Proof Theory of Graph Minors and Tree Embeddings

Martin Rudolf Arne Krombholz

Submitted in accordance with the requirements for the degree of Doctor of Philosophy

The University of Leeds

School of Mathematics

February 2018

The candidate confirms that the work submitted is his own and that appropriate credit has been given where reference has been made to the work of others.

To all those who have never dedicated a thesis to themselves.

Theorien, die einen Satz der Logik gehaltvoll erscheinen lassen, sind immer falsch.

— Ludwig Wittgenstein, *Tractatus logico-philosophicus*

Acknowledgements

Most of all, I want to thank my supervisor Michael Rathjen. Without his vast knowledge of everything proof-theoretic (and more), his guidance and support, this thesis would not have been possible. I also want to thank the University of Leeds for providing me with a scholarship, the School of Mathematics for such a nice and productive atmosphere, and the Logic Group for organizing many interesting talks and seminars. I would in particular like to thank Stan Wainer, for enriching the proof theory reading group with his contributions.

I am very grateful to the people with whom I shared my time here at Leeds, and would like to thank some of them especially. Anton Freund has been a great inspiration through his high motivation and productivity, which no doubt stimulated me in my own work, and I appreciate our many conversations. Likewise, I would also like to thank Michael Toppel for many entertaining chats. Further thanks go to Mark Carney, for helping me with UK real-life stuff. I am not sure whether I should thank or curse Cesare Galozzi for the philosophical discussions we have had.

Finally, I would like to thank Neil Robertson and Paul Seymour for writing the papers of the Graph Minors series in such a clear style. Their work strikes exactly the right balance between presenting the ideas behind their proofs and carrying out these proofs in an enjoyably precise manner.

Abstract

This thesis explores metamathematical properties of theorems appearing in the Graph Minors series. A number of these theorems have been known to have very high proof-theoretic strength, but an upper bound on many of them, including the graph minor theorem, had never been proved.

We give such upper bounds, by showing that any proofs in the Graph Minors series can be carried out within a system of Π^1_1-comprehension augmented with induction and bar-induction principles for certain classes of formulas. This establishes a narrow corridor for the possible proof-theoretic strength of many strong combinatorial principles, including the graph minor theorem, immersion theorem, theorems about patchwork containment, and various restrictions, extensions and labelled versions of these theorems. We also determine the precise proof-theoretic strength of some restrictions of the graph minor theorem, and show that they are equivalent to other restricted versions that had been considered before. Finally, we present a combinatorial theorem employing ordinal labelled trees ordered by embedding with gap-condition that may additionally have well-quasi-ordered labels on the vertices, which turns out not to be provable in the theory $\Pi^1_1 - \mathrm{CA}$. This result suggests a potential for raising the lower bounds of the immersion theorem, and the thesis concludes by outlining this possibility and other avenues for further research.

Contents

List of figures

Citation conventions

In order to improve readability and recognizability of references, we make the following conventions:

- We will refer to Robertson and Seymour (1983) as ‘Graph Minors I’,
- Robertson and Seymour (1984) as ‘Graph Minors III’,
- Robertson and Seymour (1990a) as ‘Graph Minors IV’,
- Robertson and Seymour (1986a) as ‘Graph Minors V’,
- Robertson and Seymour (1986b) as ‘Graph Minors VI’,
- Robertson and Seymour (1988) as ‘Graph Minors VII’,
- Robertson and Seymour (1990c) as ‘Graph Minors VIII’,
- Robertson and Seymour (1990b) as ‘Graph Minors IX’,
- Robertson and Seymour (1991) as ‘Graph Minors X’,
- Robertson and Seymour (1994) as ‘Graph Minors XI’,
- Robertson and Seymour (1995a) as ‘Graph Minors XII’,
- Robertson and Seymour (1995b) as ‘Graph Minors XIV’,
- Robertson and Seymour (1996) as ‘Graph Minors XV’,
- Robertson and Seymour (2003a) as ‘Graph Minors XVI’,
- Robertson and Seymour (1999) as ‘Graph Minors XVII’,
- Robertson and Seymour (2003b) as ‘Graph Minors XVIII’,
- Robertson and Seymour (2004a) as ‘Graph Minors XIX’,

- Robertson and Seymour (2004b) as ‘Graph Minors XX’, and
- Robertson and Seymour (2010) as ‘Graph Minors XXIII’.

Chapter 1

Introduction

This thesis investigates some proof-theoretic properties of the graph minor theorem, most importantly by giving new upper bounds on its proof-theoretic ordinal. As described in the standard textbook on graph theory by Diestel (2017), p. 347, the graph minor theorem is arguably the most important piece of work in graph theory:

> "Our goal in this last chapter is a single theorem, one which dwarfs any other result in graph theory and may doubtless be counted among the deepest theorems that mathematics has to offer: *in every infinite set of graphs there are two such that one is a minor of the other*. This *graph minor theorem*, inconspicuous though it may look at first glance, has made a fundamental impact both outside graph theory and within. Its proof, due to Neil Robertson and Paul Seymour, takes well over 500 pages."

This 500-page-long proof is presented in a series of 20 papers, the so-called *Graph Minors series*, which was later extended by three papers to prove Nash-William's immersion conjecture, by providing an even more general version of the graph minor theorem. Not only the length of this proof, but also its depth and complexity have been a major reason for why no upper bound on the graph minor theorem's proof-theoretic strength existed, even though the precise proof-theoretic strength of one of its restricted versions appearing in an early paper of the Graph Minors series was calibrated over thirty years ago by Friedman, Robertson, and Seymour (1987).

There are many important consequences of the graph minor theorem, the most significant being perhaps the fact that every minor-closed property can be expressed by a finite set of forbidden minors, and thus be checked computationally in cubic time. This provided for instance the existence of a polynomial-time algorithm to decide whether graph embeddings in $\mathbb{R}^3$ are knot-free, a property that was not even known to be computable before. However, the consequences of the graph minor theorem itself are not even as numerous as those of the intermediate results and concepts of its proof. The first of these is the notion of tree-decomposition, which plays a role in the above mentioned restricted form of the graph minor theorem and among others has many computational consequences. Another influential new notion that was introduced through the Graph Minors series is that of a tangle, which characterises highly connected components of a graph. Large parts of the Graph Minors series concern themselves with graph embeddings in surfaces, and one of the most central accomplishments of these investigations, the excluded minor theorem, has sparked a lot of fruitful research, as outlined by Kawarabayashi and Mohar (2007).

Another field where the graph minor theorem had a large impact is reverse mathematics. Reverse mathematics is a research program initiated and developed by Harvey Friedman and Stephen Simpson, to classify mathematical theorems in terms of the set existence axioms needed to prove it. In the words of Friedman (1974), p.235, the question that reverse mathematics asks is:

> "What are the proper axioms to use in carrying out proofs of particular theorems, or bodies of theorems, in mathematics? What are those formal systems which isolate the essential principles needed to prove them?"

The framework to study this question is that of second order arithmetic, more specifically subsystems of second order arithmetic that consist of basic arithmetic axioms and particular set existence axioms or schemata. As it turns out, most main theorems of ordinary mathematics are equivalent to one of five subsystems, called the *Big Five* of reverse mathematics. The investigation of purely combinatorial principles is of special interest, since those, in contrast to various theorems about topological or metric spaces, have a natural representation in second order arithmetic. At first glance combinatorial theorems may appear very simple from a reverse mathematics standpoint, since the statements involving them are essentially only about finite objects, and the sets used in their proofs should therefore not be too complicated. It is therefore illuminating that the graph minor theorem is not provable in the strongest system of the Big Five, $\Pi^1_1 - \mathrm{CA}_0$.

Historically, the graph minor theorem was not the first combinatorial theorem to have a surprisingly high proof-theoretic strength. The first such result was that a strengthening of Ramsey's theorem by Harrington and Paris (1977) is not provable in Peano arithmetic, and thus not in the corresponding system ACA_0 of second order arithmetic. Next, Schmidt (1979) was the first to relate tree-embeddability to ordinal numbers, which was done independently by Harvey Friedman and presented by Simpson (1985), to show that Kruskal's theorem can not be proved in the second strongest system of the Big Five, ATR_0. This work also provided an even stronger combinatorial principle — extended Kruskal's theorem — using trees labelled with natural numbers, whose proof-theoretic strength lies above that of $\Pi^1_1 - \mathrm{CA}_0$.

The exact proof-theoretic strength of Kruskal's theorem was then calibrated by Rathjen and Weiermann (1993) to lie in between ATR_0 and $\Pi^1_1 - \mathrm{CA}_0$. In terms of ordinal analysis, which roughly aims to characterize theorems in terms of the minimal ordinal that can not be proved to be well-ordered from those theorems, the proof-theoretic ordinal of Kruskal's theorem was determined in the above work to be the Ackermann or small Veblen ordinal. A related analysis by Van der Meeren, Rathjen, and Weiermann (2015), which takes up the research initiated by Diana Schmidt, recently provided the result that the maximum order type of certain well-partial orders based on trees is the big Veblen number. Rathjen (1999) gives a further overview of achievements in ordinal analysis.

Extended Kruskal's theorem mentioned above was the first combinatorial theorem shown to be stronger than $\Pi^1_1 - \mathrm{CA}_0$. However, this theorem could be seen as slightly unnatural, since the gap-condition imposed on the embedding relation is modelled after the ordinal notation system used to analyze $\Pi^1_1 - \mathrm{CA}_0$. This was remedied by Friedman, Robertson, and Seymour (1987), who showed that extended Kruskal's theorem is equivalent to the bounded graph minor theorem, a restricted version of the graph minor theorem naturally occurring in the Graph Minors series. This equivalence established the precise proof-theoretic strength of the bounded graph minor theorem, and thus a lower bound for the full graph minor theorem. An upper bound for the proof-theoretic strength of the graph minor theorem however remained elusive.

This thesis will shed some light on the upper bounds of the graph minor theorem, by establishing that it is provable in $\Pi^1_1 - \mathrm{CA}_0 + \Pi^1_2 - \mathrm{BI} + \Pi^1_3 - \mathrm{IND}$, the theory of Π^1_1-comprehension augmented with the additional principles of Π^1_2-bar induction and ordinary Π^1_3-induction. The thesis proceeds by first giving necessary definitions and basic facts from graph theory and reverse mathematics. The foundations for the investigations conducted here are then laid by presenting the construction of Friedman, Robertson, and Seymour (1987) that shows that the bounded graph minor theorem implies extended Kruskal's theorem and giving an alternative proof of its correctness. Chapter 2 continues by giving a detailed analysis of the fourth paper of the Graph Minors series, showing how to handle the combinatorial parts of the Graph Minors series in second order arithmetic and confirming rigorously some unproved claims of Friedman, Robertson, and Seymour (1987). Surfaces and graph drawings play a major role in the proof of the excluded minor theorem, the proof of which takes up almost half of the Graph Minors series. The necessary techniques for their treatment in second order arithmetic are presented in chapter 3, and applied by giving an in-depth recreation of the sixth paper of the Graph

Minors series with these techniques. Chapter 4 then gives a thorough summary of the remaining relevant papers of the Graph Minors series, examining the proof-theoretic methods used in their proofs. It also shows in detail that a Π^1_2-bar induction is sufficient for conducting the critical proofs in Graph Minors XIX and Graph Minors XXIII. Chapter 5 investigates some results for gap-embeddings of ordinal-labelled trees, and chapter 6 concludes the thesis by giving an overview over the achieved results and presenting some open questions.

1.1 Preliminaries

In this section we will introduce basic notions and notations of graph theory and reverse mathematics that are fundamental in what follows. In general, basic definitions and results for graph theory and reverse mathematics can be found in Diestel (2017) and Simpson (2009), respectively.

A graph consists of vertices and edges that connect two distinct vertices. Only one edge e can connect two vertices x and y, in this case e is also denoted xy or yx. If an edge e is directed from x to y, then x is called the tail and y the head of e, and in this case only the notation xy for e is used. Whenever loops and multiple edges between vertices are allowed, or edges are directed, this will be explicitly stated. Likewise, unless stated otherwise all graphs considered in this thesis are assumed to be finite.

A path in a graph G is a sequence of alternating vertices and edges $\langle v_0, e_1, v_1, \ldots, e_k, v_k\rangle$ of G so that e_i connects v_{i-1} and v_i, where $1 \leq i \leq k$, and so that all the v_i are distinct. If $P = \langle v_0, e_1, v_1, \ldots, e_k, v_k\rangle$ is a path in G, the vertices v_0 and v_k are called the endpoints of P. A graph is called connected if for every two distinct vertices $v_1, v_2 \in G$ there is a path in G with endpoints v_1 and v_2. A circuit or cycle in G is similarly to a path a sequence of alternating vertices and edges $\langle v_0, e_1, v_1, \ldots, e_k, v_k, e_{k+1}, v_{k+1}\rangle$, so that $\langle v_0, e_1, v_1, \ldots, e_k, v_k\rangle$ is a path in G, $v_{k+1} = v_0$ and e_{k+1} connects v_k and v_{k+1}.

A tree is then a graph which is connected and has no circuits. It follows that two vertices u, v in a tree T are connected by a unique path, which will be denoted uTv, alternatively $[u, v]_T$ or $[u, v]$ if the reference to T is clear. A rooted tree has one special vertex, called the root and denoted $root(T)$. The designation of a root induces a natural order on the vertices of the trees, so that $u \leq v$ if and only if $u \in [root(T), v]$. If $u \leq v$ then u is called a predecessor of v and v a successor of u, and in case that $u < v$ and that there is no w with $u < w < v$ the terminology immediate predecessor (or parent) and immediate successor (or child) is used. A vertex which has no successors is called a leaf of T. A rooted tree can also be defined in terms of an order as imposed by the designation of a root, so that $(T, \leq)$ is a rooted tree if there is a unique least element of T with regards to $\leq$ (which will be the root), and so that the set of predecessors of any vertex of T is linearly ordered. There is further a third definition of rooted trees in terms of ordered graphs, so that T is defined to be a rooted tree if it is a directed tree so that every vertex has at most one incoming edge, i.e. at most one parent. Then, as a tree must have exactly one more vertex than edges, it follows that there is one vertex from which every edge is directed away, which can be identified as the root of T. These three definitions are of course equivalent.

Several embedding relations between graphs are central in the context of the Graph Minors series, namely tree-embeddings, topological minors and of course minors and some variants thereof. All of these will be denoted by $\leq$ if there is no possibility of confusion, and otherwise with a suitable subscript. The first such relation is embedding between rooted trees T_1, T_2. This consists of an injective function f from the vertex set of T_1 to the vertex set of T_2 so that if $u \leq v$ in T_1, then $f(u) \leq f(v)$ in T_2, and so that if w is the infimum of u, v in T_1, denoted by $u \wedge v$, then $f(w)$ is the infimum of $f(u), f(v)$ in T_2. In other words, f has to be order and infimum preserving. A tree embedding can be generalized to the notion of a topological minor. A graph G_1 is a topological minor of a graph G_2 if there is a function f from G_1 to G_2 (called a topological expansion of G_1 in G_2) that maps vertices injectively to vertices and edges to paths connecting these vertices, so that if e_1, e_2 are distinct edges in G_1 then $f(e_1)$ and $f(e_2)$ are disjoint paths in G_2. An embedding between rooted trees is then the same as a topological minor embedding between those trees when viewed as directed trees as in the third definition above. A topological expansion can also be characterized as a sequence of subdivision of edges, where a subdivision of an edge uv consists of introducing a new vertex w, removing uv and adding the edges uw and wv.

The topological minor relation can then (in a sense) be further generalized to that of a minor, where we call G_1 a minor of G_2 if there is a function f from G_1 to G_2 that maps edges injectively to edges and vertices to connected subgraphs of G_2, so that if $v_1, v_2 \in V(G_1)$ are distinct vertices, then $f(v_1) \cap f(v_2) = \emptyset$. In this case f is called a minor expansion of G_1 in G_2. This is not quite a direct generalization of the notion of a topological minor, but nevertheless the minor relation is a generalization of the topological minor relation in the sense that if G_1 is a topological minor of G_2 then G_1 is also a minor of G_2. The reverse is known to hold for subcubic graphs. There is a second way to characterize minors, namely that G_1 is a minor of G_2 if it can be obtained from a subgraph of G_2 by contracting edges, or more explicitly if it can be obtained from G_2 by deleting edges and vertices and contracting edges. Which definition is more convenient to work with differs on the problem considered, and the two are of course equivalent, but in the following the first version will often be more useful.

The statements that the set of all finite graphs or trees is well-quasi-ordered under the corresponding relations relations possess a surprising amount of proof-theoretic strength. As mentioned in the introduction, Kruskal's theorem, which says that the set of rooted trees is well-quasi-ordered under the topological minor relation, is stronger than the theory ATR_0 of second order arithmetic. This theory is the second strongest of the "Big Five" of reverse mathematics, and most of ordinary mathematics can be developed within it. Friedman defined an extended version of Kruskal's theorem that uses rooted trees in which every vertex v is labelled with a natural number $l(v)$ which is no greater than some fixed number n, and where the embedding between two trees T_1 and T_2 is additionally subject to the *gap-condition*:

$$\forall v \in V(T_1)(l(v) = l(f(v)) \wedge \\ \wedge \forall u \leq f(v)(\neg \exists w \in V(T_1)(u \leq f(w) < f(v)) \to l(u) \geq l(v))),$$

which says that the label of v and of its image must be identical, and that any vertex that lies on the path of $f(v)$ and the image of its immediate predecessor (but is not equal to the image of its immediate predecessor) must have label greater than or equal to $l(v)$. If the statement that these trees are well-quasi-ordered under embedding with gap-condition is denoted by EKT_n, then Friedman showed that $EKT := \forall n\ EKT_n$, extended Kruskal's theorem, is not provable even in $\Pi^1_1 - \mathrm{CA}_0$, the strongest of five certain subsystems of second order arithmetic, the so-called "Big Five".

These five systems arise naturally in investigations about the proof-theoretic strength of various theorems. The goal of such investigations is most often to show that some two theorems are equivalent. This endeavor only makes sense when working within (relatively) weak base theory in which the two theorems cannot be proved, as any theorem provable in a theory T is equivalent over T to any true sentence, for example $1 = 1$. The systems of the Big Five all share some basic arithmetic axioms together with the basic induction axiom

$$\forall X(0 \in X \wedge \forall n(n \in X \to n + 1 \in X) \to \forall n(n \in X),$$

and differ mostly in what kind of set comprehension axioms they allow. The theory most often used as a base theory, and the weakest of the Big Five, is RCA_0, which allows the formation of recursive sets and a slightly stronger induction axiom. However, RCA_0 and the second weakest theory of the Big Five, WKL_0 will actually play no role in this thesis. The reason for this is that when working with stronger theorems, RCA_0 is too weak to show that central objects exist. We thus work with the base theory ACA_0, the theory which permits arithmetical comprehension, i.e. the formation of sets defined by formulas that only contain quantifiers over natural numbers. Closely related to ACA_0 is ATR_0, which permits the transfinite iteration of such arithmetical comprehension along any well-ordering. While ACA_0 is already strong enough to develop much of common mathematics, ATR_0 is much stronger and allows the classification of various theorems related directly and indirectly to countable ordinals. Finally, $\Pi^1_1 - CA_0$ is the strongest system of the Big Five. It allows comprehension of sets defined by formulas containing one set quantifier. There are few natural theorems of ordinary mathematics that are not provable in $\Pi^1_1 - CA_0$, and so it is somewhat surprising that combinatorial theorems such as extended Kruskal's theorem and the bounded graph minor theorem are not provable in $\Pi^1_1 - CA_0$.

The metamathematical unprovability results about Kruskal's theorem and extended Kruskal's theorem are due to Friedman, but were published by Simpson (1985). The idea of the proofs is to assign some (labelled) tree to an ordinal, so that if one such tree T_1 is embeddable into another such tree T_2, then the ordinal corresponding to T_1 is less than or equal to that corresponding to T_2. To give the precise definition from Simpson (1985) of this map, first define T^v to be the subtree of T with root v, i.e. the induced[1] subgraph of T with vertex set $V(T^v) := \{u \in V(T) : v \leq u\}$. Further, for a labelled tree T, set $qT = l(root(T))$. Then $o(T)$ is defined by:

I) If $|T| = 1$ define $o(T) := \Omega_{qT}$.

Otherwise, let $b_1, \ldots, b_m$ be the children of $root(T)$, indexed in such a way that $o(T^{b_1}) \geq \ldots \geq o(T^{b_m})$. Set $\beta_j := o(T^{b_j})$.

II) In case $m = 2$, $l(b_1) = qT$, $\beta_1 = \omega^{\alpha_1} + \ldots + \omega^{\alpha_{k-1}}$ and $\beta_2 = \omega^{\alpha_k}$, $\alpha_1 \geq \ldots \geq \alpha_k$, define $o(T) = \beta_1 + \beta_2$.

III) In case $m = 3$, $l(b_1) = qT$, $\beta_1 < \omega^{\beta_1}$ and $\beta_2 = \beta_3 = 0$, define $o(T) = \omega^{\beta_1}$.

IV) In case $m = 4$, $\beta_1 \in C_{qT}(\beta_1)$ and $\beta_2 = \beta_3 = \beta_4 = 0$, define $o(T) = \Psi_{qT}(\beta_1)$.

[1] A subgraph is induced if it has all possible edges of the graph it is contained in.

The sets $C_n(\alpha)$ and functions $\Psi_n(\alpha)$ above form part of an ordinal notation system for $\Pi^1_1 - \mathrm{CA}_0$. The conditions on $l(b_1)$ in II and III are imposed so that the root-label of a tree always indicates the approximate size of the corresponding ordinal. The map o is an injective, partial map on the set of all labelled trees that surjectively maps into the ordinals of the notation system for $\Psi_0(\Omega_\omega)$. Restricting it to its domain yields a bijective, total map, and it is shown in Simpson (1985) that if $T_1 \leq T_2$ and $qT_1 = 0$ then $o(T_1) \leq o(T_2)$, yielding that a bad sequence of ordinals in the proof theoretic ordinal $\Psi_0(\Omega_\omega)$ of $\Pi^1_1 - \mathrm{CA}_0$ would yield a bad sequence of labelled trees, contradicting EKT. Hence, $\Pi^1_1 - \mathrm{CA}_0$ cannot prove EKT, since it would then prove the well-foundedness of its own proof-theoretic ordinal.

A more natural result, the bounded graph minor theorem, was then found to be equivalent to EKT by Friedman, Robertson, and Seymour (1987). The central notion to define the bounded graph minor theorem is that of a tree-decomposition. A tree-decomposition of a graph G is meant to decompose it into parts arranged in a tree-like structure. It consists of a tree T and a collection of vertex sets $\langle V_t : t \in V(T)\rangle$ so that $V_t \subseteq G$ for all $t \in T$, which are called the parts of a tree-decomposition. In addition $(T, \langle V_t : t \in V(T)\rangle)$ has to satisfy three axioms:

- $\bigcup_{t \in T} V_t = V(G)$,
- every edge of G has two of its ends in some V_t,
- for every edge t_1t_2 in T, if T_1 and T_2 are the two components of $T \setminus \{t_1t_2\}$ then $V_{t_1} \cap V_{t_2}$ is a separating set between $\bigcup_{t \in T_1} V_t$ and $\bigcup_{t \in T_2} V_t$.

The third condition can be stated in various equivalent forms, such as $V_{t_1} \cap V_{t_2} = \bigcup_{t \in T_1} V_t \cap \bigcup_{t \in T_2} V_t$, or that $V_u \supseteq V_s \cap V_t$ whenever $u \in [s, t]$. The width of a tree-decomposition is defined to be one less than the maximum cardinality of any of its parts. The tree-width of a graph G is then the minimum width of all the tree-decompositions of G.

Denote by $bGMT_n$ the statement that for every sequence of graphs $\langle G_i : i \in \mathbb{N}\rangle$ so that every G_i has tree-width at most n, there are $i < j$ such that G_i is a minor of G_j. Then the statement $bGMT := \forall n\, bGMT_n$ is called the bounded graph minor theorem. The proof of Friedman, Robertson, and Seymour (1987) that the bounded graph minor theorem implies EKT gives a naturally occurring interpretation of the gap-condition as a set of pairwise disjoint paths between certain subgraphs. Their construction will be investigated in the next section.

1.2 Bounded tree-width and labelled trees

The work of Friedman, Robertson, and Seymour (1987) is the starting point of this thesis. In their paper it is shown that Friedman's extended Kruskal's theorem is equivalent to the bounded graph minor theorem, and in turn to the graph minor theorem for planar graphs (denoted $pGMT$). More precisely, theorem 4.2 of Friedman, Robertson, and Seymour (1987) summarizes its main results:

> The following are equivalent:
>
> 1) The bounded graph minor theorem,
> 2) extended Kruskal's theorem,
> 3) the well-foundedness of $\Psi_0(\Omega_\omega)$,
> 4) Π^1_1-reflection for $\Pi^1_1 - \mathrm{CA}_0$.

Note that the same theorem also claims that the full graph minor theorem is provable in $\Pi^1_1 - \mathrm{CA} + \mathrm{BI}$. This claim was later retracted by Friedman, but as will be shown it was indeed correct. Since this claim is not the only one stated without or only with a sketched proof (for example it is only claimed on page 231 of Friedman, Robertson, and Seymour (1987) that $bGMT \leftrightarrow pGMT$), it seems prudent to verify the concerned results of that paper. This will be done in the next chapter. In this chapter we present the construction that shows the implication $bGMT \rightarrow EKT$ and give a slightly adjusted proof of its correctness that however uses the same ideas as the original proof. This will lay out the fundamental ideas that are used in a similar construction in section 2.3 to show that the bounded and planar graph minor theorems are indeed equivalent.

The construction of Friedman, Robertson, and Seymour (1987) uses a small alteration to the usual definition of a labelled tree. Instead of the usual codomain of $\{1, \ldots, n\}$, n-labelled rooted trees with labels from $\{n+1, \ldots, 2n\}$ are considered, i.e. tuples (T, l), where T is a rooted tree and $l : V(T) \longrightarrow \{n+1, \ldots, 2n\}$ is called the labelling function. An n-edge-labelled rooted tree is defined analogously, with the domain of the labelling function being the edge-set of T. A tree is called k-branching if every vertex either has exactly k immediate successors or is a leaf. Denote with $p(x)$ the immediate predecessor of a vertex x, if it exists. For edge labelled trees we also use the notation $l(x)$ instead of $l(xp(x))$. Note that the gap-condition on an embedding $f : T_1 \longrightarrow T_2$ for edge-labelled trees is slightly different from the one of Simpson (1985), namely that for any edge xy of T_1 and any $e \in [f(x), f(y)]$, it has to hold that $l(e) \geq l(xy)$.

Relate with an n-edge-labelled, k-branching rooted tree T a graph G as follows: For $x \in V(T) \setminus \{root(T)\}$ let $K(x)$ be the complete graph on $2n(k+1)$ vertices, and for $x = root(T)$ let $K(x)$ be the complete graph on $2n(k+1)+1$ vertices, a distinction which will be important later. Let $K(x)$ and $K(p(x))$ intersect in exactly $l(xp(x))$ vertices, for x and y with $y \neq p(x)$, $x \neq p(y)$ let $K(x)$ and $K(y)$ be disjoint. The complete graphs $K(x)$ contain exactly enough vertices to make this construction possible, as any vertex of the tree T can have at most $k+1$ neighbours and the labels of these neighbours are no higher than $2n$.

Analogously to Friedman, Robertson, and Seymour (1987) we fix some additional notation. For $x \in V(T) \setminus \{root(T)\}$ let $Z(x) = \{z : z \geq x\}$, $Y(x) = V(T) \setminus Z(x)$ and $W(x) = V(K(x) \cap K(p(x)))$. Let furthermore $KZ(x) = \bigcup_{z \in Z(x)} K(z)$ and

$KY(x) = \bigcup_{y \in Y(x)} K(y)$. Then the following hold:

1. For any $x \neq r$, $W(x) = KZ(x) \cap KY(x)$, because no $K(y)$ for $y \neq p(x), y \neq x$ intersects both $K(x)$ and $K(p(x))$ by construction.

2. By the above point $(T, \{K(x)\}_{x \in T})$ is a tree-decomposition of G. Since $K(\mathit{root}(T))$ has size $2n(k+1)+1$, the tree-decomposition has width $2n(k+1)$ and as any complete subgraph of G must be contained in one of the parts of any tree-decomposition (see Diestel (2017), p.353, corollary 12.3.5), this must already be the tree-width of G. Hence the graphs G have bounded tree-width for fixed k and n.

3. As $W(x)$ separates $KZ(x)$ from $KY(x)$, there can be at most $l(x) = |W(x)|$ disjoint paths from any $Y \subseteq KY(x)$ to any $Z \subseteq KZ(x)$.

Following Friedman, Robertson, and Seymour (1987), we now want to produce an embedding between two labelled, rooted trees T, T', given a minor inclusion $G \leq G'$ between their associated graphs. We do this by finding, for any $x \in V(T)$, a suitable $K(x')$, $x' \in V(T')$, so that $K(x)$ and $K(x')$ correspond to each other with regard to the minor inclusion. In the following let T, T', G, G' be as above and denote the minor expansion of G in G' by f. We proceed similarly to the proof of Proposition 3.1 of Friedman, Robertson, and Seymour (1987).

Lemma 1.2.1. For any $x \in V(T)$ there exists an $x' \in V(T')$ such that

$$\forall z \in V(K(x)) : f(z) \cap K(x') \neq \emptyset.$$

Proof. Let $w \in V(T')$, $w \neq \mathit{root}(T')$, be arbitrary. As $|W(w)| \leq 2n$ and the $f(z)$ for $z \in K(x)$ are disjoint, at most $2n$ of these $f(z)$ can meet $W(w)$. As there are $|K(x)| \geq 2n(k+1)$ of them, at least $2n(k+1) - 2n = 2nk > 0$ (since $n, k > 0$) do not intersect $W(w)$. Denote the set of these vertices by $NW_x(w)$. As $f(z)$ is connected and $W(w)$ separates

$KZ(w)$ from $KY(w)$ we must therefore have $f(z) \subseteq KZ(w)$ or $f(z) \subseteq KY(w)$ for every $z \in NW_x(w)$. Furthermore, if we have two vertices z_1, z_2 in $NW_x(w)$ we cannot have $f(z_1) \subseteq KZ(w)$ and $f(z_2) \subseteq KY(w)$: As $K(x)$ is complete, there must be an edge between z_1 and z_2 and therefore an edge connecting $f(z_1)$ and $f(z_2)$ in G' by the properties of a minor inclusion. But this means that there is a path from $f(z_1) \subseteq KZ(w)$ to $f(z_2) \subseteq KY(w)$ which does not meet $W(w)$ (as $f(z_1) \cap W(w) = \emptyset = f(z_2) \cap W(w)$), in contradiction to $W(w)$ being a separating set of $KZ(w)$ and $KY(w)$.

Hence

$$f(NW_x(w)) \subseteq KZ(w) \quad \text{or} \quad f(NW_x(w)) \subseteq KY(w).$$

In the first case we say that $H(x) := f(K(x))$ is above $W(w)$, and we say that $H(x)$ is below $W(w)$ in the second case. Now give a direction to every edge $wp(w)$ of T', namely let w be the tail of $wp(w)$ if $H(x)$ is below $W(w)$, and let w be the head of $wp(w)$ if $H(x)$ is above $W(w)$.

Claim. There is a unique vertex $x' \in V(T')$ such that every edge incident with x' is directed toward x'.

Proof of Claim. First we show that at most one edge is directed away from any vertex. Assume not and let w, w_1, w_2 be distinct vertices such that both ww_1 and ww_2 are directed away from w. As w has only one predecessor, one of w_1 or w_2 must be its successor, so we can assume without loss of generality that $p(w_1) = w$. As the edge w_1w is directed toward the successor, this means that $f(NW_x(w)) \subseteq KZ(w_1)$. We now have two cases:

- If $w = p(w_2)$, then $KZ(w_1) \cap KZ(w_2) = \emptyset$, as both are children of w. But by the direction of the edge ww_2 we must have $f(NW_x(w)) \subseteq KZ(w_2)$, a contradiction as $f(NW_x(w)) \subseteq KZ(w_1)$ and $NW_x(w)$ is non-empty.

- If $w_2 = p(w)$, then as the edge ww_2 is directed toward the parent we get $f(NW_x(w)) \subseteq KY(w)$. Hence $f(NW_x(w)) \subseteq KY(w) \cap KZ(w_1)$. But $KY(w) \cap KZ(w_1) = \emptyset$ as w is a predecessor of w_1, again a contradiction.

So at most one edge is directed away from any vertex. Because T' is a tree, T' has $m-1 := |T'|-1$ edges. By directing an edge, we direct it toward one vertex and away from another. As only one edge can be directed away from any specific vertex, this means that at least $m-1$ vertices must have an incident edge that is directed away from them. On the other hand — as there are only $m-1$ edges — at most $m-1$ vertices can have an edge that is directed away from them. This means that there is exactly one vertex x' with the property that every edge that is incident with x' is directed toward x'. □(Claim)

It remains to verify that x' is as desired, i.e. that $f(z) \cap K(x') \neq \emptyset$ for every $z \in K(x)$. Assume not, so that there is $z \in K(x)$ with $f(z) \cap K(x') = \emptyset$. Then $f(z)$ must be below $W(x')$ or above $W(y')$ for some successor y' of x' (as $W(y') \subseteq K(x')$). Both are impossible since x' is the vertex that has only incoming edges, which means that in the first case $f(z)$ would have to be below $W(x')$ and in $KZ(x')$, and in the second case $f(z)$ would have to be above $W(y')$ and in $KY(x') \cup K(x')$, due to the direction of the respective edges.

□

Lemma 1.2.2. Under the above map the root of T maps to the root of T'.

Proof. Let r be the root of T, let y' be the image of r. Then $f(z) \cap K(y') \neq \emptyset$ for every $z \in K(r)$, which means that $K(y')$ contains at least $2n(k+1)+1$ vertices as all the $f(z)$ are disjoint. But there is only one such y', namely the root of T'.

□

Lemma 1.2.3. The map $x \mapsto x'$ is injective.

Proof. If $x \neq y$ then $f(z_x) \cap f(z_y) = \emptyset$ for all $z_x \in K(x), z_y \in K(y)$. Hence if x and y would map to the same vertex z', $K(z')$ would have to contain $2 \cdot 2n(k+1)$ vertices. However, such a z' does not exist since $n, k > 0$.

□

Thus, we have constructed an injective map $h : T \longrightarrow T'$ that maps the root of T to the root of T'. To confirm that h is an embedding it remains to be checked that h preserves infimums and respects the gap-condition. We will first show that h preserves infimums.

Lemma 1.2.4. If $x, y \in V(T)$ are such that $x = p(y)$ then $x' \leq y'$ in T'.

Proof. We argue by induction on the distance of y from the root r of T. If $p(y) = r$ then by the above lemma r maps to the root of T', so $r' \leq y'$.

Assume now that $p(y) = x \neq r$ and that the statement holds for all vertices that are closer to the root than y. Assume further that $x' \not\leq y'$, then $K(y') \subseteq KY(x')$. By the induction hypothesis we have that $p(x)' \leq x'$, so, as $|K(p(x)) \cap K(x)| = l(xp(x))$, there must be $l(xp(x)) \geq n + 1$ disjoint paths from $K(p(x)')$ to $K(x')$ passing through $W(x')$. But as $K(y') \subseteq KY(x')$ it must analogously hold that there are $l(yx) \geq n+1$ disjoint paths from $K(y)$ to $K(x)$ that pass through $W(x')$, which additionally must be disjoint to the former $l(xp(x))$ paths as $K(y) \cap K(p(x)) = \emptyset$. But this means that $|W(x')| \geq 2(n + 1) > 2n$, which is impossible.

□

Corollary 1.2.5. If $x \leq y$ in T then $x' \leq y'$ in T'.

Proof. If $x \leq y$ then there are $x_0, \ldots, x_n$ with $x_0 = x$, $x_n = y$ and $p(x_{k+1}) = x_k$ for all $k \in \{0, \ldots n - 1\}$. But then by the above lemma $x' = x'_0 \leq x'_1 \leq \ldots \leq x'_n = y'$.

□

Lemma 1.2.6. If $x \neq y$ with $p(x) = p(y)$ then $y' \not\leq x'$.

Proof. Assume $y' \leq x'$. By the above corollary we have $p(y)' = p(x)' \leq y'$. So $p(y)' \leq y' \leq x'$, which means that $K(p(y)') \subseteq KY(y')$ and $K(x') \subseteq KZ(y')$. Hence $W(y')$ separates $K(x')$ and $K(p(y)')$. As $K(x)$ and $K(p(x)) = K(p(y))$ have $l(x)$ vertices in common, there must be at least $l(x)$ disjoint paths linking $K(p(y)')$ and $K(x')$ in G', and by the above those paths must pass through $W(y')$. Likewise, $W(y')$ separates $K(p(y)')$ from $K(y')$ and so there must be $l(y)$ disjoint paths from $K(p(y)')$ to $K(y')$ in G. As $K(x)$ and $K(y)$ are disjoint in G those $l(x) + l(y)$ paths can be chosen to be disjoint, but this implies $|W(y')| \geq 2(n+1) > 2n$, a contradiction.

$\square$

Lemma 1.2.7. If $x \not\leq y$ then $x' \not\leq y'$.

Proof. If $y \leq x$ then we obtain $y \neq x$, and $y' \leq x'$ by corollary 1.2.5. But this means by injectivity of h that $y' < x'$, and so $x' \not\leq y'$. So we may assume $y \not\leq x$.

Consider $x \wedge y$, and let x_0, y_0 be the successors of $x \wedge y$ on the paths from $x \wedge y$ to x and y, respectively. As $x \not\leq y$ and $y \not\leq x$ we have that $x_0 \neq y_0$. Hence we are in the situation of the previous lemma and we obtain $x_0' \not\leq y_0'$ and $y_0' \not\leq x_0'$. If we had $x' \leq y'$ this would — because of $x_0' \leq x'$ by the above corollary — then imply that $x_0', y_0' \in \{z' \in T' : z' \leq y'\} =: A$. But because T' is a tree A must be a linear order, in contradiction to x_0' and y_0' being incomparable.

$\square$

With this we have established that $x \leq y$ if and only if $x' \leq y'$.

Lemma 1.2.8. $(x \wedge y)' = x' \wedge y'$.

Proof. As $x \wedge y \leq x$ and $x \wedge y \leq y$ we get $(x \wedge y)' \leq x'$ and $(x \wedge y)' \leq y'$, so $(x \wedge y)' \leq x' \wedge y'$.

If $x \leq y$ or $y \leq x$ we are done immediately by the previous lemmas, so we may assume that x and y, and hence by the previous results also x' and y', are incomparable in the tree order. Consider x_0 and y_0 as in the previous proof. As in the proof of Lemma 1.2.6 there must be $l(x_0) + l(y_0) \geq 2(n+1)$ disjoint paths from $K(x_0') \cup K(y_0')$ to $K((x \wedge y)')$. Now, if $(x \wedge y)' < x' \wedge y'$ these would have to pass through $W(x' \wedge y')$, a contradiction to $|W(x' \wedge y')| \leq 2n$. $\square$

Theorem 1.2.9. Given a minor inclusion $G \leq G'$ we can define an embedding $h : T \longrightarrow T'$ between the corresponding edge-labelled trees such that for all $x \in V(T) \setminus \{root(T)\}$ we have that if $c' \in (h(p(x)), h(x)]$ then $l(p(c')c') \geq l(xp(x))$.

Proof. Of course, we take h to be the function defined by $x \mapsto x'$, and as h is already an embedding it remains only to verify the gap condition.

Assume $x \in V(T) \setminus \{r\}$. Then there must be at least $l(xp(x))$ disjoint paths from $K(x')$ to $K(p(x)')$ in G'. Because $p(x)' \leq x'$ these must pass through all $W(c')$ with $c' \in (h(p(x)), h(x)]$. So $l(p(c')c') = |W(c')| \geq l(xp(x))$, as desired.

□

Hence the graph minor theorem for graphs of tree width $\leq 2n(k+1)$ implies a form of extended Kruskal's theorem for k-branching trees with n-edge labels. Friedman, Robertson, and Seymour (1987) go on to provide an argument to show that Kruskal's theorem for the edge-labelled trees considered above already implies extended Kruskal's theorem, thereby completing the proof that the bounded graph minor theorem implies EKT. This is done by simulating additional successors of a vertex and a stricter gap-condition by higher numbered labels. For this, four different indexed sets Q_n^1-Q_n^4 of trees and embedding relations are introduced, and it is then shown that Q_n^i being well-quasi-ordered implies that Q_n^{i+1} is well-quasi-ordered, for $1 \leq i \leq 3$.

The sets Q_n^1-Q_n^4 are defined as follows. Q_n^1 is composed of all perfect, 2-branching, n-edge labelled trees (perfect meaning that every leaf has the same distance to the root), ordered by the embedding relation considered above. Q_n^2 contains all perfect, 2-branching, n-vertex labelled trees that are ordered by embedding with a less strict than usual gap-condition, only imposing that $l(v) \geq l(f(v))$ for every embedding $f : T_1 \longrightarrow T_2$ and $v \in V(T_1)$. Q_n^3 then consists of all perfect, 4-branching, n-vertex labelled trees under the usual gap-condition, and Q_n^4 of the trees and embedding considered in extended Kruskal's theorem.

Then, if Q^1_n is well-quasi-ordered Q^2_n can be shown to be well-quasi-ordered by simulating a tree T of Q^2_n by a tree of Q^1_n, which consists of two copies of T and one root-vertex that is adjacent only to the copies of the original root vertex, and where an edge $p(y)y$ has label $l(y)$ of the tree T in Q^2_n. Next, a tree of Q^3_n can be simulated by a tree of Q^2_{2n} by transforming it into a perfect 2-branching tree through intermediate vertices and labelling those with appropriate numbers greater than n. The final simulation of a tree of Q^4_n by a tree of Q^3_{n+1} is similar to the one before, where additionally introduced vertices to obtain a perfect, 4-branching tree are labelled by $n+1$ to distinguish them from the old vertices. Since Q^4_n being well-quasi-ordered is just EKT, this completes the proof of $bGMT \rightarrow EKT$.

This section presented the proof of Friedman, Robertson, and Seymour (1987) that showed $bGMT \rightarrow EKT$. The next chapter will inspect the results of Friedman, Robertson, and Seymour (1987) that were given without or only with a sketched proof. This will firmly establish the equivalences $bGMT \leftrightarrow EKT$ and $EKT \leftrightarrow pGMT$.

Chapter 2

Metamathematics of the Graph Minor Theorem

In this chapter we will give some missing proofs of statements in Friedman, Robertson, and Seymour (1987) and a thorough proof-theoretic analysis of the proof methods used in Graph Minors IV. This will be a first step in the analysis of the proof of the graph minor theorem and provide some useful results for the analysis of further restricted forms of the graph minor theorem. The first section will show that the bounded graph minor theorem (which is proved in Graph Minors IV) is provable in the theory $\Pi^1_1 - \mathrm{CA}_0$ augmented with Π^1_1-reflection for $\Pi^1_1 - \mathrm{CA}_0$, thereby together with the results presented in the previous section establishing $bGMT \leftrightarrow EKT$, as EKT implies the well-orderedness of $\Psi_0(\Omega_\omega)$ by Simpson (1985) which is equivalent to $\Pi^1_1 - \mathrm{CA}_0 + \mathrm{RFN}_{\Pi^1_1}(\Pi^1_1 - \mathrm{CA}_0)$ over ACA_0. The first section will also give a proof-theoretic analysis of the other theorems of Graph Minors IV, establishing that all of its results can be recreated in $\Pi^1_1 - \mathrm{CA}_0$ if induction over Π^1_3-formulas is additionally allowed. This analysis will further serve the purpose of showing how the finite combinatorial parts of the Graph Minors series can be handled in ACA_0. The second section of this chapter will present some thoughts about possible ways for replacing the Π^1_3-induction with a Π^1_2-induction, and the third section will carry out the construction presented in the previous chapter with planar, trivalent graphs instead, which together with the results of section 4.3 shows that the planar graph minor theorem and the topological minor theorem for subcubic graphs are both equivalent to EKT.

2.1 Graph Minors IV

To show that the bounded graph minor theorem is provable in $\Pi^1_1 - \mathrm{CA}_0 + \mathrm{RFN}_{\Pi^1_1}(\Pi^1_1 - \mathrm{CA}_0)$, we need to recreate the proof of each instance $bGMT_n$ in $\Pi^1_1 - \mathrm{CA}_0$. Applying Π^1_1-reflection for $\Pi^1_1 - \mathrm{CA}_0$ will then yield that $\forall n\, bGMT_n$, i.e. the bounded graph minor theorem, holds. Since the proofs of the critical theorems of Graph Minors IV in $\Pi^1_1 - \mathrm{CA}_0$ will necessarily be close to the original, our focus here will be on how to code the required objects and perform the needed proof-techniques in second order arithmetic.

To carry out the proof of $bGMT_n$ in $\Pi^1_1 - \mathrm{CA}_0$, we define in second order arithmetic a graph G to consist of a code for a finite set of vertices $V(G)$ and a set of edges $E(G) \subseteq V(G) \times V(G)$. If we view G as an undirected graph we say that $u, v \in V(G)$ are connected by an edge if $\langle u, v\rangle \in E(G) \vee \langle v, u\rangle \in E(G)$, if we view G as a directed graph we say that there is an edge from u to v if $\langle u, v\rangle \in E(G)$ and we call u the tail and v the head of that edge. In this chapter we will be concerned with simple graphs, meaning graphs that have no loops (multiple edges are already ruled out by the definition), i.e. $\neg\exists x(\langle x, x\rangle \in E(G))$. For basic procedures such as coding pairs and sequences, we use the techniques of Simpson (2009).

We need some additional definitions. An undirected tree is a connected undirected graph with no circuits, i.e.

$$T \text{ is a tree} \iff T \text{ is a connected graph } \wedge \neg\exists\rho(\rho \text{ is a circuit in } T),$$

where a graph being connected abbreviates

$$\forall u, v \in V(G)\exists\rho(\rho \text{ is a path in } G \text{ from } u \text{ to } v),$$

where ρ being a path from u to v in G in turn abbreviates

$$\rho \in Seq \wedge \exists m(lh(\rho) = 2m+1) \wedge \forall i(\exists k(2k+1=i) \rightarrow \rho(i) \in E(G)) \wedge$$
$$\forall i(\exists k(2k=i) \rightarrow \rho(i) \in V(G)) \wedge \forall i \leq \frac{lh(\rho)-3}{2}(\langle \rho(2i), \rho(2i+2)\rangle = \rho(2i+1)) \wedge$$
$$\rho(0) = u \wedge \rho(lh(\rho)-1) = v \wedge \forall i,j < lh(\rho)(i \neq j \rightarrow \rho(i) \neq \rho(j)),$$

and finally ρ being a circuit in G abbreviates

$$\rho \in Seq \wedge \ \langle \rho(0), \ldots, \rho(lh(\rho)-3\rangle \text{ is a path in } G \ \wedge \rho(0) = \rho(lh(\rho-1)) \wedge$$
$$\langle \rho(lh(\rho-3)), \rho(lh(\rho-1))\rangle = \rho(lh(\rho-2)),$$

and ρ' being a path in G is shorthand for $\exists u, v(\rho'$ is a path from u to v in $G)$.

A rooted tree is then defined as in the preliminaries as a directed tree in which each vertex is the head of at most one edge. Recall that the unique path from $u \in V(T)$ to $v \in V(T)$ is denoted by uTv. We denote the projections on the first and second coordinates of a pair $\langle x, y\rangle$ by $\langle x, y\rangle_0 := x$ and $\langle x, y\rangle_1 := y$. Note that this notation is also used if a number is not explicitly written as a pair, so that for example x_0 will denote the projection on the first coordinate of x.

For the upcoming lemma we additionally need the notion of an infinite graph. An infinite graph — like a finite graph — consists of a vertex and edge set. As our language only permits the use of subsets of the natural numbers, we have to encode these two sets into one, for example by taking their direct sum: If we want to talk about a graph consisting of vertex set $V(G)$ and edge set $E(G)$, we define $G := \{\langle x, i\rangle : (i = 0 \wedge x \in V(G)) \vee (i = 1 \wedge x \in E(G))\}$. Following this idea, we say that a set G is an infinite graph if

$$G \text{ is infinite } \wedge \forall x \in G \exists i \exists j(x = \langle i, j\rangle \wedge (j = 0 \vee j = 1)) \wedge \forall x \in G(\exists i(\langle i, 1\rangle = x) \rightarrow$$
$$\exists j \exists k(x_0 = \langle j, k\rangle \wedge \langle j, 0\rangle \in G \wedge \langle k, 0\rangle \in G)).$$

We denote the edge set of G by $E(G) := \{x : \langle x, 1\rangle \in G\}$ and the vertex set by $V(G) := \{x : \langle x, 0\rangle \in G\}$. Following the terminology of Graph Minors IV, we say that a set

$X \subseteq V(G)$ is G-stable if there is no edge in G between any two elements of X or alternatively if $E(G) \restriction X \times X = \emptyset$, i.e. if $\forall x, y \in X(\langle\langle x, y\rangle, 1\rangle \notin G)$. A set $X \subseteq V(G)$ is called G-rich if it contains no infinite subset that is G-stable, i.e. if $\forall Y(\forall n(n \in Y \rightarrow n \in X) \wedge Y$ is infinite $\rightarrow$ Y is not G-stable$)$. We are now ready to prove the first lemma of Graph Minors IV, (2.1), which is together with (2.2) the main obstacle for formalizing the proof of Graph Minors IV in $\Pi^1_1 - \mathrm{CA}_0$. Lemma (2.1) condenses a minimal bad sequence argument about a sequence of trees into a combinatorial statement about an infinite graph and can thus serve as a building block of a proof of e.g. Kruskal's theorem or extended Kruskal's theorem and similar statements.

Lemma 2.1.1 ((2.1) of Graph Minors IV)**.** The following is provable in $\Pi^1_1 - \mathrm{CA}_0$:

Let $\sigma := \langle T_i : i \in \mathbb{N}\rangle$ be a sequence of trees. Let M be any infinite graph that satisfies

1) $V(M) = V(\bigsqcup_{i\in\mathbb{N}} V(T_i))$, that is formally $x \in V(M) \leftrightarrow \exists i \exists y(y \in V(\sigma(i)) \wedge x = \langle y, i\rangle)$, and

2) $E(M)$ satisfies the condition that for $i' > i \geq 1$, if $\langle\langle u, i\rangle, \langle w, i'\rangle\rangle \in E(M)$ and $v \in V(T_{i'}) \setminus \{root(T_{i'})\}$ is a vertex of $root(T_{i'})T_{i'}w$ then $\langle\langle u, i\rangle, \langle v, i'\rangle\rangle \in E(M)$, so that $E(M)$ in a sense corresponds to embeddability of the subtrees.

Now, if additionally $\{root(T_i) : i \in \mathbb{N}\}$ is M-stable then there exists an infinite M-stable set $X \subseteq V(M)$ that is minimal in a sense, i.e. such that X has at most one vertex in common with every T_i and such that $\{u : \exists v \exists i(\langle v, u\rangle \in E(T_i) \wedge v \in X)\}$ is M-rich.

Proof. The proof uses a minimal bad sequence argument. Following Graph Minors IV, define first a sequence $\langle z_i : i \in \mathbb{N}\rangle \in V(M)^\omega$ to be increasing if the corresponding sequence $\langle n_i : i \in \mathbb{N}\rangle$ is increasing, where $z_i \in T_{n_i}$. Define then a section to be an increasing sequence $\langle z_i : i \in \mathbb{N}\rangle \in (V(M) \setminus \{root(T_i) : i \in \mathbb{N}\})^\omega$ that is M-stable. The goal is to find a minimal section with regard to the subtree relation, which is made precise below. If there is no section at all, $\{root(T_i) : i \in \mathbb{N}\}$ satisfies the theorem, because it is M-stable by assumption. So we may assume that a section exists.

Let $T = \bigsqcup_i T_i$. For $v \in V(T)_0 = \{x : \exists y(\langle x, y\rangle \in V(T))\}$ define $T^v = T_i^v$, where i is such that $v \in T_i$. Robertson and Seymour then go on to define the minimal section $\langle x_i : i \in \mathbb{N}\rangle$ in a way that x_i is chosen so that:

- There exists a section that has initial segment $\langle x_1, \dots, x_i\rangle$, and
- for all $x \in V(T^{x_i}) \setminus \{x_i\}$, there exists no section that has initial segment $\langle x_1, \dots, x_{i-1}, x\rangle$.

To do this in $\Pi^1_1 - \mathrm{CA}_0$ we proceed as follows. Denote by Seq the set of all finite sequences. Let

$$R := \{\langle \sigma, n\rangle : \sigma \in Seq \wedge \exists X(\text{ } X \text{ is a section } \wedge \forall i < lh(\sigma)(\sigma(i) = X(i)) \wedge X(lh(\sigma)) = n)\}.$$

R is a set containing finite initial segments of sections together with possible continuations of the initial segment. R exists by Σ^1_1-comprehension, which is equivalent to Π^1_1-comprehension. Now define

$$R' := \{\langle \sigma, n\rangle : \langle \sigma, n\rangle \in R \wedge \forall i \forall k((\langle \sigma, k\rangle \in R \wedge n \in V(T_i)) \rightarrow k \notin V(T_i^n \setminus \{n\}))\}.$$

R' restricts R so that the extension to the sequence σ must satisfy the second condition on the minimal sequence above. Define next

$$f = \{\langle \sigma, n\rangle : (\langle \sigma, n\rangle \in R' \wedge \forall k < n(\langle \sigma, k\rangle \notin R')) \vee (\neg \exists k(\langle \sigma, k\rangle \in R) \wedge n = 0)\}.$$

f chooses for a finite sequence σ a numerically minimal possible extension, or has value 0 if there is no such extension. The default value 0 is only introduced so that f is total; as we will see it will not be of concern in the following definitions. Note that the $\forall k < n(\langle \sigma, k\rangle \notin R')$ part in the definition of f has nothing to do with the minimality of n in the sense of the second condition on the minimal sequence, this part simply picks one (namely the numerically least) of all possible minimal extensions.

Define next by primitive recursion

$$f'(0) = \langle\rangle$$
$$f'(n+1) = f'(n) \circ \langle f(f'(n)) \rangle\,,$$

where $\circ$ denotes the concatenation of sequences. The function f' is almost what we want, it gives us a sequence of sequences, the "limit" of which is our minimal section. If we simply define

$$X(n) := f'(n+1)(n),$$

then X is the desired section. To confirm this in $\Pi^1_1 - \mathrm{CA}_0$, we can easily show by arithmetical induction (which is available even in ACA_0) that $\forall n \exists k(\langle f'(n), k\rangle \in R')$: This holds for $\langle\rangle = f'(0)$ since a section exists; for $n > 0$, $\langle f'(n-1), k\rangle \in R'$ by the induction hypothesis, whence $\langle f'(n-1), f(f'(n-1))\rangle \in R'$, from which it follows that there is a section with initial segment $f'(n-1) \circ \langle f(f'(n-1))\rangle = f'(n)$, which implies that there is some k with $\langle f'(n), k\rangle \in R'$. From this it immediately follows that X is actually a section satisfying both minimality conditions.

Further we claim that X already satisfies the theorem. Write $x_i := X(i)$. To simplify notation we also write $uv \in E(G)$ instead of $\langle u, v\rangle \in E(G)$. Then

$$\forall j \forall j' \forall z((j > j' \geq 1 \wedge x_{j'}z \in E(T)) \rightarrow (x_j z \notin E(M) \wedge z x_j \notin E(M))).$$

This is because if $x_{j'} \in V(T_{i'})$ then $x_{j'} \neq o(T_{i'})$ because X is a section, and $x_{j'} \in root(T_{i'})T_{i'}z$ so by an assumption of our theorem x_j cannot be adjacent to z in M because it is not adjacent to $x_{j'}$ (since X is a section).

Let $Y = \{\langle u, i\rangle : \exists v(\langle v, u\rangle \in E(T_i) \wedge v \in X)\}$, then to conclude the proof it remains to show that Y is M-rich. First we show that no sequence in Y can be a section: Assume $\langle y_i : i \in \mathbb{N}\rangle$ is such a section in Y, then by definition of Y there is some x_k so that $x_k y_1 \in E(T)$. But then $\langle x_1, \ldots, x_{k-1}, y_1, y_2, \ldots\rangle$ is a section (using that no x_i, $i < k$, can be adjacent to any y_j, and that X and Y are both sections) and $y_1 \in V(T^{x_k} \setminus \{x_k\})$, in contradiction to X being minimal.

This implies that Y is M-rich: If there was $Y' \subseteq Y$ infinite and M-stable, then we could extract a section from Y' as follows. Define

$$Y''(0) = \mu y.(y \in Y')$$
$$Y''(n+1) = \mu y.(y \in Y' \wedge Y'(n)_1 < y_1)).$$

Then Y'' would be a section in Y.

Hence X is the desired set.

$\square$

Using the above lemma we can prove the following theorem from Graph Minors IV. The induction used in its proof would be a Π^1_3-induction, which is not available in $\Pi^1_1 - \mathrm{CA}_0$. Hence, in $\Pi^1_1 - \mathrm{CA}_0$ we need to prove the theorem by metainduction on n for every n, while it is stated universally in Graph Minors IV.

First we import some additional definitions from Graph Minors IV. For a tree T let $\phi : E(T) \longrightarrow \{0, \ldots, n\}$ be a labelling function on the edges of T. Now, $u \in V(T)$ is said to precede $v \in V(T)$ essentially if a stronger gap-condition is satisfied between u and v, that is if $u \neq root(T)$, $u \leq v$ with regard to the tree-ordering, $\phi(xu) = \phi(yv)$, where x and y are such that $xu, yv \in E(T)$, and $\phi(g) \geq \phi(xu) = \phi(yv)$ for all $g \in E(uTv)$.

To perform the induction step in the following lemma, we need the notion of a contraction onto a subset of edges of a tree T. Thus for $F \subseteq E(T)$ we call the tree S defined by $V(S) :=$ {roots of $T \setminus F$} $= \{v \in V(T) : \neg\exists u(uv \in E(T) \setminus F)\}$ and $uv \in E(S) \leftrightarrow \exists f \in F(uTv \subseteq T \setminus (F \setminus \{f\}))$, that is $uv \in E(S)$ if u and v are only separated by one edge of F in T, the contraction of T onto S.

Theorem 2.1.2 ((2.2) of Graph Minors IV)**.** The following is provable in $\Pi^1_1 - \mathrm{CA}_0$ for every n.

Let $\sigma := \langle T_i : i \in \mathbb{N}\rangle$ be a sequence of trees and $\langle \phi_i : i \in \mathbb{N}\rangle$ a corresponding sequence of codes for functions $\phi_i : E(T_i) \longrightarrow \{0, \ldots, n\}$. Note that each ϕ_i is a finite set and hence can be coded by a natural number. Let M be a graph defined by $V(M) = V(T)$, where $T = \bigsqcup_i T_i$, with the property that for $i' > i \geq 1$, if $\langle u, i\rangle, \langle w, i'\rangle \in V(T)$, $\langle u, i\rangle \langle w, i'\rangle \in E(M)$ and $v \in V(T_{i'})$ precedes w in $T_{i'}$, then $\langle u, i\rangle \langle v, i'\rangle \in E(M)$. If in addition $\{root(T_i) : i \in \mathbb{N}\}$ is M-stable then there exists an M-stable set X that contains at most one element of every T_i such that $\{u : \exists v \exists i(\langle v, u\rangle \in E(T_i) \wedge v \in V(X))\}$ is M-rich.

Proof. We argue by metainduction on n. Other than that, the proof is essentially the same as the one in Graph Minors IV.

For $n = 0$ the statement is basically just the previous lemma. For $n > 0$ define $F = \{\langle i, s\rangle : s \text{ is a code for the set } \{e \in E(T_i) : \phi_i(e) = 0\}\}$ and $S = \{\langle i, c\rangle : c \text{ is the contraction of } T_i \text{ onto } F(i)\}$. We also write F_i for $F(i)$ and S_i for $S(i)$.

Let N be the graph defined by $V(N) = \bigsqcup_i V(S_i)$ and $\langle u, i\rangle \langle v, i'\rangle \in E(N) \leftrightarrow \langle u, i\rangle \langle v, i'\rangle \in E(M)$.

Then

> For all $i' > i$: If $\langle u, i\rangle \langle w, i'\rangle \in E(N)$ and $v \neq root(S_{i'})$ with $v \in root(S_{i'}) S_{i'} w$, then $\langle u, i\rangle \langle v, i'\rangle \in E(N)$.

To show this, let $xu, yw \in E(T(i'))$, then by definition of $S(i')$ we have $\phi_{i'}(xu) = \phi_{i'}(yw) = 0$. Thus v precedes w and so $\langle u, i\rangle$ is adjacent to $\langle v, i'\rangle$ in M and hence also $\langle u, i\rangle \langle v, i'\rangle \in E(N)$, which proves the claim.

Further, since $root(T_i) = root(S_i)$, $\{root(S_i) : i \in \mathbb{N}\}$ is N-stable. Thus we can apply the previous lemma which gives us a set $X \subseteq V(N)$ such that X has at most one vertex in common with every S_i and so that $\{u : \exists v \exists i(\langle v, u\rangle \in E(S_i) \wedge v \in X)\}$ is N-rich. Applying Ramsey's theorem, we get an increasing (in the same sense as in the previous lemma) sequence $\langle z_j : j \in \mathbb{N}\rangle$ with $z_j \in X$ for all $j \in \mathbb{N}$. Viewed in the graph M, this is a sequence of roots of components R_j of $T_{i_j} \setminus F_{i_j}$, where i_j is the index so that $R_j \subseteq T_{i_j}$, such that $A := \{u \in V(T) : \exists e, j(e \in E(F_{i_j}) \wedge e_1 = u_0 \wedge e_0 \in R_j)\}$ is M-rich.

Define then $\phi'_j : E(R_j) \longrightarrow \{0, \ldots, n-1\}$ by $\phi'_j(e) = \phi_{i_j}(e) - 1$, which is possible since $F \cap R_j = \emptyset$. By the induction hypothesis we have already proved the theorem in $\Pi^1_1 - \mathrm{CA}_0$ for $n-1$, and with this we can therefore deduce that there exists an infinite M-stable set $X \subseteq \bigsqcup_i V(R_i)$ with at most one element in each T_i such that the set $B := \{\langle u, i\rangle : \exists v(\langle v_0, u\rangle \in E(R_i) \wedge v \in V(X))\}$ is M-rich.

But X already satisfies the theorem since if $C = \{\langle u, i\rangle : \exists v \exists i(\langle v_0, u\rangle \in E(T_i) \wedge v \in X)\}$, and $z \in C$ is the head of a corresponding edge $xz \in E(T)$, then either $\phi_i(xz) = 0$ and hence $z \in A$ or $\phi_i(xz) > 0$ in which case $z \in B$, where $i = z_1$. Thus $C \subseteq A \cup B$, but A and B are both M-rich and hence so is C.

$\square$

It should now be clear that the finite combinatorial parts of the Graph Minors series can be handled in ACA_0 and thus easily in $\Pi^1_1 - \mathrm{CA}_0$. The only obstructions to carry out the proof of $bGMT_n$ in $\Pi^1_1 - \mathrm{CA}_0$ are the minimal bad sequence arguments considered above. Since only a finite number of minimal bad sequence arguments can be carried out in a proof inside $\Pi^1_1 - \mathrm{CA}_0$, for the rest of the proof of each $bGMT_n$ it is important that the previous lemma is not applied in its full strength anywhere, but only for a fixed n. We will thus confirm this and give a brief summary of the remainder of Graph Minors IV.

A star is a tree where the root is adjacent to every other vertex. The idea of the remaining proof of $bGMT_n$ is to reduce the tree-decompositions to star-decompositions by collapsing the parts of the tree-decomposition not corresponding to the root, thereby essentially performing an induction on the height of the tree. For those parts of Graph Minors IV, the graphs considered may be hypergraphs, and they may have roots, which are a sequence designating special vertices in the graph. These roots of each part of a tree-decomposition will denominate the vertices of the cutset between it and the part closer to the root of the tree-decomposition. The index of a tree-decomposition is the maximum size of these cutsets (called adhesion in modern graph theoretic terminology).

Thus, the index is bounded by the maximum number of vertices in each part of the tree-decomposition.

The first theorem relying on (2.2) is theorem (3.3) (page 236 of Graph Minors IV), which essentially states that the aforementioned induction on the height of the tree-decomposition can be done with regard to minor containment in an infinite sequence. Theorem (3.3) requires that $\mathcal{S}$ be a "good" class of star-decompositions. The important criterion (in this context) of $\mathcal{S}$ being good is that it has finite index. However we cannot prove (3.3) for all such indices uniformly, since this would require the full use of (2.2). Instead, we have to restrict ourselves to one fixed index n. (3.3) can then be proved for every such fixed n in $\Pi^1_1-\mathrm{CA}_0$, using the above version of (2.2) for the same fixed n. The same issue arises again in (4.2) (page 238 of Graph Minors IV), which we can again only prove for the same fixed index n. Again the same issue arises in (5.3) on page 240, which is however already conveniently phrased without quantifying over n. (5.3) is therefore actually provable in $\Pi^1_1-\mathrm{CA}_0$ for each fixed n, since the previous lemmas are invoked for instances $\leq n+1$. Consequently, the proof of (1.5) (the bounded graph minor theorem, also page 240) can then be carried out for each tree-width n in $\Pi^1_1-\mathrm{CA}_0$. From this result, an application of Π^1_1-reflection yields the full bounded graph minor theorem.

The rest of Graph Minors IV then establishes some results about patchworks, which will only be needed later in Graph Minors VIII and Graph Minors XVIII. We present the proof-principles needed for these parts and note which theorems are again provable only for each instance in $\Pi^1_1-\mathrm{CA}_0$; this will be important in the analysis of another restricted form of the graph minor theorem in Graph Minors VIII.

First, the proofs of (4.1) and (8.2) both use an arithmetical induction. Further, the proof of (8.4) uses inductions in (4) and in the argument from (7) to (8) in the form of iterated applications of certain procedures; however these inductions are also arithmetical. Next, note that (9.1) can only be proved for each n in $\Pi^1_1-\mathrm{CA}_0$, since it relies on an instance of (5.3). Note further that unlike the bounded graph minor theorem, (9.1) is a Π^1_2-statement and is hence not provable in $\Pi^1_1 - \mathrm{CA}_0$ augmented with Π^1_1-reflection for $\Pi^1_1 - \mathrm{CA}_0$. The same is true for (9.2) which can again only be proved for each n in $\Pi^1_1 - \mathrm{CA}_0$ since it employs (9.1). Note also that (9.2) uses surfaces, which we will show not to be a problem in the next chapter. Thus, all theorems of Graph Minors IV are provable in $\Pi^1_1 - \mathrm{CA}_0$ if (2.2) is assumed to hold. Since (2.2) is a Π^1_3-statement and one induction is used to prove it, $\Pi^1_1 - \mathrm{CA}_0$ augmented with Π^1_3-induction is sufficient to prove all theorems of Graph Minors IV.

2.2 A closer look at the Π^1_3-induction

It is unlikely that the Π^1_3-induction used in theorem (2.2) of Graph Minors IV is actually necessary. The theorem essentially says that there is no sequence of trees labelled from n that is minimally bad, under an arbitrary embedding relation that satisfies (in the terms of chapter 5) that if $T_1 \leq T_2$ and T_2 is a gap-subtree of T_3 and $qT_2 = qT_3$, then $T_1 \leq T_3$. The proof of (2.2) uses only (2.1), a version of (2.2) for unlabelled trees, and in the induction step itself for trees labelled from a smaller number. It is very likely that instead one could perform an induction on the theorems where (2.2) is actually used, using (2.1) in the base case of this induction and previous instances of the theorem in the induction step. (2.2) itself is only used in (3.3). In the following we give a proof of (3.3) without applying (2.2). Since this section only makes sense if read in conjunction with Graph Minors IV, and a disproportionately large number of definitions and notations are necessary to state it in this context, we refer to Graph Minors IV for the definitions of a "good" set of star-decompositions (p. 235), the "tips" of a star-decomposition (p.234), "linked" tree-decomposition (p. 235), "branching" of a tree-decomposition (p. 235), "simulation" (p.235), axioms 1 and 2 of simulation (p. 235), "well-simulated" (p. 235), "order" of an edge (p.234), and the notations $\langle T, \tau \rangle$ for tree-decompositions (p. 234) and $\tau \times T$ (p.234).

Theorem 2.2.1 ((3.3) of Graph Minors IV)**.** Let $\mathcal{S}$ be a good set of star-decompositions, and let $\langle\langle T_i, \tau_i\rangle : i \in \mathbb{N}\rangle$ be a sequence of linked tree-decompositions so that every branching of any T_i is in $\mathcal{S}$. Then $\tau_i \times T_i$ is simulated in $\tau_{i'} \times T_{i'}$ for some $i < i'$, $i, i' \in \mathbb{N}$.

Proof. The proof is by induction on the number n of distinct indices of members of $\mathcal{S}$. In the case of $n = 1$ we can use the proof of (3.3) of Graph Minors IV, with an application of (2.1) instead of (2.2). Otherwise, we use ideas from the proof of (2.2) and apply them in this context. Assume thus that $\langle T_i, \tau_i\rangle$ is a bad sequence, each branching of which is in $\mathcal{S}$. Let M be defined as having vertex-set $V(\bigsqcup_i T_i)$, and for $u \in T_i$, $v \in T_{i'}$, $i < i'$, let $\langle u, v\rangle \in E(M)$ if and only if $\tau_i \times T_i^u$ is simulated in $\tau_{i'} \times T_{i'}^v$. Define a labelling function on the edges of every T_i by letting $\phi_i(e)$ be the order of e. Let n_0 be the minimum index of members of $\mathcal{S}$. Let $F_i = \{e \in E(T_i) : \phi_i(e) = n_0\}$. Let T_i' be the contraction of T_i onto F_i, and let N be defined by restricting M to $\bigcup_i T_i'$.

Then (1) of the proof of (3.3) of Graph Minors IV holds, and using that we get (1) of the proof of (2.2) of Graph Minors IV, which gives that for all $i < i'$, if $uw \in E(N)$ for some $u \in T_i'$ and $w \in T_{i'}'$, then $uv \in E(N)$ as well for every $v \in (root(T_{i'}'), w]$. Since $\langle T_i, \tau_i\rangle$ was assumed to be bad, $\{root(T_i') : i \in \mathbb{N}\}$ is N-stable. Thus we can apply (2.1) of Graph Minors IV to get (2) of the proof of (2.2), saying that there is an increasing, M-stable sequence $\langle z_j : j \in \mathbb{N}\rangle \subseteq M^\omega$ of roots of corresponding components $R_j \subseteq T_{i_j} \setminus F_{i_j}$ (where $i_j \neq i_{j'}$ can be assumed if $j \neq j'$), so that the set $\{v \in \bigcup_i V(T_i) : \exists j \exists w \in R_j (wv \in F_{i_j})\}$ is M-rich.

Define a tree-decomposition $\langle R_j, \rho_j\rangle$ on $\tau_{i_j} \times T_{i_j}^{w_j}$ by letting $\rho_j(v) = \tau_{i_j}(v) \cup \bigcup_{w \in Succ_0(v)} \tau_{i_j} \times T_{i_j}^w$, where $Succ_0(v)$ is the set of all successors w of v so that $\phi_{i_j}(\langle v, w\rangle) = n_0$. Let $\mathcal{S}'$ be the set of all branchings of the $\langle R_j, \rho_j\rangle$, we want to show that $\mathcal{S}'$ is good. Conditions (a) and (b) of "good" are trivially satisfied. To show condition (c) let $\langle S_i : i \in \mathbb{N}\rangle$ be a sequence in $\mathcal{S}'$ so that the set of all its tips is well-simulated, then without loss of generality we may assume that $\langle S_i : i \in \mathbb{N}\rangle$ is increasing, in the sense that the indices of the corresponding T_i (or equivalently R_j) are strictly increasing. We want to show that the corresponding set of tips, using τ_i instead of ρ_i, is well-simulated.

This set of tips can be partitioned into two sets: The first corresponding to edges that have label $> n_0$ and are thus the same for τ_i and ρ_i, and the second those corresponding to edges of label n_0, which exist for τ_i but not for ρ_i. The first set was assumed to be well-simulated, and the second set is well-simulated because the set of all heads of all edges with tails in $\bigcup_j R_j$ was shown to be M-rich above. Because the union of two well-simulated sets is again well-simulated, the set of tips with regards to τ_i is well-simulated. Hence there must be $i < j$ so that $\sigma_i \times S_i$ is simulated in $\sigma_j \times S_j$, because $\mathcal{S}$ was assumed to be good and every branching of any T_i in $\mathcal{S}$. Thus condition (c) of "good" holds for $\mathcal{S}'$. Since by construction n_0 is not an index of any member of $\mathcal{S}'$ and all its other indices are indices of $\mathcal{S}$, we can apply the induction hypothesis to $\mathcal{S}'$ and $\langle\langle R_j, \rho_j\rangle : j \in \mathbb{N}\rangle$, to get that there are $j_1 < j_2$ so that $\rho_{j_1} \times R_{j_1}$ is simulated in $\rho_{j_2} \times R_{j_2}$. But then $\tau_{j_1} \times T_{j_1}^{z_{j_1}}$ is also simulated in $\tau_{j_2} \times T_{j_2}^{z_{j_2}}$, a contradiction to $\langle z_i : i \in \mathbb{N}\rangle$ being M-stable.

$\square$

Thus, (2.2) is not actually necessary to prove (3.3). However, condition (c) of "good" is a Π^1_2-statement, and so (3.3) is itself a Π^1_3-statement. Shifting the induction from (2.2) to (3.3) does hence not immediately achieve an improvement in proof-theoretic upper bounds. Examining the use of (3.3) throughout the Graph Minors series, it can be seen that it is only used in (4.2) of Graph Minors IV, and in (6.6) in Graph Minors XVIII. (4.2) of Graph Minors IV, being again a Π^1_3-statement, is only used in (5.3) of Graph Minors IV. (5.3) is then used in the proof of the bounded graph minor theorem, but more importantly it is used in the more general theorem (9.1) of Graph Minors IV. Thus, two theorems are ultimately affected by (3.3), (6.6) of Graph Minors XVIII and (9.1) of Graph Minors IV. It seems possible that the induction of (3.3) (or originally (2.2)) might be shifted to these theorems directly. Alternatively, the theorem (3.3) is only used for two concrete instances of classes of star-decompositions (for each index n). Proving (3.3) only for these instances, it might be possible to do without the assumption that $\mathcal{S}$ be good, and thus simplify the induction formula. This remains to be verified, however.

2.3 The planar graph minor theorem

As claimed in Friedman, Robertson, and Seymour (1987), the construction that shows that the bounded graph minor theorem implies extended Kruskal's theorem can be carried out with planar graphs instead. This is the aim of this section. As the planar graph minor theorem follows without proof-theoretic difficulties from the bounded graph minor theorem, this will show that extended Kruskal's theorem is equivalent over ACA_0 even to the planar graph minor theorem for cubic graphs of bounded tree-width, or to the topological graph minor theorem for these graphs.

For subcubic graphs, the existence of a minor embedding between graphs G_1 and G_2 implies the existence of a topological minor embedding between G_1 and G_2, as shown for example in Friedman, Robertson, and Seymour (1987), p. 235. As the construction below uses subcubic graphs, it is thus possible to work with topological minors, greatly simplifying some of the arguments.

First, define the $m \times n$ grid is to be the graph with vertex-set the coordinates $\{\langle i, j\rangle : 1 \leq i \leq m, 1 \leq j \leq n\}$, where $\langle i, j\rangle$ and $\langle k, l\rangle$ are adjacent if and only if $i = k \wedge |j - l| = 1$ or $j = l \wedge |i - k| = 1$. The $m \times n$ wall is then obtained from the $2m \times n$ grid by alternatingly leaving out every second horizontal edge in a sense, i.e. by removing the edges between $\langle i, j\rangle$ and $\langle i + 1, j\rangle$ for odd i and even j, and for even i and odd j.

Instead of the complete graphs of Friedman, Robertson, and Seymour (1987), we use enclosures in our construction. Enclosures are obtained as follows: Starting from an $m \times n$ wall, identify the vertices $(1, k)$ with $(2m, k)$, $1 \leq k \leq n$, where the coordinates stem from the wall being viewed as a subgraph of the $2m \times n$ grid. The resulting graph is called an enclosure with n circles (the induced subgraphs of the enclosure with vertex set $\{(l, k) : 1 \leq l \leq 2m\}$) and m spokes (the induced subgraphs of the enclosure with vertex set $\{(l,k) : 1 \leq k \leq n\} \cup \{(l+1,k) : 1 \leq k \leq n\})$. One example of an enclosure is depicted in figure 2.1.

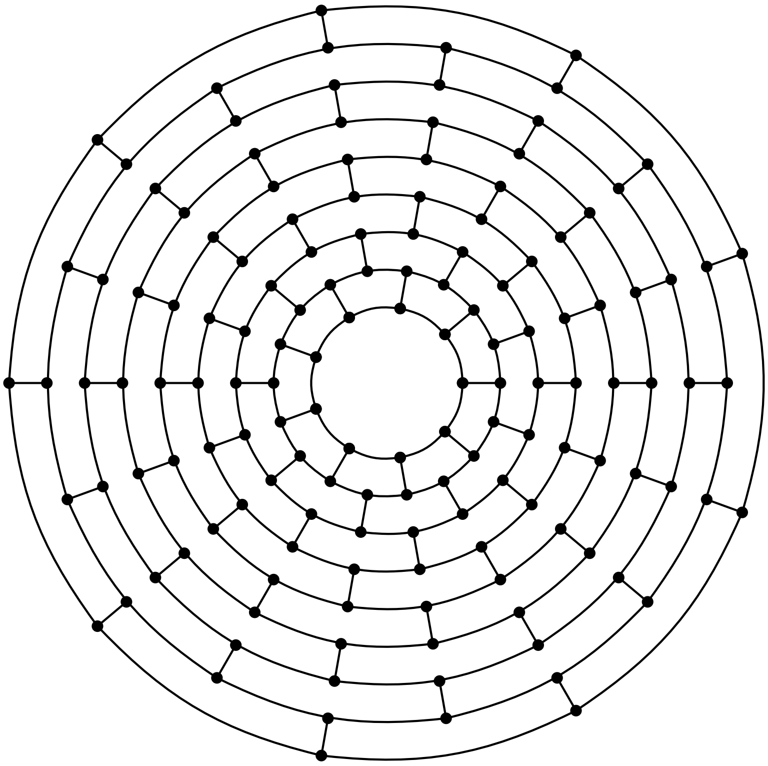

Figure 2.1: An enclosure

The circle consisting of the vertices with second coordinate 1, respectively n, is called the inner, respectively outer, circle of the enclosure.

While the complete graphs used to prove the implication $bGMT \to EKT$ were pasted together at vertices, we will use edges to connect enclosures in a tree-like shape. This is necessary to make sure that our final construction is still a cubic graph, and because under a topological minor embedding the edges are turned into paths, and paths are the objects that correspond naturally the gap-condition under embedding relations.

By the same arguments as for the usual construction, it suffices to show the result for 2-branching trees. This means that in the following, enclosures with with $3n$ spokes and

$2n+1$ circles will be used ($3n$ spokes and $3n+2$ circles for the enclosure corresponding to the root), where n is the maximum label number occurring in the labelled trees. To connect two enclosures, each edge on the outer circles of both enclosures is first subdivided. Then, if an edge with label i is meant to be simulated, if $v_1, \ldots, v_i$ are successive new vertices obtained by the subdivision in clockwise order around the first enclosure and $w_1, \ldots, w_i$ are successive new vertices obtained by the subdivision in counter-clockwise order around the second enclosure, then a new edge between v_j and w_j is added, for $1 \leq j \leq i$. One example of this process is pictured in figure 2.2.

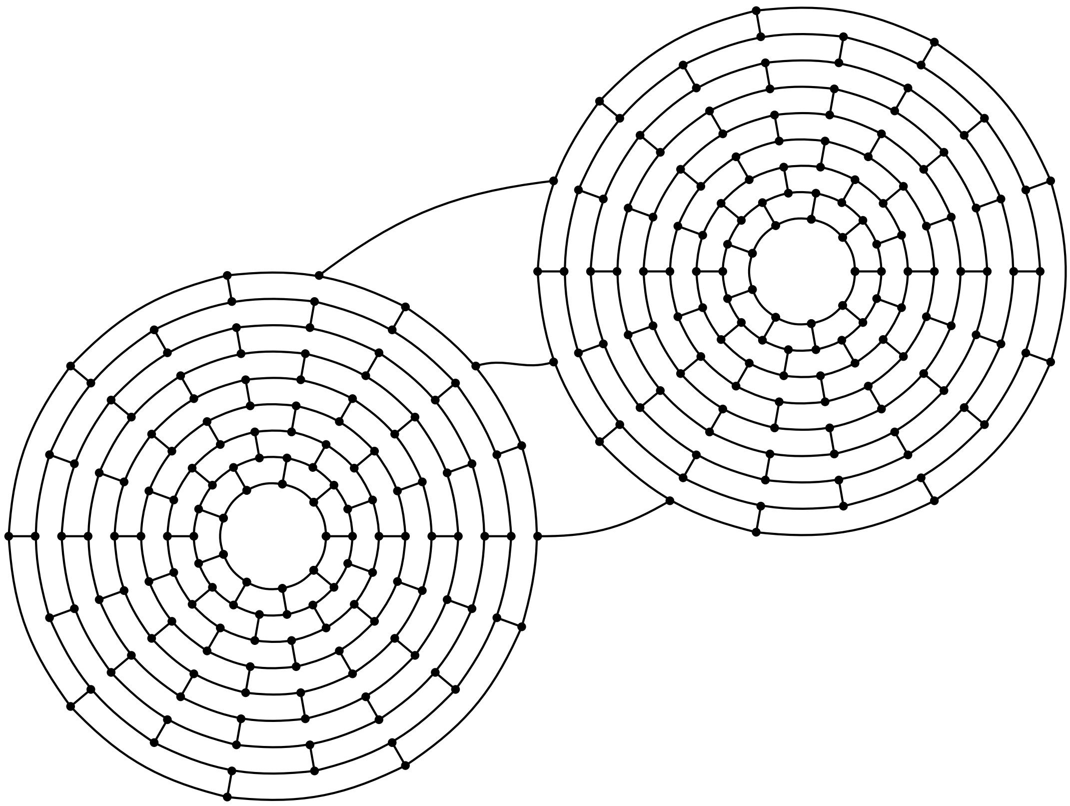

Figure 2.2: Connecting two enclosures

Note that between two adjacent enclosures there are i paths between the respective inner circles, where i is the number of edges connecting them; these paths are obtained by going along the spokes corresponding to the connecting edges in each enclosure. The goal is now to show that under a topological embedding between two constructions G_1, G_2 simulating trees T_1, T_2, the image of the inner circle of an enclosure in G_1 is completely contained in an enclosure in G_2. From this it will follow immediately that the gap-condition holds, and the other tree embedding properties follow as in the Friedman-Robertson-Seymour construction using complete graphs. For the rest of the section, topological expansions will simply be called expansions.

Lemma 2.3.1. For any topological expansion of an enclosure E of G_1 in G_2 there is a unique enclosure E' of G_2 that intersects every circle of E.

Proof. Assume $G_1 \leq G_2$ and consider any two circles C_1, C_2 of an enclosure in G_1. Their expansions must again be circles. Since they are connected by $3n$ disjoint paths in G_1 they are also connected by $3n$ disjoint paths in G_2.

Hence (as $3n > n$) some of their vertices must lie in one and the same enclosure in G_2. This is true for any two such circles. Consider now a third circle C_3 and its expansion, and the three vertices v_{12}, v_{23}, v_{13} in T_2 corresponding enclosures $E_{ij}, 1 \leq i < j \leq 3$ containing vertices of both the expansions of C_i and C_j. Then the enclosures corresponding to vertices on the path between v_{12} and v_{13} in T_2 also contain vertices of the expansion of C_1, and similar for the other circles. A standard property of trees now says that the paths between any three vertices in a tree meet in one (unique) vertex, which means that the expansions of C_1, C_2, C_3 have some of their vertices in one common enclosure of G_2. Arguing inductively, we can show that all expansions of circles of E have some of their vertices in one common enclosure in G_2. This shows the existence of E'.

To show uniqueness, consider the expansion of an enclosure of G_1 as a subdrawing of G_2. By a theorem of Whitney (1932) any drawing of a 3-connected simple graph in the sphere is unique, up to homeomorphisms of the sphere. If we consider the outer region special in a drawing of a graph in the plane, any drawing of a 3-connected graph in the plane can be obtained by considering the unique drawing in the sphere and designating an arbitrary region as the outer region (i.e. projecting it from a point inside that region onto the plane). This clearly holds for subdivisions of such graphs, and thus for expansion of such graphs, as well. Due to symmetry, effectively the only ways to draw the expansion of such an enclosure are hence in two "piles" of concentric circles, an example for such a drawing is given in figure 2.3.

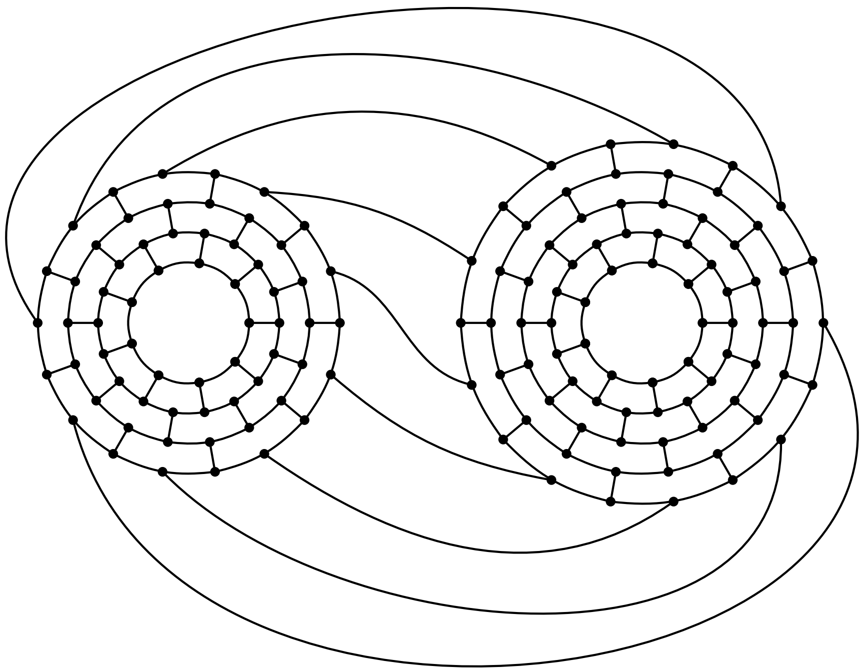

Figure 2.3: One example of a possible drawing

Since the piles must be drawn in concentric circles and two enclosures are connected by at most n paths in G_2, at most $n/2$ of the outermost circles of a pile can have vertices in an enclosure other than E'.

Since the total number of circles is $2n+1 > 2 \cdot n/2$ (or $3n+2 > 2 \cdot n/2$ for the enclosure corresponding to the root, respectively) this means that E' is unique.

□

The argument above also shows that for E corresponding to $root(T_1)$, E' corresponds to $root(T_2)$, since at least $3n + 2 - 2 \cdot n/2 = 2n + 2 > 2n + 1$ concentric circles must be contained in E'. For E an enclosure, let E' be as in the lemma above.

Lemma 2.3.2. The expansion of the inner circle of an enclosure E of G_1 in G_2 is contained in E' (assuming $G_1 \neq E$).

Proof. We may assume E does not correspond to the root of T_1, as then E' would correspond to the root of T_2 and the argument that follows can be analogously applied to this case.

So let E be an enclosure with $2n+1$ circles. By the argument about drawings of enclosures in the plane in the proof of the previous lemma, either the expansion of the outer circle or the expansion of the inner circle has at least n concentric circles around it. So either the expansion of the inner or the outer circle must be contained in E'.

Assume first that the expansion of the outer circle lies in the pile with at least $n+1$ circles, and is hence contained in E'. Because $G_1 \neq E$, the outer circle is connected to some other enclosure. Since the expansion of the outer circle is contained in E', the expansion of that enclosure has to be drawn inside the expansion of the outer circle.

But this means that E' contains at least the full expansion of one enclosure and the expansion of the outer circle, a contradiction since all enclosures have the same size. Hence the expansion of the inner circle must be completely contained in E'.

□

Theorem 2.3.3. A topological minor inclusion $G_1 \leq G_2$ induces an embedding of the corresponding labelled trees T_1, T_2, where the vertex corresponding to an enclosure E is mapped to the vertex corresponding to the enclosure E'.

Proof. We have to show that the embedding map is injective, preserves the ordering of the tree and infima, and satisfies the gap-condition.

- Injectivity: If E does correspond to $root(T_1)$, E' corresponds to $root(T_2)$ by a remark above. Then it cannot happen that for another enclosure E_* the enclosure E'_* coincides with E': The proof of the lemmas above shows that $3n+2-n = 2n+2$ circles would otherwise have to be contained in E' due to E, and that $2n+1-n = n+1$ additional circles would have to be contained in E' due to E_*. But these $2n+2+n+1 = 3n+3$ circles are more than can be accommodated in E', a contradiction. Similarly, for E not corresponding to $root(T_1)$ a minimum number of $2n+1-n = n+1$ circles have to be fully contained in E'. Since E' contains only $2n+1$ such circles, no two expansions E_1, E_2 in G_1 can have the same expansion E' in G_2 contain their inner circles.

- Ordering and infima: By the arguments above, because the enclosures corresponding to the root vertices have $3n+2$ circles we can ensure that the root is mapped to the root. By taking labels from $\{n+1, \ldots 2n\}$ instead of $\{1, \ldots, n\}$, ordering and infima are preserved by the arguments as for the usual construction.

- Gap condition: Any two adjacent inner circles E_u and E_v in G_1 are connected by k disjoint paths, where k is the edge label in the corresponding tree between u and v. Since these disjoint paths simply turn into disjoint paths under a topological expansion, there must also be at least k disjoint paths between E'_u and E'_v in G_2, meaning that every edge label on the path between the images of u and v is $\geq k$.

□

Thus the topological minor theorem for subcubic graphs, denoted by $tGMT_{\leq 3}$, implies extended Kruskal's theorem. As noted above, this means that the planar graph minor theorem implies EKT as well. This chapter thus rigorously established the equivalence

$$bGMT \leftrightarrow EKT \leftrightarrow pGMT \leftrightarrow tGMT_{\leq 3} \leftrightarrow WO(\Psi_0(\Omega_\omega)) \leftrightarrow \Pi^1_1\text{-}\mathrm{CA}_0 + \mathrm{RFN}_{\Pi^1_1}(\Pi^1_1\text{-}\mathrm{CA}_0)$$

originally stated by Friedman, Robertson, and Seymour (1987). In the next chapter we will present a way to handle surfaces in second order arithmetic, which through our analysis of Graph Minors VIII in section 4.5 will ultimately add another restricted version of the graph minor theorem to this list.

Chapter 3

Graph Minors VI

This chapter will present several methods of encoding surfaces in second order arithmetic as natural numbers. By focusing on the method that allows proofs in second order arithmetic to stay as close as possible to the original proofs, it will be shown that the contents of Graph Minors VI can be easily recreated even in ACA_0. Together with the previous chapter, this shows how to recreate most arguments of the Graph Minors series (aside from some infinitary proof techniques) in $\Pi^1_1 - \mathrm{CA}_0$, or even ACA_0, since these arguments are mostly concerned with finite combinatorial objects and surfaces. It is worth mentioning that the excluded minor theorem and many other results of the Graph Minors series are actually first order statements, and that the methods presented here allow a formalization of these results in Peano arithmetic.

In the first section of this chapter we will thus present three methods of handling surfaces in second order arithmetic. In the subsequent two sections we will then choose one of these methods to recreate the results of Graph Minors VI. It should be noted that due to the aim of recreating the proofs of Graph Minors VI in ACA_0, the contents of sections 3.2 and 3.3 are necessarily close to the original, Graph Minors VI.

3.1 Surfaces in second order arithmetic

There are two main areas of mathematics that concern themselves with expressing graphs drawn on surfaces combinatorially. The first one is of course the area of topological graph theory, which has strong interaction with certain results of the Graph Minors series such as the excluded minor theorem. The second one is computational topology, which tries to give efficient algorithms for topological problems. We will first present a method of Mohar and Thomassen (2001) that essentially only uses graphs themselves together with some additional information to encode surfaces, and then one method using fundamental polygons and graph drawings as explicit coordinates in these polygons, of which variants are used in computational topology. It should be noted that all the surfaces appearing in this thesis are compact.

We present the first method, due to Mohar and Thomassen (2001), using graphs and π-walks to represent surfaces. Firstly, to encode embeddings in orientable surfaces we choose an arbitrary, possibly edge-crossing embedding of a graph G, which for every vertex $v \in V(G)$ induces a cyclic, clockwise permutation π_v of the edges that have v as one of their endpoints. Set $\pi = \{\pi_v : v \in V(G)\}$. A π-walk is then defined to be a sequence $v_1 e_1 v_2 \dots v_k e_k v_1$ of alternating vertices and edges of minimal length with the property that $\pi_{v_{i+1}}(e_i) = e_{i+1}$, $1 \leq i \leq k$, where $v_{k+1} := v_1$ and $e_{k+1} := e_1$; note that e_1 and e_{k+1} are traversed in the same direction with respect to the π-walk, in the sense that the edge is traversed from v_1 to v_2. We then draw each such π-walk as a polygon in the plane, which will correspond to a region of the graph drawing. Then we glue together corresponding edges (i.e. edges of the polygons which were the same edge in the original graph) of all the polygons obtained in this way, where the edges are directed in the direction they were traversed in. This results in a surface in which our graph G is 2-cell embedded, which means that every region of G is isomorphic to a disk. Using this method, any 2-cell embedding of a graph in an orientable surface can be encoded by a natural number.

Likewise, if G is 2-cell embedded in a possibly non-orientable surface, observe first that we can still talk about a clockwise ordering of edges around a vertex v, since there is a neighbourhood of v that is isomorphic to the plane, as long as v does not have degree 2. To carry out this construction, vertices of degree 2 thus have to be suppressed; however this is no restriction to generality since vertices of degree 2 do not pose an obstacle to embeddability. As above, a clockwise ordering of edges around v gives rise to a permutation π_v of edges incident with v; set $\pi = \{\pi_v : v \in V(G)\}$. Furthermore, for each edge $e = uv$ we can check whether the chosen clockwise orderings around u and v represent the same direction or not, again because there is a neighborhood around e, u and v which is planar. If the orderings agree we set $\lambda(e) = 1$, if not then we set $\lambda(e) = -1$. We can now construct a walk and polygons as in the orientable case, if we modify our procedure so that we use π^{-1} after having traversed an edge e with $\lambda(e) = -1$. Finally, to get an encoding of the original graph without suppressed vertices, we can subdivide edges of G until we obtain the original graph, setting $\lambda(e) = 1$ for one of the new edges obtained by any subdivision and retaining the old value of λ for the other.

The drawback of this method is that it can only encode drawings of graphs in surfaces that are 2-cell. This problem could be solved by introducing additional edges and vertices to obtain a 2-cell drawing of a graph that is embedded, but not 2-cell embedded in a surface, and then to label these additional vertices and edges to signify that they are not part of the original graph. In general, a direct translation of the proofs of the Graph Minors series into second order arithmetic using this method would be rather complicated. It should however be noted that this method would likely be fully capable of carrying out the proofs of the Graph Minors series, as for example a proof of the excluded minor theorem (originally proved in Graph Minors XVI) using this method is given by Mohar and Thomassen (2001). Nonetheless, in order to simplify the translation of the proofs of the Graph Minors series, we examine a second method.

The second method works as follows: We encode the surface by its fundamental polygon, and edges drawn as general curves are replaced by polygonal chains. Using polygonal chains causes no loss of generality, as shown for example in Mohar and Thomassen (2001), Lemma 2.1.2. This will make it possible to encode the surface and the drawing of the graph by a natural number, thus greatly simplifying the arguments in second order arithmetic.

To do this we work with the unit square $[0,1]^2$ (which produces codes that are easier to handle than for example the unit circle) and divide it into (labelled) sections corresponding to the fundamental polygon of the surface that is to be represented. For example, the torus $aba^{-1}b^{-1}$ would then lead to a division of $[0,1]^2$ into $\{\langle 0,x\rangle : x \in [0,1]\}$, $\{\langle x,1\rangle : x \in [0,1]\}$, $\{\langle 1,x\rangle : x \in [0,1]\}$ and $\{\langle x,0\rangle : x \in [0,1]\}$, labelled with a, b, a^{-1} and b^{-1} respectively. In general, we obtain a division from polygon $a_1 \dots a_n$ by dividing $[0,4]$ into n sections of equal length and then mapping it to the unit square via

$$x \longmapsto \begin{cases} \langle 0,x\rangle & \text{if } 0 \leq x \leq 1, \\ \langle x-1,1\rangle & \text{if } 1 \leq x \leq 2, \\ \langle 1,3-x\rangle & \text{if } 2 \leq x \leq 3, \\ \langle 4-x,0\rangle & \text{if } 3 \leq x \leq 4. \end{cases}$$

We also make one further convention: If a polygonal chain goes through identified sides, we require that an additional vertex be added on each of those sides at the point where it intersects the polygonal chain. Furthermore, we require that all vertices of our graphs and polygonal chains have rational coefficients, which is possible without loss of generality.

Then a graph drawn on a surface can be encoded by the sequence of sides of the fundamental polygon of the surface together with their directions, and the (rational) coordinates of the vertices of the graph and the polygonal chains.

The most important tools for handling surfaces used in the Graph Minors series are:

- Deciding whether two surfaces are homeomorphic,

- pasting two surfaces together at two cuffs,
- cutting surfaces along a graph (an operation very frequently used in the Graph Minors series), and
- deciding whether two paths on a surface are homotopic.

We will now show how to handle these operations in second order arithmetic.

Deciding whether two surfaces are homeomorphic can be done by just looking at the sequences of the sides of the fundamental polygon, as in Lefschetz (1954), page 75 onward. The technique presented there is only for connected surfaces without boundary. Deciding whether two surfaces with boundary are homeomorphic is then achieved by counting the number of cuffs, if they are equal and the surfaces with boundary filled in are homeomorphic, then the surfaces with boundary are homeomorphic. For surfaces with multiple components, to decide whether two surfaces are homeomorphic it suffices to pair the components of the two surfaces so that each pair is homeomorphic.

Since a surface with multiple components can be encoded as a sequence or set of connected surfaces, we will in the following restrict ourselves to connected surfaces, noting that everything generalizes without problems to disconnected surfaces.

To carry this out formally, for $\sigma = \langle x_1, \ldots, x_n \rangle$ denote by σ^{-1} the sequence $\langle x_n^{-1}, \ldots, x_1^{-1} \rangle$. Then, to decide whether the closed surfaces encoded by two sequences σ_1, σ_2 are homeomorphic, written $\sigma_1 \sim \sigma_2$, can be decided by repeated application of the following four rules, see Lefschetz (1954), page 75 onward:

1) For any cyclic permutation σ' of σ, $\sigma \sim \sigma'$.

2) If $\sigma = \sigma_1 \circ \langle x, x^{-1} \rangle \circ \sigma_2$, then $\sigma \sim \sigma_1 \circ \sigma_2$.

3) If $\sigma = \sigma_1 \circ \langle x \rangle \circ \sigma_2 \circ \sigma_3 \circ \langle x^{-1} \rangle$, then $\sigma \sim \sigma_1 \circ \langle y \rangle \circ \sigma_3 \circ \sigma_2 \circ \langle y^{-1} \rangle$, where y does not appear in σ.

4) If $\sigma = \sigma_1 \circ \langle x \rangle \circ \sigma_2 \circ \sigma_3 \circ \langle x \rangle$, then $\sigma \sim \sigma_1 \circ \langle y \rangle \circ \sigma_3 \circ \langle y \rangle \circ \sigma_2^{-1}$, where y does not appear in σ.

However, simply deciding whether two surfaces are homeomorphic is not sufficient when there are graphs drawn on these surfaces. In this case we need to do the actual surface manipulations corresponding to the rules above by performing the corresponding cutting and pasting procedures. We will now show how to achieve this.

Pasting along two sides can be done as follows: First "elevate" the side to paste, by introducing a new segment on the polygonal chain going straight up (without loss of generality the side contains a segment of $\{\langle x, 1 \rangle : x \in [0, 1]\}$, otherwise rotate the unit square), then fold the other sides on the same side on the square to the left and right, respectively, then paste the two polygons together at the side where they should go together. Then rescale the square and then the sides, to make them match with our definitions. This procedure is depicted in figure 3.1.

For the above procedure we need to specify how to scale sides to intended lengths. This is needed firstly for pasting because the sides to be pasted need to be in the correct position with regard to their respective squares, and they need to have the same length. Secondly, after pasting the two squares together, the sides need to be rescaled in order to comply with our definition. This will also be necessary when cutting the surface along a graph. When scaling a side, there are two cases. The first and easy one is that the side and its scaled end-version lie completely inside the same side of the square. In this case, we can lift all vertices by some small rational ϵ away from the side, so that no vertex not lying on one of the sides contained in the same side of the square is within ϵ distance of the side of the square, and so that lifting the vertices does not cause any crossings in the drawing of the graph. We can then scale the side, and connect the lifted vertices with the points corresponding to their original position on the rescaled side, depicted in figure 3.2. The second one is when a side is moved around a corner. In that case, we need to first "stagger" the vertices on the side within an ϵ distance, then connect them straight to

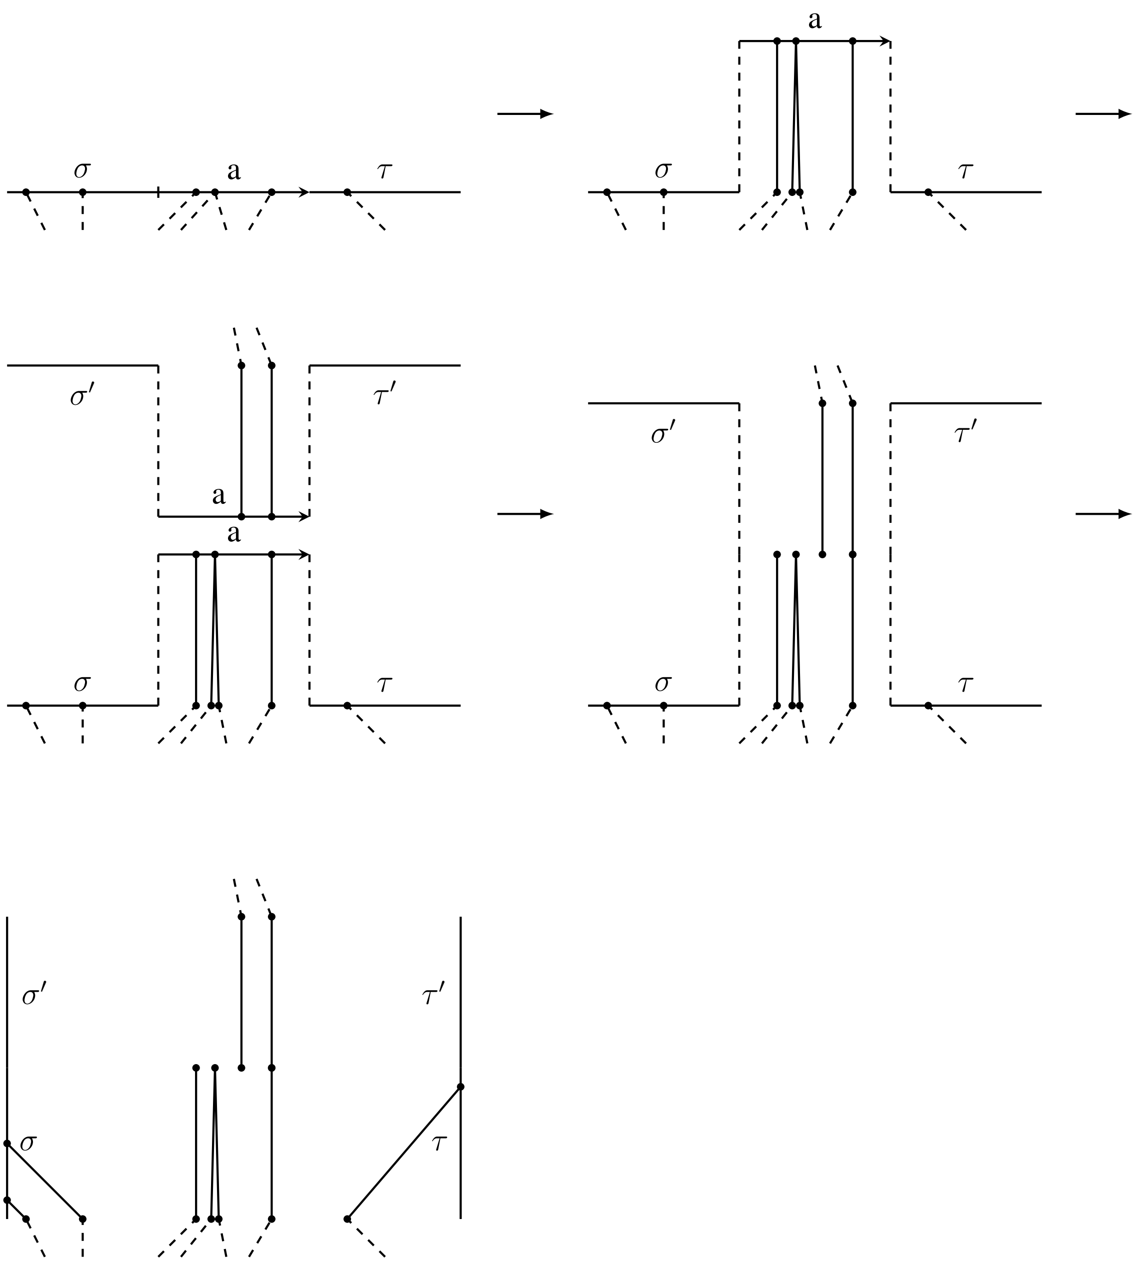

Figure 3.1: Pasting two polygons. The two sides to be pasted are elevated so that they do not interfere with the rest of the polygons. When pasting the sides together, corresponding vertices of the graph drawings are identified. Finally, the two fundamental polygons are connected to create the new surface.

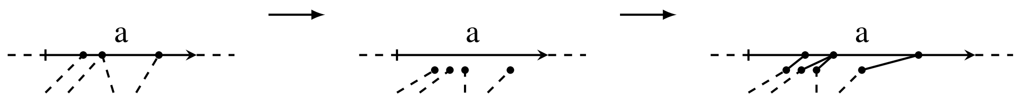

Figure 3.2: Scaling a side

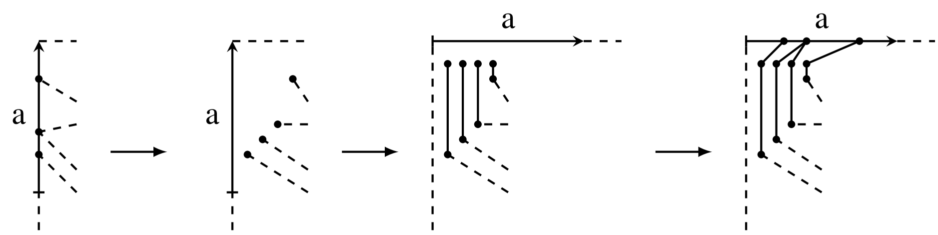

Figure 3.3: Moving a side around a corner. In order to preserve the graph drawing, vertices originally lying on the side have to be connected to their intended new location. Since edge crossings need to be avoided during this process, the vertices have to be moved away from the side by differing distances.

vertices situated some ϵ' away from the side the edge is moved to, then connect them to the appropriate points on the new side. This procedure is illustrated in figure 3.3. If only parts of a side move around a corner, we can split the side into two parts, the one that has to move around the corner and the one that does not, then apply the appropriate steps of first moving the second part away from the corner, then moving the first part around the corner and then merging the two parts again.

For cutting, it suffices to investigate cutting along a closed non-crossing polygonal chain (i.e. polygonal chains homeomorphic to a circle), since cutting along any other graph can be reduced to subsequent cutting along such chains. However, some considerations need to be made to carry out this procedure properly. If the polygonal chain cuts through a side of the fundamental polygon, then the two sides so created and their possible equivalents have to receive new labels in order to identify them correctly. Then, the components finally obtained by the cutting procedure have to be pieced together again by identifying some corresponding sides, and finally the sides have to be reordered or rearranged to

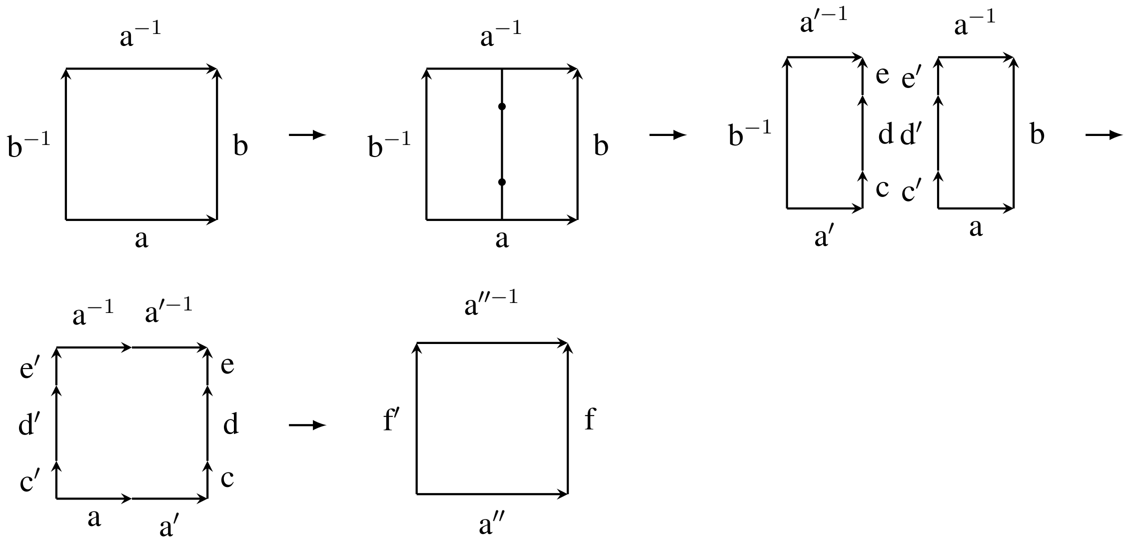

Figure 3.4: Cutting the torus along a circle

comply with the standard form of a surface. The whole procedure is depicted for the simple example of the torus in figure 3.4. The polygonal chain will in general not be neatly aligned with the unit square, so we also need to make explicit how to move the newly obtained cuff to the side of the square, in order for the fundamental polygon to match our definition. This is done by projecting the sides of any "triangle" until we reach the side of the square, depicted in figure 3.5.

If the chain does not intersect the square, then we have to draw an extra line toward the boundary of the square, split the surface along it and identify the sides of the extra line introduced in this way. By taking these sides to not be perfectly horizontal or vertical, we can use the above procedure of projecting triangles to then bring the new surface into normal form. This piece of the cutting procedure is pictured in figure 3.6.

One further tool is needed for applications in for example Graph Minors VI, namely to decide whether two paths on a surface are homotopic. Conveniently, a simple linear time algorithm for this problem due to Erickson and Whittlesey (2013) already exists. The algorithm (given only for orientable surfaces without cuffs which are not the torus) works

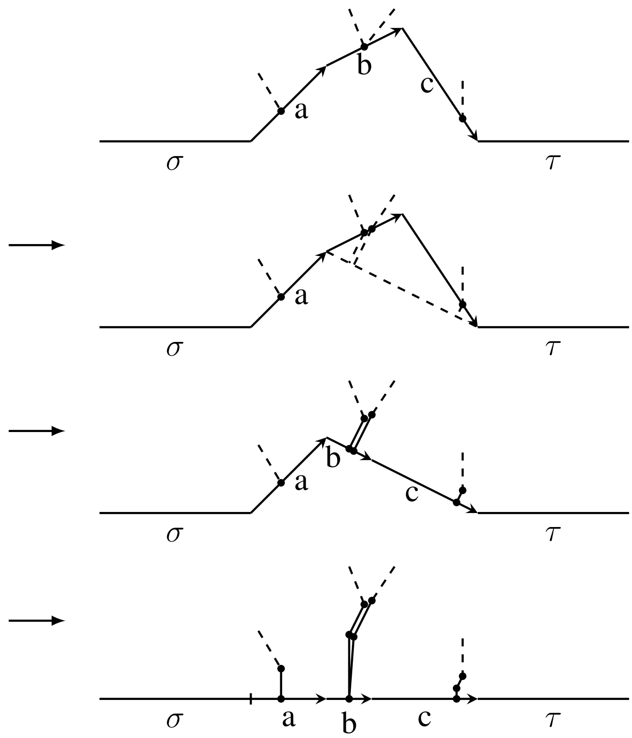

Figure 3.5: Projecting a cut toward the sides. Since the fundamental polygon is required to coincide with the boundary of $[0,1]^2$, this is a necessary step in order to retain the normal form after a cut. The simplest way to remove any "indention" made by a cut is to project two adjacent sides toward the third side of the triangle formed by their endpoints until all sides lie on the boundary of $[0,1]^2$.

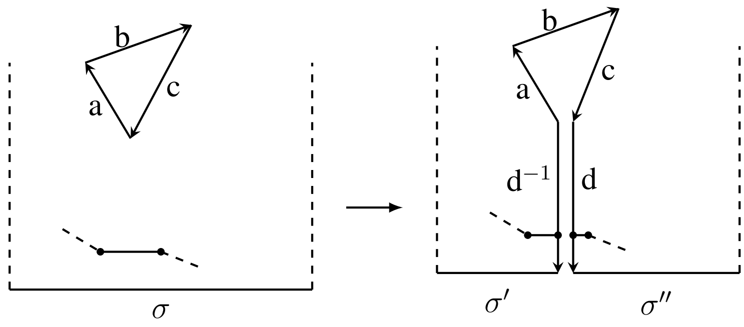

Figure 3.6: Connecting a cut to the sides of the square. In order to perform this procedure, the surface has to be divided along a line toward the sides of the fundamental polygon. Identifications have to be made whenever an edge is intersected by this division. After that, the procedure is the same as in figure 3.5.

as follows: Assume we are given a surface as a fundamental polygon and a closed walk w on that surface, with its vertices on the corners of the fundamental polygon; we want to decide whether the walk is null-homotopic. Then to decide whether two paths are homotopic, we look at whether their concatenation (with one of the paths reversed) is a null-homotopic walk. The algorithm of Erickson and Whittlesey (2013) proceeds as follows.

First, we transform our surface into a system of quads, that is we add a new vertex z in the middle of our polygon and one edge from z to each corner of the polygon. The closed walk w is transformed into a walk w' on this system of quads by replacing every edge $e = uv$ with the two edges uz and zv. Since v and u lie on the boundary of the polygon, there are actually two possible choices for uz and zv, any choice works and we choose the one that is further to the left and then top of the unit square. Clearly w is null-homotopic if and only if its transformation is, since the path is merely changed in a planar neighbourhood.

We import a number of definitions and notations from Erickson and Whittlesey (2013). The turn between two subsequent edges uz and zv (and on the vertex where these edges meet) of w' is defined as the number of corner points encountered when going around the fundamental polygon from u to v in clockwise direction. The turn sequence of w' is then the sequence of turns of pairs of subsequent edges of w'. Further, define the notations $\bar{t} := -t \mod p$, where p is the number of corners of our polygon, and $\langle a^n, b^m, \ldots\rangle := \langle \underbrace{a, \ldots, a}_{n \text{ times}}, \underbrace{b, \ldots, b}_{m \text{ times}}, \ldots\rangle$. The following reductions then serve to eliminate certain features appearing in a turn sequence, and a cycle is null-homotopic if and only if its turn sequence can be reduced to the empty sequence by these; see Erickson and Whittlesey (2013) for a proof of correctness and further explanation of the reductions.

$$\tau \circ \langle x, 0, y\rangle \circ \tau' \to \tau \circ \langle x + y\rangle \circ \tau'$$
$$\langle 0, 0\rangle \to \langle\rangle$$
$$\tau \circ \langle x, 1, 2^k, 1, y\rangle \circ \tau' \to \tau \circ \left\langle x - 1, \bar{2}^k, y - 1\right\rangle \circ \tau'$$
$$\langle x, 1, 2^k, 1\rangle \to \left\langle x - 2, \bar{2}^k\right\rangle$$
$$\langle 1, 2^k\rangle \to \left\langle \bar{3}, \bar{2}^{k-2}\right\rangle$$
$$\tau \circ \left\langle x, \bar{1}, \bar{2}^k, \bar{1}, y\right\rangle \circ \tau' \to \tau \circ \langle x + 1, 2^k, y + 1\rangle \circ \tau'$$
$$\left\langle x, \bar{1}, \bar{2}^k, \bar{1}\right\rangle \to \langle x + 2, 2^k\rangle$$
$$\left\langle \bar{1}, \bar{2}^k\right\rangle \to \langle 3, 2^{k-2}\rangle$$

Testing whether a turn sequence can be reduced to the empty sequence can be done in linear time. However, since we are not concerned with runtime, we use a conceptually even simpler, but similar algorithm. We partition the surface by choosing as above a vertex z lying in the middle of our polygon, then connecting it to the corners of the sides of the polygon and (as opposed to Erickson and Whittlesey (2013)) also connecting the corners via the sides of the polygon themselves to obtain a triangulation of the surface instead of a system of quads. Then we identify paths that go in the same direction around a triangle, while accounting for identification of sides of the fundamental polygon. Then two paths are homotopic if and only if they can be transformed into each other by some sequence of such identifications. An example of this procedure is presented in section 3.3, see figure 3.7.

3.2 Paths on a disk

Our goal is now to show how to recreate the contents of Graph Minors VI in ACA_0, using the techniques presented in the previous section. Graph Minors VI deals with the problems of determining whether it is possible to find k non-crossing paths between some specific vertices in a graph drawn on the disk or cylinder. This section will examine the parts of Graph Minors VI that deal with the disk case, and the following section those that deal with the cylinder case. It should be stated again that the aim of this section and the next is to recreate the proofs of Graph Minors VI as true to the original as possible, so as to establish that this is mostly straightforward and need not be done in detail for the later papers of the Graph Minors series. We start off with some necessary definitions of Graph Minors VI.

A division in the context of Graph Minors VI is a partition of some vertex set, i.e. a finite set of sets of vertices $\Delta := \{\delta_1, \ldots, \delta_n\}$ where $\delta_i \neq \emptyset$ and $\delta_i \cap \delta_j = \emptyset$ for $i \neq j$. The idea is that every element δ of Δ is meant to represent a tree T_δ in G connecting the vertices of δ, and that $T_\delta \cap T_{\delta'} = \emptyset$ if $\delta \neq \delta'$. If it is possible to find such trees, we obtain a forest with $|\Delta|$ components in G, where distinct components of the forest connect distinct vertex sets $\delta \in \Delta$. Such a forest is called a realization of Δ in G, and if such a realization exists, Δ is called feasible in G.

The objective of Graph Minors VI is to show that the problem of deciding whether a division Δ in a graph G is realizable has a polynomial time algorithm, in the case where G is drawn on a disk or cylinder and $\bigcup \Delta$ is subset of the boundary of the disk or cylinder. For our purposes, only the solution to this problem matters, whether the algorithm runs in polynomial time or not is irrelevant.

Graph Minors VI first examines the case of divisions on a circle, that is when G is embedded in a disc and $\bigcup \Delta$ is a subset of the boundary of the disc. A disc in our model is just a fundamental polygon that has only one boundary component C, say. Let C be such a circle, i.e. the boundary of the disc. Since C is then the boundary of the unit square, it is straightforward to define when for a vertex set $V = \{v_1, \dots, v_n\}$ on the circle, one v_i comes after some v_j when going around the circle in clockwise order. In most cases, it suffices to map $[0, 1]^2$ to $[0, 4]$ as detailed in the previous section, then determine whether the image of v_i lies to the left (i.e. is less than) the image of v_j and there is no image of another v_k in between the two. There is one special case which this procedure does not cover, namely if v_j is such that there is no vertex to the left of its image on $[0, 4]$. In this special case we can paste another copy of our interval $[0, 4]$ to the left of $[0, 4]$, and then decide whether the image v_i in the left copy comes directly before the image of v_j in the right copy. If $V \subseteq C$ then we call each part of C between clockwise subsequent elements of V a segment with respect to V. Now, if Δ is a division so that Δ is a partition of V, then we define an additional set U as follows. If $V = \emptyset$ then also $U := \emptyset$, if $V \neq \emptyset$ then let $\mathcal{S}$ be the set of segments of C with respect to V, and for each $S \in \mathcal{S}$ choose some $u(S) \in S$ (for example let $u(S)$ be the midway point of the segment S), then let $U = \{u(S) : S \in \mathcal{S}\}$.

For vertices a, b, c, d lying on the boundary of our circle, $\{a, c\}$ is defined to cross $\{b, d\}$ if, when starting at a and going around the boundary in clockwise direction, a, b, c, d are encountered exactly in this order. Then the lines ac and bd intersect inside the circle, and since a, b, c, d lie on the boundary of the circle, every polygonal chain from a to c must cross every polygonal chain from b to d.

For Δ a division and $u, u' \in U$, where U is as above, write $u \sim u'$ if and only if $\{u, u'\}$ is not crossed by any $\{v, v'\}$ where v and v' lie in the same $\delta \in \Delta$. Since $U \cap V = \emptyset$, the relation $\sim$ is transitive and hence an equivalence relation. Then $U/\sim$, the set of equivalence classes of U under $\sim$, is a division Δ^* on U. Analogously to Δ^*, we can define a division Δ^{**} on V by starting with Δ^* and U instead of Δ and V. The relation $\sim$ can be thought of as follows: If $u \sim u'$, then for any $\delta \in \Delta$, $\bigcup \delta$ is either completely contained in (u, u') or in (u', u), where (u, u') denotes the part of c that is obtained by going clockwise from u to u' around c. Thus, for $w, w' \in \delta^* \in \Delta^*$, we can find a path in C between w and w' that does not cross any of the trees corresponding to the $\delta \in \Delta$.

We expand the definition of crossing from two-element sets to arbitrary sets A, B, and say that A crosses B if and only if there are $a, a' \in A$, $b, b' \in B$ such that $\{a, a'\}$ crosses $\{b, b'\}$. A division Δ is called cross-free if none of its members cross. If G is a graph drawn on a disc and Δ a feasible division in G then Δ has to be cross-free by the remarks above.

We also make some further definitions of frequent notions in the context of graph drawings. First, we make the convention that for the rest of this section all graphs G are drawn in such a way that edges intersect the boundary C only in their endpoints. Denote by $R(G)$ the set of regions of G (inside the disk). The geometrical dual G^* of a connected graph G can then be obtained by taking the vertex set of G^* to be $R(G)$, with two vertices of G^* connected by an edge if they are adjacent as regions in the drawing of G. Then there is a natural drawing of G^*, with the vertices of G^* inside their corresponding regions, edges of G^* crossing the edge that separates the two regions corresponding to its endpoints, and with vertices that correspond to regions incident with C drawn on C. A well-known fact is that $G^{**} = G$.

The following lemmas of Graph Minors VI are straightforward to prove in our model, since they do not actually require the use of a surface at all. Thus, we only reproduce the proof of the first such lemma, and give proofs where they have not been given in Graph Minors VI.

Lemma 3.2.1 ((2.1) of Graph Minors VI)**.** For all Δ, Δ^* is cross-free.

Proof. Suppose we have $u_1, u_3 \in \delta^*$, $u_2, u_4 \in \delta^{*'}$, $\delta^*, \delta^{*'} \in \Delta^*$, so that $\{u_1, u_3\}$ crosses $\{u_2, u_4\}$. Suppose $\delta^* \neq \delta^{*'}$, then since δ^* and $\delta^{*'}$ are different equivalence classes with regard to $\sim$, there are $v_1, v_2 \in \delta \in \Delta$ such that $\{u_1, u_2\}$ is crossed by $\{v_1, v_2\}$. But then $\{v_1, v_2\}$ must cross either $\{u_1, u_3\}$ or $\{u_2, u_4\}$ as well, a contradiction.

$\square$

Lemma 3.2.2 ((2.2) of Graph Minors VI)**.** Whenever Δ is cross-free, Δ and Δ^{**} coincide.

Lemma 3.2.3 ((2.3) of Graph Minors VI)**.** If $U \cup V \neq \emptyset$ then there is $\delta \in \Delta \cup \Delta^*$ with $|\delta| = 1$.

For $x, y \in U \cup V$, let $\Delta(x, y)$ be the cardinality of the set

$$\{\delta \in \Delta : \delta \cap \{x, y\} \neq \emptyset \vee \exists u, v \in \delta(\{u, v\} \text{ crosses } \{x, y\})\}.$$

$\Delta^*(x, y)$ is then defined similarly, using Δ^* instead of Δ in the set above.

If we view each $\delta \in \Delta$ as a tree T_δ in G connecting certain vertices on the boundary of C, $\Delta(x, y)$ can be interpreted as the number of such trees every polygonal chain from x to y inside C has to intersect. Likewise, if we view $\delta^* \in \Delta^*$ as a tree-like set of polygonal chains connecting members of U without intersecting any T_δ, $\Delta^*(x, y)$ denotes the number of such sets any polygonal chain from x to y in C has to intersect. Thus, intuitively, for $x, y \in V$, if a polygonal chain from x to y intersects as few members of Δ as possible, say n, then it intersects exactly $n - 1$ members of Δ^*, and similarly for $x, y \in U$ with the roles of Δ and Δ^* interchanged. Of course, this interpretation is only possible if Δ is feasible in G, but the next lemma tells us that basically the same result holds simply when Δ is cross-free.

Lemma 3.2.4 ((2.4) of Graph Minors VI)**.** Let Δ be cross-free, $x, y \in U \cup V$. Then Δ and Δ^* are related as follows:

$$\Delta(x, y) - \Delta^*(x, y) = \begin{cases} 1 & \text{in case both } x \in V \text{ and } y \in V, \\ -1 & \text{in case both } x \in U \text{ and } y \in U, \\ 0 & \text{in case } x \in V, y \in U \text{ or } x \in U, y \in V. \end{cases}$$

Graph Minors VI now turns toward providing an algorithm to decide whether an arbitrary division Δ is feasible in a graph G, and toward giving an equivalent criterion of feasibility. Let G be a graph drawn on a disc D, let Δ be a division with $\bigcup \Delta = V(G) \cap bd(\Sigma)$.

Lemma 3.2.5 ((3.2) of Graph Minors VI)**.** As above, let D be a disc with fundamental polygon C, and assume Δ is a cross-free division of the vertices of G lying on C. Without loss of generality assume further that the set vertices of G^* lying on C coincides with U, i.e. that Δ^* is a division of these vertices.

Then there is a realization of Δ in G if and only if there is one of Δ^* in G^*.

Proof. Assume Δ is realizable in G. By the discussion above, an element δ^* of Δ^* corresponds to a maximal set of vertices of $V(G^*) \cap C$ which can be connected by a tree-shaped set of polygonal chains P without crossing any tree T_δ, $\delta \in \Delta$, of the realization of Δ. Denote by R_P the set of regions of G through which any element of P passes, and by E_P the set of edges of G which are crossed by an element of P. Then $T_{\delta^*} := \langle R_P, E_P \rangle$ is a tree in G^* connecting the vertices of δ^*, and no two such trees intersect. Hence $\{T_{\delta^*} : \delta^* \in \Delta^*\}$ is a realization of Δ^* in G^*.

Since G is connected, $G^{**} = G$, and so by the same argument as above if Δ^* is feasible in G^* then Δ is feasible in G. This finishes our proof.

□

Graph Minors VI then presents some facts which can immediately seen to be correct and that help reduce the problem of realizability to smaller subproblems, and with the results so far give an algorithm to decide whether a division Δ is feasible in a graph G.

Lemma 3.2.6 ((3.3) of Graph Minors VI)**.** The following hold:

i) If there is $\delta \in \Delta$, $v \in G$ such that $\delta = \{v\}$, then there is a realization of Δ in G if and only if there is one of $\Delta \setminus \{\delta\}$ in $G \setminus v$. Hence we can simplify our problem by deleting appropriate vertices.

ii) If there is a component $G' \subseteq G$ with $G' \cap C = \emptyset$, then there is a realization of Δ in G if and only if there is one of Δ in $G \setminus G'$, so we can delete such components and obtain a simpler problem.

iii) If there are components $G', G'' \subseteq G$ with $G' \neq G''$ and $\delta \in \Delta$ with $\delta \cap G' \neq \emptyset \neq \delta \cap G''$ then there can be no realization of Δ in G.

iv) If $G_1, \ldots, G_n$, $n \geq 2$, are the components of G and the preceding two reductions can not be applied, let $\Delta_1, \ldots, \Delta_n \subseteq \Delta$ be chosen so that $\bigcup \Delta_i = V(G_i) \cap C$. Then there is a realization of Δ in G if and only if there is a realization of Δ_i in G_i for all i.

The lemma above effectively gives an algorithm to decide whether a division is feasible in a graph.

Lemma 3.2.7 ((3.4) of Graph Minors VI)**.** The following algorithm decides whether a division Δ is feasible in a graph G.

1) First verify whether Δ is cross-free, if not return no.

2) Apply the reductions corresponding to lemma 3.2.6 until they cannot be applied anymore. If 3.2.6.iii is applicable at any point, return no.

3) If G now still consists of at least one vertex then G must be connected since 3.2.6.iv cannot be applied. Since 3.2.6.i cannot be applied to G either, it can be applied to G^* by 3.2.3. We can thus replace G by G^* and Δ by Δ^* and go back to step (2).

4) If G is null return yes.

Graph Minors VI now looks at what can go wrong if a division is not realizable in D. A vertex of G is called peripheral if it lies on the boundary of D, and a region of G is called peripheral if it is incident with the boundary of D. Let V, U and Δ^* be as above, and let $x, y \in V \cup U$. A sequence $\langle A_1, \ldots, A_k \rangle$, $k \geq 1$ of vertices and regions is called an (x, y)-chain if it satisfies the following:

i) each A_i, A_{i+1} is an adjacent vertex-region pair,

ii) $A_1 = x$ or $x \in A_1$, depending on whether A_1 is a vertex or region, and similar for A_k,

iii) only A_1 and A_k are peripheral.

If $\langle A_1, \ldots, A_k \rangle$ is an (x, y)-chain, define its length $l(\langle A_1, \ldots, A_k \rangle)$ to be $|\{i : 1 \leq i \leq k \wedge A_i \text{ is a vertex}\}|$, and its Δ-redundancy by $l(\langle A_1, \ldots, A_k \rangle) - \Delta(x, y)$. Δ-redundancy can be understood as a kind of "unnecessary detours" that an (x, y)-chain takes to get from x to y.

Lemma 3.2.8 ((3.5) of Graph Minors VI)**.** Let G be a connected graph embedded in a disk. Then for every (x, y)-chain in G^* we can find a corresponding (x, y)-chain in G. If in addition Δ is cross-free then the redundancies of the two chains are equal.

Proof. Because G is connected, no two distinct vertices of U are contained in some region R of G. Hence we can draw G^* so that $V(G^*) \cap c = U$, and then any (x, y)-chain in G^* is an (x, y)-chain in G, where of course vertices in G^* are regions in G and vice versa.

For the second part, consider first the case where x and y are regions in G^* and thus vertices in G. An (x, y)-chain in G^* then contains one vertex less than the corresponding (x, y)-chain in G, so its length is one less in G^* than it is in G. On the other hand, because Δ is cross-free we can apply lemma 3.2.4, and hence $\Delta(x, y) - \Delta^*(x, y) = 1$. Their Δ/Δ^*-redundancy respectively is thus the same, and an analogous argument proves the case where x and y are both vertices of G^* and the case where one of x, y is a region and the other a vertex.

□

From the results so far, it is clear that if Δ is cross-free it is realizable in some graph G', for example by taking G' to be tree-like sets of polygonal chains, each set connecting the vertices of one $\delta \in \Delta$. The only other obstacle to realizability is thus a lack of edges in G to accommodate the necessary trees disjointly. This is made precise in the following lemma.

Lemma 3.2.9 ((3.6) of Graph Minors VI)**.** Let G and Δ be as above, then there exists a realization of Δ in G if and only if

i) no two elements of Δ cross, and

ii) for every $x, y \in U \cup V$ and (x, y)-chain A, it holds that $l(A) \geq \Delta(x, y)$.

Proof. That (i) is necessary is obvious by the remarks and results so far. That (ii) is necessary follows because for every $\{u, v\}$ that crosses $\{x, y\}$ there has to be a path from u to v which necessarily has at least one vertex in common with every (x, y)-chain. Because the length of any (x, y)-chain is the number of vertices in it, by Menger's theorem this means that this number is at least the number of δ that contain $\{u, v\}$ that cross $\{x, y\}$, plus the number of δ that contain x or y, otherwise there would not be enough paths in G linking u and v. Thus the length of any (x, y)-chain has to be greater than or equal to $\Delta(x, y)$.

Sufficiency is proved as in Graph Minors VI, by induction on $|V(G)| + |R(G)|$. In short, we apply 3.2.6 if possible and are done by the induction hypothesis (3.2.6.iii cannot apply because of assumption (ii)), and if no case of 3.2.6 can be applied, we may consider G^* and Δ^* instead by 3.2.1 and 3.2.8, then apply 3.2.6.i to G^* and Δ^* by 3.2.3, and conclude by using the induction hypothesis and 3.2.5.

□

3.3 Paths on a cylinder

This section handles the parts of Graph Minors VI that deal with deciding whether there are a certain number of disjoint paths across a cylinder. This case is more complicated than the disk case since two paths on a cylinder are not necessarily homotopic. Again, it should be noted that the aim of this section is to reproduce the remaining contents of Graph Minors VI as close as possible to the original in second order arithmetic.

In the following it will be important that the homotopy type of a path from one cuff of the cylinder to the other is determined by its winding number. To show this we use our method of triangulating surfaces presented in section 3.1.

Let C denote the cylinder. In the fundamental polygon C, label the cuffs with c_1 and c_2 and the identified sides with a and a^{-1}, respectively. Assume that c_1 and c_2 point in the same direction, say to the right, and a points upwards. Introduce a new vertex z in the middle of the polygon and connect it to all the corner points. Denote the corner points by the two letters of the sides (in counterclockwise order) that meet there, e.g. ac_1 for the corner where a meets c_1. Denote the edges incident with z by ac_1z, zc_1a, zc_2a and ac_2z in counterclockwise order starting in the bottom left corner, their direction as indicated by the letters. Denote the inverse of an edge e by $\overline{e}$. Apart from the identifications $a \sim a^{-1}$, $ac_1 = c_1a$ and $ac_2 = c_2a$ due to the fundamental polygon, we make the following additional identifications to conform with the triangulation of the fundamental polygon depicted in figure 3.7:

$$\langle ac_1z, zc_1a\rangle \sim c_1$$

$$\langle ac_2z, zc_2a\rangle \sim c_2$$

$$\langle ac_1z, zc_2a\rangle \sim a^{-1}$$

$$\langle \overline{zc_1a}, \overline{zac_2}\rangle \sim a$$

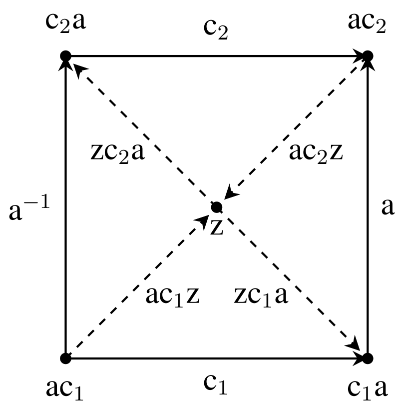

Figure 3.7: Triangulation of the fundamental polygon of the cylinder

Let p be any path from cuff c_1 to cuff c_2. By cutting the cylinder along the appropriate line we may assume without loss of generality that the endpoints of p are ac_1 and c_2a. From the above identifications it can be seen that any such path ultimately takes the form $\langle c_1^{y_1}, a, c_2^{y_2}\rangle$, $y_1, y_2 \in \mathbb{Z}$. Furthermore $\langle c_1, a\rangle \sim \langle a, c_2\rangle$ by applying the above identifications consecutively, and thus any such path takes a final form $\langle a, c_2^y\rangle$, $y \in \mathbb{Z}$. Hence any path is solely determined by this integer y, which we call the winding number of the path p. In the following we will concern ourselves with multiple paths across the cylinder, and so there may be some offset between the endpoints of a path. To take this offset into account, if $x = \langle 0, x_2\rangle$ and $y = \langle 1, y_2\rangle$ are the endpoints of p, we also subtract or add the difference $x_2 - y_2$ from the integer winding number above. This rational winding number will be denoted by $\theta(p)$.

Graph Minors VI now investigates the problem of ROUTED FOREST CONTAINMENT: For a graph G and a forest H embedded in a cylinder C such that the vertices of H lying on the cuffs of C are also vertices of G, decide whether G contains a forest that is homotopic to H.

For this, we first have to define when two forests H_1 and H_2 in C are homotopic.

Definition. For H_1 and H_2 as above, H_1 is homotopic to H_2 if

i) The vertices of H_1 and H_2 lying on the cuffs of C coincide, and

ii) any two such vertices s and t lie in the same component of H_1 if and only if they lie in the same component of H_2, and the unique paths between them in H_1 and H_2 are homotopic.

Note that this does not mean that every path in H_1 must have a corresponding homotopic path in H_2 and vice versa; we are only considering paths with endpoints on the cuffs.

Graph Minors VI then gives a number of reductions similar to those presented in the previous section, which essentially reduce the problem to whether for vertices $s_1, \ldots, s_n$ lying on c_1 and $t_1, \ldots, t_n$ on c_2 there are n disjoint paths of a given homotopy class connecting s_i with t_i, $1 \leq i \leq n$. Since the arguments for these reductions do not really differ when using our model of a cylinder, we only present the reductions and the necessary definitions and focus and replicating the proof of the reduced case mentioned above.

For u and v lying both on the same cuff c_i, we can define a segment S between u and v as in the previous section. We remark that there are two possibilities for choosing S, and that any path from u to v in C is homotopic to one of these choices. If the the endpoints u and v of a segment S are vertices of a graph G drawn in C and there is a path in G homotopic to S, then denote by $P(u, v, S)$ the unique path that bounds the regions of G incident with S. The following lemma gives criteria for infeasibility, similar to 3.2.6.iii.

Lemma 3.3.1 ((4.1) of Graph Minors VI)**.** Let C, G, H be as above and let S be a segment between $u, v \in V(H)$ so that $[u, v] \subseteq H$ is homotopic to S. If one of the following applies, H is not homotopic to a forest contained in G, and ROUTED FOREST CONTAINMENT can thus be answered negatively:

i) there is no path homotopic to S in G,

ii) $P(u, v, S)$ exists but contains some peripheral vertices of a component of H that does not contain u and v, or

iii) $P(u, v, S)$ exists and for some vertex $w \in P(u, v, S) \cap bd(C) \cap H$, the path from w to v in H is not homotopic to the path from w to v in $P(u, v, S)$.

We can now introduce the reductions of Graph Minors VI.

Lemma 3.3.2 (Reductions 1 and 2 of chapter 4 of Graph Minors VI)**.** The following operations transform a graph G and a forest H drawn on C into a graph G' and a forest H' so that G has a subgraph homotopic to H if and only if G' has a subgraph homotopic to H'.

1) Let H^* be a component of H and $V = V(H^*) \cap bd(C)$. Let G^* consist of the components of G containing the vertex set V. If G^* has a subgraph that is homotopic to H^*, then we can reduce G to $G' := G \setminus V$ and H to $H' := H \setminus H^*$.

2) Let S be a segment between u and v so that S is homotopic to a path in H. Then:

   i) If w lies inside a region of G incident with S we can reduce G to $G' := G \setminus \{w\}$.

   ii) Let $W = P(u, v, S) \cap (G \setminus bd(C))$. If $W \neq \emptyset$ then we can reduce G to G' which is obtained by contracting one of the edges of $P(u, v, S)$ that have w as one their endpoints, for every $w \in W$.

   iii) If case (i) can not be applied and $P(u, v, S) = \langle u, e, v \rangle$, then we can reduce G to G' obtained by contracting e so that u remains in place, and reduce H to H' obtained by likewise contracting e so that u remains in place, and if a circuit is introduced by this process, also deleting the edge originally incident with v.

There is one case left when neither lemma 3.3.1 nor lemma 3.3.2 apply.

Lemma 3.3.3 ((4.2) of Graph Minors VI)**.** For C a cylinder with cuffs c_1, c_2, G a graph drawn in C and H a forest drawn in C, if neither 3.3.1 nor 3.3.2 can be applied, then H consists of paths $p_1, \ldots, p_n$ joining vertices $s_i \in c_1$ and $t_i \in c_2$, $1 \leq i \leq n$.

Due to lemma 3.3.3 we need to investigate under which conditions there are k disjoint paths from c_1 to c_2 in G drawn in C, with the endpoints of the paths coinciding with the vertices of G lying on c_1 and c_2. We can assume without loss of generality that the endpoints of the paths are spaced out equidistantly on the cuffs on the cylinder. Thus, analogously to Graph Minors VI let

$$M_i = \left\{ \left\langle i, \frac{j}{k} \right\rangle : 1 \leq j \leq k \right\}, \qquad (i = 1, 2)$$

where the coordinates in $\mathbb{Q}^2$ are points on the cuffs in our fundamental polygon C. A set L of k pairwise disjoint paths in C from M_1 to M_2 is called a linkage. Since the paths in a linkage do not cross and their endpoints are spaced out evenly, they must have the same winding number, which we will refer to as $\theta(L)$, or just θ if the context is clear.

Since a linkage is determined by its winding number, we can focus on the winding number instead and ask whether a number $\theta = p + j/k$, $1 \leq j \leq k$, $p \in \mathbb{Z}$, is feasible in G, meaning whether there exists a linkage as a subgraph of G that has this winding number. From the geometric intuition, it should be the case that if $\{p_1, \ldots, p_k\}$ and $\{p'_1, \ldots, p'_k\}$ are linkages with winding numbers $0 \leq \theta = p + j/k < \theta'$, respectively, and the initial vertex of p_i and p'_i being $\langle 0, i/k \rangle$, then also $p + (j+1)/k$ should be feasible: Assuming $p + (j+1)/k < \theta'$, by following the path p_i until it is crossed by the path p'_{i-1}, then following p'_{i-1} until it crosses p_{i+1} and then following p_{i+1} to its endpoint (where all indices are modulo k, as they will be for the rest of this section), we create a linkage with winding number $p + (j+1)/k$. Extending this to non-negative winding numbers by symmetry, we see that the feasible winding numbers should constitute an interval, and it will be shown in the following that this is indeed correct.

To find the endpoints of such an interval we need to decide whether we can enlarge it, that is decide whether in a graph G the number $\theta + 1/k$ is feasible if θ is known to be feasible. To do this we first require some further reductions. Let $\{p_1, \ldots, p_k\}$ be a linkage as above, and denote by s_i the first vertex of p_i. Let C_i be the area in C in between p_i and p_{i+1} that contains no other p_j, and in G let R_i be the region which is incident with the segment between s_i and s_{i+1}.

Reduction 1 (Reductions 1-4 of chapter 5 of Graph Minors VI)**.** The following reductions will be used to simplify our problem.

1) If $e \in E(G)$ has s_i as one of its endpoints and $e \subseteq C_{i-1}$, pass to $G \setminus e$.

2) If $\langle G_1, G_2 \rangle$, $G_1, G_2 \neq G$ is a separation of G that is disjoint or intersects only in one vertex s_i, and all paths $p_1, \ldots, p_k$ are contained in one G_j, $j = 1, 2$, replace G by G_j.

3) If there is $i \leq k$ such that s_i is only adjacent with one vertex v and $v \cap c_2 = \emptyset$, then contract $s_i v$ so that s_i stays in place.

4) If there is p_i and $v \in V(p_i)$ so that v lies on the boundary of R_i but not all of $[s_i, v]_{p_i} \subseteq p_i$ does, then let t_i be the endpoint of p_i lying on c_2 and let r be the path bounding R_i. Then change p_i into $[s_i, v]_r \circ [v, t_i]_{p_i}$.

Note that 4 above is actually valid since the newly defined p_i does not intersect p_{i+1}. Also note that in reduction 1.4 we change the paths P_i and therefore all other objects associated with them, most importantly the C_i.

The goal is now to show that our current problem is equivalent on the graph G' obtained by employing any number of reductions from Reduction 1. Recall that the problem is not whether a number θ is feasible in G (which obviously holds for G' if it does for G), but to decide whether $\theta + 1/k$ is feasible if θ is.

Lemma 3.3.4 ((5.1) of Graph Minors VI)**.** Assume that θ is feasible in G but that it is not possible to apply any reduction of Reduction 1, then either there is $i \leq k$ so that s_i is only incident with t_i, or there is a linkage in G with winding number $\theta + 1/k$.

Proof. First note that if some s_i has only one neighbour v in G, then $v = t_i$, since otherwise Reduction 1.3 could be applied. Thus, assume that every s_i has more than one neighbour, we need to show that $\theta + 1/k$ is feasible in G.

As s_i has more than one neighbour, there is an edge e incident with s_i but not in p_i. Because Reduction 1.1 cannot be applied, e must be contained in C_i. Since Reduction 1.4 cannot be applied, e is incident with p_i only in s_i. Finally, as Reduction 1.2 cannot be applied, e must be the first edge in a path q_i from s_i to p_{i+1}. Hence, e must be part of the path bounding R_i, which thus intersects p_i only in s_i. Since further Reduction 1.1 cannot be applied to s_{i+1}, the path q_i must intersect p_{i+1} in some $v_i \neq s_{i+1}$. This holds for every s_i, since they all have more than one neighbour.

Let $p'_i = q_i \circ [v_i, t_{i+1}]_{p_{i+1}}$. Then p'_i has winding number $\theta + 1/k$, $i = 1, \ldots, k$, whence $\{p'_1, \ldots, p'_k\}$ is a linkage with winding number $\theta + 1/k$, and so $\theta + 1/k$ is feasible in G.

$\square$

Let f be any polygonal chain in C, i.e. not necessarily $f \subseteq G$. Then f is called G-normal if $f \cap \bigcup E(G) = \emptyset$. For a G-normal f let $L(f)$ be the number of times f crosses G, i.e. let $L(f) = |f \cap \bigcup V(G)|$. For $j = 1, 2$ let N_j be the set of points in between the set of end- or start-vertices of the paths p_j,

$$N_j = \left\{ \left\langle j, \frac{i}{k} + \frac{1}{2k} \right\rangle : 1 \leq i \leq k \right\}.$$

A G-normal polygonal chain f is called a helix in G if it has initial vertex in N_1 and terminal vertex in N_2.

Lemma 3.3.5 ((5.2) of Graph Minors VI)**.** A helix f approximates a feasible θ by the inequalities

$$\theta(f) - k^{-1}L(f) \leq \theta \quad \text{and} \quad \theta(f) + k^{-1}L(f) \geq \theta,$$

or equivalently

$$L(f) \geq k \cdot |\theta(f) - \theta| .$$

Proof. Let $\{p_1, \ldots, p_k\}$ be a linkage in G with winding number θ. We show $L(f) \geq k \cdot |\theta(f) - \theta|$ by induction on $k \cdot |\theta(f) - \theta|$. Note that our fundamental polygon, the (drawing of the) graph G and all linkages and helices are completely parametrized by rational numbers, and can thus be encoded as a natural number. Note also that any occurring θ is a rational number, and so the statement of the lemma is arithmetical. The induction is thus an arithmetical induction, which can be performed in ACA_0.

The claim is obvious for $|\theta(f) - \theta| = 0$, so assume $k \cdot |\theta(f) - \theta| = n > 0$. By symmetry (i.e. reflecting our fundamental polygon with regard to the line $\{\langle x, 1/2\rangle : x \in \mathbb{R}\}$), we may assume that $\theta(f) > \theta$. Let n_i be the initial and n'_j be the terminal vertex of f, so that $n_i \in C_i$ and $n'_j \in C_j$. Since $\theta(f) \neq \theta$, $f \not\subseteq C_i$, and so f must cross p_{i+1} (as $\theta(f) > \theta$) in a vertex v. Subdivide the edge $e \subseteq C_{i+1}$ incident with v, obtaining a new vertex $v' \in C_{i+1}$. By connecting v' to $n_{i+1} := n_i + \left\langle 0, \frac{1}{k}\right\rangle$ (where the addition is pointwise) we obtain a new polygonal chain f', and we may without loss of generality assume that $v'n_{i+1}$ does not intersect G. Thus f' is a helix in G, and $k \cdot |\theta(f') - \theta| = n-1$. Thus $L(f') \geq k \cdot |\theta(f') - \theta|$ by the induction hypothesis, and so $L(f) \geq (n-1) + 1 = n = k \cdot |\theta(f) - \theta|$.

□

For convenience, define $\phi(f) = \theta(f) + k^{-1}L(f)$.

Lemma 3.3.6 ((5.4) of Graph Minors VI)**.** Let G be a graph with linkage L and G' be a graph with linkage L'. Assume G' and L' are obtained from G and L by applying one of Reduction 1.1-4. Then there is a linkage in G with winding number $\theta(L) + 1/k$ if there is one with winding number $\theta(L') + 1/k$ in G'.

Proof. If G', L' were obtained by Reduction 1.1-4, let $\{p'_1, \ldots, p'_k\}$ be a linkage in G' with winding number $\theta + 1/k$. If G' was obtained by contracting $s_i v$ (i.e. in the case of Reduction 1.3), replace p'_i by $p''_i := \langle s_i, v\rangle \circ p'_i$ in G, where the first vertex of p'_i is changed to v. Then $\{p''_1, \ldots, p''_k\}$ is a linkage in G with winding number $\theta + 1/k$. All other cases are trivial, since then $\{p'_1, \ldots, p'_k\}$ is already a linkage in G. □

Again, it is becoming clear that many of the arguments can be phrased in the same way as in Graph Minors VI. The main difference lies in using polygonal chains instead of arbitrary continuous functions, and in working in the fundamental polygon instead of the actual surface. To conclude our analysis, we will thus only give the proofs of those lemmas directly relevant to the final result where other aspects of our model are of relevance.

Lemma 3.3.7 ((5.5) of Graph Minors VI)**.** If one of Reduction 1.1-4 is applied to graph G and linkage L and results in graph G' and linkage L', then for every helix f' in G' we can find a helix f in G so that $\phi_G(f) \leq \phi_{G'}(f')$.

Lemma 3.3.8 ((5.3) of Graph Minors VI)**.** Assume that there is a linkage in G with winding number θ, but none with winding number $\theta + 1/k$. Then $\phi(f) = \theta$ for some helix f.

Proof. The lemma is proved by induction on the number of vertices and edges of G. Note again that all our objects are codable by a natural number, and the induction is thus arithmetical.

Assume first that one of Reduction 1.1-3 can be applied and results in graph G' and linkage L'. By lemma 3.3.6, $\theta + 1/k$ is not feasible in G', and since G' has one vertex or edge fewer than G we can apply the induction hypothesis to G' to get a helix f' with $\phi_{G'}(f') = \theta$. Then by 3.3.7 there is a helix f in G with $\phi_G(f) \leq \phi_{G'}(f') = \theta$. But by lemma 3.3.5 $\theta \leq \phi_G(f)$, and hence $\phi_G(f) = \theta$.

Assume next that Reduction 1.4 can be applied, then we can use one of Reduction 1.1-3 on the resulting graph: Let v, r and i be as in Reduction 1.4. If the first edge of $[s_i, v]_{p_i}$ is not an edge of $[s_i, v]_r$ then it will not be part of p_i in L', and so it will lie in C_{i-1} and Reduction 1.1 can be applied. If the first edge of $[s_i, v]_{p_i}$ is an edge of $[s_i, v]_r$ and Reduction 1.1 can not be applied then there are two cases. First assume that s_i has more than one neighbour. Then since Reduction 1.1 can not be applied this additional neighbour

v' must lie completely in R_i, and in fact the whole subgraph of $G \setminus \{s_i v'\}$ that does not contain s_i must lie in R_i. Thus this subgraph cannot contain any edges or vertices of the paths p_j, and so Reduction 1.2 can be applied. Second, if s_i has only one neighbour, then this neighbour cannot lie in c_2 since Reduction 1.4 can be applied. Hence Reduction 1.3 can be applied.

So after an application of Reduction 1.4 one of the other reductions can be used on G' and L'. Using the induction hypothesis and 3.3.6, 3.3.7 and 3.3.5 twice as above, it follows that there is a helix f in G with $\phi(f) = \theta$

Finally, assume no reduction from Reduction 1 can be used on G. Employing lemma 3.3.4, we get that there is an s_i that has only one neighbour t which lies in c_2, and thus $t = t_j$ for some j and $p_i = \langle s_i, s_i t_j, t_j \rangle$. We want to define a polygonal chain f from $n_i \in N_1$ to $n'_{j-1} \in N_2$ that intersects G only in t_j. To do this, first find a point x so that $n_i x$ and $x t_j$ do not intersect G, where the edge $n_i x$ is drawn as a single straight line (possibly crossing the side a of the fundamental polygon) and $x t_j$ is drawn in close proximity to $s_i t_j$. Note that such a point x exists since Reduction 1.4 cannot be applied to G and since p_i consists only of one edge and there are thus no paths from p_{i+1} to p_i that do not have terminal vertex t_j. Then denote by $t_j n'_{j-1}$ the straight line from t_j to n'_{j-1} in C (again, possibly crossing a), and set $f = n_i x \circ x t_j \circ t_j n'_{j-1}$.

Then f is a helix in G which only intersects G in t_j, so $L(f) = 1$ and $\theta(f) = \theta - 1/k$ since f runs close to $s_i t_j$ from x to t_j, but has endpoints n_i and n'_{j-1}. So $\phi(f) = \theta - \frac{1}{k} + \frac{1}{k} = \theta$, and we are done.

□

From this lemma follow a series of corollaries that finally provide an algorithm for checking which θ are feasible in a graph G. Since their proofs are simply repeated applications of the previous lemmas and thus trivial to reproduce in our model, we state the corollaries and algorithms without proofs.

Corollary 3.3.9 ((5.6) of Graph Minors VI)**.** If one of Reduction 1.1-4 is applied to graph G with linkage L that has winding number θ and results in graph G' and linkage L', then there is a linkage with winding number $\theta + 1/k$ in G' if and only if there is such a linkage in G.

The following algorithm now determines whether $\theta + 1/k$ is feasible in a graph G, assuming that θ is.

Algorithm 1 ((5.7) of Graph Minors VI)**.**

1) Apply reductions 1.1-4 as often as possible. The resulting problem is equivalent by the above corollary. If it is not possible to apply any reduction anymore, go to step 2.

2) Determine whether some s_i has as only neighbour t_j. If this is the case, by 3.3.4 $\theta + 1/k$ is not feasible in G, so return 'no', otherwise return 'yes'.

We also get a first characterization of when a number $\theta + 1/k$ is feasible in a graph G.

Corollary 3.3.10 ((5.8) of Graph Minors VI)**.** Assume that there is a linkage L in G with winding number θ. Then the following are equivalent:

i) There exists some helix f so that $\phi(f)$ coincides with θ,

ii) there is no linkage L' in G with $\theta(L') = \theta + 1/k$,

iii) for all $\theta' > \theta$, there is no linkage L' in G with winding number θ'.

This implies the following.

Corollary 3.3.11 ((5.9) of Graph Minors VI)**.** Let $\theta_1 \leq \theta_3$ so that θ_1 and θ_3 are both feasible in G. Let z be an integer so that $k \cdot \theta_1 \leq z \leq k \cdot \theta_3$. Then $\theta_2 := z/k$ is feasible in G.

These corollaries lead to a characterization of feasibility, and this characterization can then be applied to provide an algorithm for deciding when some θ is feasible in G.

Lemma 3.3.12 ((5.10) of Graph Minors VI)**.** Let θ be a rational number, G be a graph drawn in C. Then there is a linkage L in G with $\theta(L) = \theta$ if and only if

i) $k \cdot \theta \in \mathbb{Z}$,

ii) there exists a linkage L' in G, not necessarily $\theta(L') = \theta$, and

iii) all helices f in G satisfiy $L(f) \geq k \cdot |\theta - \theta(f)|$.

We can now give the final algorithm of Graph Minors VI that determines whether any given θ is feasible in a graph G.

Algorithm 2 ((5.11) of Graph Minors VI)**.**

1) Determine whether 3.3.12.i holds, if not then output 'no'.

2) Determine whether 3.3.12.ii holds, which can be done effectively by producing a linkage using Menger's theorem. If not, output 'no'.

3) Take the linkage L obtained in step 2. If $\theta(L) > \theta$ flip C along $\{\langle x, \frac{1}{2}\rangle : x \in \mathbb{R}\}$. Then $\theta(L) \leq \theta$ holds; proceed to step 4.

4) If $\theta(L) = \theta$ return 'yes'. Otherwise check whether $\theta(L) + 1/k$ is feasible by applying algorithm 1. If not return 'no', otherwise substitute a linkage L' with winding number $\theta + 1/k$ for L, and return again to step 4.

The running time of the algorithms plays no part in the proof of the graph minor theorem or in the Graph Minors series in general, and hence this concludes our analysis of Graph Minors VI.

This chapter showed how surfaces and graph drawings in surfaces appearing in the Graph Minors series can be handled in second order arithmetic. It presented a model in which (compact) surfaces and operations on surfaces can be described by arithmetical statements and thus be easily carried out in ACA_0. The use of this model was then demonstrated on the example of Graph Minors VI. The next chapter will give a summary of the remaining papers of the Graph Minors series, where this method is implicitly applied when examining the proof methods used.

Chapter 4

Proof Methods of the Graph Minors Series

Chapters 2 and 3 gave a proof-theoretic treatment of Graph Minors IV and VI, respectively. This chapter will give a summary of the remaining relevant papers of the Graph Minors series, pointing out any proof methods that are notable from a metamathematical point of view and presenting more important results of the Graph Minors series in greater detail. For most of the papers, this means that we will argue that their proofs can be carried out in ACA_0, given the theorems of earlier papers. It will turn out that in addition to the principles of Graph Minors IV only one major proof principle will be needed to carry out the proofs of the Graph Minors series, namely Π^1_2-bar induction. This will establish that the graph minor theorem and the immersion theorem can be proved in the theory $\Pi^1_1 - \mathrm{CA}_0 + \Pi^1_2 - \mathrm{BI} + \Pi^1_3 - \mathrm{IND}$. Good overall summaries and surveys of the steps in the proof of the graph minor theorem and its most important intermediate results can be found in e.g. Diestel (2017) and Kawarabayashi and Mohar (2007). It should also be noted that each of the papers in the Graph Minors series give a good summary of their objectives and the overall structure of their proofs in their introductions.

4.1 Graph Minors I

Graph Minors I gives a result similar to the excluded minor theorem for planar graphs, namely that if $G_1 \not\leq G_2$, where G_1 is a forest, then G_2 has bounded path-width with the bound depending only on G_1. Here path-width is defined similar to tree-width, where only paths instead of general trees are allowed in a decomposition. All objects occurring in Graph Minors I are finite and combinatorial, and consequently only arithmetical inductions are used. Thus, it is not surprising that Graph Minors I can be carried out in ACA_0. In particular, the following inductions are used in Graph Minors I.

- The proof of (2.3) uses induction on a natural number, and since (2.3) is an arithmetical statement, the induction is arithmetical.
- Likewise, Step 2 of the proof of (2.5) uses an arithmetical induction on the length of a sequence.
- An arithmetical induction is used in the proof of (2.6) to show that certain minima exist.
- The proof of (2.7) again uses arithmetical induction via a minimal counterexample.
- The same is the case for the proof of (2.8).
- Part 1 of the proof of (3.1) (after the statement of (3.7)) also uses arithmetical induction.
- Part 3 of the same proof uses again an arithmetical induction.

4.2 Graph Minors III

While Graph Minors III is not essential for the proof of the graph minor theorem and its main theorem is a weaker version of that of Graph Minors V, some of its results are used in Graph Minors VIII for an important restricted version of the graph minor theorem. The main theorem of Graph Minors III is the second of the type of an excluded minor theorem in the Graph Minors series, and states that if H and G are both planar graphs so that $H \not\leq G$, then $tw(G) \leq k_H$, where $t(G)$ denotes the tree-width of G and the number k_H depends only on H and is thus the same for all such G.

The only objects appearing in Graph Minors III are finite graphs embeddable in the plane, such as cylinders (a certain kind of graph, not the surface), which are similar to the enclosures considered in chapter 2, constructions involving them, such as sleeve unions, and subgraphs thereof. Thus again, all of Graph Minors III can be carried out in ACA_0, as only arithmetical inductions are used:

- The induction used to prove (2.2) is arithmetical.
- The same is the case for (4.4).
- Finally, (5.2) features another arithmetical induction.

4.3 Graph Minors V

The main theorem of Graph Minors V is similar to that of Graph Minors III. It states that with all graphs H that are planar we can associate a natural number k_H so that for all graphs G (where G does not have to be planar), if $tw(G) > k_H$ then already $H \leq G$. This adds another form of excluded minor theorem similar to that of Graph Minors III where only one graph of the sequence needs to be planar, since together with the bounded graph minor theorem this result implies that any sequence of graphs that contains a planar graph is good. Since, as described below, the proofs of Graph Minors V do not contain any advanced proof-theoretic methods, the bounded graph minor theorem in particular implies the planar graph minor theorem.

The objects used in Graph Minors V are finite graphs and finite sets of them, such as tree-decompositions and various notions concerning subgraphs of a given graph, such as paths and separations. Two other such central objects of Graph Minors V are so-called webs and spiders. A web can be viewed as a grid-like subgraph of a graph, in that it is a set of paths that is arranged somewhat like a grid. A spider is a subgraph that is disjoint from the edges of either the "vertical" or "horizontal" paths of a web, and intersects each of the other type of paths in exactly one vertex that is only connected by one edge to the rest of that subgraph. These objects can be easily encoded by natural numbers.

The second part of Graph Minors V gives a more general version of a theorem originally proved by Erdős and Pósa (1965). The main theorem of this part is that the graphs H possessing the Erdős-Pósa property are exactly the planar graphs, where H possessing the Erdős-Pósa property means that any graph G either contains a certain number of disjoint minor expansions of H or after removal of a certain other number of vertices there can be no expansion of H in G, where these two numbers depend on each other. The proofs of this part also use orientable surfaces.

We give a brief overview of the structure of the proof of the main theorem of Graph Minors V.

- Firstly, it suffices to prove the result when H is a grid, since every planar graph is a minor of some grid.
- It is then proved that if G does not have H as a minor, then there is no (θ_2, θ_2)-web in G ($\theta_1, \ldots, \theta_9$ are natural numbers depending on H).
- This result is then refined by showing that if $H \not\leq G$ then there is no (θ_5, θ_6)-mesh in G, where a mesh is a generalization of a web.
- This is then combined with the existence of a certain separation in G to obtain the conclusion that G has tree-width not higher than θ_9.

The following proof principles are used in Graph Minors V.

- The proof of (3.2) employs an arithmetical induction.
- This is also the case in (5.1).
- The proof of (6.2) uses an induction on the number of vertices of a graph. The statement of 6.2 is arithmetical since it only involves finite graphs and the function w is finite as well. This is hence a simple arithmetical induction.
- Since the set $\mathcal{B}$ in the proof of (6.3) and the set of all such $\mathcal{B}$ are finite (since $\mathcal{B}$ is a set of subgraphs of G of fixed size), the lexicographically least such $\mathcal{B}$ can be obtained using arithmetical induction.
- The proof of (7.3) again employs an induction, which is arithmetical since $G \in \mathcal{F}_\theta$ is the statement that G does not have the θ-grid as a minor.
- (8.5) is again proved using arithmetical induction.
- The proof of (8.12) uses induction on finite sequences of natural numbers, ordered lexicographically (via assuming that there is a certain surface for which the genus is least with regard to that ordering). This would correspond to a transfinite induction of length ω^ω. Since the induction is performed over an arithmetic formula, it is a transfinite arithmetical induction, which is available in ACA_0.

Thus, any proof in Graph Minors V can be carried out in ACA_0.

4.4 Graph Minors VII

Graph Minors VII generalizes the results of Graph Minors VI. It investigates the question of when there are k disjoint paths between given vertices in a graph G that is embedded in an arbitrary surface Σ. For this it uses some techniques similar to the ones in Graph Minors VI, in particular generalized helices. Since G is now drawn in an arbitrary surface, the additional condition that G have large parts of its drawing across every handle and crosscap is also needed, so that the drawing of G makes proper use of the space the surface provides. The first preliminary version of this generalized result is (4.6), which is a statement about the feasibility of forests in G, and is then refined in (5.9) and further in (6.1), where it is shown that conditions on the number of intersections of G with $bd(\Sigma)$ are not necessary. The final version with additional conditions on the cuffs removed is then (7.5).

More interesting for a restricted version of the graph minor theorem in Graph Minors VIII are a series of theorems, (9.1) – (9.5), which are applications of the main theorem, and roughly say that for every graph H embedded in Σ, there is a number k depending on H so that if G is a graph drawn in Σ that has more than k of its parts drawn across every handle or crosscap as above, then H is a minor of G. Finally, an algorithm to check feasibility is given in a similar manner to those of Graph Minors VI, which basically consists of applying the main theorem (7.5) or the algorithms of Graph Minors VI if possible or reducing the problem to smaller subproblems.

Note that graphs are technically defined as topological spaces in this paper. This is not necessary however and only done so that Graph Minors VII does not need to distinguish between graphs and drawings of graphs, and essentially a graph in Graph Minors VII is defined to be a drawing of a graph. The methods outlined in our treatment of Graph Minors VI can thus be applied without problems. Further, note that chapter 11, which deals with various topological results, is not needed for our purposes since our objects are defined to be codes in second order arithmetic and not actual topological spaces. These results are thus built-in into our actual definition of surfaces, paths, homeomorphisms etc.. In particular, (11.8) which is proved using Zorn's Lemma is trivial in our model, since "flat and arc-connected" and "simply-connected" are simply the same thing in this model.

The following proof methods are used in Graph Minors VII.

- The proof of (3.3) uses induction to show that there are only a finite number of certain drawings of graphs up to homeomorphism. Since a homeomorphism is a finite function in our model, the induction used is arithmetical.
- (4.4) also uses arithmetical induction.
- Likewise, the proof of (6.1) also uses arithmetical induction.
- (7.3) is also proved using arithmetical induction.
- Further arithmetical inductions are used to show (on pages 247 and 248 of Graph Minors VII) that the algorithm given on page 247 is correct.

4.5 Graph Minors VIII

Graph Minors VIII is not essential for the proof of the graph minor theorem, but it proves an interesting restricted form of the graph minor theorem. The main theorem of Graph Minors VIII is that for a fixed surface Σ, if we have a sequence $\langle G_1, G_2, \ldots\rangle$ of graphs that can all be embedded into Σ, then $G_i \leq G_j$ for some $i < j$, $i, j \in \mathbb{N}$. In other words, the graph minor theorem holds for graphs embeddable into a fixed surface Σ. It is even proved that this holds if the vertices that are drawn on the cuffs of Σ stay fixed after each contraction; for this it is needed that all the graphs in the sequence have the same number of vertices drawn on the boundary of Σ. Denote this version of the graph minor theorem by $\Sigma - GMT$, and denote the statement that this holds for every surface by $\forall\Sigma - GMT$.

As it turns out, $\Sigma - GMT$ and $\forall\Sigma - GMT$ are both equivalent to the planar graph minor theorem and hence also to the bounded graph minor theorem over ACA_0. This can be shown by recreating the proof the main theorem of Graph Minors VIII in $\Pi^1_1 - \mathrm{CA}_0$

augmented with Π^1_1-reflection for $\Pi^1_1 - \mathrm{CA}_0$. This is for the most part straightforward, however one result of Graph Minors VIII references (9.2) from Graph Minors IV. Here it is important to check that only instances of (9.2) are used in the proof, as otherwise a Π^1_3-induction would be required. Further, a Π^1_2-induction is used in the proof of (13.3), and hence it also needs to be checked that only instances of (13.3) suffice to prove each instance of the cylinder case (17.3) and the final theorem (18.3). To do this, we need to trace the uses of these theorems throughout Graph Minors VIII.

The proof of the graph minor theorem for general surfaces is first given for the case where Σ is a disk, which interestingly is the most difficult case. Up to (10.1), which uses a result from Graph Minors IV, the proofs can be reformulated in ACA_0. As noted above, (10.1) references (9.2) from Graph Minors IV, and thus can only be proved for each tree-width w individually in $\Pi^1_1 - \mathrm{CA}_0$. Since (10.1) is further a Π^1_2-statement, it can also only be proved for every instance of w in $\Pi^1_1 - \mathrm{CA}_0 + \mathrm{RFN}_{\Pi^1_1}(\Pi^1_1 - \mathrm{CA}_0)$. The only place where (10.1) is applied is in the proof of an instance of (6.1), the disc case, on page 273, section 12 of Graph Minors IV. There, an induction is performed on the index (essentially the number of vertices on the cuff) of a graph, invoking (10.1) for a certain, fixed tree-width w in the induction step. Let D denote the disk. Then for each fixed index D-GMT is thus provable in $\Pi^1_1 - \mathrm{CA}_0$, and since D-GMT is a Π^1_1-statement it then holds in its full strength in $\Pi^1_1 - \mathrm{CA}_0 + \mathrm{RFN}_{\Pi^1_1}(\Pi^1_1 - \mathrm{CA}_0)$.

Next, the proof is given for the case where Σ is a cylinder in (17.3), the other main problematic case, and concluded for the general case in (18.3). One of the theorems leading up to (17.3), namely (13.3), uses a Π^1_2-induction. Since (17.3) is a Π^1_1-statement, we hence need to again use Π^1_1-reflection for $\Pi^1_1 - \mathrm{CA}_0$ to prove it in its full strength. Since (13.3) is also used in the final theorem (18.3), which is again a Π^1_1-statement, it also needs to be shown that (18.3) can be proved for every index k using only instances of (13.3).

Lemma (13.3) is applied in (14.4), (15.4), (16.1) and (18.3). Further, the disk case —

which is only provable for each fixed index in $\Pi^1_1 - \mathrm{CA}_0$ — is also used in the proof of (15.4). Consequences of all these theorems can thus only be proved for a fixed index in $\Pi^1_1 - \mathrm{CA}_0$, and overall it needs to be shown that only instances of the following theorems are required to prove instances of the desired theorems (17.3) and (18.3):

- As noted above, (14.4), (15.4) and (16.1) cannot be used in their full form.
- (14.4) is used in (14.5), which is in turn used (only) in (15.4).
- (15.4) is used in (16.1), which is used in (17.3), the cylinder case.
- Finally, (17.3) is used in the proof of (18.3) as the base case.

We will now show that only instances of these theorems are needed in the proofs of all instances of (17.3) and (18.3) in $\Pi^1_1 - \mathrm{CA}_0$. It has been shown already that the disk case is provable in $\Pi^1_1 - \mathrm{CA}_0$ for each fixed index. To prove an instance of (15.4), a fixed instance of the disk case for index $k + r(r-1) + 2s$ is first applied (k, r, s arbitrary, fixed natural numbers), which is combined with an application of (13.3) for the same index to show that a certain set $\mathcal{B}'$ that has index at most $p := k + r(r-1)$ is well-rooted. Then (14.5) is applied for this index p, which requires an application of (14.4) for the same index p, which in turn requires and application of (13.3) for indices $3p$ and $2p$. Thus, each instance of (15.4) is provable in $\Pi^1_1 - \mathrm{CA}_0$. To show the same for (16.1), we note that to prove (16.1) for a fixed index k one must apply (13.3) twice, once for index $3k-1$ and once for index $2k-2$. Then, to prove the cylinder case (17.3) for index k, (16.1) needs to be applied for index $k-1$, whence the cylinder case is provable with Π^1_1-reflection for $\Pi^1_1 - \mathrm{CA}_0$. Finally, in the proof of (18.3), first note that the two performed inductions do not depend on each other. We can thus fix an index k, a surface Σ and a number N (which depends on the surface and the first term of a sequence) and prove (18.3) for each such index k, surface Σ and number N by a metainduction on these numbers. In this proof, first (17.3) is applied in the case that Σ is the sphere, disk or cylinder. In the induction step, (13.3) is applied twice, once for index $k+2N+2$ and once for index $3k-2$. Thus, (18.3) can be proved for all such k, Σ and N in $\Pi^1_1-\mathrm{CA}_0$. Applying Π^1_1-reflection for $\Pi^1_1-\mathrm{CA}_0$, (18.3) can be seen to be true.

To summarize, the following proof-theoretic principles are used.

- (10.1) uses (9.2) from Graph Minors IV, and in $\Pi^1_1-\mathrm{CA}_0$ is hence only provable for every fixed tree-width.

- The conclusion of the proof of (6.1) for the disc case (on page 273, chapter 12) can then be proved using Π^1_1-reflection for $\Pi^1_1-\mathrm{CA}_0$.

- The proof of (13.3) uses a Π^1_2-induction. Only fixed instances of (13.3) are thus provable in $\Pi^1_1-\mathrm{CA}_0$. As pointed out above, using Π^1_1-reflection for $\Pi^1_1-\mathrm{CA}_0$ one can still prove the important theorems of Graph Minors VIII.

- Claim (3) of (14.2) employs arithmetical induction, since it does not require the assumption that $\mathcal{A}$ be well-rooted, as it is essentially a general statement about simulation of grafts in path-decomposition and uses only related techniques.

- Another arithmetical induction is used in the proof of (15.2).

- The cylinder case, (17.3), can be proved in $\Pi^1_1-\mathrm{CA}_0+\mathrm{RFN}_{\Pi^1_1}(\Pi^1_1-\mathrm{CA}_0)$ as detailed above.

- The final result (18.3) uses a nested induction, each of the two on a natural number over a Π^1_1-formula. These ultimately rely on (13.3) and (17.3) as the base case and can thus not be performed in $\Pi^1_1-\mathrm{CA}_0$, but can however be circumvented through Π^1_1-reflection for $\Pi^1_1-\mathrm{CA}_0$.

4.6 Graph Minors IX

Graph Minors IX introduces some lemmas about vortices that are used later in the Graph Minors series. Vortices are local non-planar subsets of a graph, and will be a major component of the structure that can be attributed to a graph that does not contain a minor of another graph, the so-called excluded minor theorem.

As laid out in the abstract of Graph Minors IX, informally the main theorem states the following, which is made precise in (6.1). For any cyclic order on a subset V' of $V(G)$, where G is sufficiently highly connected, one of two things must hold:

i) Either G can be embedded inside a disk with the vertices of V' on the boundary in the same clockwise order as the cyclic order on V' except for crossings in a locally constrained region, or

ii) V' can be partitioned into two segments on the boundary of a disk with the rest of G drawn inside, and then there is a high number of vertex-disjoint (but possibly edge-crossing) paths from one segment to the other so that every such path lies completely between only two other paths or is crossed by another path.

More precisely, the main theorem states that there can be no large crooked transaction in a rural neighbourhood that has sufficiently high sophistication. The converse of this theorem turns out to hold as well, but only up to 3-separations. The main objects used in the proofs of Graph Minors IX are thus transactions, which are sets of vertex-disjoint paths in a graph that is drawn (again, possibly with edge-intersections) in a disk with the vertices on the boundary in a fixed order (a so-called society), rural neighbourhoods, which are graphs drawn in a cylinder (without crossings) with their vertices on the boundaries in the respectively correct orders, and the graphs used to define sophistication, which are obtained by glueing together a rural neighbourhood and a society without a large transaction. The main objects of Graph Minors IX are thus finite graphs drawn in a plane, disk or cylinder, which pose no problems for our model.

The second half of Graph Minors IX makes some results more precise or brings them into a more convenient form. For instance, it is shown that in any large crooked transaction it is possible to find another large crooked transaction that has one out of three forms: a crosscap, a leap or a doublecross. Another noteworthy result is that if a society S has no large transaction, then it is possible to find a linear decomposition of S that has small depth.

The following proof theoretically relevant techniques are used in Graph Minors IX:

- Arithmetical induction is used in (2.4) to show that a certain society is rural.
- The main theorem (6.1) also uses arithmetical induction in the form of a least counterexample.
- (10.3) uses another arithmetical induction to prove that a crooked transaction can be found in a subset of a rural neighbourhood, if the subset is large enough with regard to the breadth (or radius in a sense) of that neighbourhood.
- A minimal consolidation needed for the proof of (11.11) can be found using arithmetical comprehension.

4.7 Graph Minors X

Graph Minors X introduces the very important concept of a tangle. Roughly, a tangle is a classification of low-order separations in a graph (or hypergraph) into a big and a small side. It is a dual concept to that of a tree-decomposition because a graph has large tree-width if and only if it contains a tangle of large order, as proved in (5.2) of Graph Minors X, and it can be shown that any graph has a tree-decomposition whose parts correspond to its tangles. Since by Graph Minors V large tree-width is correlated with the existence of a large grid minor, the same can be deduced with respect to the existence of a tangle of large order. Tangles can also be used to control tree-decompositions in a certain sense which is the subject of section 11 of Graph Minors X. Further, it can be shown that if $H \leq G$ and $\langle G_1, G_2 \rangle$ is a low-order separation occurring in a tangle, then exactly one of the two sides (the "big" one) contains an expansion of a vertex of H.

Graph Minors X thus deals with finite graphs and the finite sets of separations in such graphs. These and all other concepts used in Graph Minors X can be obtained using arithmetical comprehension. The following proof principles of note are used.

- Arithmetical inductions and comprehension are used in (3.3), (3.4), (3.6), (4.1), (4.2) and (4.5).
- Another arithmetical induction is used in the proof of (7.2), a lemma for the proof that there is a tangle of order θ in the θ-grid.
- The proof of the existence of a tiebreaker in (9.2) requires
$$\binom{|L(G)|+1}{|L(G)|-1} = \binom{|L(G)|+1}{2} = \frac{(|L(G)|+1)(|L(G)|)}{2}$$
rationally independent real numbers (the number of 2-combinations with repetitions from $|L(G)|$ elements), where $L(G) = V(G) \cup E(G)$. Since the square roots of distinct prime numbers are rationally independent, and they exist even in RCA_0 (for example, the Babylonian method provides a Cauchy sequence), it is possible to use those. Since further $(\Lambda, <)$ from (9.2) effectively consists of only finitely many elements (i.e. the range of λ is finite), after having found a tiebreaker we can discard the real numbers occurring in it and replace $(\Lambda, <)$ with any ordering of the same order type involving only rational (or integer) numbers.
- In (6) of the proof of (10.3) another arithmetical induction on the order of a distinction is used.

- One more arithmetical induction is used in (11.3) to prove that a tree-decomposition in a graph can be found that corresponds to a certain design under the right circumstances.

4.8 Graph Minors XI

Graph Minors XI shows any graph G drawing on a surface that represents the surface sufficiently well induces a tangle in G, and similarly that such a graph also induces a metric on the surface, which will play a role in later papers of the Graph Minors series. The distance function needed to define the metric is introduced at the very end of Graph Minors XI in section 9, roughly speaking in terms of the number of regions one has to traverse to get from one point to the other.

Another important result of Graph Minors XI is that the interior of any curve which does not meet the graph often enough (in the above sense) is homeomorphic to a disk, which combined with the above distance function gives the result that every point has a planar neighbourhood, the size of which corresponds to the representativeness of the graph G.

The central tools for achieving these results are pretangles and slopes, where a slope defines the inside of every region in G. It is proved that pretangles and slopes are in 1-1 correspondence, by way of showing that a slope can be derived from a pretangle and the pretangle this slope induces is then again the original slope, which is made precise in (6.3). It should be noted that the purely topological results of sections 7 and 8 are not needed for the purpose of recreating the proof in second order arithmetic.

The proof of Graph Minors XI can thus be formalized in ACA_0, and the following notable proof-theoretic techniques are employed.

- Arithmetical inductions are used to prove (4.4), (5.4), (5.5), (6.1), (6.2) and (6.4).

- Technically, (7.1) and (8.8) use further inductions, but as noted above the results are largely obvious with our treatment of surfaces.

4.9 Graph Minors XII

Graph Minors XII expands on the metric defined in the previous Graph Minors paper. It explores what impact small changes in the drawing of the graph on a surface or even the surface itself have on this metric. The metric is used to reformulate in (3.2) the main theorem of Graph Minors VII, so that it roughly states that there are disjoint connected subgraphs connecting vertices on certain disks in a graph if

i) such subgraphs are topologically feasible,

ii) the disks are sufficiently far apart with regard to the metric, and

iii) the vertices lying on the boundaries of a disk are free, for every disk,

where a subset of a graph is free if it is small with respect to the tangle order, but not contained in any small side of a separation that has even smaller order than the size of the subset.

Graph Minors XII also introduces the notion of when a tangle controls a minor L of a graph G. Roughly, this is when the tangle number is bigger than the size of L and no inflation of any vertex is contained in a small part of a separation in the tangle that has order smaller than the size of L, i.e. the minor lies mostly in the large parts of the tangle. An important application of the above reformulation is that for any graph L embedded on a surface Σ it is possible to find a natural number such that any graph G which can be drawn 2-cell on Σ such that it contains a tangle of order higher than that number controls an L minor with that tangle, which is shown in (4.3).

Section 5 of Graph Minors XII then investigates when a tangle can be extended to surfaces and graphs of higher genus while keeping some of its properties. A tangle is called respectful if for any closed curve that does not meet the graph too many times, the curve bounds a disk and the subgraph contained in that disk is the small side in a separation of the tangle. Assume that two surfaces Σ_1 and Σ_2 with graphs G_1, G_2 drawn on them, respectively, are glued together along a number disks (the most interesting cases are when Σ_2 is a cylinder or Möbius band and there are two or one disks, i.e. when a handle or crosscap is added), such that the graphs naturally connect at these disks. Theorem (5.3) of Graph Minors XII then says that a respectful tangle in Σ_1 can be extended to a respectful tangle in the resulting surface if

i) the disks are far enough apart,

ii) any subset of vertices lying on one of the disks is free with respect to the tangle, and

iii) G_2 represents Σ_2 well enough with respect to the order of the tangle.

The remainder of Graph Minors XII then analyzes what happens with a metric or respectful tangle if a planar neighbourhood of the drawing is altered, i.e. a certain drawing of a graph in a disk is added, or a subgraph contained in a small disk deleted.

To summarize, Graph Minors XII deepens the results about tangles from the earlier papers of the Graph Minors series and mainly uses the same concepts, and the inductions occuring in Graph Minors XII are arithmetical inductions in (2.2), (7.4), (7.5) and (7.6). Its proofs can thus be reproduced in ACA_0.

4.10 Graph Minors XIV

Graph Minors XIV continues where Graph Minors XII left off, namely with the extending of embeddings in surfaces. It studies which obstructions have to occur in a graph G so that an already existing embedding of a subgraph H in some surface cannot be extended to one of G. If H represents the surface well enough (as defined in earlier papers of the Graph Minors series) then Graph Minors XII established that every one of its vertices must be surrounded by a large planar region, so locally this problem reduces to drawings in disks. To obtain results that are more easily applicable to later papers in the Graph Minors series, only "rigid" H, that have a somewhat unique drawing in such disks, are considered. Further, some of the utilized results about societies from Graph Minors IX only hold up to 3-separations, and thus it is also necessary to consider extensions to embeddings of G only up to 3-separations. It then turns out that there are only two obstructions to extending an embedding, which is made precise in (10.1) of Graph Minors XIV.

The proofs of Graph Minors XIV can be recreated in ACA_0 like the earlier papers of the Graph Minors series leading up to the excluded minor theorem, and the following types of induction are used.

- The proof of (5.4) uses arithmetical induction where (5.3) is applied in the induction step, and similarly for (5.5).
- Likewise, (7.3) and (8.4) use arithmetical induction.
- Another arithmetical induction is used in (10.1), the main theorem of Graph Minors XIV.

4.11 Graph Minors XV

Graph Minors XV lays the foundation for the excluded minor theorem, to be proved in Graph Minors XVI. It investigates which features a graph G must have in order to not have a drawing which is an extension of a graph H embedded in a surface Σ. The question was answered in Graph Minors XIV for the case where G was not allowed to have any vortices, in Graph Minors XV it is answered for any bounded number of vortices. Graph

Minors XV shows that if G can not be extended then there are three possible obstructions, made precise in (1.1) of Graph Minors XV.

The excluded minor theorem is used in the proof of the graph minor theorem roughly in the following way. Given a sequence $\langle G_i : i \in \mathbb{N} \rangle$, it can be assumed that G_0 is not a minor of any other G_i, otherwise the graph minor theorem holds. But then the excluded minor theorem can be applied to all G_i, $i \geq 1$, to infer that these graphs have a certain structure, from which it can then be concluded that some G_i must be a minor of some G_j, for $1 \leq i < j$. As outlined in Graph Minors XV, the excluded graph minor theorem gives stronger results if G_1 is a complete graph. But G_1 can be replaced by the complete graph on $|V(G_1)|$ vertices in the above argument, since no G_i for $i \geq 1$ can have that complete graph as a minor either, and this stronger form may be used.

The introduction of Graph Minors XV gives the following ideas of how its main theorem will be used. By the above, assume that G has no K_p minor. Assume that $H \subseteq G$ is drawn in Σ so that it is representative of that surface and no subgraph of G is representative of a surface of higher genus. Then one of four cases holds:

i) A small region of H can be replaced by a subgraph of G, where a crosscap is also added in this small region and the resulting graph is represents the new surface well; this however is inconsistent with the maximality of the genus of Σ, or

ii) H can be altered in such a way that it is drawn in Σ up to a fixed number of edge-crossings which however lie far apart in the surface; then it can be shown that K_n must be a minor of G, also a contradiction, or

iii) allowing a number of vortices, G can be embedded in Σ in the intended way, as in Graph Minors XIV up to 3-separations, or

iv) after drawing a small part of H with slight alterations, G can be drawn in Σ in such a way that it has a path between two vertices of H that lie far apart in the surface, but which apart from that does not intersect H.

Thus, in the situation outlined above, only (iv) needs to be dealt with. The main theorem is made precise in (1.1) for any graph G and proved in a series of very technical lemmas, then extended in (8.2) and (8.4) to the special form for G having no K_p minor. It should be noted that (i)-(iii) above correspond to (i)-(iii) of (1.1), and (iv) above corresponds to the condition of (1.1) that a Σ-span be (λ, μ)-flat.

As for previous papers leading up to the excluded minor theorem, the proofs of Graph Minors XV can be recreated in ACA_0, and the following types of induction are used in Graph Minors XVI.

- (2.1) uses arithmetical induction to prove that a certain segregation is central with regard to a certain tangle.
- The subproof of (4) in the proof of (8.3) uses arithmetical induction through a least counterexample.

4.12 Graph Minors XVI

Graph Minors XVI proves the excluded minor theorem, arguably the most important result of the Graph Minors series after the graph minor theorem itself. It generalizes the result from Graph Minors V that said that if $H \not\leq G$ for H planar, then G has a tree-decomposition into parts with $< k_H$ vertices, where k_H is the same number for all such G and only depends on H. The excluded minor theorem now says that if H is any graph (not necessarily planar) and $H \not\leq G$, then there is a tree-decomposition of G into parts which can be "nearly" drawn in a surface Σ in which H cannot be drawn.

This idea can be formalized as follows. A graph X being (d, r, w)-nearly embeddable into Σ means that:

- X can have up to d r-rings. That is, if $t_1, \ldots, t_n$ are the vertices around a region of X' (where X' is embedded into Σ), and $X_1, \ldots, X_n$ is a path decomposition (or ring decomposition) of width $\leq r$ with $t_i \in V(X_i)$ where $\bigcup_{1 \leq i \leq n} X_i$ intersects X' only in $\{t_1, \ldots, t_n\}$, then $X' \cup \bigcup_{1 \leq i \leq n} X_i$ is said to have one r-ring.

- There can be up to w apex vertices $v_1, \ldots, v_w \in V(X)$, meaning that $X \setminus \{v_1, \ldots, v_w\}$ is embeddable into Σ with rings as above allowed.

The excluded minor theorem, (1.3) of Graph Minors XVI, then says:

> For every graph H there are natural numbers d, r and w such that if $H \not\leq G$, then there is a tree-decomposition of G into parts which are (d, r, w)-nearly embeddable into a surface into which H is not embeddable.

Note that up to homeomorphism, there are only finitely many surfaces into which any fixed graph H is not embeddable, so the excluded minor theorem provides an effective structural characterization of graphs of which H is not a minor.

In Graph Minors XIX, roughly it will be proved that any sequence of parts of such a tree-decomposition is good. Hence it is convenient to focus on one part of the tree-decomposition, and it is possible to do this by investigating the tangle centered on that part, assuming that the parts of the tree-decomposition are as small as possible, since almost the whole part will be located on the big side of the tangle. Then, roughly speaking, it is also possible to replace the rings from above by vortices of bounded depth.

The excluded minor theorem is then rephrased in (3.1) into the following:

> For every graph H, there are d, r and w such that for any graph G and any tangle in G that does not control an H-minor in G, after removing w vertices from G there is a drawing of G with at most d vortices of depth $\leq r$ in a surface in which H cannot be embedded.

The proof of (3.1) then proceeds by reducing it to simpler forms. A Σ-span in a graph G is roughly a rigid subgraph onto which a tangle restricts respectfully. Then the first reduction results in (4.1), which roughly says:

For every surface Σ, $\phi \geq 1$ and $p \in \mathbb{N}$ there are d, r, w and $\theta \geq 1$ so that for every tangle in G and Σ-span with order of at least θ one of the following possibilities holds.

- For one of the surfaces Σ' constructed by attaching a handle or crosscap to Σ it is possible to find a Σ'-span that has order at least ϕ, or
- G can almost be embedded into Σ (with conditions as in (3.1)), or
- G contains the complete graph on p vertices as a minor.

Note that the second case is the desired one. In the application to the proof of the graph minor theorem, case 3 will not be possible because G will be assumed to have no K_p minor, and case 1 will not apply because Σ will be assumed to be of maximal genus.

The next step in the proof of the excluded minor theorem is to introduce animals with horns and hairs. An animal is more or less a Σ-span as above, horns are paths outside this animal with a common endpoint, and hairs correspond to paths outside the animal with endpoints far apart. The strength of an animal is related to the order of the tangle associated with the span, the length of its hairs and the length and breadth of its horns. It is then possible to provide a further reduction, (5.1), which roughly postulates that in addition to the possibilities in the previous theorem it is also possible that the number of horns of the Σ-span increases (provided the animal has no hairs).

For every surface Σ, $\phi, \psi \geq 1$ and $p, \tau, \chi \in \mathbb{N}$ there are σ, d, r, w and $\theta \geq 1$ so that for every tangle in G and every animal with no hairs, χ horns and strength at least (θ, σ) one of the following possibilities holds.

- For one of the surfaces Σ' constructed by attaching a handle or crosscap to Σ it is possible to find a Σ'-span that has order at least ϕ, or
- G contains an animal with no hairs, $\chi + 1$ horns and strength at least (ψ, τ), or

- G can almost be embedded into Σ (with conditions as in (3.1)), or
- G contains the complete graph on p vertices as a minor.

This result can then be used to increase the number of hairs instead of horns in the second case, which is done in (6.1):

> For every surface Σ, $\phi, \psi \geq 1$ and $p, \tau, \chi \in \mathbb{N}$ there are σ, d, r, w and $\theta \geq 1$ so that for every tangle in G and every animal with δ hairs, χ horns and strength at least (θ, σ) one of the following possibilities holds.
>
> - For one of the surfaces Σ' constructed by attaching a handle or crosscap to Σ it is possible to find a Σ'-span that has order at least ϕ, or
> - G contains an animal with $\delta + 1$ hairs, χ horns and strength at least (ψ, τ), or
> - G can almost be embedded into Σ (with conditions as in (3.1)), or
> - G contains the complete graph on p vertices as a minor.

The next reduction, (7.1), then uses the results from Graph Minors XV in order to alter the third and fourth possibilities above, so that the possibilities of near embeddings and existence of K_p minors are replaced with level Σ-spans instead, which are Σ-spans that can be drawn flat on the surface in a certain sense. Note also that the first possibility now only involves handles and no crosscaps.

> For every surface Σ, $\phi, \psi \geq 3$, $\theta' \geq 1$ and $\delta, \lambda, \tau, \chi \in \mathbb{N}$ there are $\sigma \geq 0$ and $\theta \geq 1$ so that for every tangle in G and every animal with δ hairs, χ horns and strength at least (θ, σ) one of the following possibilities holds.
>
> - For one of the surfaces Σ' constructed by attaching a handle to Σ it is possible to find a Σ'-span that has order at least ϕ, or

- G contains an animal with $\delta + 1$ hairs, χ horns and strength at least (ψ, τ), or
- after removal of no more than $\chi + \frac{1}{2}\delta^2\phi^2$ apex vertices from G there is a $(\lambda, 2\psi)$-level Σ-span that has order at least θ'.

The final reduction (8.1) then replaces the requirement that a Σ-span be level with the existence of large disjoint zones so that the rest of the graph not contained in these zones contains only a bounded number of long disjoint paths between the zones. The possibility of certain animals existing is replaced with the possible existence of a path that avoids a subgraph, the endpoints of which lie far apart, and one of those endpoints also being far from any of the zones.

> For every surface Σ, $\phi, \psi \geq 3$ and $\delta \in \mathbb{N}$ there are $\gamma \geq 0$ and $\theta > (4\gamma + 2)\delta$ so that for every tangle in G, Σ-span H that has order at least θ, and $Y \subseteq H$ of cardinality δ such that two distinct points in Y are at least θ apart in the metric induced by the tangle, one of the following possibilities holds.
>
> - For one of the surfaces Σ' constructed by attaching a handle to Σ it is possible to find a Σ'-span that has order at least ϕ, or
> - there is path in G with endpoints in, but otherwise disjoint from, H that are at least ψ apart so that at least one endpoint is also ψ apart from every $y \in Y$, or
> - there are pairwise disjoint γ-zones around any $y \in Y$ so that the drawing of H outside of any of them is rigid, and after removing at most $\frac{1}{2}\delta^2\phi^2$ apex vertices in G that do not both lie in any of these zones and in H, there are no paths with endpoints in, but otherwise disjoint from, H' where these endpoints are also farther than 2ψ apart under the metric of the $(4\gamma + 2)\delta$-compression in H' of the tangle of G, where H' consists of the parts of H that do not lie in one of the zones.

In section 9 of Graph Minors XVI these possibilities are then proved to be the only ones, completing the proof of the excluded minor theorem.

Again, all results and proofs of Graph Minors XVI only use finite graphs drawn in surfaces, and the proof of the excluded minor theorem in Graph Minors XVI is thus formalizable in ACA_0. Graph Minors XVI uses the following types of induction.

- The proof that (4.1) implies (3.1) uses an implicit arithmetical induction in constructing a certain sequence of numbers. The same is true for the proof that (5.1) implies (4.1), that (6.1) implies (5.1) and that (4) in the proof that (8.1) implies (7.1) holds (in the form of δ applications of (6.2)).
- One further arithmetical induction is used in the proof of (5.5).
- The final proof, that of (8.1), also uses arithmetical induction in (3) on page 72 to prove that a certain set is free with regard to a number of tangles.

4.13 Graph Minors XVII

Graph Minors XVII investigates how to encode vortices and apex vertices as labels of a quasi-order in order to transform near-embeddable graphs into graphs that are fully embeddable into a surface, but now do have labels from the mentioned quasi-order. More precisely, after deletion of the apex vertices, the almost embeddable graph will be converted into an edge-labelled hypergraph where every edge is adjacent with either 2 or 3 vertices.

To this end, Graph Minors XVII introduces paintings, which are the same as drawings, except that instead of lines paintings utilize disks to connect vertices (in a sense "fat" lines), which are called cells and can have at most three ends. Another difference is that paintings must contain the boundary of the surface they are drawn in, i.e. if the boundary circles were filled, the resulting disks would be regions of the painting, and every cell bordering the boundary can only have two ends.

A portrayal of a hypergraph G is then defined as a painting that encodes G in a specific way. The cells encode subhypergraphs of G, and the nodes the intersections of these subhypergraphs. For internal vertices (those not lying on the boundary), the intersections can only contain one vertex. For border vertices, the subhypergraphs corresponding to the border cells can be arranged in a circle in the same way as the corresponding cells, so that their intersections can be encoded as subhypergraphs corresponding to the common end of two cells (note again that border cells have only two ends).

The portrayals of Graph Minors XVII are often assumed to be in some sense minimal. For this purpose, the warp of a border cell c connecting n_1 and n_2 is defined as the least p so that the subhypergraphs corresponding to n_1 and n_2 have size at most $p + 1$, and if there are $p + 1$ disjoint paths betweens them through the 1-skeleton (which is a simple graph as similar as possible to the hypergraph) of the hypergraph encoded by c then one of these paths must have two certain fixed endpoints which are encoded by the portrayal. The warp of a portrayal is then the maximum over the warps of its cells. For a tangle $\mathcal{T}$ in the hypergraph corresponding to a portrayal, the portrayal is called $\mathcal{T}$-central if every subhypergraph corresponding to a cell is a small side of a separation in $\mathcal{T}$. Graph Minors XVII then shows that up to the deletion of a few apex vertices, every $\mathcal{T}$-central portrayal can be transformed into another, better connnected $\mathcal{T}$-central portrayal, provided that the portrayal has minimal warp in a sense and that it is as simple as possible, i.e. using as few internal edges with many ends as possible, and as few vertices as possible, with internal vertices weighted more strongly.

Graph Minors XVII then applies these results in (14.2) with a version of the excluded minor theorem in terms of such portrayals, where the graph that the minor is excluded from is allowed to be a hypergraph. (14.2) roughly states that:

> Given an ordinary graph H, there exist $p, q, z \in \mathbb{N}$ and $\theta > z$ so that if G is

any arbitrary hypergraph containing a tangle $\mathcal{T}$ of order at least θ, one of two possibilities holds.

- Either H is a minor of the 1-skeleton of G, or
- up to removal of at most z apex vertices, there is a $\mathcal{T}$-central portrayal of G that has warp at most p in a surface in which H is not embeddable and that has at most q boundary components, and further that even up to removal of $2p+7$ vertices, the portrayal is minimal in the same sense as above.

Overall, there are no obstructions to carrying out the proofs of Graph Minors XVII in ACA_0, and the following types of induction are used.

- At the end of the proof of (4.3), an arithmetical induction is used.
- (10.2) and (10.3) use arithmetical induction via a minimal counterexample.
- Arithmetical induction is used in the proof of (11.1) to show that a certain subgraph is a circuit which has length 6.
- (12.1) uses arithmetical induction to show that under certain assumptions, any complete subgraph is contained in a subgraph corresponding to just one border cell.
- Another arithmetical induction is used in the proof of (12.6), which establishes a connection between portrayals of a graph and its 1-skeleton.
- One more arithmetical induction is used in the proof of (13.3), a lemma which is applied to obtain portrayals of higher redundancy, i.e. where a higher number of vertices can be deleted with the portrayal staying minimal.

4.14 Graph Minors XVIII

Graph Minors XVIII introduces another restricted form of the graph minor theorem, namely the graph minor theorem for graphs that already have minor-well-quasi-ordered central components that in addition can be separated out by every tangle, which extends the results of Graph Minors IV in a sense. This is a central result, necessary in the proofs of Graph Minors XIX and the graph minor theorem in Graph Minors XX, and even in the proof of Nash-Williams' immersion conjecture in Graph Minors XXIII. Graph Minors XVIII also begins the intensive use of patchworks (briefly used before in Graph Minors IV and VIII), which are a central tool in the following Graph Minors papers and used for defining and working with minors of hypergraphs.

The theorem above, theorem (6.7) of Graph Minors XVIII, is formulated in terms of Ω-patchworks, hearts and θ-isolation. Since the result is so important for the graph minor theorem itself, it make sense to roughly (technically, in Graph Minors XVIII they are defined in terms of rooted hypergraphs and simulations instead) define these terms here.

A patch is a collection of disconnected cliques on a given vertex set. A patchwork then consists of a triple $\langle G, \Delta, \mu \rangle$ with G a hypergraph, and Δ and μ functions on the edges of that hypergraph. The function μ is defined only on a subset of all edges of G and assigns to each edge an ordering of its endpoints. The function Δ is defined on all edges of G and assigns to every edge a patch on the endpoints of that edge, with the additional condition that the patch contains all possible sets of disconnected cliques if the edge is not in the domain of μ.

Next, a location $\mathcal{L}$ is a set of separations $\langle A_i, B_i \rangle$, $i = 1, \ldots, n$, in a patchwork $P = \langle G, \mu, \Delta \rangle$ so that $A_i \subseteq B_j$ if $i \neq j$. It can then be shown that $H := \bigcap\{B_i : 1 \leq i \leq n\}$ contains all the separating vertices $A_i \cap B_i$. This subgraph H is defined to be the heart of $\langle P, \mathcal{L} \rangle$ if it has additional edges e_i with endpoints $A_i \cap B_i$, $1 \leq i \leq n$, and is imbued with functions μ' and Δ', where μ' is an extension of μ to H and Δ' is an extension of Δ that captures the connectedness of the vertices of $A_i \cap B_i$ in A_i.

An Ω-patchwork then is a patchwork with an additional function ϕ that assigns to each edge an element of the quasi-order Ω. A location $\mathcal{L}$ θ-isolates a tangle $\mathcal{T}$ if all its separations have order smaller than θ and are also in $\mathcal{T}$, and if for every other tangle $\mathcal{T}'$ that has order at least θ, any $(C, D) \in \mathcal{L}$ that distinguishes $\mathcal{T}$ from $\mathcal{T}'$ is bigger than the $(\mathcal{T}, \mathcal{T}')$-distinction[1] (A, B), in the sense that $A \subseteq C$ and $D \subseteq B$ hold.

The main theorem (6.7) of Graph Minors XVIII now roughly says that for any well-quasi-order Ω and for any set of Ω-patchworks (in which the patches contain 2-cliques for every possible subset of the endpoints of an edge), if for all tangles in any of these patchworks there exists a location such that the hearts (with regard to a tiebreaker which is allowed to depend on the patchwork) of these locations are well-quasi-ordered under the minor relation and the location θ-isolates the tangle, then the patchworks are well-quasi-ordered as well.

Section 7 then concludes with a helpful lemma which shows that certain locations θ-isolate tangles, provided that the elements (A, B) of the location are minimal with respect to a tie-breaker, in the sense that if $A \subseteq A'$ and $B' \subseteq B$ for any (A', B') in the tangle, then (A', B') has higher order with respect to the tie-breaker than (A, B).

Aside from the use of quasi-orders, the proofs of Graph Minors XVIII are largely finitely combinatorial. Up to section 6, some care must be taken when considering the used inductions, since the theorems now include proper set quantifiers due to the quasi-orders. In section 6, results from Graph Minors IV are used. Since these ultimately rely on a Π^1_3-induction, they are mentioned explicitly below, together with the other inductions used.

- The proofs of (3.4) and (4.1) use arithmetical induction.
- A Π^1_1-induction is used in (5.7) to show that a minor relation exists between certain

[1]The $(\mathcal{T}, \mathcal{T}')$-distinction with respect to a tiebreaker is the least (A, B) with respect to that tiebreaker so that $(A, B) \in \mathcal{T}$ but $(A, B) \notin \mathcal{T}'$.

Ω-patchworks.

- (6.2) uses (9.1) from Graph Minors IV, the most important result from Graph Minors IV aside from the bounded graph minor theorem.
- (6.4) uses (4.1) from Graph Minors IV, which is however completely self-contained and does not rely on the Π^1_3-induction in (2.2).
- (6.6) uses (3.3) from Graph Minors IV, which crucially employs (2.2) and thus a Π^1_3-induction in its proof.
- (6.7) uses (6.1) from Graph Minors IV which is however again self-contained.

4.15 Graph Minors XIX

Graph Minors XIX is from a proof theoretic perspective the most important paper of the Graph Minors series. It features a Π^1_2-transfinite induction which cannot be dealt with by the proof principles considered so far. This thus opens up the serious possibility that the graph minor theorem is strictly stronger than the bounded graph minor theorem.

The main theorem of Graph Minors XIX is (2.1), and similar to the main result of Graph Minors VIII it is a statement about the well-quasi orderedness of graphs drawable in a fixed surface under a kind of minor relation, where the edges on the boundary must stay in place and may however be labelled by a well-quasi order. Another difference to Graph Minors VIII is that the graphs considered in Graph Minors XIX are actually hypergraphs that can have edges with 2 or 3 ends, and that the endpoints of edges are ordered and the minor relation has to respect this order.

More formally, a painting in the context of Graph Minors XIX is the same as a painting in Graph Minors XVII, with the main difference being the additional condition that it also has a function γ associated with it that assigns to each edge an order of the ends of that edge. An inflation of a painting Γ_1 in another painting Γ_2 is essentially a minor embedding σ of the graphs underlying Γ_1 and Γ_2 with the additional condition that it also has to respect the order of the endpoints of an edge. The notion of a linear inflation is introduced to make σ preserve some features of the surface, that is to impose that

- σ maps border edges to border edges bordering the same cuff and internal edges to internal edges, and
- if e' is a border edge of Γ_2 but not in the image $\sigma(E(\Gamma_1))$ then it must be contained in the expansion of some vertex of Γ_1, and
- it preserves the cyclic order of edges around each cuff.

A painting Γ is internally 3-connected if any non-null homotopic curve that is Γ-normal (i.e. only intersects vertices of Γ) intersects Γ at least 3 times.

Theorem (2.1) is then stated as follows:

> For any sequence $\langle \Gamma_1, \Gamma_2, \ldots \rangle$ of paintings in a surface Σ that are edge-labelled from a well-quasi-order Ω via ϕ_i and further internally 3-connected, there are $i < j$ so that there exists a linear inflation $\sigma : \Gamma_i \longrightarrow \Gamma_j$ with the additional property that $\phi_i(e) \leq \phi_j(\sigma(e))$ for every edge $e \in E(\Gamma_1)$.

The statement of the theorem is then augmented slightly to obtain another theorem (3.1), in the proof of which the aforementioned Π^1_2-transfinite induction is used. To state this augmented version (3.1), define a frame to be a drawing Φ of which the underlying graph has directed edges, that coincides with the boundary of Σ and where additionally the edges are classified as long and short so that long edges are only adjacent with short edges. A painting Γ fits a frame Φ if the frame is in a sense a subpainting of Γ, that is if its vertices are also vertices of Γ, its short edges are also edges of Γ and every other border edge of Γ is contained in a long edge of Φ, where border edges of Γ have the same order on their endpoints as they have in Φ. For technical reasons, Γ is also required to be internally 3-connected, and for any edge with three endpoints and any region incident with one of the segments f of $bd(e) \setminus V(e)$ (where $bd(e)$ denotes the boundary of the disk e and $V(e)$ the endpoints of e), either the region is incident with less than 3 vertices of Γ or at least one of the endpoints incident with f lies on the boundary of Σ. An inflation σ of Γ_1 in Γ_2 is said to respect a frame Φ if border edges on every cuff are preserved under σ and any such border edge and its image are contained in the same edge of Φ.

A colour scheme χ is then composed of some surface Σ and a frame Φ for that surface, well-quasi orders $\Omega(2)$ and $\Omega(3)$ intended for the internal edges of size 2 and size 3 of paintings fitting the frame, and for every edge S of Φ one well-quasi-order $\Omega(S)$, where the well-quasi-orders corresponding to the short edges are distinct one-element orders (and thus trivial). Further, a χ-coloured painting is a painting Γ fitting the frame Φ together with a function ϕ that assigns to each edge of Γ an element of the corresponding well-quasi order, so an element of $\Omega(2)$ or $\Omega(3)$ for internal edges of size two or three respectively, and an element of $\Omega(S)$ for edges contained in S. A linear inflation between χ-colored paintings is then a linear inflation between the underlying paintings that also respects the frame Φ and the labels of the well-quasi-orders. Theorem (3.1) of Graph Minors XIX, the augmented version of the main theorem that uses the Π^1_2-bar induction, can then be stated as:

> The set of χ-coloured paintings is well-quasi-ordered under the relation induced by linear inflation, for any colour scheme χ.

4.15.1 The bar induction

The proof of (3.1) actually requires the use of four nested inductions, two of which are Π^1_2-bar inductions. These are stated with conditions S_1 to S_4 on the colour scheme χ, that say that χ is a minimal colour scheme that is bad in a certain sense.

To state these conditions, some further definitions are needed. First, a well-quasi-order Ω is an initial ideal of a well-quasi-order Ω', denoted $\Omega \preceq \Omega'$, if it is contained in Ω' and closed downward with regard to Ω', that is

$$\Omega \preceq \Omega' \longleftrightarrow \forall x \in \Omega(x \in \Omega' \wedge \forall x' \leq_{\Omega'} x(x' \in \Omega)).$$

A colour scheme χ' is then called a refinement of a colour scheme χ if:

- The underlying surfaces of χ and χ' are isomorphic and the well-quasi orders corresponding to the edges of size 2 and 3 are identical for χ and χ', and
- there is a function f (not necessarily injective or surjective) from the long edges of the frame associated with χ' (which is just called a long edge of χ') to the long edges of the frame associated with χ that satisfies:
  - for long edges R of χ', $\Omega_{\chi'}(R) \preceq \Omega_{\chi}(f(R))$, and
  - if R_1 and R_2, $R_1 \neq R_2$, are long edges of χ' that are mapped to the same long edge in χ, then $\Omega_{\chi'}(R_i) \prec \Omega_{\chi}(f(R_i))$, $i = 1, 2$, and
  - if f is bijective and none of the initial ideal relations in the first condition is strict, then the number of short sides of χ' is strictly less than that of the short sides of χ.

A colour scheme χ is called orientedly bad if there is a bad sequence of χ-coloured paintings and if either there is no orientation of Σ_{χ} at all, or there exists an orientation of Σ_{χ} so that for every edge with three ends in the sequence of χ-coloured paintings, the cyclic order of the endpoints of the edge induced by the χ-coloured painting is the same as the cyclic order induced by the orientation. We denote by $X^{<\omega}$ the set of finite sequences on X, by $[X]^{<\omega}$ the set of finite subsets of X, and by $\Omega \oplus \Omega'$ the direct sum of two well-quasi-orders Ω, Ω'. The conditions S_1 to S_4 of Graph Minors XIX on χ are then stated as follows:

S_1 If χ' is a bad colour scheme then it cannot be the case that after filling up the cuffs of both surfaces, Σ_χ can be derived by adding a number of handles and crosscaps to $\Sigma_{\chi'}$.

S_2 If χ' is an orientedly bad colour scheme so that Σ_χ has the same number of handles and crosscaps as $\Sigma_{\chi'}$, then it cannot be the case that $\Omega_{\chi'}(2)\oplus\Omega_{\chi'}(3) \prec \Omega_\chi(2)\oplus\Omega_\chi(3)$.

S_3 If χ' is an orientedly bad colour scheme then it cannot be the case that both $\Omega_{\chi'}(2) = \Omega_\chi(2)$, $\Omega_{\chi'}(3) = \Omega_\chi(3)$ hold, Σ_χ and $\Sigma_{\chi'}$ have the same number of handles and crosscaps, and that simultaneously $\Sigma_{\chi'}$ has fewer cuffs than Σ_χ.

S_4 If χ' is an orientedly bad colour scheme then it cannot be a refinement of χ.

By essentially asserting that χ is a least counterexample, S_1 corresponds to a Π^1_2-induction (since all the quasi-orders are allowed to change in previous instances of the induction), S_2 to a Π^1_2-bar induction (since $\Omega(S)$ is not specified for long edges S), S_3 to another Π^1_2-induction (because again $\Omega(S)$ is not specified for long edges) and S_4 to a Π^1_2-induction and another Π^1_2-bar induction. It may be beneficial to elucidate this last point. Since the definition of "refinement" is split into essentially two cases, one of which being that the refined colour scheme has less short sides, it makes sense to also split the induction on the corresponding formula into one bar-induction when at least one of the initial ideal inclusions is strict, and a standard induction otherwise. However, aside from the case where the refined colour scheme has fewer short edges, there is another case where none of the initial ideal inclusions has to be strict. This is when f is injective but not surjective, and the refined colour scheme has fewer long edges. Thus, two inductions are used to deal with S_4, one standard induction on the number of long and short edges, and one bar-induction corresponding to the relation defined by refinement in the case where the refined colour scheme has neither fewer short nor fewer long edges.

Before the inductions can be carried out in second order arithmetic, it is however necessary to address one further issue. The standard form of bar induction is the following:

$$\forall X(WF(X) \to \forall j(\forall i <_X j\varphi(i) \to \varphi(j)) \to \forall n \in X\varphi(n)).$$

The kind of induction used in Graph Minors XIX however is different, namely:

$$\forall X(WQO(X) \to (\forall X' \prec X(\forall X'' \prec X'\varphi(X'') \to \varphi(X')) \to \varphi(X))).$$

It is not clear whether this scheme is implied by the usual bar-induction scheme, because not every subset of the natural numbers (or of any arbitrary $X \subseteq \mathbb{N}$) can be encoded by a natural number. But the actual relation used in the proofs is a weaker one, namely:

$$X' \prec_w X :\Leftrightarrow \exists y \in X\forall x'(x' \in X' \leftrightarrow x' \in X \wedge x' \not\geq y),$$

so only sets determined by one element $y \in X$ would have to be considered. However this relation is actually not transitive, and when showing that the usual bar induction scheme implies the intended initial ideal induction scheme it is convenient to instead set

$$X' \prec_1 X :\Leftrightarrow \exists \langle y_1, \dots, y_n \rangle \in X^{<\omega}\forall x'(x' \in X' \leftrightarrow x' \in X \wedge \forall i < n(x' \not\geq y_i)),$$

which is suitable, as will be shown. Define $X^{y_1,\dots,y_n} := \{x \in X : \forall i < n(x \not\geq y_i)\}$.

The aim is to provide a well-founded relation on the natural numbers that corresponds to the above relation on well-quasi-orderings. First, define a relation $\leq_1$ on $[X]^{<\omega}$ by

$$\{a_1, \dots, a_n\} \leq_1 \{b_1, \dots, b_m\} \leftrightarrow \forall j \in \{1, \dots, m\}\exists i \in \{1, \dots, n\}a_i \leq b_j.$$

The quantifiers might seem backwards at first glance, but they are as required because $\prec_1$ essentially cuts out segments above certain elements. It needs to be shown that this relation is actually well-founded. To do this it is convenient to first prove a lemma from Graph Minors XIX, (4.3), which says that there can be no infinite descending sequence of well-quasi-orders under the (full) initial ideal relation.

Lemma 4.15.1 ((4.3) of Graph Minors XIX)**.** The following is provable in RCA_0. Let Ω^* be a well-quasi-order. Then there is no infinite sequence $\langle \Omega_i : i \in \mathbb{N} \rangle$ with $\Omega_i \preceq \Omega^*$ and $\Omega_{i+1} \prec \Omega_i$ for all $i \in \mathbb{N}$.

Proof. Assume there is such an infinite descending sequence, i.e. a set Ω with $\Omega_{i+1} \prec \Omega_i$ for all i, where $\Omega_l = \{n : \langle n, l\rangle \in \Omega\}$. Let F be defined by $F(n) = \mu k.\, \langle k, n\rangle \in \Omega \wedge \langle k, n+1\rangle \notin \Omega$. Then for all $n \geq 0$ and $k \geq 1$, $F(n) \not\leq F(n+k)$: Since $F(n) \in \Omega_n$ but $F(n) \notin \Omega_{n+1}$, also $F(n) \notin \Omega_{n+k}$ as $\Omega_{n+k} \preceq \Omega_{n+1}$. Assume $F(n) \leq F(n+k)$. Then as $\Omega_{n+k} \prec \Omega_n$ and $F(n+k) \in \Omega_{n+k}$ it follows from the definition of $\prec$ that $F(n) \in \Omega_{k+n}$, a contradiction.

Thus F is a bad sequence in Ω^*, a contradiction.

□

Next, it needs to be shown that $\leq_1$ is actually a well-founded relation. This is proved in the following lemma, but see e.g. Forster (2003) for an alternative proof of this fact.

Lemma 4.15.2. Let X be well-quasi-ordered. Then there is no infinite descending $\leq_1$-sequence in $[X]^{<\omega}$.

Proof. First it is shown that if $\{a_1, \dots, a_n\} \leq_1 \{b_1, \dots, b_m\}$, then $X^{a_1,\dots,a_n} \preceq X^{b_1,\dots,b_m}$.

Let $\{a_1, \dots, a_n\} \leq_1 \{b_1, \dots, b_m\}$. Then $x \in X^{a_1,\dots,a_n} \implies \forall i < n : x \not\geq a_i \implies \forall j < m : x \not\geq b_j \implies x \in X^{b_1,\dots,b_m}$, so $X^{a_1,\dots,a_n} \subseteq X^{b_1,\dots,b_m}$. But $X^{a_1,\dots,a_n}$ is also closed downward with regard to X and hence with regard to $X^{b_1,\dots,b_m}$, so $X^{a_1,\dots,a_n} \preceq X^{b_1,\dots,b_m}$. But then from a descending $<_1$-sequence one could construct a descending $\prec$-sequence, which is not possible by lemma 4.15.1.

□

Then for well-quasi-ordered X (note that X indeed needs to be well-quasi-ordered for $\prec_1$ and $\leq_1$ to be well-founded) the Π^1_k-bar induction scheme for $\leq_1$ implies the Π^1_k-bar induction scheme for $\prec_1$, for any $k \geq 1$.

Lemma 4.15.3. If for every well-quasi-ordered set X^* and every Π^1_k-formula $\varphi'(n)$

$$\forall j(\forall i <_1 j \varphi'(i) \rightarrow \varphi'(j)) \rightarrow \forall n \in [X^*]^{<\omega} \varphi'(n),$$

then for every well-quasi-ordered set X and every Π^1_k-formula $\varphi(Y)$

$$(\forall X' \prec_1 X(\forall X'' \prec_1 X'\varphi(X'') \to \varphi(X')) \to \varphi(X)).$$

Proof. Assume $WQO(X)$ and let $\hat{X} = X \cup \{\top\}$ where $\top$ with $\top > x$ for all $x \in X$ is a new element. Further assume that the usual bar induction scheme for Π^1_k-formulas with regard to $[\hat{X}]^{<\omega}$ and $\leq_1$ holds. Let $\varphi(X)$ be a Π^1_k-formula. The aim is to show the $\prec_1$-bar induction scheme for φ. So assume progression for φ with respect to $\prec_1$, i.e. $\forall X' \prec_1 X(\forall X'' \prec_1 X'\varphi(X'') \to \varphi(X'))$. Then it needs to be proved that $\varphi(X)$ holds. To do this, define $\varphi'(\{a_1, \ldots, a_n\})$ to be $\varphi(\{x \in \hat{X} : \forall i < n : x \not\geq a_i\})$, i.e.

$$\varphi'(i) := \forall Y(i = \{a_1, \ldots, a_n\} \to (\forall x(x \in Y \leftrightarrow x \in \hat{X} \wedge \forall j < n : x \not\geq_1 a_j) \to \varphi(Y))).$$

Note that such a set Y exists, so φ' is actually the intended statement, and that $\varphi'(i)$ is still a Π^1_k-formula. Now, the idea is to utilize Π^1_k-bar induction for φ' to show that $\varphi(X)$ holds. For this it needs to be shown that φ' is progressive. So assume that $\forall i <_1 j\varphi'(i)$ (where i and j are codes for finite subsets of $\hat{X}$), then $\varphi'(j)$ needs to be proved. That $\forall i <_1 j\varphi'(i)$ implies $\forall X'' \prec_1 X^j\varphi(X'')$:

$X'' \prec_1 X^j$ means that $X'' = X^{b_1,\ldots,b_m,c_1,\ldots,c_k}$ for some $c_1, \ldots, c_k$, where $j = \{b_1, \ldots, b_m\}$, and trivially $\{b_1, \ldots, b_m, c_1, \ldots, c_k\} \leq \{b_1, \ldots, b_m\}$, where in fact strict inequality holds since $X'' \prec_1 X^j$. Let $i = \{b_1, \ldots, b_m, c_1, \ldots, c_k\}$ then $\varphi'(i)$ holds by assumption, and hence (since $X^i = X''$) $\varphi(X'')$ holds.

So $\forall X'' \prec_1 X^j\varphi(X'')$. Since progressiveness for φ with regard to $\prec_1$ was assumed, this gives $\varphi(X^j)$, and hence $\varphi'(j)$. Hence φ' is progressive. Applying Π^1_k-bar induction thus results in $\forall x \in [\hat{X}]^{<\omega}\varphi'(x)$. In particular $\varphi'(\{\top\})$ holds, which implies $\varphi(X)$ and thus completes the proof.

□

Bar induction for $\prec_1$ corresponds to condition S_2. Another bar induction is needed for a relation corresponding to refinement. On direct sums of well-quasi orders, define

$$\begin{aligned}
&X_1 \oplus \ldots \oplus X_n \prec_2 Y_1 \oplus \ldots \oplus Y_m \\
\leftrightarrow\ &\exists f : \{1, \ldots, n\} \longrightarrow \{1, \ldots, m\}(\forall i \leq n(X_i \preceq_1 Y_{f(i)}) \wedge \\
&\exists i \leq n(X_i \prec_1 Y_{f(i)}) \wedge \forall i, j(i \neq j \wedge f(i) = f(j) \rightarrow X_i \prec_1 Y_{f(i)})).
\end{aligned}$$

To perform bar induction on this relation, an analogous relation is needed for natural numbers. Define analogously on $([X]^{<\omega})^{<\omega}$ a relation

$$\begin{aligned}
&\langle \omega_1, \ldots, \omega_n \rangle <_2 \langle \omega'_1, \ldots, \omega'_m \rangle \\
\leftrightarrow\ &\exists f : \{1, \ldots, n\} \longrightarrow \{1, \ldots, m\}(\forall i \leq n(\omega_i \leq_1 \omega'_{f(i)}) \wedge \\
&\exists i \leq n(\omega_i <_1 \omega'_{f(i)}) \wedge \forall i, j(i \neq j \wedge f(i) = f(j) \rightarrow \omega_i <_1 \omega'_{f(i)})).
\end{aligned}$$

Analogous to $<_1$, it needs to be shown that $<_2$ is well-founded.

Lemma 4.15.4. Let X be well-quasi-ordered. Then there is no infinite descending $<_2$-sequence in $([X]^{<\omega})^{<\omega}$.

Proof. Since X is well-quasi-ordered, by lemma 4.15.2 $[X]^{<\omega}$ is well-founded with regard to $<_1$. If $\langle \omega_1, \ldots, \omega_n \rangle <_2 \langle \omega'_1, \ldots, \omega'_m \rangle$ via f, say that ω'_j branches into $\omega_{i_1}, \ldots, \omega_{i_{m_j}}$ if $f^{-1}(j) = \{i_1, \ldots, i_{m_j}\}$ and $\omega_{i_1} <_1 \omega'_j$ (this last condition is only necessary in case that the preimage of j contains only one element).

Now assume that there is a descending $<_2$-sequence $s := \langle \langle \omega^i_1, \ldots, \omega^i_{n_i} \rangle : i \in \mathbb{N} \rangle$, and let $\langle f_i : i \geq 2 \rangle$ be the sequence of functions witnessing the $<_2$ relations in that sequence. Interpret each ω^i_k as a term, and identify two such terms transitively if $\omega^{i+1}_k = \omega^i_l$ and $f_{i+1}(k) = l$. Let $S = \{\omega^i_k : i \in \mathbb{N} \wedge k \leq n_i\}$, and for $\omega, \omega' \in S$ say that ω' is a successor of ω if (the element underlying) ω branches into (the element underlying) ω' at some point in s. Note that every ω branches only once, and that it branches only into finitely many successors. This branching relation thus defines a forest on S, which is infinite since s is bad and in which every tree is finitely branching. Since this forest consists of

n_1, and thus finitely many, trees, there must be one such tree which is infinite. Applying König's Lemma to this tree yields an infinite, strictly decreasing $<_1$-sequence in $[X]^{<\omega}$, a contradiction to $[X]^{<\omega}$ being well-founded.

$\square$

Then the bar induction scheme for $<_2$ implies the bar induction scheme corresponding to the refinement relation.

Lemma 4.15.5. If for every well-quasi-ordered set X^* and every Π^1_k-formula $\varphi'(n)$

$$\forall j(\forall i <_2 j\varphi'(i) \to \varphi'(j)) \to \forall n \in ([X^*]^{<\omega})^{<\omega}\varphi'(n),$$

then for every direct sum of well-quasi-ordered sets $X := X_1 \oplus \ldots \oplus X_n$ and every Π^1_k-formula $\varphi(Y)$

$$(\forall X' \prec_2 X(\forall X'' \prec_2 X'\varphi(X'') \to \varphi(X')) \to \varphi(X)).$$

Proof. The proof is analogous to that of lemma 4.15.3. Let $\varphi(X)$ be a Π^1_k-formula, for $1 \leq i \leq n$ define $\hat{X}_i := X_i \cup \{\top_i\}$, where $\top_i$ is a new element with $\top_i > x_i$ for all $x_i \in X_i$, and let $X^* = \hat{X}_1 \oplus \ldots \oplus \hat{X}_n$. Assume regular bar induction for Π^1_k-formulas with regard to $([X^*]^{<\omega})^{<\omega}$ and $\leq_2$, and also assume progression for $\prec_2$, i.e. $\forall X' \prec_2 X(\forall X'' \prec_2 X'\varphi(X'') \to \varphi(X'))$, the aim is to show $\varphi(X)$. To do this, define $\varphi'(\langle\omega_1, \ldots, \omega_m\rangle)$ as $\varphi(X^{\omega_1}_{k_1} \oplus \ldots \oplus X^{\omega_m}_{k_m})$ where $k_i \in \{1, \ldots, n\}$, i.e.

$$\begin{aligned}\varphi'(l) := \quad & \forall Y(l = \langle\omega_1, \ldots, \omega_m\rangle \to (\forall i(i \in Y \leftrightarrow (i = \langle a_1, \ldots, a_m\rangle \wedge \\ & \forall k \leq m\exists k' \leq n(a_k \in X_{k'}) \wedge \forall k \leq m\forall j \leq |\omega_k| : a_k \ngeq \omega_{kj}) \to \varphi(Y))).\end{aligned}$$

Then the goal is to show that φ' is progressive. So assume that $\forall i <_2 j\varphi'(i)$, then $\varphi'(j)$ needs to be proved. That $\forall i <_2 j\varphi'(i)$ implies that $\forall X'' \prec_2 X^j\varphi(X'')$ since the two relations are defined completely analogously. Since by assumption progression for φ holds, this gives $\varphi(X^j)$, hence $\varphi'(j)$. So φ' is progressive, Π^1_k-bar induction thus gives

$\varphi'(x)$ for every $x \in ([X^*]^{<\omega})^{<\omega}$. In particular $\varphi'(\langle\{\top_1\},\ldots,\{\top_n\}\rangle)$, which implies $\varphi(X)$.

$\square$

Now the proof of Graph Minors XIX that requires the Π^1_2-bar induction can be carried out in second order arithmetic. First note that "χ is a colour scheme" is a Π^1_1-formula, since it involves the statement that a certain quasi-order is a well-quasi-order. Next, note that "χ is good" and "χ is orientedly good" are also Π^1_1-statements, since they assert that there is no bad sequence of χ-coloured paintings.

The first induction is performed on a formula corresponding to S_1, namely the formula

$$\varphi_1 = \forall\,\langle n,k\rangle\,\forall\chi(\chi \text{ is a colour scheme with } n \text{ handles and } k \text{ crosscaps} \rightarrow\ \chi \text{ is good}).$$

The pairs $\langle n,k\rangle$ here are ordered by $\langle n',k'\rangle \leq \langle n,k\rangle$ if and only if $n' \leq n \wedge k' \leq k$. For a colour scheme χ, write $\chi(X,Y,Z)$ to make the quasi-orders $X = \Omega(2)$, $Y = \Omega(3)$ and Z as the direct sum of the $\Omega(S)$ explicit. The formula corresponding to S_2 is

$$\varphi_2 = \forall X\forall Y\forall Q(\chi(X,Y,Q) \text{ is a colour scheme} \rightarrow \chi(X,Y,Q) \text{ is orientedly good}).$$

Two nested Π^1_2-bar inductions are then performed with regard to $\prec_1$, the first on X and the second on Y. Alternatively, a single Π^1_2-bar induction can be performed on the disjoint union of X and Y. The formula corresponding to S_3 is

$$\varphi_3 = \forall k\forall\chi(\chi \text{ is a colour scheme with } k \text{ cuffs} \rightarrow\ \chi \text{ is orientedly good}).$$

Here the induction will be a simple Π^1_2-induction on the number of cuffs. Finally, we split S_4 into two formulas, one referring to the number of long and short edges, and one corresponding to $\prec_2$. The first formula is

$$\begin{aligned}\varphi_4 = \forall n\forall k\forall\chi(&\chi \text{ is a colour scheme with } n \text{ long edges and } k \text{ short edges}\\ &\rightarrow\ \chi \text{ is orientedly good}),\end{aligned}$$

where two nested Π^1_2-inductions are carried out. The formula corresponding to $\prec_2$ is

$$\begin{aligned}\varphi_5 = \forall S^*\forall Q_1\forall Q_2(\chi(Q_1, Q_2, S^*) \text{ is a colour scheme}\\ \rightarrow \chi(Q_1, Q_2, S^*) \text{ is orientedly good}),\end{aligned}$$

where a Π^1_2-bar induction is carried out with regard to $\prec_2$ on the disjoint sum of quasi-orders corresponding to the long sides, S^*. Note that if $\chi(Q_1, Q_2, S^*)$ is a colour scheme and every colour scheme $\chi'(Q_1, Q_2, S'^*)$ satisfies φ_4 and φ_5, then in particular every refinement of χ is orientedly good, and thus S_4 holds for χ.

Made precise, the proof of Graph Minors XIX that (4.1) implies (3.1), on the bottom of page 334, can be carried out as follows. Since it is used in the proof, note that (4.2) of Graph Minors XIX states that every bad colour scheme induces an orientedly bad colour scheme on the same surface.

Theorem 4.15.6 (Proof of (3.1) assuming (4.1), p.334 of Graph Minors XIX)**.** If every colour scheme that satisfies S_1-S_4 is orientedly good, then there is no bad sequence of χ-coloured paintings, for every colour scheme χ.

Proof. Assume that every colour scheme satisfying S_1-S_4 is orientedly good, and let χ be an arbitrary colour scheme. First, perform Π^1_2-induction on φ_1 and hence assume that φ_1 holds for all colour schemes with surfaces of lesser genus, i.e. assume that there is no bad sequence of χ'-coloured paintings, where χ' has lesser genus than χ.

Next, perform a Π^1_2-bar induction on φ_2, and hence additionally assume that there is no orientedly bad sequence of χ'-coloured paintings where the well-quasi-orders $\Omega_{\chi'}(k)$, $k = 2, 3$, are $\prec_1$-predecessors of $\Omega_\chi(k)$, $k = 2, 3$.

Then, analogously continue arguing inductively for φ_3, φ_4 and φ_5. By the inductive hypothesis for φ_1, χ satisfies S_1. By the inductive hypotheses for φ_2 and φ_3, χ satisfies S_2 and S_3, respectively. As noted the discussion above, by the inductive hypotheses for φ_4 and φ_5, χ satisfies S_4. Since χ satisfies S_1-S_4 it is orientedly good by the assumption of our theorem. By (4.2) of Graph Minors XIX it is thus also good. This proves the theorem. □

To prove that it is correct to restrict the $\prec$-relations of Graph Minors XIX to $\prec_1$- and $\prec_2$ relations, it remains to show that actually only those are used in the proofs of Graph Minors XIX. Hence, all the applications of S_2 and S_4 need to be looked at.

The first use of S_4 is in lemma (10.5). Here, removing a freedom-flaw results in a painting that has several fewer short edges and/or at least one fewer long edge than the previous one, which corresponds to φ_4. The other use of S_4 is in the proof of lemma (11.1) (on page 361 of Graph Minors XIX). Here one long edge S of a painting is replaced by several long edges X_i with corresponding quasi-orders $\{x \in \Omega(S) : x \not\geq \omega_i'\}$ (here ω_i' is just an element of $\Omega(S)$, not a set of elements as in our lemmas above), which is in accordance with φ_5. Finally, S_2 is used in the proof of (14.1) (on the bottom of page 337 and top of page 338 of Graph Minors XIX). Here, $\Omega_{\chi'}(k)$ of the new painting is defined to be $\{x \in \Omega_\chi(k) : x \not\geq \omega_0\}$, $k = 2, 3$, which is in accordance with φ_2. These are all the applications of S_2 and S_4.

4.15.2 The rest of Graph Minors XIX

The rest of Graph Minors XIX is then devoted to proving (4.1), which says that every colour scheme that satisfies $S_1 - S_4$ is orientedly good. The idea is roughly that if a graph is not a minor of another, it must have certain structural flaws. Conditions S_1-S_4 are then used to deal with these flaws.

First, the definitions and terms used in Graph Minors XIX are made to match with those of earlier papers of the Graph Minors series. For this, conditions S_1 and S_3 are used, among other things to be able to pass from inflations to linear inflations. One important intermediate result is (9.2), which says that it is possible to focus on locations in the paintings in order to produce a good sequence. Since it can be assumed that the first painting of a given sequence of paintings does not admit an inflation in any other painting of the sequence, this non-existence of an inflation gives the result that any tangle in the paintings must have some kind of flaw; conditions S_1-S_4 are then used to show that these flaws give rise to locations that can be used to apply (9.2).

The first such flaw is a representativeness flaw, meaning that the representativeness of a tangle (defined in terms of the number of intersections of null-homotopic curves with the painting) is lower than that of the painting and than its order, which is dealt with by S_1. The next kind of flaw, called distance flaw, meaning that the distance between some two poles (points representing disks cut out of the surface) is less than the representativeness of a tangle, is dealt with by S_3. After that there is the case of a freedom flaw, meaning that there exists a circuit of certain maximum length containing a long edge or a large number of vertices, and hence short edges, of the frame. The condition of S_4 corresponding to φ_4 deals with this flaw.

Investigated next are flaws due to the labelling on the edges; in essence this happens if sections above one element of the well-quasi-order are not used as labels. S_4 deals with such insufficient generality on the long sides by its condition corresponding to φ_5. S_2 then similarly deals with the same kind of flaw on internal edges, i.e. edges labelled from $\Omega(2)$ and $\Omega(3)$.

Also note that the "well-behaved" sets $\mathcal{C}_1$-$\mathcal{C}_5$, constructed in (10.3), (10.4), (10.5), (11.1) and (14.1) respectively, which are used in the final proof, correspond to the formulas φ_1-φ_5, namely $\mathcal{C}_1$ to φ_1, $\mathcal{C}_2$ to φ_3, $\mathcal{C}_3$ to φ_4, $\mathcal{C}_4$ to φ_5 and $\mathcal{C}_5$ to φ_2.

The inductions used in the proof that (4.1) implies (3.1) are the only noteworthy proof principles used in Graph Minors XIX for the purpose of determining the proof-theoretic strength of the graph minor theorem. Another induction, that is however arithmetical, is used in the proof of (5.2).

4.16 Graph Minors XX

Graph Minors XX finishes the proof of the graph minor theorem by combining the results of previous papers of the Graph Minors series. As outlined already earlier in the Graph Minors series, the proof of the graph minor theorem is concluded as follows. First, it is noted that given an infinite sequence $\langle G_1, G_2, \ldots \rangle$ of graphs, it may be assumed that G_1 is not a minor of any other G_i, otherwise the graph minor theorem holds. But then the results of Graph Minors XVII show that all the G_i for $i \geq 2$ can be decomposed into parts which can be nearly embedded in a surface in which G_1 cannot be embedded. To show that a sequence consisting of such graphs must contain G_i and G_j, $i < j$, such that G_i is a minor of G_j, the results from Graph Minors XVIII, saying that this holds if the hearts of the associated patchworks are well-quasi-ordered, and Graph Minors XIX, which gives the graph minor theorem for labelled hypergraphs embedded on a fixed surface, are combined. Most of Graph Minors XX is devoted to removing technical obstructions to combining these theorems. However, Graph Minors XX finishes with a stronger theorem about the well-quasi-orderedness of patchworks, which is somewhat similar to that of Graph Minors XVIII and is used in the proof of Nash-Williams' immersion conjecture.

The following proof-theoretic principles are used in Graph Minors XX:

- The proof of (4.5) uses a Π^1_2-induction to show that one set of patchworks is well-quasi-ordered if another is.
- The same is the case for (4.8), (4.9) and (4.13).

4.17 Graph Minors XXIII

Graph Minors XXIII proves two important theorems, the first in a sense a strengthening of the graph minor theorem to hypergraphs (stated in (1.2) of Graph Minors XXIII), the second Nash-Williams' immersion conjecture (stated in (1.1) of Graph Minors XXIII). An immersion is somewhat similar to a topological minor, but less strict in that the paths may not necessarily be vertex disjoint, i.e. an immersion of one graph into another maps vertices injectively to vertices and edges to edge-disjoint paths connecting the images of their endpoints.

The main proof-theoretic principle used in Graph Minors XXIII is another Π^1_2-bar induction, similar to the one in Graph Minors XIX. However this one is now easier to handle, since it can be dealt with by the same methods developed for the Π^1_2-bar induction in Graph Minors XIX. Some additional care must be taken however because of differing definitions of Graph Minors XIX and Graph Minors XXIII. It is defined on page 187 of Graph Minors XXIII that effectively only subsets of a well-quasi-order Ω generated by the labels on the edges of a sequence are considered, that is subsets of the form $\bigcup_{i\in\mathbb{N}}\{x \in \Omega : \exists e \in S_i(x \leq l_i(e))\}$ for some subset $S_i \subseteq E(G_i)$ and where $l_i(e)$ denotes the label of e. These are used in the definition of a shadow of some set of partial Ω-patchworks, which is a finite sequence consisting of subsets of a well-quasi-order, a natural number and two finite sets, which are ordered lexicographically by the $\prec$-relation, the $<$-relation and the subset relation (which is clearly well-founded for finite sets), respectively. However, it would also be possible to treat a shadow like a colour scheme from Graph Minors XIX, and decree that the respective edges be labelled from the associated well-quasi-order. A shadow of Graph Minors XXIII would then be the minimal "shadow-colour scheme" from which a certain set could be labelled. Since the codomains of functions do effectively not play a role in any of the proofs of Graph Minors XXIII, this does not change any of the proofs and is only a change in terminology.

However, in the treatment of second order arithmetic this change of terminology allows the use of $\prec_1$ instead of $\prec$, and thus the formalization of the proof with Π^1_2-bar induction. The bar-induction appears in the discussion at the end of section 3 (on page 187 of Graph Minors XXIII), where it is employed to show that one shadow is good if all those below it with regard to the lexicographic ordering above are, which corresponds to bar-induction on a Π^1_2-statement. To show that $\prec_1$ can be indeed used instead of $\prec$, all the instances where a shadow is reduced to a smaller one due to a $\prec$-relation have to be considered. This is the case in 6.1 and 8.1 of Graph Minors XXIII. In 6.1 at most some finite number of edges are allowed to have label $\geq \xi$ for some $\xi \in \Omega_\omega$ (note that Ω_ω in this context is a certain well-quasi-order, not to be confused with the ordinal Ω_ω), these are then removed so that no edge has label $\geq \xi$, or in other words, so that every edge is labelled from Ω^ξ_ω. In 8.1 the same is the case and no vertices need to be removed, but a case distinction has to be made between partial Ω-patchworks and full Ω-patchworks.

To summarize, the main theorem of Graph Minors XXIII is (2.1), a theorem about Ω-patchworks from which the immersion theorem (1.1) and the extended version of the graph minor theorem (1.2) follow. As laid out at above, the proof of (2.1) consists of a Π^1_2-bar induction to show that every shadow is good. For the induction hypothesis, it needs to be shown that there can be no minimally bad shadow. This occupies most of Graph Minors XXIII, and consists of first showing that a certain number of conditions are sufficient for one patchwork to be contained in another, and then showing that if any of these conditions fail, the last theorem of Graph Minors XX can be applied to yield a good sequence of patchworks. By assuming that the first term of a sequence of patchworks is not simulated in any later term, this thus yields the desired result.

The Π^1_2-bar induction is the strongest principle used in these proofs. Another induction is used in claim (2) in the proof of 7.2, but this induction is arithmetical.

This chapter gave a summary of the proof methods used in the Graph Minors series. It showed that $\Sigma - GMT$ (and its uniform version) and the bounded graph minor theorem are equivalent over ACA_0, thus adding $\Sigma - GMT$ to the equivalence chain

$$\Sigma - GMT \leftrightarrow bGMT \leftrightarrow EKT \leftrightarrow pGMT \leftrightarrow tGMT_{\leq 3} \leftrightarrow WO(\Psi_0(\Omega_\omega)).$$

Furthermore, this chapter established that the graph minor theorem, the immersion theorem and a generalized version of the graph minor theorem ((1.2) of Graph Minors XXIII) are all provable in $\Pi^1_1 - \mathrm{CA}_0 + \Pi^1_2 - \mathrm{BI} + \Pi^1_3 - \mathrm{IND}$. The proof-theoretic strength of these three theorems thus lies between $\Pi^1_1 - \mathrm{CA}_0 + \mathrm{RFN}_{\Pi^1_1}(\Pi^1_1 - \mathrm{CA}_0)$ and $\Pi^1_1 - \mathrm{CA}_0 + \Pi^1_2 - \mathrm{BI} + \Pi^1_3 - \mathrm{IND}$. Since labelled trees with gap-embedding are often used to show lower bounds on similar combinatoric theorems, the next chapter will give some results about such ordinal-labelled trees, demonstrating which principles can not be used for such a task and which ones possibly might.

Chapter 5

Generalized Kruskal theorems

This final chapter investigates the proof-theoretic strength of versions of well-quasi-orderedness statements for labelled trees under gap-condition, to possibly provide some combinatorial principle that has proof-theoretic lower bounds in the vicinity of, but higher than that of the graph minor theorem or immersion theorem, and to show which principles can not be used for this purpose. A general ordinal analysis of trees labelled from arbitrary ordinals has been given in Gordeev (1993) and Gordeev (1990). However, these papers use non-standard ordinal notation systems and the proofs are rather long and involved. Therefore a shorter proof using standard methods is given in the first two sections of this chapter, which will show that trees labelled from an ordinal α are closely related to iterated Π^1_1-comprehension of length α. Since these results are already known, the proofs will be highly formal to keep them short. The third section of this chapter will give a new combinatorial principle employing ordinal labelled trees that have additional labels from an arbitrary well-quasi-order, and for which slightly better lower bounds can be shown.

A labelled rooted tree T for the purposes of this chapter is a rooted tree together with a function $l : V(T) \longrightarrow \alpha$, where α is an arbitrary ordinal. In this case T is said to have labels from α. A tree is ordered if the set of immediate successors of any vertex has a linear order associated with it, so that we can speak of the i-th immediate successor of a vertex. All trees considered in the following are rooted, ordered and labelled.

Recall the following definitions and notations. For vertices $u, v \in T$ we write $[u, v]$ for the unique path from u to v in T. If v lies on $[w, root(T)]$ then w is a successor of v, also written $v \leq w$. A vertex w is an immediate successor of v if w is a successor of v and there is no vertex third vertex u on $[v, w]$. An embedding between rooted labelled trees T_1 and T_2 is a function $f : T_1 \longrightarrow T_2$ that is infimum-preserving (with regard to the tree-ordering $\leq$), order-preserving (with regard to the ordering of immediate successors) and satisfies the gap-condition:

> If w is an immediate successor of v in T_1 and $u \in (f(v), f(w)]$, then $l(u) \geq l(w)$. If $u \leq f(root(T_1))$ then also $l(u) \geq l(root(T_1))$.

The existence of such a function is denoted by $T_1 \leq T_2$. A sequence $s = \langle T_i : i \in \mathbb{N} \rangle$ of trees is bad if there are no $i, j \in \mathbb{N}$, $i < j$, such that $T_i \leq T_j$. GKT$_\alpha$ is the statement that there is no bad sequence of trees labelled from α, and GKT that $\forall \alpha$GKT$_\alpha$.

We denote by qT the root-label of a tree T, and by $T' \ll T$ that T' is a gap-subtree of T, i.e. that T' is a proper subtree of T and that $\min\{l(c) : c \in [root(T), root(T')]\}$ equals either qT' or qT. Sequences of trees may not be defined everywhere, so as to increase readability in the following and simplify notation. Accordingly, a sequence s is thus a function from an infinite subset Ds of $\mathbb{N}$ to the set of labelled trees. A sequence s is said to be regular if $qs(i) \leq qs(j)$ whenever $i < j$, $i, j \in Ds$.

Let Bad be the set of all bad sequences, RegBad be the set of all regular bad sequences. As these sets do not actually exists in second order arithmetic, in the following they are thus to be understood as statements about the existence of a bad sequence or about all bad sequences. For h a sequence let Sub(h) be the set (which again, does not exist in second order arithmetic) of all sequences s such that $\forall i \in Ds(s(i) \ll h(i))$.

Our metamathematical analysis of ordinal labelled trees now proceeds in two parts. First we show the upper bounds on the proof-theoretic strength of GKT$_\alpha$ by augmenting the original proof of Kriz (1989) so that it works in a theory of transfinitely iterated Π^1_1-comprehension, which will be the subject of the next section. Then we use techniques from Simpson (1985) to establish that those bounds are sharp in certain cases, which is done in the second section.

5.1 Provability of GKT$_\alpha$

This section will concern itself with the upper bounds of GKT$_\alpha$ and GKT, by analyzing and augmenting the original proof of these principles. The idea of the original proof in Kriz (1989) is as follows: Assuming that a bad sequence exists, we want to find a minimal (with regard to the gap-subtree relation) regular bad sequence h of trees, i.e. a sequence which satisfies the condition that there is no bad sequence s with $Ds \subseteq Dh$ and $s(i) \ll h(i)$ for all $i \in Ds$. The following lemma then corresponds to 1.7 of Kriz (1989).

Lemma 5.1.1. If we can construct a minimal sequence h as above from the assumption that a bad sequence h^* exists, then GKT holds.

Proof. Let S be the set of trees that can be obtained by removing the root from some tree in h. Observe that any such tree obtained from some $h(i)$ is trivially a gap-subtree of $h(i)$. Since h is minimal S must be well-quasi-ordered with respect to embeddability. By Higman's theorem $S^{<\omega}$ is well-quasi-ordered, too. This means that we can find $i < j$ such that $\langle T_i^1, \dots, T_i^{n_i} \rangle \leq \langle T_j^1, \dots, T_j^{n_j} \rangle$, where $T_k^1, \dots, T_k^{n_k}$ (in order) are the trees obtained by deleting the root from $h(k)$. We can then map $root(h(i)) \mapsto root(h(j))$ and the immediate subtrees as given by the above embeddings to obtain an embedding of $h(i)$ into $h(j)$. Note that the gap condition is satisfied by this embedding because h is regular. But $h(i) \leq h(j)$ is in contradiction to h being bad, so there can be no bad sequence h^* in the first place, i.e. GKT holds. $\square$

We remark that the above argument holds for trees with vertices that are additionally labelled from a well-quasi-order — which will be used in section 5.3 — by choosing an infinite subsequence that has increasing well-quasi-order labels on the roots. The following arguments work unchanged for such labelled trees as well, since Facts 1 and 2 below still hold in that case. We now need to construct a minimal bad sequence h as above.

The construction in Kriz (1989) produces an $\aleph_0 \times \aleph_1$-sized tableau where each row is a regular bad sequence of gap-subtrees of the previous rows; in successor rows a prefix of the previous sequence has to be added to the next sequence so that the tableau is not empty after limit steps. The root-labels of this tableau are weakly increasing in the rows and columns. The problem is that the tableau is uncountable and hence too big to be represented in second order arithmetic.[1] We thus want to make the construction terminate earlier; for this we need that the root-labels are strictly increasing in the columns (if not part of an added prefix). Then at each step one possible root-label is ”eliminated”, meaning that if our trees are labelled from α, after α steps there can be no eligible trees left to construct the next row and our construction terminates.

Like Dershowitz and Tzameret (2003) we present the construction as an algorithm to improve readability. Note however that the algorithm presented there is not suited for our purpose since its root-labels are not increasing in the columns due to the limit step.

Our algorithm works as follows. The subroutine lex is used to construct the rows of our tableau. It chooses a subsequence of its input h that in some sense has lexicographically minimal labels: There is no sequence b that has the same initial segment as lex(h) (say, up to k) but has lower root-label on its next value, i.e. $qb(k_b) < qh(k)$, where k_b is the least number greater than k on which b is defined. Doing only this would result in the same construction as in Kriz (1989) (aside from the limit step which in our version has an extra application of lex to speed up our algorithm). However, among all such sequences we also always take one that is defined as early as possible, so that in the successor step ”small”

[1] However, changing the argument slightly one can carry out the proof in Π^1_2-CA_0.

root-labels cannot persist in the appended prefix of the $\alpha+1$-st row $h_{\alpha+1}$ when they could have been removed. Further, we also always take the minimal gap-subtree possible so that we cannot have the same root-label in subsequent steps of the construction.

In short, there are three measures that make sure that each stage the least root-label in a bad sequence is eliminated:

1) We choose a subsequence that has minimal possible root-labels,

2) we choose such a subsequence that is always defined as early as possible, and

3) we choose a sequence whose elements are not a gap-subtrees of any other sequence satisfying the two conditions above.

For i a natural number and s a sequence, let i_s be the first element of Ds that is greater than or equal to i. Then lex is defined as follows.

lex(h)	K:= RegBad $\cap$ Sub(h)
	
	$i := 0$
	
	while $i < \infty$ **do**
	
	    $M :=$ arg min$\{qs(i_s) : s \in K\}$
	
	    $t :\in$ arg min$\{i_s : s \in M\}$
	
	    $K := \{s \in K : i_s = i_t \wedge qs(i_s) = qt(i_t)\}$
	
	    $K := \{s \in K : \neg\exists s' \in K(s'(i_t) \ll s(i_t))\}$
	
	    $i := i_t$
	
	$k :\in K$
	
	return k

Let lex_0 be the same algorithm as lex, except that it does not take any input and the first line is replaced by $K :=$ RegBad.

The construction of a minimal bad sequence is as follows, where the limit in stage λ is meant to be taken pointwise:

Stage 0	$h_0 := g_0 := \text{lex}_0$
Stage $\alpha + 1$	**if** Bad $\cap$ Sub$(h_\alpha) = \emptyset$ **return** h_α $g_{\alpha+1} :=$ lex(h_α) **for** $i = 0, \ldots, \infty$ **do** $h_{\alpha+1}(i) := \begin{cases} h_\alpha(i), & \text{if } i < Dg_{\alpha+1} \\ g_{\alpha+1}(i), & \text{otherwise} \end{cases}$
Stage λ	$h := \lim_{\gamma \to \lambda} h_\gamma$ **if** Bad $\cap$ Sub$(h) = \emptyset$ **return** h $h_\lambda := g_\lambda :=$ lex(h)

Note that in our construction h_α corresponds to a_α, $g_{\alpha+1}$ to a'_α (both defined on pages 220 and 221 of Kriz (1989)) and lex to an application of lemma 2.2 of Kriz (1989).

For the correctness of the construction, note that by Ramsey's theorem Bad $\cap$ Sub$(h) = \emptyset$ if and only if RegBad $\cap$ Sub$(h) = \emptyset$, so in the first line of lex K is non-empty. That the definition of h in the limit stage is possible follows because if $\langle h_\gamma(i) : \gamma < \lambda \rangle$ would not eventually become stationary for each i, we would produce an infinite sequence of pairwise distinct proper subtrees of the finite tree $h_0(i)$, a contradiction.

In order to prove that our construction works as intended, we need the following two facts from lemma 1.6 of Kriz (1989).

Fact 1. If $T_1 \leq T_2$, $T_2 \ll T_3$ and $qT_1 \leq qT_3$ then also $T_1 \leq T_3$.

Fact 2. If $T_1 \ll \ldots \ll T_n$ with $\min\{qT_i : 1 \leq i \leq n\} = qT_j$ where $j = 1$ or $j = n$, then also $T_1 \ll T_n$.

The following theorem is the main step in showing that our construction of a minimal bad sequence terminates as quickly as possible, by demonstrating that the root-labels of the trees are strictly increasing in the columns wherever possible. To prove the theorem we are technically carrying out a simultaneous transfinite induction on the statement of the theorem and the following two corollaries. Since we only need the corollaries in the limit step and they would clog up the proof unnecessarily otherwise, they are stated separately instead.

Theorem 5.1.2. Let $\alpha > 0$ be any ordinal and h be the sequence from which g_α was obtained by an application of lex. Then for all $i \in Dg_\alpha$, $qh(i) < qg_\alpha(i)$.

Proof. The proof is by induction on α. Assume for a contradiction that there is an i such that $qh(i) \geq qg_\alpha(i)$; let k be the least such i.

Assume first that α is a successor, $\alpha = \beta + 1$, hence $h = h_\beta$. There are two cases:

- If $\beta = 0$, define h' by
$$h' := h \restriction k \cup g_\alpha \restriction (\mathbb{N} \setminus k).$$

  Then h' is bad: Let $i < j$, we want to show $h'(i) \not\leq h'(j)$. Since h and g_α are themselves bad the only interesting case is $i < k \leq j$. Assume for a contradiction that $h'(i) \leq h'(j)$. Then we have $h(i) = h'(i) \leq h'(j) = g_\alpha(j) \ll h(j)$ and $qh(i) \leq qh(j)$ by the regularity of h. But then $h(i) \leq h(j)$ by Fact 1, a contradiction to h being bad.

  h' is regular: Again, let $i < k \leq j$. Then $qh'(i) = qh(i) \leq qg_\alpha(i) \leq qg_\alpha(j) = h'(j)$, where the first inequality holds by the minimality of k.

The existence of h' stands in immediate contradiction to h being the output of lex_0 by either the second (in case $\min Dh > k$) or third (in case $\min Dh \leq k$) conditions on our sequence.

- If $\beta > 0$, let h^* be the sequence from which g_β was obtained by an application of lex. Define g'_β by

$$g'_\beta := g_\beta \upharpoonright k \cup g_\alpha \upharpoonright (\mathbb{N} \setminus k).$$

  Assume first that $\min Dg_\beta \leq k$.

  g'_β is bad: Let $i < k \leq j$, and assume $g'_\beta(i) \leq g'_\beta(j)$. Then we have $g_\beta(i) \leq g_\alpha(j) \ll g_\beta(j)$ and $qg_\beta(i) \leq qg_\beta(j)$, whence $g_\beta(i) \leq g_\beta(j)$ by Fact 1, a contradiction.

  g'_β is regular: If $i < k \leq j$ then $qg'_\beta(i) = qg_\beta(i) \leq qg_\alpha(i) \leq qg_\alpha(j)$, where the first inequality holds once again by minimality of k.

  For all $i \in Dg'_\beta$, $g'_\beta(i) \ll h^*(i)$: This is clear for $i < k$, so let $k \leq i$. Then we have $g'_\beta(i) = g_\alpha(i) \ll g_\beta(i) \ll h^*(i)$. By the induction hypothesis $qg_\beta(i) > qh^*(i)$, so $\min\{qg_\alpha(i), qg_\beta(i), qh^*(i)\} \in \{qg_\alpha(i), qh^*(i)\}$. By Fact 2 this implies $g_\alpha(i) \ll h^*(i)$.

  Because of $g'_\beta(k) \ll g_\beta(k)$, $qg'_\beta(k) \leq qg_\beta(k)$, the above are in contradiction to g_β being the output of lex(h^*) by the third condition on the construction of lex.

  Assume now that $\min Dg_\beta > k$. Similar to the case above we get that $g_\alpha \in$ RegBad $\cap$ Sub(h^*), again a contradiction to g_β being the output of lex(h^*) by our second condition on the construction of lex.

Assume now that α is a limit ordinal, then $h = \lim_{\gamma \to \alpha} h_\gamma$. Let γ be the stage after which $\langle h_\gamma(k) : \gamma < \alpha \rangle$ becomes stationary (i.e. so that $h_\gamma(k) = h(k)$). Then $\min Dg_{\gamma+1} > k$. We want to show $g_\alpha(i) \ll h_\gamma(i)$ for all $i \in Dg_\alpha$; then $qg_\alpha(\min Dg_\alpha) \leq qg_\alpha(k) \leq$

$qh(k) = qh_{\gamma+1}(k) \le qg_{\gamma+1}(\min Dg_{\gamma+1})$, a contradiction to the choice of $g_{\gamma+1}$ as the output of lex by the third condition on lex.

So let $i \in Dg_\alpha$. Because the theorem is true by the induction hypothesis for all $\xi < \alpha$, we may apply the following corollary 5.1.3 (recall that we are performing a simultaneous induction on these corollaries) to get $qh_\gamma(i) \le qh(i)$ for all $i \in Dh$. Corollary 5.1.4 after that gives us $h_\gamma(i) \gg h(i) \vee h_\gamma(i) = h(i)$ for all $i \in Dh$. An application of Fact 2 together with $Dg_\alpha \subseteq Dh$ therefore gives $h_\gamma(i) \gg g_\alpha(i)$ for all $i \in Dg$.

This completes the proof.

□

The theorem implies the following for every α, using transfinite induction:

Corollary 5.1.3. For all $\beta < \alpha$, $i \in Dh_\alpha$: $qh_\beta(i) \le qh_\alpha(i)$ and $qh_\beta(i) < qg_\alpha(i)$.

Together with Fact 2 and another transfinite induction this gives us for every α:

Corollary 5.1.4. For all $\beta < \alpha$, $i \in Dh_\alpha$: $h_\beta(i) = h_\alpha(i)$ or $h_\beta(i) \gg h_\alpha(i)$. In particular, for all $i \in Dg_\alpha$: $h_\beta(i) \gg g_\alpha(i)$.

We now get the following lemma, which gives us that if the labels on the trees are all strictly less than α, a minimal bad sequence can be found after α steps.

Lemma 5.1.5. For all $i \in Dg_\alpha$, $qg_\alpha(i) \ge \alpha$.

Proof. By transfinite induction on α.

For $\alpha = 0$ the statement is trivial.

Let $\alpha > 0$ and assume the induction hypothesis for all $\beta < \alpha$. Assume for a contradiction that $qg_\alpha(i) = \beta$ for some $\beta < \alpha$. By induction hypothesis $qg_\beta(j) \ge \beta$ for all $j \in Dg_\beta$, so if $i \ge \min Dg_\beta$ we would have $qg_\alpha(i) \le qg_\beta(i)$, a contradiction. So $i < \min Dg_\beta$. Let h be the sequence from which g_β was obtained by lex (if $\beta = 0$ we are immediately done by the definition of lex_0), then $g_\alpha(j) \ll h(j)$ for all $j \in Dg_\alpha$, but $qg_\alpha(\min Dg_\alpha) \leq \beta \leq qg_\beta(\min Dg_\beta)$, a contradiction to g_β being the output of lex(h).

□

We can thus make sure that after stage α there are no root-labels $< \alpha$ remaining. For α a limit (or 0) this follows from the lemma since then $h_\alpha = g_\alpha$; for α a successor we change the construction so that at stage α we simply set $h_\alpha := g_\alpha$. Hence, if we only consider trees with labels from α our construction stops after stage α. Since the output of lex (and lex_0) is just a lexicographically minimal bad sequence, it can be constructed via Π^1_1-comprehension (see, for example, Marcone (1996) or chapter 2 of this thesis). This means that $\alpha + 1$ iterated Π^1_1-comprehensions are enough to prove GKT_α, i.e. $(\Pi^1_1 - \text{TR})_{\alpha+1} \vdash \text{GKT}_\alpha$. For α limit, this means $(\Pi^1_1 - \text{TR})_{<\alpha} \vdash \text{GKT}_{<\alpha}$, where $\text{GKT}_{<\alpha} := \forall\beta < \alpha(\text{GKT}_\beta)$. Also, $\Pi^1_1 - \text{TR}_0 \vdash \text{GKT}$.

5.2 Lower bounds on ordinal labelled trees

In this section we will prove lower bounds on GKT_α and GKT. To do this, we proceed similar to Simpson (1985). For this we need an ordinal representation system for the proof-theoretic ordinals of the systems $(\Pi^1_1 - \text{TR})_{<\alpha}$ and a function relating these ordinals to labelled trees. Let $\Omega_0 := 0$, $\Omega_n = \aleph_n$ for $n \in \omega \setminus \{0\}$ and $\Omega_\xi := \aleph_{\xi+1}$ for $\xi \geq \omega$. Let φ denote the Veblen function.

Definition. $C_\xi(\alpha)$ and $\psi_\xi(\alpha)$ are defined as follows:

- $[0, \Omega_\xi] \subseteq C_\xi(\alpha)$,
- $\beta, \gamma \in C_\xi(\alpha) \Longrightarrow \beta + \gamma, \varphi\beta\gamma \in C_\xi(\alpha)$,
- $\eta \in C_\xi(\alpha) \Longrightarrow \Omega_\eta \in C_\xi(\alpha)$,

- $\gamma < \alpha, \beta, \gamma \in C_\xi(\alpha) \Longrightarrow \psi_\beta(\gamma) \in C_\xi(\alpha)$,
- $\psi_\xi(\alpha) := \min\{\beta : \beta \notin C_\xi(\alpha)\}$.

Definition. Normal forms for ordinals in this system are:

- $\eta =_{\text{NF}} \eta_1 + \ldots + \eta_n \Longleftrightarrow \eta_1 = \omega^{\xi_1} \wedge \ldots \wedge \eta_n = \omega^{\xi_n} \wedge \eta_1 \geq \ldots \geq \eta_n$ for some $\xi_1, \ldots, \xi_n$.
- $\eta =_{\text{NF}} \varphi\alpha\beta \Longleftrightarrow \alpha, \beta < \varphi\alpha\beta$.
- $\eta =_{\text{NF}} \psi_\mu(\gamma) \Longleftrightarrow \gamma \in C_\mu(\gamma)$.

The ordinal notation system obtained in this manner is then an ordinal notation system for $|\Pi^1_1 - \text{TR}_0|$, and only allowing Ω_ξ and ψ_ξ for $\xi < \alpha$, α limit, yields an ordinal notation system for $|(\Pi^1_1 - \text{TR}_0)_{<\alpha}|$, see Rathjen (1989). We also need the following facts:

Fact 3. $\eta =_{\text{NF}} \eta_1 + \ldots + \eta_n, \eta \in C_\xi(\alpha) \Longrightarrow \eta_i \in C_\xi(\alpha)$ for all $i = 1, \ldots, n$.

Fact 4. $\eta =_{\text{NF}} \varphi\alpha\beta, \eta \in C_\xi(\alpha) \Longrightarrow \alpha, \beta \in C_\xi(\alpha)$.

Fact 5. $\eta =_{\text{NF}} \psi_\mu(\gamma), \eta \in C_\xi(\alpha), \mu \geq \xi \Longrightarrow \gamma \in C_\xi(\alpha)$.

We now define the ordinal $o(T)$ associated with a tree T as follows (note that this definition requires the tree to be labelled "correctly" in the sense of the next lemma; if T has not the required form then $o(T)$ is not defined):

- If $|T| = 1$ define $o(T) = \Omega_{qT}$.

Otherwise let $b_1, \ldots, b_n$ be the children of $root(T)$ in order. Let $\beta_1, \ldots \beta_n$ be the already associated ordinals.

- If $n = 2$, β_1 is an additive principal number, $\beta_1 + \beta_2 =_{\text{NF}} \beta_1 + \eta_1 + \ldots + \eta_n$ where $\beta_2 =_{\text{NF}} \eta_1 + \ldots + \eta_n$, and $qT = l(b_1)$, define $o(T) = \beta_1 + \beta_2$.

- If $n = 3$, $\beta_3 = 0$, $\varphi\beta_1\beta_2 =_{\text{NF}} \varphi\beta_1\beta_2$ and $qT = \max\{l(b_1), l(b_2)\}$, define $o(T) = \varphi\beta_1\beta_2$.
- If $n = 4$, $\beta_2 = \beta_3 = \beta_4 = 0$ and $\psi_{qT}(\beta_1) =_{\text{NF}} \psi_{qT}(\beta_1)$, define $o(T) = \psi_{qT}(\beta_1)$.

Lemma 5.2.1. If $qT = \xi$ then $\Omega_\xi \leq o(T) < \Omega_{\xi+1}$.

Proof. By induction on $|T|$. If $|T| = 1$ then $o(T) = \Omega_\xi$.

Otherwise $o(T)$ has successors $b_1, \ldots, b_n$, $n = 2, 3$ or 4.

- If $o(T)$ has 2 successors then $o(T) = \beta_1 + \beta_2$ with $\Omega_\xi \leq \beta_1 < \Omega_{\xi+1}$ by the induction hypothesis, hence also $\Omega_\xi \leq o(T) < \Omega_{\xi+1}$.
- If $o(T)$ has 3 successors then $o(T) = \varphi\alpha\beta$ with $\alpha, \beta < \Omega_{\xi+1}$ and $\Omega_\xi \leq \alpha$ or $\Omega_\xi \leq \beta$. Now $\alpha, \beta \leq \varphi\alpha\beta$, so $\Omega_\xi \leq o(T)$. On the other hand, if $\alpha, \beta < \Omega_{\xi+1}$ then also $\varphi\alpha\beta < \Omega_{\xi+1}$ (because $\Omega_{\xi+1}$ is regular), hence $o(T) < \Omega_{\xi+1}$.
- If $o(T)$ has 4 successors then $o(T) = \psi_\xi(\alpha)$ and we are done.

$\square$

Lemma 5.2.2. If $f : T_1 \longrightarrow T_2$ is an embedding between ordered trees, $\xi = qT_1$ and $o(T_2) \in C_\xi(\gamma)$, then also $o(T_1) \in C_\xi(\gamma)$ holds, for all γ. Furthermore, $o(T_1) \leq o(T_2)$.

Proof. Let $a = root(T_1)$, $\mu = qT_2$. The proof is by induction on the distance of $root(T_2)$ to $f(a)$.

If $root(T_2) = f(a)$ we need a second induction on $|T_2|$.

- If $|T_2| = 1$ then $o(T_1) = \Omega_\xi \in C_\xi(\gamma)$. Also, $o(T_1) = \Omega_\xi \leq o(T_2)$ by the previous lemma.

(In the following the trivial case $o(T_1) = \Omega_\xi$ is omitted.)

Otherwise let $b_1, \ldots, b_m$, $c_1, \ldots, c_n$ be the successors of a, $root(T_2)$ in order. Let $\beta_1, \ldots, \beta_m$, $\gamma_1, \ldots, \gamma_n$ be the ordinals of the corresponding subtrees. There are three cases.

- If $o(T_2) = \gamma_1 + \gamma_2$ then $o(T_1) = \beta_1 + \beta_2$ and $T_1^{b_1} \leq T_2^{c_1}, T_1^{b_2} \leq T_2^{c_2}$. We also have $\gamma_1, \gamma_2 \in C_\xi(\gamma)$ by Fact 3 and hence by the induction hypothesis (using that $l(b_1) = \xi$) we get $\beta_1 \in C_\xi(\gamma)$. Also, if $l(b_2) < \xi$ then by lemma 5.2.1 $\beta_2 < \Omega_\xi$ so $\beta_2 \in C_\xi(\gamma)$. If $l(b_2) = \xi$ then by the induction hypothesis $\beta_2 \in C_\xi(\gamma)$.

  So we have $\beta_1, \beta_2 \in C_\xi(\gamma)$. Hence also $o(T_1) = \beta_1 + \beta_2 \in C_\xi(\gamma)$.

  By the induction hypothesis $\gamma_1 \geq \beta_1, \gamma_2 \geq \beta_2$, so $o(T_2) = \gamma_1 + \gamma_2 \geq \beta_1 + \beta_2 = o(T_1)$.

- If $o(T_2) = \varphi\gamma_1\gamma_2$ we have again $\gamma_1, \gamma_2 \in C_\xi(\gamma)$. Note that if $m = 2$ we have $\beta_2 > 0$ and hence $T_1^{b_2} \leq T_2^{c_2}$, so in both cases $m = 2$ or $m = 3$ we have $T_1^{b_1} \leq T_2^{c_1}$ and $T_1^{b_2} \leq T_2^{c_2}$. This gives us $\beta_1, \beta_2 \in C_\xi(\gamma)$ as above (again with the case distinction on the label of b_2).

  Hence, no matter whether $m = 2$ or $m = 3$, we get $o(T_1) \in C_\xi(\gamma)$.

  Also, $\varphi\gamma_1\gamma_2 \geq \varphi\beta_1\beta_2$ and $\varphi\gamma_1\gamma_2 \geq \beta_1 + \beta_2$, hence $o(T_2) \geq o(T_1)$.

- If $o(T_2) = \psi_\mu(\gamma_1)$ then $m = 3$ or $m = 4$ and $\beta_2 = 0 \in C_\xi(\gamma)$. Also $\gamma_1 \in C_\xi(\gamma)$ using Fact 5 and the fact that $\mu \geq \xi$ by the gap-condition.

  If $m = 3$ we may assume $T_1^{b_1} \leq T_2^{c_1}$ as otherwise trivially $\beta_1 = 0 \in C_\xi(\gamma)$ and $o(T_1) = \varphi 00 \in C_\xi(\gamma)$. But then $\beta_1 \in C_\xi(\gamma)$ by the induction hypothesis (since $\beta_2 = 0$ and hence b_1 must have label ξ) and hence also $o(T_1) = \varphi\beta_1 0 \in C_\xi(\gamma)$.

  To show $o(T_1) \leq o(T_2)$ we may again assume $l(b_1) = \xi$ or else $\beta_1 = 0$ and $o(T_1) = 1$. If $\mu > \xi$ then $o(T_1) < \Omega_{\xi+1} \leq \psi_\mu(\gamma_1)$. If $\xi = \mu$ then $\gamma_1 \in C_\xi(\gamma_1) = C_\mu(\gamma_1)$, hence by induction hypothesis $\beta_1 \in C_\xi(\gamma_1)$. Hence $o(T_1) = \varphi\beta_1 0 \in C_\xi(\gamma_1)$. But also $o(T_1) < \Omega_{\xi+1}$, so $o(T_1) < \psi_\xi(\gamma_1) = o(T_2)$, since $\psi_\xi(\gamma_1) = \Omega_{\xi+1} \cap C_\xi(\gamma_1)$.

So now assume $m = 4$, i.e. $o(T_1) = \psi_\xi(\beta_1)$ and $\beta_1 \in C_\xi(\beta_1)$. Then $o(T_2) \in C_\xi(\gamma)$, but $o(T_2) \notin C_\mu(\gamma_1) \supseteq C_\xi(\gamma_1)$ (since $o(T_2) = \psi_\mu(\gamma_1)$), hence $\gamma > \gamma_1 \geq \beta_1$, where the second inequality holds by the induction hypothesis. Since $\beta_1 \in C_\xi(\beta_1)$ this implies that also $\beta_1 \in C_\xi(\gamma)$. But if $\beta_1 < \gamma$, $\beta_1 \in C_\xi(\gamma)$ then also $o(T_1) = \psi_\xi(\beta_1) \in C_\xi(\gamma)$. Also, $\psi_\xi(\beta_1) \leq \psi_\mu(\gamma_1)$ since $\beta_1 \leq \gamma_1, \xi \leq \mu$.

We proceed with our initial induction. Assume that $root(T_2) \neq f(a)$ and let $c_1, \ldots, c_n$ be the successors of $root(T_2)$ in order. Let i be such that $T_1 \leq T_2^{c_i}$

- If $n = 2$ then $o(T_2) = \gamma_1 + \gamma_2 \in C_\xi(\gamma)$ and hence $\gamma_1, \gamma_2 \in C_\xi(\gamma)$. By the induction hypothesis (no matter whether $i = 1$ or 2) we get $o(T_1) \in C_\xi(\gamma)$, and obviously $o(T_1) \leq \gamma_i \leq o(T_2)$.

- If $n = 3$ then $o(T_2) = \varphi\gamma_1\gamma_2$ and as above $\gamma_1, \gamma_2 \in C_\xi(\gamma)$. Again as above we get $o(T_1) \in C_\xi(\gamma)$ (if $i = 3$ then $o(T_1) = 0$), and again obviously $o(T_1) \leq \gamma_i \leq o(T_2)$.

- If $n = 4$ then $o(T_2) = \psi_\mu(\gamma_1)$ and $\gamma_1 \in C_\mu(\gamma_1)$. Now, if $o(T_2) = \psi_\mu(\gamma_1) \in C_\xi(\gamma)$, then because of $\mu \geq \xi$ also $\gamma_1 \in C_\xi(\gamma)$, and so we get $o(T_1) \in C_\xi(\gamma)$ by the induction hypothesis (or because $o(T_1) = 0$).

  If $\mu > \xi$ we have $o(T_1) < \Omega_{\xi+1} \leq \psi_\mu(\gamma_1) = o(T_2)$, if $\mu = \xi$ then $\gamma_1 \in C_\xi(\gamma_1)$, hence $o(T_1) \in C_\xi(\gamma_1)$ by the induction hypothesis, hence (because $o(T_1) < \Omega_{\xi+1}$) $o(T_1) < \psi_\xi(\gamma_1) = o(T_2)$.

□

The above lemma gives the desired result that $T_1 \leq T_2 \Rightarrow o(T_1) \leq o(T_2)$. Hence the statement that the ordered trees with labels from α are well-quasi-ordered under gap-embedding implies that the ordinal $\psi_0(\Omega_\alpha)$ is well-founded (provided that the ordinal notation system consisting of $+, \varphi, \psi_\xi, \Omega_\xi$ is sufficient to represent $\psi_0(\Omega_\alpha)$). Together with our upper bounds and the ordinal analysis of Rathjen (1989) this gives $\psi_0(\Omega_\alpha) = |(\Pi^1_1 - \text{TR})_{<\alpha}| = |\text{GKT}_{<\alpha}|$ for $\alpha < \Gamma^\Omega_0$ limit, and $\psi_0(\Gamma^\Omega_0) = |\Pi^1_1 - \text{TR}_0| = |\text{GKT}|$, where Γ^Ω_0 denotes the first fixed point of $\alpha \mapsto \aleph_\alpha$.

5.3 Lower bounds on trees with well-quasi-ordered labels

The previous sections showed that Kruskal theorems for trees labelled from $\omega + 1$ are already stronger than the graph minor theorem. For possibly increasing the lower bounds of the graph minor theorem and related principles, it would be desirable to find some combinatorial statement that has lower bounds above but not too far above for example $\Pi^1_1 - \text{CA}$. This section gives such a combinatorial principle by looking at trees labelled from ω with additional labels from a well-order on the leaves which are independent from the gap-condition.

A suitable ordinal notation system for this endeavour is presented in Rathjen and Thompson (n.d.). Let X be a well-order (interpreted as an ordinal) and for $x \in X$ let $\overline{x} = \Omega_\omega \cdot (1 + x)$. Then $C^X_m(\alpha)$, $m \in \mathbb{N}$, and $\psi_m\alpha, m \in \mathbb{N}$, are defined by induction on α as follows. First, let $C^X_m(\alpha)$ be the least set $C \supseteq \Omega_m \cup \{\Omega_i : i \in \mathbb{N}\} \cup \{\overline{x} : x \in X\}$ so that:

- If $\xi, \eta \in C \cap \Omega_\omega$ then $\xi + \eta \in C$ and $\omega^\xi \in C$,
- if $\alpha \in C \cap \Omega_\omega$ and $x \in X$ then $\overline{x} + \alpha \in C$, and
- if $\gamma \in C \cap \alpha$ then $\psi_n\gamma \in C$, for all $n \in \mathbb{N}$.

Then $\psi_m(\alpha) := \min\{\xi : \xi \notin C^X_m(\alpha)\}$. The following results from Rathjen and Thompson (n.d.) give the necessary facts and intuition about ψ_m.

Lemma 5.3.1 (Lemma 2.3 of Rathjen and Thompson (n.d.))**.** The following hold:

i) $\psi_m\alpha \leq \psi_n\beta$ and $C_m^X(\alpha) \subseteq C_n^X(\beta)$ whenever $m \leq n$ and $\alpha \leq \beta$.

ii) If $\omega^{\xi_1} + \ldots + \omega^{\xi_k}$ is the Cantor normal form of $\xi \in C_m^X(\alpha) \cap \Omega_\omega$, then also $\xi_1, \ldots, \xi_k \in C_m^X(\alpha)$.

iii) For all $x \in X$, if $\overline{x} + \beta \in C_m^X(\alpha)$ then also $\beta \in C_m^X(\alpha)$.

iv) $\Omega_m \leq \psi_m(\alpha)$ and $\psi_m(\alpha) = \Omega_{m+1} \cap C_m^X(\alpha)$.

v) For all $m \in \mathbb{N}$, all α and β with $\alpha \in C_n^X(\alpha)$, $\beta \in C_n^X(\beta)$, the following hold:

$$\psi_m\alpha < \psi_m\beta \Leftrightarrow \alpha < \beta, \text{ and}$$

$$\psi_m\alpha = \psi_m\beta \Leftrightarrow \alpha = \beta$$

As for denoting normal forms, we write $\alpha =_{NF} \omega^{\alpha_1} + \ldots + \omega^{\alpha_k}$ if $\alpha > \alpha_1 \geq \ldots \geq \alpha_k$ and $\alpha =_{NF} \psi_m\alpha_1$ if $\alpha_1 \in C_m^X(\alpha_1)$.

Definition. The ordinal notation system $OT(\Omega_\omega \cdot X)$ is defined by Rathjen and Thompson (n.d.) as follows.

- $\{\Omega_m : m \in \mathbb{N}\} \cup \{\overline{x} : x \in X\} \subseteq OT(\Omega_\omega \cdot X)$.
- For all $x \in X$ and $\beta \in OT(\Omega_\omega \cdot X)$ with $0 < \beta < \Omega_\omega$, also $\overline{x} + \beta \in OT(\Omega_\omega \cdot X)$.
- If $\alpha_1, \ldots, \alpha_n \in OT(\Omega_\omega \cdot X) \cap \Omega_\omega$ and $\alpha =_{NF} \omega^{\alpha_1} + \ldots + \omega^{\alpha_n}$ then also $\alpha \in OT(\Omega_\omega \cdot X)$.
- If $\alpha_1 \in OT(\Omega_\omega \cdot X)$ and $\alpha =_{NF} \psi_m\alpha_1$ then $\alpha \in OT(\Omega_\omega \cdot X)$.

In this ordinal notation system, $|\Pi_1^1 - \mathrm{CA}| = \psi_0(\Omega_\omega \cdot \varepsilon_0)$.

We now need to define the trees to represent this ordinal notation system. The trees considered are vertex-labelled, ordered trees that have two distinct labels per vertex. One is a natural number (under embedding subject to the gap-condition), denoted by $l(v)$ for a vertex v, and the other one from a well-quasi-order, denoted by $l_Q(v)$. This well-quasi-order consists of a well-order W and additional elements $\psi, +, \omega^\cdot$ which are incomparable to all other elements. The well-order W has a special element -1, so that $-1 < w$ for all other $w \in W$, and so that $\psi_n(-1) = \Omega_n$.[2] Interpreting $W \setminus \{-1\}$ as an ordinal, we set $\psi_n(w + \beta) := \psi_n(\overline{w} + \beta)$ for all $w \in W \setminus \{-1\}$ and $\beta \in OT(\Omega_\omega \cdot X) \cap \Omega_\omega$.

An embedding between two such trees is then an embedding f that satisfies the gap-condition with regard to l and so that $l_Q(v) \leq l_Q(f(v))$ for all vertices v. Denote by $\mathrm{GKT}_\omega(W)$ the statement that the trees defined in this manner are well-quasi-ordered under the embedding relation. Since we are only concerned with trees corresponding to ordinal notations, some restrictions are made on the trees used in the following. First, any vertex labelled with $+$ must have exactly two successors, any vertex labelled with ψ or $\omega^\cdot$ must have exactly one successor and any vertex labelled from W must be a leaf or have exactly one successor. Further, only trees corresponding to terms in normal form will be allowed, and if v has label $+$ or $\omega^\cdot$ then its natural number label must be the maximum of its successor labels.

An ordinal $o(T)$ corresponding to a tree T is then defined as follows:

- If $l_Q(root(T)) \in W$ and $root(T)$ has no successor, then $o(T) = \psi_n w$, where $n = qT$ and $w = l_Q(root(T))$.
- If $l_Q(root(T)) \in W \setminus \{-1\}$ and $root(T)$ has one successor v, then $o(T) = \psi_n(w + o(T^v))$, where $n = qT$ and $w = l_Q(root(T))$.
- If $l_Q(root(T)) = +$ let v_1 and v_2 be the successors of $root(T)$, in order. Then $o(T) = o(T^{v_1}) + o(T^{v_2})$.

[2] Of course, -1 is just the least element of W and would more naturally correspond to 0, but considering the definition of $\psi(-1)$ the designation -1 makes more sense to bring the trees and ordinal notation system in accordance.

- If $l_Q(root(T)) = \omega^{\cdot}$ let v be the successor of $root(T)$. Then $o(T) = \omega^{o(T^v)}$.
- If $l_Q(root(T)) = \psi$ let again v be the successor of $root(T)$. Then $o(T) = \psi_n o(T^v)$, where $n = qT$.

Then for every $\alpha \in OT(\Omega_\omega \cdot W)$ there is a labelled tree T with $o(T) = \alpha$. By 5.3.1.iv we also immediately get:

Lemma 5.3.2. For a labelled tree T as above, if $qT = n$ then $\Omega_n \leq o(T) < \Omega_{n+1}$.

The following theorem then establishes the desired lower bound on $\mathrm{GKT}_\omega(W)$. In the following we always treat $\alpha \in W$ as $\overline{\alpha}$, and write $C_m(\alpha)$ for $C_m^W(\alpha)$.

Theorem 5.3.3. If W is a well-ordering and $\mathrm{GKT}_\omega(W)$ holds, then $OT(\Omega_\omega \cdot W)$ is well-ordered.

Proof. We show by induction on $|[f(root(T_1)), root(T_2)]|$ and by subsidiary induction on $|T_1|$ that if $T_1 \leq T_2$ via f then $o(T_1) \leq o(T_2)$, and that if $o(T_2) \in C_k(\gamma)$ with k, γ arbitrary, then $o(T_1) \in C_k(\gamma)$ as well. Let $m := qT_1$, $n := qT_2$, $\alpha := o(T_1)$, $\beta := o(T_2)$.

First assume that $|[f(root(T_1)), root(T_2)]| = 1$, i.e. that $f(root(T_1)) = root(T_2)$. If $|T_1| = 1$ then $l_Q(root(T_1)), l_Q(root(T_2)) \in W$. Let $l_Q(root(T_1)) = \lambda$, $l_Q(root(T_2)) = \tau$. Then $\lambda \leq \tau$ and hence $o(T_1) = \psi_m\lambda \leq \psi_n\tau = o(T_2)$ by 5.3.1.i. If further $o(T_2) \in C_k(\gamma)$, then either $o(T_1) \leq o(T_2) \leq \Omega_k$ and hence $o(T_1) \in C_k(\gamma)$ as well, or $\lambda \leq \tau < \gamma$ and so $o(T_1) \in C_k(\gamma)$ in this case as well.

Now, if $|T_1| > 1$ let a_i and b_i be the successors of $root(T_1)$ and $root(T_2)$, respectively, where $i = 1$ or $i = 1, 2$ depending on $l_Q(root(T_1))$. Let α_i, β_i be the ordinals associated with the corresponding subtrees. Hence we have by the induction hypothesis that $\alpha_i \leq \beta_i$ and $\alpha_i \in C_k(\gamma)$ whenever $\beta_i \in C_k(\gamma)$.

If $l_Q(root(T_1)) \in W$, let λ, τ be as above. Then $o(T_1) = \psi_m(\lambda + \alpha_1) \leq \psi_n(\tau + \beta_1) = o(T_2)$ by 5.3.1.i. If further $o(T_2) \in C_k(\gamma)$, then as above there are two cases. The first is that $o(T_1) \leq o(T_2) \leq \Omega_k$ and hence $o(T_1) \in C_k(\gamma)$ as well, and since this case is identical in all steps of this proof, it will not be covered in the following. The second case is that $\lambda + \alpha_1 \leq \tau + \beta_1 < \gamma$ and $\tau + \beta_1 \in C_k(\gamma)$. Then by 5.3.1.iii $\beta_1 \in C_k(\gamma)$ and so by the induction hypothesis $\alpha_1 \in C_k(\gamma)$. Hence also $o(T_1) = \psi_m(\lambda + \alpha_1) \in C_k(\gamma)$. If $l_Q(root(T_1)) = \psi$ we can argue completely analogously, and if $l_Q(root(T_1)) = +$ or $l_Q(root(T_1)) = \omega^{\cdot}$ then the claims follow immediately from the induction hypothesis and 5.3.1.ii.

Assume now that $|[f(root(T_1)), root(T_2)]| > 1$. If $l_Q(root(T_2)) = +$ or $l_Q(root(T_2)) = \omega^{\cdot}$ the claims follow again immediately by the induction hypothesis and 5.3.1.ii. So assume that $l_Q(root(T_2)) \in W \setminus \{-1\}$ or $l_Q(root(T_2)) = \psi$. Let b_1, β_1 be as above, and assume first that $l_Q(root(T_2)) = \psi$. Then $o(T_2) = \psi_n\beta_1$ and $\beta_1 \in C_n(\beta_1)$ by the normal form condition. If $n > m$ we immediately get $o(T_1) \leq o(T_2)$ by 5.3.2, so we may assume $m = n$. Since also $T_1 \leq T_2^{b_1}$ we get $o(T_1) \in C_n(\beta_1)$ by the induction hypothesis, but then $o(T_1) \leq o(T_2)$ by 5.3.1.iv. For the second claim suppose $o(T_2) \in C_k(\gamma)$, then as above we may assume $\beta_1 \in C_k(\gamma)$, but then we get $o(T_1) \in C_k(\gamma)$ immediately by the induction hypothesis. If $root(T_2) \in W \setminus \{-1\}$ we can argue analogously, the only difference being that we have to use 5.3.1.iii to obtain $\beta_1 \in C_n(\tau + \beta_1)$ and $\beta_1 \in C_k(\gamma)$.

This finishes the induction. Thus $o(T_1) \leq o(T_2)$ whenever $T_1 \leq T_2$, and so $\mathrm{GKT}_\omega(W)$ implies the well-orderedness of $OT(\Omega_\omega \cdot W)$.

□

Since the proof-theoretic ordinal for $\Pi^1_1 - \mathrm{CA}$ is $\psi_0(\Omega_\omega \cdot \varepsilon_0)$, taking $W := \varepsilon_0$ in the above lemma shows that $\mathrm{GKT}_\omega(W)$ can not be provable in $\Pi^1_1 - \mathrm{CA}$. This thus presents a combinatorial theorem that is not provable in $\Pi^1_1 - \mathrm{CA}$ but provable in $(\Pi^1_1 - \mathrm{TR})_\omega$ by the first section of this chapter.

One approach for applying the ideas presented in this section to graph minor theory would be to attempt a similar construction using tree-like graphs with multiple edges ordered

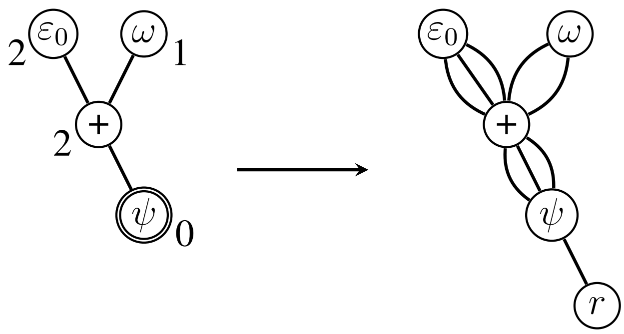

Figure 5.1: Simulating a labelled tree by a graph with multiple edges. The labels from the well-quasi-order are drawn inside the nodes, while the natural number labels of the tree are drawn next to the nodes. The root of the tree is marked by a node with double borders.

by vertex-label preserving immersion (introduced in section 4.17). The corresponding graphs would be defined as follows. To simulate a tree T as considered for the statement GKT$_\omega(W)$, add first one more possible label r for vertices, where r is incomparable to all other elements of $W \cup \{\psi, +, \omega^\cdot\}$. The idea is to encode the natural number labels as multiple edges of the graph constructed, which under an immersion expansion will correspond to disjoint paths. However, since the graphs used in the immersion theorem cannot have roots in the same sense that trees do, a new label r is needed to simulate the root. Thus, to carry out the construction to simulate T, copy the vertex set of T with the natural number labels discarded but the labels from W kept. Then introduce a new vertex and give it label r. Connect the vertex labelled with r with $qT + 1$ parallel edges[3] to the original root of T (one additional edge is needed to ensure that the two vertices are connected), and likewise for any vertex v corresponding to a $v' \in V(T) \setminus \{root(T)\}$ connect v and the vertex corresponding to pv' by $l(v') + 1$ parallel edges. One simple example of such a construction is depicted in figure 5.1.

A graph G constructed in this way corresponds thus to a tree of GKT$_\omega(W)$ by designating

[3]Alternatively, if simple graphs are preferred for some reason, any parallel edge can be subdivided and any vertices so introduced labelled with a new well-quasi-order label $\circ$, say.

the only neighbour of the vertex labelled with r as the root (denoted by v_r), giving the vertices natural number labels corresponding to the amount of parallel edges (minus one) that connect them to the vertex closer to v_r and finally deleting v_r and all parallel edges of G. Under an immersion expansion of G in G', the parallel edges of G would correspond to edge-disjoint paths in G' and thus satisfy a kind of gap-condition. However, aside from v_r mapping to the corresponding vertex v'_r of G' labelled with r, the tree structure that would be respected by an embedding between trees is not preserved under immersion embedding. However, whenever the immersion violates infima or the order of the corresponding trees, the labels of the vertices in the region where the violation occurs are significantly increased due to the number of edge-disjoint paths that have to be present in the corresponding graph G'. Thus, an immersion between two such graphs G and G' would correspond to an embedding between trees that does not have to respect infima or the tree-order (and is thus not an embedding at all), but with a stronger gap condition that requires the summation over labels on paths where such a violation occurs. Thus, even though there is not a perfect match between the graph immersions and tree embeddings considered, due to lemma 5.2.1 it might still be possible that $o(T_1) \leq o(T_2)$ whenever there is an immersion between the corresponding graphs G_1 and G_2. However, when attempting such a proof, additional induction hypotheses of the form "if $o(T_2) \in C_k(\gamma)$ with k, γ arbitrary, then $o(T_1) \in C_k(\gamma)$" as in theorem 5.3.3 can not be applied (as the corresponding embedding is not order-preserving), and thus the question whether this or a similar construction is suitable for increasing the lower bounds on the immersion theorem (with quasi-ordered labels allowed) is open.

Finally, it should be noted that this problem would disappear if directed edges and immersions could be used in the construction above, so that the edges are directed away from the vertex labelled with r; this would make the corresponding tree embedding order preserving. That the immersion theorem holds for such directed graphs can be seen since it is implied by $\mathrm{GKT}_\omega(W)$, which however has much higher upper bounds than the ordinary immersion theorem for undirected graphs. It should further be noted that the immersion theorem is known to not extend to directed graphs in general, and that at this point the only major positive result for partial classes of directed graphs seems to be for directed tournaments, see Chudnovsky and Seymour (2011).

To summarize, this final chapter presented an alternative metamathematical analysis of the combinatorial principles GKT and GKT_α, thus illustrating that proof methods beyond $\Pi^1_1{-}\mathrm{CA}_0{+}\Pi^1_2{-}\mathrm{BI}{+}\Pi^1_3{-}\mathrm{IND}$ are used in currently existing proofs of even GKT_ω, and that $\mathrm{GKT}_{\omega+1}$ is already stronger than any theorems considered in the Graph Minors series. The third section of this chapter introduced a new combinatorial principle $\mathrm{GKT}_\omega(W)$ which was proved to be stronger than $\Pi^1_1 - \mathrm{CA}$, and presented a possible option to employ a similar construction to raise the upper bounds of the immersion theorem with vertex labels in the future.

Chapter 6

Conclusion

This thesis investigated proof-theoretic aspects of theorems of the Graph Minors series and well-quasi-ordering theorems of labelled trees. Results postulated in Friedman, Robertson, and Seymour (1987) — namely that the bounded graph minor theorem is equivalent to extended Kruskal's theorem, Vázsonyi's conjecture $tGMT_{\leq 3}$ and thus the planar graph minor theorem — were confirmed rigorously. It was then shown that the bounded graph minor theorem is also equivalent to the graph minor theorem for graphs drawn on an arbitrary, fixed surface, and to the statement that this surface graph minor theorem holds uniformly for all surfaces, which altogether established the following equivalences:

$$EKT \leftrightarrow bGMT \leftrightarrow pGMT \leftrightarrow tGMT_{\leq 3} \leftrightarrow \Sigma - GMT \leftrightarrow \forall\Sigma - GMT.$$

The strongest result of this thesis is that 1.6 of Graph Minors XXIII, a generalized version of the graph minor theorem for hypergraphs with labels from a well-quasi-order on the edges and an ordering on the endpoints of edges with at most a fixed number of vertices allowed, is provable in $\Pi^1_1 - \mathrm{CA}_0 + \Pi^1_2 - \mathrm{BI} + \Pi^1_3 - \mathrm{IND}$. This gives the same upper bound for the graph minor theorem, more generally the graph minor theorem with well-quasi-ordered labels on vertices and edges and directed edges allowed, and for the immersion theorem with well-quasi-ordered labels on the vertices allowed. The proof-theoretic strength of these theorems thus lies in the interval between $\Pi^1_1-\mathrm{CA}_0+\mathrm{RFN}_{\Pi^1_1}(\Pi^1_1-\mathrm{CA}_0)$ and $\Pi^1_1-\mathrm{CA}_0+\Pi^1_2-\mathrm{BI}+\Pi^1_3-\mathrm{IND}$. The corresponding theorems about Ω-patchworks from the Graph Minors series consequently also fall in this range, which are (9.1) of Graph Minors IV, 6.7 of Graph Minors XVIII, 7.3, 10.4 and 11.2 from Graph Minors XX, and 2.1 of Graph Minors XXIII. The well-quasi-ordering theorem about patchworks of bounded tree-width, (9.1) of Graph Minors IV, was shown to be provable even in $\Pi^1_1-\mathrm{CA}_0+\Pi^1_3-\mathrm{IND}$. Further, a simpler and more refined analysis of the proof-theoretic strength of generalized Kruskal's theorem for trees labelled with arbitrary ordinals under gap-condition was given. This was extended to an analysis of the corresponding theorem for trees with well-quasi-ordered labels on the vertices. The main result here is that the generalized Kruskal's theorem for trees labelled from ω and additional well-quasi-ordered labels on the vertices allowed is not provable in $\Pi^1_1-\mathrm{CA}$, but provable in $(\Pi^1_1-\mathrm{TR})_\omega$.

The most obvious further work arising from these investigations would be to determine the exact proof-theoretic strength of the graph minor theorem and related theorems mentioned above. It is somewhat unlikely that the Π^1_3-induction in Graph Minors IV is really needed for the proofs, and so one possibly relatively easy reduction in the upper bound would be to remove this induction, for which a broadly sketched strategy is outlined in section 2.2. Removing the Π^1_2-bar induction could be attempted by replacing the well-quasi-orders used by the classes of graphs that these well-quasi-orders are meant to encode. Strategies for possibly raising the lower bounds on the graph minor theorem are harder to devise; however as made precise in section 5.3, the immersion theorem with well-quasi-ordered labels might be used to conduct an ordinal analysis similar to that for ordinal labelled trees with well-quasi-ordered labels. This would raise the lower bounds of the immersion theorem with well-quasi-ordered labels, and thus also that of the most general version of the graph minor theorem (1.6 of Graph Minors XXIII), above $\Pi^1_1-\mathrm{CA}$. In general, it would be interesting to investigate whether versions of the theorems above with well-quasi-ordered labels admitted are proof-theoretically stronger than those without such labels. It seems plausible that they might be equal for theorems that have proof-theoretic strength in the vicinity of $\Psi_0(\Omega_\omega)$, like the bounded graph minor theorem and (9.1) of Graph Minors IV, since any additional strength that an instance of these theorems gains would only apply to some segment between some $\Psi_0(\Omega_n)$ and $\Psi_0(\Omega_{n+1})$, and would be nullified in going to the limit towards $\Psi_0(\Omega_\omega)$. Other theorems which may allow a more uniform correspondence to the ordinal notation system for $\Psi_0(\Omega_\omega)$, such as the immersion conjecture, might however be strictly stronger in the forms with well-quasi-ordered labels. Finally, a hypothetical "reification" similar to Rathjen and Weiermann (1993) might be achieved by considering a canonical tree-decomposition, like for example those of Carmesin, Diestel, Hamann and Hundertmark (2016a, 2016b), and associating it to a labelled tree. There are some immediate problems with this approach however, firstly that it is not clear that such an association would be injective, secondly that the tree-decompositions produce only underlying unrooted trees, thirdly that it is not clear whether an embedding between the trees would induce a minor relation between the corresponding graphs, and lastly that those canonical tree-decompositions depend on a prior arbitrarily chosen number k.

To conclude, although this thesis has made some progress in establishing proof-theoretic bounds on the graph minor theorem and related theorems of the Graph Minors series, there is still a lot to investigate in the metamathematics of graph minor theory.

Bibliography

Carmesin, Johannes et al. (2016a). “Canonical tree-decompositions of finite graphs I. Existence and algorithms”. In: *Journal of Combinatorial Theory, Series B* 116, pp. 1–24.

– (2016b). “Canonical tree-decompositions of finite graphs II. Essential parts”. In: *Journal of Combinatorial Theory, Series B* 118, pp. 268–283.

Chudnovsky, Maria and Paul Seymour (2011). “A well-quasi-order for tournaments”. In: *Journal of Combinatorial Theory, Series B* 101, pp. 47–53.

Dershowitz, Nachum and Iddo Tzameret (2003). “Gap Embedding for Well-Quasi-Orderings”. In: *Electronic Notes in Theoretical Computer Science* 84, pp. 80–90.

Diestel, Reinhard (2017). *Graph Theory*. 5th ed. Springer.

Erdős, Paul and Lajos Pósa (1965). “On independent circuits contained in a graph”. In: *Canadian Journal of Mathematics* 17, pp. 347–352.

Erickson, Jeff and Kim Whittlesey (2013). “Transforming Curvers on Surfaces Redux”. In: *SODA ’13 Proceedings of the Twenty-Fourth Annual ACM-SIAM Symposium on Discrete Algorithms*, pp. 1646–1655.

Forster, Thomas (2003). “Better-quasi-orderings and coinduction”. In: *Theoretical Computer Science* 309, pp. 111–123.

Friedman, Harvey (1974). “Some Systems of Second Order Arithmetic and Their Use”. In: *Proceedings of the International Congress of Mathematicians*. Ed. by Ralph James. Vol. 1. American Mathematical Society, pp. 235–242.

Friedman, Harvey, Neil Robertson, and Paul Seymour (1987). "Metamathematics of the Graph Minor Theorem". In: *Logic and Combinatorics*. Ed. by Stephen Simpson. Contemporary Mathematics 65. American Mathematical Society, pp. 229–261.

Gordeev, Lev (1990). "Generalizations of the Kruskal-Friedman Theorems". In: *The Journal of Symbolic Logic* 55.1, pp. 157–181.

– (1993). *Strong Well-Quasi-Ordering Tree Theorems*. URL: http://www.wsi.uni-tuebingen.de/lehrstuehle/logik-und-sprachtheorie/personen/ehemalige-und-sonstige-mitarbeiter/pd-dr-lev-gordeev/publications.html (visited on 08/05/2016).

Harrington, Leo and Jeff Paris (1977). "A Mathematical Incompleteness in Peano Arithmetic". In: *Handbook of Mathematical Logic*. Ed. by Jon Barwise. Studies in Logic and the Foundations of Mathematics. Elsevier Science Publishers (North Holland), pp. 87–117.

Kawarabayashi, Ken-ichi and Bojan Mohar (2007). "Some Recent Progress and Applications in Graph Minor Theory". In: *Graphs and Combinatorics* 23.1, pp. 1–46.

Kriz, Igor (1989). "Well-Quasiordering finite Trees with Gap-Condition. Proof of Harvey Friedman's conjecture". In: *Annals of Mathematics* 130.1, pp. 215–226.

Lefschetz, Solomon (1954). *Introduction to Topology*. Princeton University Press.

Marcone, Alberto (1996). "On the logical strength of Nash-Williams' theorem on transfinite sequences". In: *Logic: From Foundations to Applications: European Logic Colloquium*. Clarendon Press, pp. 327–351.

Mohar, Bojan and Carsten Thomassen (2001). *Graphs on Surfaces*. The John Hopkins University Press.

Rathjen, Michael (1989). "Untersuchungen zu Teilsystemen der Zahlentheorie zweiter Stufe und der Mengenlehre mit einer zwischen Δ^1_2-CA und Δ^1_2-CA+BI liegenden Beweisstärke". PhD thesis. Westfälische Wilhelms-Universität Münster.

– (1999). "The Realm of Ordinal Analysis". In: *Sets and Proofs*. Ed. by S. Barry Cooper and John Truss. Cambridge University Press, pp. 219–279.

Rathjen, Michael and Ian Alexander Thompson (n.d.). "Well ordering principles, ω-models and Π^1_1-comprehension". forthcoming.

Rathjen, Michael and Andreas Weiermann (1993). "Proof-theoretic investigations on Kruskal's theorem". In: *Annals of Pure and Applied Logic* 60.1, pp. 49–88.

Robertson, Neil and Paul Seymour (1983). "Graph Minors. I. Excluding a Forest". In: *Journal of Combinatorial Theory, Series B* 35.1, pp. 39–61.

– (1984). "Graph Minors. III. Planar Tree-Width". In: *Journal of Combinatorial Theory, Series B* 36.1, pp. 49–64.

– (1986a). "Graph Minors. V. Excluding a Planar Graph". In: *Journal of Combinatorial Theory, Series B* 41.1, pp. 92–114.

– (1986b). "Graph Minors. VI. Disjoint Paths across a Disc". In: *Journal of Combinatorial Theory, Series B* 41.1, pp. 115–138.

– (1988). "Graph Minors. VII. Disjoint Paths on a Surface". In: *Journal of Combinatorial Theory, Series B* 45.1, pp. 212–254.

– (1990a). "Graph Minors. IV. Tree-Width and Well-Quasi-Ordering". In: *Journal of Combinatorial Theory, Series B* 48.1, pp. 227–254.

– (1990b). "Graph Minors. IX. Disjoint Crossed Paths". In: *Journal of Combinatorial Theory, Series B* 49.1, pp. 40–77.

– (1990c). "Graph Minors. VIII. A Kuratowski Theorem for General Surfaces". In: *Journal of Combinatorial Theory, Series B* 48.1, pp. 255–288.

– (1991). "Graph Minors. X. Obstructions to Tree-Decomposition". In: *Journal of Combinatorial Theory, Series B* 52.1, pp. 153–190.

– (1994). "Graph Minors. XI. Circuits on a Surface". In: *Journal of Combinatorial Theory, Series B* 60.1, pp. 72–106.

– (1995a). "Graph Minors. XII. Distance on a Surface". In: *Journal of Combinatorial Theory, Series B* 64.1, pp. 240–272.

– (1995b). "Graph Minors. XIV. Extending an Embedding". In: *Journal of Combinatorial Theory, Series B* 65.1, pp. 23–50.

Robertson, Neil and Paul Seymour (1996). “Graph Minors. XV. Giant Steps”. In: *Journal of Combinatorial Theory, Series B* 68.1, pp. 112–148.

– (1999). “Graph Minors. XVII. Taming a Vortex”. In: *Journal of Combinatorial Theory, Series B* 77.1, pp. 162–210.

– (2003a). “Graph Minors. XVI. Excluding a non-planar graph”. In: *Journal of Combinatorial Theory, Series B* 89.1, pp. 43–76.

– (2003b). “Graph Minors. XVIII. Tree-decompositions and well-quasi-ordering”. In: *Journal of Combinatorial Theory, Series B* 89.1, pp. 77–108.

– (2004a). “Graph Minors. XIX. Well-quasi-ordering on a surface”. In: *Journal of Combinatorial Theory, Series B* 90.1, pp. 325–385.

– (2004b). “Graph Minors. XX. Wagner’s conjecture”. In: *Journal of Combinatorial Theory, Series B* 92.1, pp. 325–357.

– (2010). “Graph Minors XXIII. Nash-William’s immersion conjecture”. In: *Journal of Combinatorial Theory, Series B* 100.1, pp. 181–205.

Schmidt, Diana (1979). *Well-partial orderings and their maximal order types*. Habilitationsschrift.

Simpson, Stephen (1985). “Nonprovability of certain combinatorial properties of finite trees”. In: *Harvey Friedman’s Research on the Foundations of Mathematics*. Studies in Logic and the Foundations of Mathematics. Elsevier Science Publishers (North Holland), pp. 87–117.

– (2009). *Subsystems of Second Order Arithmetic*. 3rd ed. Cambridge University Press.

Van der Meeren, Jeroen, Michael Rathjen, and Andreas Weiermann (2015). “Proof-theoretic investigations on Kruskal’s theorem”. In: *Archive for Mathematical Logic* 54.1-2, pp. 193–230.

Whitney, Hassler (1932). “Congruent Graphs and the Connectivity of Graphs”. In: *American Journal of Mathematics* 54.1, pp. 150–168.