(P)$; note that such an edge $\vec{f}$ exists since $\vec{e}$ is a suitable choice. The above observation then implies $\vec{f} \notin \vec{E}(P^-)$, and we hence have $f \in A(P, v) \setminus (A \cup E(P))$. Thus, there exists a v–$\overleftarrow{f}$ trail T' in $A(P, v) \cup E(P)$ and, as above, we let $\vec{f}'$ be the first edge on T' not contained in $\overleftrightarrow{A} \cup \vec{E}(P)$. Again, we have $\vec{f}' \in A(P, v) \setminus (A \cup E(P))$ by the above observation. By this construction, $\vec{f}'T'\overleftarrow{f}$ is a (P, v, A)-quasi-ear trail in B.

Now let T be a (P, v, A)-quasi-ear trail in B. By definition, T is a (P, v, A)-ear trail if and only if $E(T) \cap (A \cup E(P)) = \emptyset$. So suppose that T is not itself a (P, v, A)-ear trail, and let $\vec{f}T\vec{f}'$ be a maximal subtrail of T in $A \cup E(P)$. Since the first and the last edge of T are not in $\overleftrightarrow{A} \cup \overset{\leftrightarrows}{E}(P)$, there exist an edge $\vec{g}$ preceding $\vec{f}$ on T and an edge $\vec{g}'$ succeeding $\vec{f}'$ on T. We now claim that at least one of $T\vec{g}$ and $\vec{g}'T$ is again a (P, v, A)-quasi-ear trail. This then concludes the proof: since both $T\vec{g}$ and $T\vec{g}'$ contain strictly fewer edges in $A \cup E(P)$, we can recursively apply the above step until we obtain a (P, v, A)–ear trail.

Both $\vec{g}$ and $\vec{g}'$ are not in $\overleftrightarrow{A} \cup \overset{\leftrightarrows}{E}(P)$ by construction. So by the definition of $\mathrm{Bon}(P, v, A)$, the claim holds if $\vec{f} \in \overleftrightarrow{A} \cup \vec{E}(P^-)$ or $\vec{f}' \in \overleftrightarrow{A} \cup \vec{E}(P)$. Thus, we can assume that $\vec{f} \in \vec{E}(P) \setminus \overleftrightarrow{A}$ and $\vec{f}' \in \vec{E}(P^-) \setminus \overleftrightarrow{A}$ since $f, f' \in A \cup E(P)$. Let $\vec{h}$ be the first edge of $\overset{\leftrightarrows}{E}(\vec{f}T\vec{f}')$ along P. If we have $\vec{h} \in \vec{E}(T)$, then $P\vec{h}T\vec{f}'$ is a trail. This contradicts our above observation since $\vec{f}' \in \vec{E}(P^-) \setminus \overleftrightarrow{A}$. Otherwise, $\vec{h} \in \vec{E}(T^-)$ and then $P\vec{h}T^-\overleftarrow{f}$ is a trail. This contradicts our above observation since $\overleftarrow{f} \in \vec{E}(P^-) \setminus \overleftrightarrow{A}$. $\square$

Corollary 9.29. *There exists a polynomial time algorithm that, given a path P with startvertex v in a bidirected graph B, computes its appendage $A(P, v)$.*

Proof. The algorithm starts with the (P, v)-admissible set $A = \emptyset$ and extends A recursively to (P, v)-admissible sets until $A = A(P, v)$. Let A be any (P, v)-admissible set.

The algorithm first checks whether there exists an edge $e \in E(P) \setminus A$ such that $A \cup \{e\}$ is again (P, v)-admissible. Such an edge e exists if and only if there exists a v–$\vec{e}$ trail in $A \cup E(P)$ where $\vec{e} \in \vec{E}(P^-)$. This can be checked in polynomial time by Corollary 9.27. If such an edge e exists, then the algorithm restarts the recursion with $A \cup \{e\}$.

If there is no such edge, then the algorithm looks for a (P, v, A)-ear trail T in B. This can again be done in polynomial time by Corollary 9.27 since there exists such a (P, v, A)-ear trail if and only if there is a non-trivial trail in $B - (A \cup E(P))$ that starts and ends in $\mathrm{Bon}(P, v, A)$. If such T exists, then the algorithm restarts the recursion with $A \cup E(T)$ which is again (P, v)-admissible by the definition of (P, v, A)-ear trail.

If there is also no (P, v, A)-ear trail, then Lemma 9.28 implies $A = A(P, v)$, and the algorithm terminates. Since the number of recursion steps is bounded by $|E(B)|$, this algorithm runs in polynomial time. $\square$

With all these polynomial time algorithms at hand, we are now ready to prove Theorem 9.23.

Proof of Theorem 9.23. Our algorithm follows the construction in the proof of Theorem 9.17 which thus yields the correctness of the algorithm. Each construction step of the recursive procedure in the proof of Theorem 9.17 can be done in polynomial time by Lemma 9.25, Corollary 9.27 and Corollary 9.29. The number of recursion steps is at most the sum of the lengths of the paths P_i and hence bounded by $|E(B)|$. Altogether, this yields a polynomial running time. $\square$

Appendix

Chapter 10

Supplementary Material

The proof of Theorem 3.3 in Chapter 3 is done with the help of a computer algorithm. The files containing the algorithm, named `K_4_game.py` and `functions.py`, are provided with the print of this thesis. The files can also be found online, e.g. at [BG23b].

Chapter 11

Summaries

Summary of Results (English)

Games on Graphs

In Part I we investigate (G, H)-games for different graphs G and H: two players alternately claim edges of a graph G with the aim to have a copy of H contained in the claimed edges. In the strong version of the game it is the aim of both players, in the Maker-Breaker version it is just the aim of Maker, while Breaker tries to prevent this. The main result of Chapter 3 is that there is a winning strategy for the first player in the strong $(K^{\aleph_0}, K^4)$-game.

In Chapter 4 we consider the Maker-Breaker $(K^{\aleph_0}, K^{\aleph_0})$-game with a number of different additional structural requirements. We prove that in the basic version of the game there is a winning strategy for Maker.

A possible structural property is to colour the vertices of the board and require that Maker's target graph H reflects the colouring of the board. We show that if there are only finitely many colours then Maker can obtain a complete subgraph in which all colours appear infinitely often, but that Breaker can prevent this if there are infinitely many colours. Even when there are infinitely many colours, we show that Maker can obtain a complete subgraph in which infinitely many of the colours each appear infinitely often.

Another possible structural property is to enumerate the board with the rational numbers and require that the vertex set of Maker's complete infinite graph with the induced ordering on the vertex set is order-isomorphic to $\mathbb{Q}$. We prove that there is a winning strategy for Maker in this game in Section 4.4. We also prove that there is a winning strategy for Breaker in the game where Maker must additionally make the vertex set of her complete graph dense in $\mathbb{Q}$.

In Chapter 5 we investigate Maker-Breaker games with boards of size $\aleph_1$ in which Maker's goal is to build a copy of the host graph. In contrast to the smaller boards, set-theoretic considerations come into play and a firm dependence of the outcome of the game on the axiomatic framework is established. We prove that there is a winning strategy for Maker in the $(K^{\omega,\omega_1}, K^{\omega,\omega_1})$-game under ZFC+MA+¬CH and a winning strategy for Breaker under ZFC+CH. We prove a similar result for the $(K^{\omega_1}, K^{\omega_1})$-game where Maker has a winning strategy under ZF+DC+AD, while Breaker has one under ZFC+CH again.

Directed and Bidirected Graphs

In Part II we study directed and bidirected graphs. We touch on different subjects in the field and prove a number of results.

In Chapter 6 we investigate ubiquity in directed graphs by considering two simple classes of directed graphs: digraphs whose underlying undirected graph is a ray or a double ray. We prove that a digraph whose underlying undirected graph is a ray is ubiquitous if and only if it has only finitely many vertices of in-degree 2 as well as only finitely many vertices of out-degree 2. For a digraph whose underlying undirected graph is a double ray that is not a directed double ray we prove that it is ubiquitous if and only if the sum of the number of vertices of in-degree 2 and vertices of out-degree 2 is odd.

In Chapter 7 we generalise a result by László Lovász. He proved that for any finite digraph D with a root r there is a spanning subdigraph L such that for every vertex $v \in V(D) \setminus \{r\}$ some maximal set of paths from r to v in L covers all incoming edges of v in D while one may choose from each path of such a maximal set of paths precisely one edge or internal vertex in such a way that the resulting set separates r from v in the original graph D. Attila Joó proved the corresponding statement for countably infinite digraphs. We prove the corresponding result for digraphs of size $\aleph_1$.

In Chapter 8 we investigate 1-separations in (possibly infinite) directed graphs. For this purpose we introduce a relation for tight set partitions, correspondence, and a novel structure called torsoids. We show that being in correspondence is an equivalence relation on a certain class of tight set partitions and obtain a good insight into how tight sets relate to torsoids. We further prove that if there is a torsoid corresponding to a tight set partition, it is unique. In summary, we show that torsoids are a global tool in the sense that they capture all tight cut contractions of a given digraph independent of the precise choice of a tight cut family.

Lastly in Chapter 9 we investigate connectivity in bidirected graphs, a generalisation of directed graphs. An important theorem in graph theory due to Karl Menger asserts that in a graph the maximum number of disjoint paths connecting two vertex sets X and Y in a graph is the same as the minimum number of vertices separating X and Y. Unfortunately this statement can not be moved verbatim to bidirected graphs, which we demonstrate. We also prove that there is a corresponding statement in bidirected graphs for which we introduce an additional property, cleanness. We also prove that there is a Menger type statement for edge disjoint paths in bidirected graphs and that both statements imply Menger's Theorem in directed and undirected graphs. Our statement also gives a polynomial time algorithm to calculate a set of paths and a separating set of vertices.

Zusammenfassung der Ergebnisse (Deutsch)

Spiele auf Graphen

In Teil I untersuchen wir (G, H)-Spiele für verschiedene Graphen G und H: zwei Spieler reklamieren abwechselnd Kanten eines Graphen G für sich, mit dem Ziel eine Kopie von H in den von ihnen eingenommenen Kanten zu enthalten. In der starken Variante des Spiels ist dies das Ziel beider Spieler, in der Erbauer-Zerstörer Variante ist dies lediglich das Ziel des Erbauers, während es das Ziel des Zerstörers ist, den Erbauer vom Erreichen seines Ziels abzuhalten. Das Hauptergebnis von Kapitel 3 ist, dass es eine Gewinnstrategie für den ersten Spieler in der starken Variante des $(K^{\aleph_0}, K^4)$-Spiels gibt.

In Kapitel 4 betrachten wir das Erbauer-Zerstörer $(K^{\aleph_0}, K^{\aleph_0})$-Spiel mit verschiedenen zusätzlichen strukturellen Anforderungen. Wir beweisen, dass es in der Grundversion des Spiels eine Gewinnstrategie für den Erbauer gibt.

Eine mögliche strukturelle Eigenschaft ist, die Ecken des Spielfelds zu färben und zu verlangen, dass der Zielgraph H des Erbauers die Färbung des Spielfelds widerspiegelt. Wir zeigen, dass, wenn es nur endlich viele Farben gibt, der Erbauer einen vollständigen Teilgraphen erhalten kann, in dem alle Farben unendlich oft auftreten, aber dass der Zerstörer dies verhindern kann, wenn es unendlich viele Farben gibt. Für die Variante in der es unendlich viele Farben gibt, zeigen wir, dass der Erbauer einen vollständigen Teilgraphen erhalten kann, in dem unendlich viele verschiedene Farben jeweils unendlich oft auftreten.

Eine weitere mögliche strukturelle Eigenschaft ist, das Spielfeld mit den rationalen Zahlen aufzuzählen und zu verlangen, dass die Eckenmenge des vollständigen unendlichen Graphen des Erbauers mit der induzierten Ordnung auf der Eckenmenge isomorph zu $\mathbb{Q}$ ist. Wir beweisen in Abschnitt 4.4, dass es eine Gewinnstrategie für den Erbauer in diesem Spiel gibt. Wir beweisen ebenfalls, dass es eine Gewinnstrategie für den Zerstörer in der Variante des Spiels gibt, in der die Eckenmenge des vollständigen Graphen des Erbauers dicht in $\mathbb{Q}$ sein muss.

In Kapitel 5 untersuchen wir Erbauer-Zerstörer Spiele auf Spielfeldern der Kardinalität $\aleph_1$, in denen es das Ziel des Erbauers ist, eine Kopie des Spielfelds zu erstellen. Im Gegensatz zu den kleineren Spielfeldern sind hier mengentheoretische Aspekte relevant. Wir stellen fest, dass es eine starke Abhängigkeit des Spielausgangs vom axiomatischen Rahmen gibt. Wir beweisen, dass es eine Gewinnstrategie für den Erbauer im $(K^{\omega,\omega_1}, K^{\omega,\omega_1})$-Spiel unter ZFC+MA+$\neg$CH und eine Gewinnstrategie für den Zerstörer unter ZFC+CH gibt. Wir beweisen ein ähnliches Ergebnis für das $(K^{\omega_1}, K^{\omega_1})$-Spiel, wo der Erbauer eine Gewinnstrategie unter ZF+DC+AD hat, während der Zerstörer eine unter ZFC+CH hat.

Gerichtete und Bigerichtete Graphen

In Teil II untersuchen wir gerichtete und bigerichtete Graphen. Wir schneiden diverse Themen auf diesem Gebiet an und beweisen Ergebnisse in diesen verschiedenen Bereichen.

In Kapitel 6 untersuchen wir Allgegenwart in gerichteten Graphen, indem wir zwei einfache Klassen von gerichteten Graphen betrachten: gerichtete Graphen, deren zugrundeliegender ungerichteter Graph ein Strahl oder ein Doppelstrahl ist. Wir beweisen, dass ein gerichteter Graph, dessen zugrunde liegender ungerichteter Graph ein Strahl ist, allgegenwärtig ist, genau dann, wenn er endlich viele Ecken mit Eingangsgrad 2 und endlich viele Ecken mit Ausgangsgrad 2 hat. Für einen gerichteten Graphen dessen zugrunde liegender ungerichteter Graph ein Doppelstrahl aber kein gerichteter Doppelstrahl ist, beweisen wir, dass er allgegenwärtig ist, genau dann, wenn die Summe der Anzahl der Ecken mit Eingangsgrad 2 und Ecken mit Ausgangsgrad 2 ungerade ist.

In Kapitel 7 verallgemeinern wir ein Ergebnis von László Lovász. Er bewies, dass es für jeden endlichen gerichteten Graphen D mit designierter Wurzel r einen aufspannenden gerichteten Teilgraphen L gibt, so dass es für jede Ecke $v \in V(D) \setminus \{r\}$ eine maximale Menge von Wegen von r zu v in L gibt, die alle eingehenden Kanten von v in D überdeckt und, dass man aus jedem Weg einer solchen maximalen Menge von Wegen genau eine Kante oder interne Ecke auswählen kann, sodass die resultierende Menge im ursprünglichen Graph D die Wurzel r von der Ecke v trennt. Attila Joó bewies die entsprechende Aussage für abzählbar unendliche gerichtete Graphen. Wir beweisen das entsprechende Ergebnis für gerichtete Graphen der Kardinalität $\aleph_1$.

In Kapitel 8 untersuchen wir 1-Separationen in (möglicherweise unendlichen) gerichteten Graphen. Zu diesem Zweck führen wir eine Relation zwischen Partitionen von knappen Mengen ein, welche wir Korrespondenz nennen, und wir definieren eine neue Struktur namens Torsoid. Wir zeigen, dass „in Korrespondenz stehen“ eine Äquivalenzrelation auf einer bestimmten Klasse von Partitionen von knappen Mengen ist und erhalten einen guten Einblick in die Interaktion von knappen Mengen mit Torsoiden. Wir beweisen außerdem, dass wenn ein Torsoid existiert, der in Korrespondenz zu einer Partition von knappen Mengen steht, dieser eindeutig ist. Zusammenfassend zeigen wir, dass Torsoide in dem Sinne ein globales Werkzeug sind und, dass sie alle Kontraktionen knapper Schnitte eines gegebenen gerichteten Graphen erfassen, unabhängig von der genauen Wahl einer Familie von knappen Schnitten.

Zuletzt untersuchen wir in Kapitel 9 Zusammenhang in bigerichteten Graphen, einer Verallgemeinerung von gerichteten Graphen. Ein wichtiges Theorem der Graphentheorie von Karl Menger besagt, dass in einem Graphen die maximale Anzahl disjunkter Wege, die zwei Eckenmengen X und Y in einem Graphen verbinden, gleich der minimalen Anzahl von Ecken ist, die X und Y trennen. Wir beweisen, dass diese Aussage nicht direkt auf bigerichtete Graphen übertragen werden kann. Wir beweisen aber eine dazu passende Aussage in bigerichteten Graphen, für die wir eine zusätzliche Eigenschaft einführen: Sauberkeit. Wir zeigen dann eine mengerartige Aussage für kantendisjunkte Wege in bigerichteten Graphen für Eckenmengen, die sauber sind. Wir beweisen ebenfalls, dass sowohl die eckendisjunkte als auch die kantendisjunkte Aussage die jeweiligen nach Menger benannten Theoreme in gerichteten und ungerichteten Graphen implizieren. Zuletzt zeigen wir, dass unsere Aussage einen Polynomialzeitalgorithmus zur Berechnung einer Wegemenge und einer trennenden Eckenmenge liefert.

Chapter 12

Publications Related to this Work

Games on Graphs

(1) Chapter 3 is based on [BG23b], which is based on [Gut17].

(2) In Chapter 4, Section 4.4 is based on [BG23a]. The other sections are based on [BEG23], which are based on [Gut20].

(3) Chapter 5 is based on [BGJP22].

Directed and Bidirected Graphs

(4) In Chapter 6, Section 6.1 is based on [GKR22] and Section 6.2 is based on [GKR23].

(5) Chapter 7 is based on [GJ23].

(6) Chapter 8 is based on [BGH$^+$23].

(7) Chapter 9 is based on [BGG$^+$23].

Chapter 13

Declaration of my Contribution

Games on Graphs

Chapter 3 is based on [BG23b] and is the result of research that I conducted under the supervision of Nathan Bowler. I drafted the chapter and the associated algorithm found in the files `K_4_game.py` and `functions.py`. The topic of this chapter is the same as the topic of my bachelor thesis [Gut17], thus the proof ideas are the same but the proof presented in Chapter 3 is completely new, especially the computer algorithm was drafted within my time as a doctoral student.

Chapter 4 is based on [BEG23] and [BG23a]. As with Chapter 3, Section 4.4 is the outcome of research that I conducted under the supervision of Nathan Bowler and I drafted the section. [BEG23] was researched by Marit Emde and I under the supervision of Nathan Bowler. The research for Section 4.1 was done by Marit Emde under supervision, while the research for Section 4.2 and Section 4.3 was done by me under supervision. I drafted Section 4.1, Section 4.2 and Section 4.3. A number of ideas in this chapter have their root in my master thesis, in particular [BEG23] has overlap with my master thesis [Gut20].

Chapter 5 is based on [BGJP22] and the outcome of research conducted together with Nathan Bowler, Attila Joó and Max Pitz following a talk about Chapter 4 that I gave in the weekly research seminar of the discrete maths group of the University of Hamburg. A first draft of the paper was written by Attila Joó.

Directed and Bidirected graphs

Chapter 6 is the result of research conducted in close collaboration with Thilo Krill and Florian Reich and based on [GKR22] and [GKR23]. We undertook the research together and wrote the papers together.

Chapter 7 is based on [GJ23] and arose in collaboration with Attila Joó. As Attila Joó published a paper on the corresponding result in countable graphs [Joó19b] and first introduced me to elementary submodels, he provided the key ideas. We wrote the paper together.

Chapter 8 is the result of a project by Nathan Bowler, Meike Hatzel, Ken-ichi Kawarabayashi, Irene Muzi, Florian Reich and me within the context of a research stay in Tokyo, Japan. As a result the research was conducted in close collaboration with all authors. Subsection 8.1.1 was drafted by Meike Hatzel, Section 8.3 was drafted by Nathan Bowler, Section 8.5 was drafted by Florian Reich, while Section 8.2 and Section 8.4 were drafted by Florian Reich and me in close collaboration.

Chapter 9 is the result of a research seminar that I organised, leading to a publication [BGG$^+$23] by Nathan Bowler, Ebrahim Ghorbani, Raphael W. Jacobs, Florian Reich and me. We conducted the research together in regular meetings and drafted Section 9.1 together. Section 9.2 and Section 9.4 were drafted by Nathan Bowler, Section 9.3 and Section 9.6 were drafted by Florian Reich while I drafted Section 9.5.

Acknowledgement

First and foremost I want to thank my advisor Nathan Bowler without whom this thesis would not have been written. I am thankful for your advice and guidance throughout my studies.

Florian Reich who joined the supervision of Nathan Bowler shortly after me has become a reliable research partner and through collaborative work and even sharing a flat in Japan became a friend. Likewise, I am happy Florian and I shared an office with Thilo Krill. Thank you to both of you for making office hours an insightful and simultaneously fun time.

Furthermore, I am thankful to all of my co-authors and many people from the maths department at the University of Hamburg, especially the discrete mathematics group. I particularly want to mention Ann-Kathrin Elm, Raphael W. Jacobs, Attila Joó, Paul Knappe, Deniz Sarikaya, and Lucas Wansner. Thank you for discussions, constructive criticism and time spent together.

Finally, to Hanna Niehaus and my family, Friederike, Petra and Michael: surely I would not be where I am today without your loving support. Thank you for always being there for me.

Eidesstattliche Versicherung

Hiermit versichere ich an Eides statt, die vorliegende Dissertation selbst verfasst und keine anderen als die angegebenen Hilfsmittel benutzt zu haben. Darüber hinaus versichere ich, dass diese Dissertation nicht in einem früheren Promotionsverfahren eingereicht wurde.

Ort, Datum	Florian Gut

Bibliography

[AB09] R. Aharoni and E. Berger, *Menger's Theorem for Infinite Graphs*, Inventiones mathematicae **176** (2009), no. 1, 1–62.

[AFN96] K. Ando, S. Fujishige, and T. Nemoto, *Decomposition of a Bidirected Graph into Strongly Connected Components and its Signed Poset Structure*, Discrete Applied Mathematics **68** (1996), no. 3, 237–248.

[AK06] G. Appa and B. Kotnyek, *A Bidirected Generalization of Network Matrices*, Networks **47** (2006), no. 4, 185–198.

[And02] T. Andreae, *On Disjoint Configurations in Infinite Graphs*, Journal of Graph Theory **39** (2002), no. 4, 222–229.

[And13] T. Andreae, *Classes of Locally Finite Ubiquitous Graphs*, Journal of Combinatorial Theory, Series B **103** (2013), no. 2, 274–290.

[BCP15] N. Bowler, J. Carmesin, and J. Pott, *Edge-Disjoint Double Rays in Infinite Graphs: A Halin Type Result*, Journal of Combinatorial Theory, Series B **111** (2015), 1–16.

[Bec02] J. Beck, *Ramsey Games*, Discrete Mathematics **249** (2002), no. 1, 3–30. Combinatorics, Graph Theory, and Computing.

[Bec08] J. Beck, *Combinatorial Games: Tic-Tac-Toe Theory*, Vol. 114, Cambridge University Press, 2008.

[BEE$^+$22] N. Bowler, C. Elbracht, J. Erde, J. P. Gollin, K. Heuer, M. Pitz, and M. Teegen, *Ubiquity in Graphs I: Topological Ubiquity of Trees*, Journal of Combinatorial Theory, Series B **157** (2022), 70–95.

[BEE$^+$18] N. Bowler, C. Elbracht, J. Erde, J. P. Gollin, K. Heuer, M. Pitz, and M. Teegen, *Ubiquity in Graphs II: Ubiquity of Graphs With Nowhere-Linear End Structure*, arXiv: 1809.00602 (2018).

[BEE$^+$20] N. Bowler, C. Elbracht, J. Erde, J. P. Gollin, K. Heuer, M. Pitz, and M. Teegen, *Ubiquity in Graphs III: Ubiquity of Locally Finite Graphs with Extensive Tree-Decompositions*, arXiv: 2012.13070 (2020).

[BEG23] N. Bowler, M. Emde, and F. Gut, *The $K^{\aleph_0}$-Game: Vertex Colouring*, Mathematika **69** (2023), no. 3, 584–599.

[BG23a] N. Bowler and F. Gut, *The Rational Number Game*, arXiv: 2309.05526 (2023).

[BG23b] N. Bowler and F. Gut, *The K^4-Game*, arXiv: 2309.02132 (2023).

[BGG$^+$23] N. Bowler, E. Ghorbani, F. Gut, R. W. Jacobs, and F. Reich, *Menger's Theorem in Bidirected Graphs*, arXiv: 2306.01921 (2023).

[BGH01] T. Böhme, F. Göring, and J. Harant, *Menger's Theorem*, Journal of Graph Theory **37** (2001), no. 1, 35–36.

[BGH+23] N. Bowler, F. Gut, M. Hatzel, K.-i. Kawarabayashi, I. Muzi, and F. Reich, *Decomposition of (Infinite) Digraphs Along Directed 1-Separations*, arXiv: 2305.09192 (2023).

[BGJP22] N. Bowler, F. Gut, A. Joó, and M. Pitz, *Maker–Breaker Games on K_{ω_1} and K_{ω,ω_1}*, The Journal of Symbolic Logic **88** (2022), no. 2, 697–703.

[BJG08] J. Bang-Jensen and G. Z. Gutin, *Digraphs: Theory, Algorithms and Applications*, Springer Science & Business Media, 2008.

[BMMN06] M. Bhattacharjee, D. Macpherson, R. G. Möller, and P. M. Neumann, *Notes on Infinite Permutation Groups*, Springer, 2006.

[BZ06] E. D. Bolker and T. Zaslavsky, *A Simple Algorithm that Proves Half-Integrality of Bidirected Network Programming*, Networks **48** (2006), no. 1, 36–38.

[Car23] J. Carmesin, *On Andreae's Ubiquity Conjecture*, Journal of Combinatorial Theory, Series B **162** (2023), 68–70.

[CK23] J. Carmesin and J. Kurkofka, *Canonical Decompositions of 3-Connected Graphs*, arXiv: 2304.00945 (2023).

[Cal78] G. Calvillo Vives, *Optimum Branching Systems*, Ph.D. Thesis, 1978.

[DH21] F. Dorais and D. Hathaway, *Bounding 2D Functions by Products of 1D Functions*, arXiv: 1601.05454 (2021).

[Die17] R. Diestel, *Graph Theory*, 5th ed., Vol. 173, Springer, 2017.

[DM41] B. Dushnik and E. W. Miller, *Partially Ordered Sets*, American Journal of Mathematics **63** (1941), no. 3, 600–610.

[Edm65] J. Edmonds, *Paths, Trees, and Flowers*, Canadian Journal of Mathematics **17** (1965), 449–467.

[EGJ21] J. Erde, J. P. Gollin, and A. Joó, *Enlarging Vertex-Flames in Countable Digraphs*, Journal of Combinatorial Theory, Series B **151** (2021), 263–281.

[ELP82] J. Edmonds, L. Lovász, and W. R. Pulleyblank, *Brick Decompositions and the Matching Rank of Graphs*, Combinatorica **2** (1982), no. 3, 247–274.

[Erd19] J. Erde, *personal communication*, 2019.

[GG55] R. E. Greenwood and A. M. Gleason, *Combinatorial Relations and Chromatic Graphs*, Canadian Journal of Mathematics **7** (1955), 1–7.

[GJ23] F. Gut and A. Joó, *Large Vertex-Flames in Uncountable Digraphs*, Israel Journal of Mathematics (2023).

[GKR22] F. Gut, T. Krill, and F. Reich, *Ubiquity of Oriented Rays*, arXiv: 2210.06952 (2022).

[GKR23] F. Gut, T. Krill, and F. Reich, *On the Ubiquity of Oriented Double Rays* (2023), arXiv: 2310.09857.

[Gut17] F. Gut, *Ramseyspiele auf einem unendlichen Spielfeld: Betrachtung einfacher Graphen*, Bachelor Thesis, Universität Hamburg, 2017.

[Gut20] F. Gut, *Unendliche Graphentheorie: Rekonstruktion und das $K_{\aleph_0}$-Spiel*, Master Thesis, Universität Hamburg, 2020.

[GW21] A. C. Giannopoulou and S. Wiederrecht, *The Flat Wall Theorem for Bipartite Graphs with Perfect Matchings*, Extended Abstracts EuroComb 2021, 2021, pp. 429–435.

[Hal65] R. Halin, *Über die Maximalzahl fremder unendlicher Wege in Graphen*, Mathematische Nachrichten **30** (1965), no. 1–2, 63–85.

[Hal70] R. Halin, *Die Maximalzahl fremder zweiseitig unendlicher Wege in Graphen*, Mathematische Nachrichten **44** (1970), 119–127.

[HKN$^+$17] D. Hefetz, C. Kusch, L. Narins, A. Pokrovskiy, C. Requilé, and A. Sarid, *Strong Ramsey Games: Drawing on an Infinite Board*, Journal of Combinatorial Theory, Series A **150** (2017), 248–266.

[HKSS14] D. Hefetz, M. Krivelevich, M. Stojaković, and T. Szabó, *Positional Games*, Springer, Basel, 2014.

[HRW19] M. Hatzel, R. Rabinovich, and S. Wiederrecht, *Cyclewidth and the Grid Theorem for Perfect Matching Width of Bipartite Graphs*, Graph-Theoretic Concepts in Computer Science, 2019, pp. 53–65.

[Jec03] T. Jech, *Set Theory*, Vol. 14, Springer, 2003.

[Jec68] T. Jech, ω_1 *Can be Measurable*, Israel Journal of Mathematics **6** (1968), no. 4, 363–367.

[Joó19a] A. Joó, *The Complete Lattice of Erdős-Menger Separations*, arXiv: 1904.06244 (2019).

[Joó19b] A. Joó, *Vertex-Flames in Countable Rooted Digraphs Preserving an Erdős-Menger Separation for Each Vertex*, Combinatorica **39** (2019), no. 6, 1317–1333.

[Joó21] A. Joó, *Greedoids from Flames*, Journal of Graph Theory **98** (2021), no. 1, 49–56.

[Kan08] A. Kanamori, *The Higher Infinite: Large Cardinals in Set Theory from their Beginnings*, Springer Science & Business Media, 2008.

[KK15] K.-i. Kawarabayashi and S. Kreutzer, *The Directed Grid Theorem*, Proceedings of the forty-seventh annual ACM symposium on Theory of Computing, 2015, pp. 655–664.

[Kot59a] A. Kotzig, *Z teórie konečných grafov s lineárnym faktorom. I.*, Matematicko-fyzikálny časopis **09** (1959), no. 2, 73–91.

[Kot59b] A. Kotzig, *Z teórie konečných grafov s lineárnym faktorom. II.*, Matematicko-fyzikálny časopis **09** (1959), no. 3, 136–159.

[Kot60] A. Kotzig, *Z teórie konečných grafov s lineárnym faktorom. III.*, Matematicko-fyzikálny časopis **10** (1960), no. 4, 205–215.

[Kun11] K. Kunen, *Set Theory*, Vol. 34, College Publications, London, 2011.

[Kur21] J. Kurkofka, *Ubiquity and the Farey Graph*, European Journal of Combinatorics **95** (2021).

[Lea08] I. B. Leader, *Lecture on Hypergraph Games*, 2008.

[Lov73] L. Lovász, *Connectivity in Digraphs*, Journal of Combinatorial Theory, Series B **15** (1973), no. 2, 174–177.

[Lov87] L. Lovász, *Matching Structure and the Matching Lattice*, Journal of Combinatorial Theory, Series B **43** (1987), no. 2, 187–222.

[LP09] L. Lovász and M. D. Plummer, *Matching Theory*, Reprint of the 1986 original published by North-Holland, Providence, RI: AMS Chelsea Publishing, 2009.

[McC00] W. McCuaig, *Even Dicycles*, Journal of Graph Theory **35** (2000), no. 1, 46–68.

[Men27] K. Menger, *Zur allgemeinen Kurventheorie*, Fundamenta Mathematicae **10** (1927), 96–115.

[Mos09] Y. N. Moschovakis, *Descriptive Set Theory*, 2nd ed., American Mathematical Society, 2009.

[NT06] S. Norine and R. Thomas, *Minimal Bricks*, Journal of Combinatorial Theory, Series B **96** (2006), no. 4, 505–513.

[NT07] S. Norine and R. Thomas, *Generating Bricks*, Journal of Combinatorial Theory, Series B **97** (2007), no. 5, 769–817.

[NW60] C. S. J. A. Nash-Williams, *Decomposition of Graphs into Closed and Endless Chains*, Proceedings of the London Mathematical Society **3** (1960), no. 1, 221–238.

[Per68] H. Perfect, *Applications of Menger's Graph Theorem*, Journal of Mathematical Analysis and Applications **22** (1968), no. 1, 96–111.

[Pym69] J. S. Pym, *The Linking of Sets in Graphs*, Journal of the London Mathematical Society **1** (1969), no. 1, 542–550.

[RST99] N. Robertson, P. D. Seymour, and R. Thomas, *Permanents, Pfaffian Orientations, and Even Directed Circuits*, Annals of Mathematics **150** (1999), no. 3, 929–975.

[Sie33] W. Sierpiński, *Sur un problème de la théorie des relations*, Annali della Scuola Normale Superiore di Pisa - Classe di Scienze **2** (1933), no. 3, 285–287.

[Sou11] L. Soukup, *Elementary Submodels in Infinite Combinatorics*, Discrete Mathematics **311** (2011), no. 15, 1585–1598.

[Tut66] W. T. Tutte, *Connectivity in Graphs*, University of Toronto Press, Toronto, 1966.

[Wie22] S. Wiederrecht, *Matching Minors in Bipartite Graphs*, Ph.D. Thesis, Technische Universität Berlin, 2022.