Finiteness properties of fibre products

Benno Kuckuck

Wolfson College

University of Oxford

A thesis presented for the degree of Doctor of Philosophy

Michaelmas Term, 2012

Abstract

A group Γ is of type F_n for some $n \geq 1$ if it has a classifying complex with finite n-skeleton. These properties generalise the classical notions of finite generation and finite presentability.

We investigate the higher finiteness properties for fibre products of groups. The central conjecture is this: If $N_1 \hookrightarrow \Gamma_1 \overset{\pi_1}{\twoheadrightarrow} Q$ and $N_2 \hookrightarrow \Gamma_2 \overset{\pi_2}{\twoheadrightarrow} Q$ are short exact sequences with N_1 of type F_n, Γ_1 and Γ_2 of type F_{n+1} and Q of type F_{n+2}, then the fibre product $P = \{(\gamma_1, \gamma_2) \in \Gamma_1 \times \Gamma_2 \mid \pi_1(\gamma_1) = \pi_2(\gamma_2)\}$ is of type F_{n+1}. This is called the n-$(n+1)$-$(n+2)$ Conjecture.

The fibre product is the basic construction that all subgroups of direct products can be built from inductively. The conjecture therefore has implications for the finiteness properties of these subgroups. In particular, we show that the n-$(n+1)$-$(n+2)$ Conjecture for nilpotent Q implies another open conjecture, the Virtual Surjections Conjecture: A subgroup $P \leq \Gamma_1 \times \cdots \times \Gamma_n$ of a direct product of groups of type F_k is itself of type F_k if it virtually surjects to k-tuples of factors, i.e. if the image of the projection $P \to \Gamma_{i_1} \times \cdots \times \Gamma_{i_k}$ has finite index for all $i_1, \ldots, i_k$ with $1 \leq i_1 < \cdots < i_k \leq n$.

We prove the n-$(n+1)$-$(n+2)$ Conjecture in many cases, such as when the quotient Q is abelian or when the epimorphism π_2 splits. We prove a form of the conjecture where the F_n-properties are replaced by weaker homological finiteness conditions. We prove a converse which states (under slightly stronger assumptions) that the fibre product can *only* be of type F_{n+1} if the quotient Q is of type F_{n+2}. We describe an explicit map in group homology whose kernel gives an obstruction to the n-$(n+1)$-$(n+2)$ Conjecture. We prove that the general conjecture can be reduced to the special case where Γ_2 is finitely generated free and for this case we construct a concrete finite, $(n+1)$-dimensional complex with fundamental group P, which is in a certain sense close to being n-aspherical (and might actually be n-aspherical).

Our results on the n-$(n+1)$-$(n+2)$ Conjecture have consequences for the Virtual Surjections Conjecture. For example, we show that any subgroup of a direct product that virtually surjects to k-tuples, where k is greater than half the number of factors, is of type F_k. We also prove that a subgroup that virtually surjects to k-tuples satisfies a finiteness condition on k-dimensional homology.

In order to establish these result, we present different approaches to the main problem, employing a range of methods from topology and homological algebra, each of which sheds light on the conjecture from a different angle.

As my work on this thesis draws to a close it is my obligation and great pleasure to acknowledge the invaluable contributions other people have made to it.

First and foremost, I am deeply indebted to my supervisor, Martin Bridson, whose guidance and advice made this work possible and whose infectious enthusiasm for mathematics and eagerness to share his knowledge and ideas made it exciting.

Among the many people who have taught me mathematics in the past, Robert Bieri and Richard Weidmann deserve my special gratitude not only for introducing me to this subject area but also for their continued support over the years.

My studies in Oxford were funded by the German Academic Exchange Service and the Engineering and Physical Sciences Research Council, for which I am very grateful indeed.

I also would like to thank my friends and fellow students who provided much encouragement, motivation and comfort.

Finally, I would hardly have been able to start this work, much less finish it, without the constant and unconditional support from my family.

Contents

1

Introduction

❦

Ever tried. Ever failed. No matter.
Try again. Fail again. Fail better.

—Samuel Beckett

This thesis is about the higher finiteness properties of fibre products and of subgroups of direct products of groups. Our account is built around a central theme:

The n-$(n+1)$-$(n+2)$ Conjecture. *Let*

$$N_1 \hookrightarrow \Gamma_1 \overset{\pi_1}{\twoheadrightarrow} Q$$
$$N_2 \hookrightarrow \Gamma_2 \overset{\pi_2}{\twoheadrightarrow} Q$$

be two short exact sequences of groups and

$$P := \{(\gamma_1, \gamma_2) \in \Gamma_1 \times \Gamma_2 \mid \pi_1(\gamma_1) = \pi_2(\gamma_2)\}$$

their fibre product.
If N_1 is of type F_n, both Γ_1 and Γ_2 are of type F_{n+1} and Q is of type F_{n+2}, then P is of type F_{n+1}.

The aim of this introduction then is to explain, first, the motivations behind and consequences of this conjecture — and hence, why we should be interested in it. And second, to give a synopsis of the progress we have made on resolving it, the task with which, in one way or another, we shall be concerned with for the remainder of the thesis.

1.1 Background

Subgroups of direct products of free groups The origins of the n-$(n+1)$-$(n+2)$ Conjecture lie in the study of subgroups of direct products of free groups, surface groups and limit groups. Historically, subgroups of direct products of free groups were the first source of examples for groups with exotic higher finiteness properties. Stallings in [Sta63] discovered the first example of a finitely presented group with non-finitely generated integral homology in dimension 3 as a subgroup of a direct product of three free groups.

Over time, the systematic study of such subgroups of direct products has revealed that higher finiteness properties play a central role in understanding this class of groups. In particular, the literature has brought to light an apparent dichotomy between those subgroups which satisfy strong finiteness properties and those which do not — a contrast first observed by Baumslag and Roseblade in [BR84]. They showed that although in a direct product of two finitely generated free groups one can find uncountably many isomorphism classes of subgroups, the only finitely presented subgroups are virtually direct products of two free groups.

The "wildness" of general subgroups of a direct product of free groups is further exemplified by a number of unsolvability results: In a direct product $F \times F$ of two free groups on two generators

- there is a finitely generated subgroup $L \leq F \times F$ with unsolvable membership problem [Mih66] and unsolvable conjugacy problem [Mil92],
- the isomorphism problem among finitely generated subgroups is unsolvable [Mil92],
- there are uncountably many subgroups with non-isomorphic abelianisation [BM09], and
- there is no algorithm to compute the rank of the abelianisation of an arbitrary finitely generated subgroup [BM09].

The other side of Baumslag and Roseblade's dichotomy was greatly expanded on in a series of papers by Baumslag, Bridson, Howie, Miller and Short [BBMS00, BHMS02, BH07b, BH07a, BHMS09, BHMS12]. In particular, they generalised the result from [BR84] to direct products of any (finite) number of free groups, or more generally, surface groups or limit groups[1]. This generalisation says that a subgroup $P \leq \Gamma_1 \times \cdots \times \Gamma_n$ in a direct product of finitely generated free (surface, limit) groups $\Gamma_1, \ldots, \Gamma_n$, which is of type F_n, is itself virtually a direct product of n or fewer finitely generated free (surface, limit) groups [BHMS02, Theorem A], [BHMS09, Theorem A]. In fact, if the intersection of P with each factor Γ_i is non-trivial, then the product

[1] The latter are most easily defined as the finitely generated *fully residually free* groups, i.e. those finitely generated groups Γ such that for every finite subset $X \subset \Gamma$ there is a free group F and a homomorphism $f : \Gamma \to F$ which is injective on X. This class includes free groups, free abelian groups and fundamental groups of orientable surfaces.

of these intersections will have finite index in P [BHMS02, Theorem B], [BHMS09, Theorem C]. So, virtually, the only subgroups of type F_n are the "obvious" ones.

In both results the higher finiteness condition is essential: Stallings' example mentioned earlier and its generalisations by Bieri [Bie81] provide finitely presented subgroups in a direct product of three or more free groups (and intersecting each factor non-trivially), which do not satisfy the conclusions of the theorems.

This indicates that a thorough study of the higher finiteness properties of subgroups of direct products will be necessary in order to gain a full understanding of this class of groups. In particular, there is great interest in comprehending the class of finitely presented subgroups of direct products of limit groups, as these are just the finitely presented residually free groups. A systematic study of these groups was started in [BHMS12].

In light of the work cited above, one might think of the subgroups of a direct product of n free (or limit) groups as arranged along a spectrum according to which finiteness properties they satisfy. At the lower end, the class of all finitely generated subgroups is hopelessly complicated. At the upper end, the only subgroups of type F_n are the obvious ones. But if we aim to understand all the finitely presented groups we need to find out what happens as one imposes successively stronger finiteness conditions. The hope is that the stark contrast between "finitely generated" and "type F_n" will, to some extent, be resolved by considering the classes in between.

There are some promising signs that this is not a futile endeavour. The examples by Stallings and Bieri are subgroups of type F_{n-1} in a direct product of n free groups. And while they are certainly not as well-behaved as the subgroups of type F_n, they are not anywhere as arbitrary as finitely generated subgroups either. In fact, the original examples are kernels of maps from the direct product to a finitely generated abelian group, and it turns out that any subgroup of type F_{n-1} has to be (virtually) of that type ([BM09, Theorem E] for the case of three factors; see our Corollary 3.4 for a stronger statement).

On the other end of the spectrum, the work of [BHMS12] has demonstrated that the finitely presented subgroups in a direct product of limit groups, while still a very large class of groups, do not exhibit quite as much "wildness" as the finitely generated subgroups. For example, in contrast to the unsolvability results cited above, the finitely presented subgroups do have solvable conjugacy and membership problems [BHMS12, Theorems J & K].

The Virtual Surjections Conjecture In [BHMS12] a general criterion was given to recognise finitely presented subgroups of a direct product: If $P \leq \Gamma_1 \times \cdots \times \Gamma_n$ is a subgroup of the direct product of finitely presented groups $\Gamma_1, \ldots, \Gamma_n$ such that the image of P under each projection $P \to \Gamma_i \times \Gamma_j$ to a pair of factors has finite index in $\Gamma_i \times \Gamma_j$ then P is finitely presented. The latter condition has been termed "virtual surjection to pairs" and the criterion has been used to great effect to study the finitely presented subgroups of free and limit groups (yielding some of the results we cited above).

Given the utility of this criterion in studying finitely presented subgroups and in view of our proclaimed goal to understand what the subgroups of type F_k in a direct product are for any k, the following conjecture suggests itself:

> The Virtual Surjections Conjecture. *Let $\Gamma_1, \ldots, \Gamma_n$ be groups of type F_k and $P \leq \Gamma_1 \times \cdots \times \Gamma_n$ a subgroup of their direct product. Assume that P* virtually surjects to k-tuples of factors, *i.e. for any $i_1, \ldots, i_k$ with $1 \leq i_1 < \cdots < i_k \leq n$ the image of P under the projection $P \to \Gamma_{i_1} \times \cdots \times \Gamma_{i_k}$ is of finite index. Then P is of type F_k.*

This conjecture appears in the literature in various forms, e.g. [Koc10, Section 5] or [Dis08, Section 12.5]. While the conjecture makes no assumptions on the factor groups, beyond the finiteness condition F_k, it is especially relevant for direct products of free groups (or limit groups), where a converse is known to hold.

To state this, let us introduce two pieces of terminology: If $P \leq \Gamma_1 \times \cdots \times \Gamma_n$ is a subgroup of a direct product, we shall call P a *subdirect product* of $\Gamma_1, \ldots, \Gamma_n$ if all the projection maps to the factors $P \to \Gamma_i$ are onto. P is called a *full* subdirect product if, in addition, the intersections $P \cap \Gamma_i$ are non-trivial for all i.

It is often possible to restrict attention to subdirect products or full subdirect products, by means of two simple observations: If $P \leq \Gamma_1 \times \cdots \times \Gamma_n$ is any subgroup then P is a subdirect product of $\Gamma_1', \ldots, \Gamma_n'$ where Γ_i' is the image of P under the projection $P \to \Gamma_i$.

And if P is a subdirect product, such that the intersection of P with, say, Γ_1 is trivial, then the projection to the remaining factors $P \to \Gamma_2 \times \cdots \times \Gamma_n$ has trivial kernel, so we can embed P as a subgroup into a direct product with fewer factors. We can repeat this procedure if necessary — one by one projecting away from the factors that intersect P trivially — until we end up with an embedding of P as a full subdirect product of some subset of the factors $\Gamma_1, \ldots, \Gamma_n$.

In particular, in most cases of interest, restricting to full subdirect products does not impose a real restriction on the class of groups considered.

For subdirect products of non-abelian free (or limit) groups the following is conjectured to hold:

> Conjecture. *A full subdirect product $P \leq \Gamma_1 \times \cdots \times \Gamma_n$ of the non-abelian limit groups $\Gamma_1, \ldots, \Gamma_n$ is of type F_k if and only if it virtually surjects to k-tuples of factors.*

One direction of this is, of course, contained in the Virtual Surjections Conjecture. The other direction — that any full subdirect product of type F_k virtually surjects to k-tuples — has been established in [BHMS12, Theorem D] for $k = 2$ and in [Koc10, Theorem D] in the general case. Proving the Virtual Surjections Conjecture would thus complete the picture and give a full characterisation of the subdirect products of type F_k of free groups.

One of the main reasons why we care about the n-$(n+1)$-$(n+2)$ Conjecture is that it implies the Virtual Surjections Conjecture (Theorem 3.14).

Higher finiteness properties Another reason why we believe the n-$(n+1)$-$(n+2)$ Conjecture deserves to be studied is because of its relation to the theory surrounding higher finiteness properties of groups.

While the higher finiteness properties are defined as natural generalisations of the fundamental properties of finite generation and finite presentability and have been studied for a long time (they first appear in [Wal65] and [Ser71]), they are still in some ways poorly understood. The interest in studying subgroups of direct products, where, as we explained above, finiteness conditions do have palpable effects on the group structure, goes both ways: Finiteness conditions might serve as a tool for understanding subdirect products, but conversely, understanding subdirect products might also shed more light on the nature of the higher finiteness properties. Viewed from this angle, the n-$(n+1)$-$(n+2)$ Conjecture supplies a novel way of establishing these properties.

Let us mention two connections in particular.

Amalgamated free products First, it is worth mentioning that, from the viewpoint of category theory, the fibre product is just the pullback in the category of groups, and thus dual to the amalgamated free product. For the latter operation, an analogous statement to the n-$(n+1)$-$(n+2)$ Conjecture exists: If $N \hookrightarrow \Gamma_1$ and $N \hookrightarrow \Gamma_2$ are monomorphisms such that N is of type F_{n-1} and Γ_1 and Γ_2 are of type F_n then the associated amalgamated product $\Gamma_1 *_N \Gamma_2$ is of type F_n. For $n = 2$, where the converse holds as well, this is a remark by Baumslag [Bau62] and in general it follows, for example, from the construction of classifying complexes for amalgamated products in [SW79] (see also [Bie81, Theorem 2.13]).

Clearly, the conjecture for fibre products is decidedly more subtle, as it involves finiteness conditions not just on the factor groups and the quotient but also on one of the kernels in an essential way (for more on the necessity of the conditions in the conjecture, see also the end of Section 3.3). Remarkably, the proof of the statement for amalgamated products is not difficult either: Essentially, the classifying complex for the amalgamated product constructed in the naïve way will have the required finiteness properties (e.g., in the case $n = 2$ it suffices to consider the "usual" presentation for the amalgamated product). The same, alas, cannot be said for the fibre product, but in any case, as a complement to the basic theorem on finiteness properties of amalgamated products, the n-$(n+1)$-$(n+2)$ Conjecture seems worth striving for. For more on this relationship, see also the "Further remarks" in Section 5.4.

Bestvina-Brady groups and Σ-invariants While it is generally a hard problem to determine the finiteness properties of subgroups — even for the properties of finite generation and finite presentability — there is one class where a far-reaching theory is available: For kernels of maps to abelian groups, the theory of Σ- or Bieri-Neumann-Strebel-Renz invariants, provides great insights and a very useful set of tools for establishing higher finiteness properties (we will use them for this purpose in Chapter 8). In [BB97] Bestvina and Brady used their "synthetic Morse theory" to determine the precise finiteness properties of kernels of maps from right-angled Artin groups to $\mathbb{Z}$, essentially by determining the Σ-invariants of general right-angled Artin groups (see also [MMV98, BG99]). These kernels — the Bestvina-Brady groups — have been studied extensively since and have proved a rich source of examples of groups with various exotic properties.

Of course, direct products of free groups are examples of right-angled Artin groups and the results of Bestvina and Brady (or the earlier work by Meinert in [Mei94]) can be used to obtain information on the finiteness properties of subgroups containing the commutator subgroup. The n-$(n+1)$-$(n+2)$ Conjecture and the Virtual Surjections Conjecture would complement this picture by supplying information on the higher finiteness properties of arbitrary subgroups that are out of reach of methods using Σ-invariants.

Viewed from the angle of Σ-theory, an answer to the Virtual Surjections Conjecture would be especially interesting because the subdirect products that virtually surject to k-tuples of factors contain some term of the lower central series, though usually not the commutator subgroup (see also Section 3.1). There have been attempts to generalise some aspects of Σ-invariants to give information on such subgroups (e.g. [BG95, BG00]), but progress on this problem has been slow and a general theory, possibly encompassing the higher-dimensional invariants, has yet to emerge. Though this is speculative at present, the n-$(n+1)$-$(n+2)$ Conjecture and Virtual Surjections Conjecture would provide an alternative approach to deriving higher finiteness properties of subgroups contatining a term of the lower central series in one particular setting, possibly giving some clues on what a more general theory of Σ-invariants might look like.

1.2 Results

While we have not yet succeeded in proving the n-$(n+1)$-$(n+2)$ Conjecture (and thus the Virtual Surjections Conjecture) in full generality, we have made substantial progress, employing a variety of methods to establish certain special cases, rule out potential counterexamples, prove homological analogues or shed light on the general case.

Theorems 5.6 & 8.3. *The n-$(n+1)$-$(n+2)$ Conjecture holds*

- *when the second epimorphism $\pi_2 : \Gamma_2 \twoheadrightarrow Q$ splits, or*
- *when Q is abelian.*

The latter result is particularly interesting in light of the following:

Theorem 3.14. *If the n-$(n+1)$-$(n+2)$ Conjecture holds for nilpotent Q then the Virtual Surjections Conjecture holds.*

In particular, the result on abelian quotients implies the following special case of the Virtual Surjections Conjecture:

> Corollaries 8.5 & 8.6. *Let $\Gamma_1, \dots, \Gamma_n$ be groups of type F_k. A subdirect product $P \le \Gamma_1 \times \cdots \times \Gamma_n$ that virtually surjects to k-tuples is of type F_k if $\Gamma_i/(P \cap \Gamma_i)$ is abelian for all i (previously obtained in [Koc10, Theorem 34]); this holds, in particular, if $k > \frac{n}{2}$.*

Furthermore, we show that the general n-$(n+1)$-$(n+2)$ Conjecture can be reduced to a special case:

> Proposition 5.7. *If the n-$(n+1)$-$(n+2)$ Conjecture is true whenever Γ_2 is finitely generated free, then it holds in general.*

We consider the n-$(n+1)$-$(n+2)$ Conjecture for the case where Γ_2 is free in Section 7.3. In that case we manage to produce — assuming a certain fact, for which we have strong evidence, but which we cannot currently prove — an $(n+1)$-dimensional complex W, which has the fibre product P as its fundamental group, and satisfies $\pi_k(W) = 0$ for $2 \le k \le n-1$ and $H_0(N_2, \pi_n(W)) = 0$. We believe that there is a good chance that, in fact, $\pi_n(W) = 0$ which would prove the n-$(n+1)$-$(n+2)$ Conjecture.

In the literature on subdirect products, another finiteness property has proved particularly useful. A group Γ is said to be of type wFP_n if the integral homology groups $H_k(\Gamma', \mathbb{Z})$ are finitely generated for any $k \le n$ and any finite index subgroup $\Gamma' \le \Gamma$. In many of the results we quoted from the literature earlier, the F_n-assumptions can in fact be weakened to wFP_n. We show a variant of the n-$(n+1)$-$(n+2)$ Conjecture and the Virtual Surjections Conjecture for these finiteness properties:

> Theorem 6.8. *If*
>
> $$N_1 \hookrightarrow \Gamma_1 \twoheadrightarrow Q$$
> $$N_2 \hookrightarrow \Gamma_2 \twoheadrightarrow Q$$
>
> *are short exact sequences with N_1 of type wFP_n, Γ_1 of type wFP_{n+1}, Γ_2 of type FP_{n+1} and Q of type FP_{n+2}, then the associated fibre product is of type wFP_{n+1}.*

> Corollary 6.9. *If $\Gamma_1, \dots, \Gamma_n$ are groups of type F_k and $P \le \Gamma_1 \times \cdots \times \Gamma_n$ is a subgroup that virtually surjects to k-tuples then P is of type wFP_k.*

In conjunction with Kochloukova's result, cited earlier, this shows

> Corollary 6.11. *A full subdirect product $P \le \Gamma_1 \times \cdots \times \Gamma_n$ of non-abelian free (or limit) groups virtually surjects to k-tuples (where $k \ge 2$) if and only if it is of type wFP_k.*

We also obtain a result that can be considered a kind of converse to the n-$(n+1)$-$(n+2)$ Conjecture:

THEOREM 6.20. *If*

$$N_1 \hookrightarrow \Gamma_1 \twoheadrightarrow Q$$
$$N_2 \hookrightarrow \Gamma_2 \twoheadrightarrow Q$$

are short exact sequences with N_1 *of type* FP_n, Γ_1, Γ_2 *of type* FP_{n+2} *and the associated fibre product is of type* FP_{n+1}, *then* Q *is of type* FP_{n+2}.

Finally, we produce an obstruction to the Conjecture in the form of a kernel (which might always vanish) of a particular explicit map between homology groups.

PROPOSITION 6.3. *Let*

$$N_1 \hookrightarrow \Gamma_1 \twoheadrightarrow Q$$
$$N_2 \hookrightarrow \Gamma_2 \twoheadrightarrow Q$$

be short exact sequences and P *their associated fibre product. Then there is a map*

$$\phi : H_0(\Gamma_2, H_n(N_1, \prod \mathbb{Z}P)) \to H_0(Q, H_n(N_1, \prod \mathbb{Z}\Gamma_1)).$$

If this map is injective then the n-$(n+1)$-$(n+2)$ *Conjecture holds for the sequences above.*

1.3 Summary

We shall give a brief rundown of the main part of the thesis.

In Chapter 2, we recall some of the fundamental properties of fibre products and their relation to subdirect products of groups.

The main result of Chapter 3 is that the n-$(n+1)$-$(n+2)$ Conjecture implies the Virtual Surjections Conjecture. We also prove the 0-1-2 Lemma, which is a standard result, and provide a strengthening of the Virtually Nilpotent Quotients Theorem from [BM09] and [BHMS12], which we use later to prove Corollary 8.6, quoted earlier. The methods in this section are mostly elementary.

In Chapter 4, we provide a proof of the 1-2-3 Theorem, which was first established in [BBMS00] and [BHMS08]. While the results are not original, our method differs from the proof in the literature. We provide this proof largely because it serves as a guide to the generalisations in the following chapters. The tools we use are based around crossed modules.

In Chapter 5, we pursue a topological approach to the problem, trying to establish higher finiteness properties by constructing classifying complexes directly. Here we prove the Split n-$(n+1)$-$(n+2)$ Theorem and the reduction to the case where the second factor is finitely generated free. The main tools we use are from the theory of "stacks" as defined by Geoghegan [Geo08].

In Chapter 6, we establish the wFP_n-version of the two conjectures, as quoted above. The approach in this chapter is much more combinatorial, mainly using tools from homological algebra, such as the Lyndon-Hochschild-Serre spectral sequence. We also establish the "converse" alluded to above, employing the Bieri-Eckmann characterisation of FP_n-properties. Furthermore this viewpoint allows us to formulate the obstruction given above.

In Chapter 7, we attack the Conjecture by trying to build suitable resolutions for groups. The goal is to develop an algebraic approach which mirrors Chapter 5, but allows for explicit computation and is at the same time more powerful than the one based on homology groups from the previous chapter.

In Chapter 8, we prove the n-$(n+1)$-$(n+2)$ Conjecture for abelian quotients and the corresponding case of the Virtual Surjections Conjecture. We also give a brief analysis of the situation for nilpotent quotients. The methods we use here are based on the theory of Bieri-Neumann-Strebel-Renz- or Σ-invariants.

Finally, in an appendix, we collect some (mostly) standard results on the finiteness properties F_n and FP_n, which we use throughout the main text, for the convenience of the reader.

2

Fibre products and subdirect products

We shall start off by giving the basic definitions and recalling some fundamental facts about the construction of fibre products and how that basic construction can be used to iteratively build up subdirect products of groups.

Definition 2.1. A subdirect product of the groups $\Gamma_1, \dots, \Gamma_n$ is a subgroup $P \leq \Gamma_1 \times \cdots \times \Gamma_n$ of the direct product such that the canonical projections $P \to \Gamma_i$ are surjective for all i.

One way of obtaining a subdirect product of two groups is via the fibre product construction: Let

$$N_1 \hookrightarrow \Gamma_1 \overset{\pi_1}{\twoheadrightarrow} Q$$
$$N_2 \hookrightarrow \Gamma_2 \overset{\pi_2}{\twoheadrightarrow} Q$$

be two short exact sequences. Then we call

$$P := \{(\gamma_1, \gamma_2) \in \Gamma_1 \times \Gamma_2 \mid \pi_1(\gamma_1) = \pi_2(\gamma_2)\}$$

the fibre product associated to these sequences (or to π_1 and π_2).[1]

A crucial observation is that to any such fibre product we can associate two exact sequences of the form $N_1 \hookrightarrow P \twoheadrightarrow \Gamma_2$ and $N_2 \hookrightarrow P \twoheadrightarrow \Gamma_1$. More precisely:

[1]If the two sequences are identical, i.e. $\pi_1 = \pi_2$, P is often called a "symmetrical" or "untwisted" fibre product (associated to π_1), but we will make little use of this concept.

Lemma 2.2. *The fibre product above fits into a commutative diagram*

$$\begin{tikzcd}
 & 1 \times N_2 \arrow[r, equal] \arrow[d, hook] & N_2 \arrow[d, hook] \\
N_1 \times 1 \arrow[r, hook] \arrow[d, equal] & P \arrow[r, two heads, "p_2"] \arrow[d, two heads, "p_1"] & \Gamma_2 \arrow[d, two heads, "\pi_2"] \\
N_1 \arrow[r, hook] & \Gamma_1 \arrow[r, two heads, "\pi_1"] & Q
\end{tikzcd}$$

with exact rows and columns, where $p_1 : P \to \Gamma_1$ and $p_2 : P \to \Gamma_2$ are the restrictions of the canonical projections $\Gamma_1 \times \Gamma_2 \to \Gamma_i$ and all the injections are inclusion homomorphisms.

Proof. The lower right square commutes simply by the definition of the fibre product. If $(\gamma_1, \gamma_2) \in P$ is in the kernel of p_2, then $\gamma_2 = 1$ and $\pi_1(\gamma_1) = \pi_2(\gamma_2) = 1$, so $\gamma_1 \in \ker \pi_1 = N_1$. Obviously, $N_1 \times 1$ is sent isomorphically to $N_1 \le \Gamma_1$ under the projection p_1. The rest follows by symmetry. □

We will usually identify $N_1 \times 1$ with N_1 and $1 \times N_2$ with N_2.

It is a basic observation that every subdirect product of two groups is a fibre product. Precisely, we have the following:

Lemma 2.3. *Let $P \le \Gamma_1 \times \Gamma_2$ be a subdirect product of Γ_1 and Γ_2. Let $N_1 := P \cap \Gamma_1$ and $N_2 := P \cap \Gamma_2$, where Γ_1 denotes the subgroup $\Gamma_1 \times 1 \le \Gamma_1 \times \Gamma_2$ and similarly for Γ_2 (i.e. we have identified each Γ_i with its image under the natural inclusion $\Gamma_i \hookrightarrow \Gamma_1 \times \Gamma_2$). Then*

$$\Gamma_1/N_1 \cong P/N_1 \times N_2 \cong \Gamma_2/N_2$$

and P is the fibre product associated to two short exact sequences

$$N_1 \hookrightarrow \Gamma_1 \twoheadrightarrow Q$$
$$N_2 \hookrightarrow \Gamma_2 \twoheadrightarrow Q.$$

Proof. Let $p_i : P \to \Gamma_i$ for $i = 1, 2$ denote the restriction of the canonical projection $\Gamma_1 \times \Gamma_2 \to \Gamma_i$. Clearly, N_1 is the kernel of p_2, so N_1 is normal in P. Symmetrically, N_2 is normal in P, as well. Therefore $N_1 \times N_2$ is normal in P and the quotient map $P \to P/(N_1 \times N_2)$ factors over both projections p_i, giving a commutative square

$$\begin{tikzcd}
P \arrow[r, "p_2"] \arrow[d, "p_1"'] & \Gamma_2 \arrow[d, "\pi_2"] \\
\Gamma_1 \arrow[r, "\pi_1"] & P/N_1 \times N_2
\end{tikzcd}$$

The kernel of π_1 is $p_1(N_1 \times N_2) = N_1$, and similarly $\ker \pi_2 = N_2$. Denoting $P/(N_1 \times N_2)$ by Q, we thus have two short exact sequences:

$$N_1 \hookrightarrow \Gamma_1 \overset{\pi_1}{\twoheadrightarrow} Q$$
$$N_2 \hookrightarrow \Gamma_2 \overset{\pi_2}{\twoheadrightarrow} Q.$$

Let P' denote their fibre product. Commutativity of the above square means that $P \leq P'$. Conversely, if $\gamma' = (\gamma_1', \gamma_2') \in P'$, set $q := \pi_1(\gamma_1') = \pi_2(\gamma_2')$. Then pick a $\gamma := (\gamma_1, \gamma_2) \in P$ with
$$\gamma(N_1 \times N_2) = q.$$
Then $\pi_i(\gamma_i) = q = \pi_i(\gamma_i')$ for $i = 1, 2$ by definition of the maps π_i. And hence $\gamma\gamma'^{-1} \in N_1 \times N_2 \subset P$ so $\gamma' \in P$.

This shows that $P = P'$, so P is the fibre product associated to the two short exact sequences. □

Indeed, not only is every subdirect product of two groups a fibre product, but we can build up any subdirect product of three or more groups by iterating the fibre product construction. The following notation will be frequently useful, when dealing with subgroups of direct products of more than two groups:

Notation. Let $\Gamma_1, \ldots, \Gamma_n$ be groups and $P \leq \Gamma := \Gamma_1 \times \cdots \times \Gamma_n$ a subgroup of their direct product. For $J = \{i_1, \ldots, i_m\} \subset \{1, \ldots, n\}$ with $i_1, \ldots, i_m$ pairwise distinct, we will denote by Γ_J the group $\Gamma_{i_1} \times \cdots \times \Gamma_{i_m}$, which we will also, via the canonical injection, consider a subgroup of Γ. Furthermore, we will denote by p_J the canonical projection map $P \to \Gamma_J$ and by N_J the intersection $P \cap \Gamma_J$. Sometimes we will omit curly braces to unclutter the notation, e.g. we will write $p_{1,\ldots,n-1}$ for the projection $P \to \Gamma_1 \times \cdots \times \Gamma_{n-1}$, in place of $p_{\{1,\ldots,n-1\}}$.

Now let $P \leq \Gamma_1 \times \cdots \times \Gamma_n$ be a subdirect product. Using the notation introduced above, set
$$T := p_{1,\ldots,n-1}(P) \leq \Gamma_1 \times \cdots \times \Gamma_{n-1}.$$
Then P is a subdirect product of T and Γ_n, so by Proposition 2.3 it is the fibre product associated to short exact sequences

$$N_{1,\ldots,n-1} \hookrightarrow T \twoheadrightarrow Q$$
$$N_n \hookrightarrow \Gamma_n \twoheadrightarrow Q.$$

Note that T, in turn, is a subdirect product of $\Gamma_1, \ldots, \Gamma_{n-1}$. This allows us to argue about subdirect products of any number of groups by using induction on the number of factors and dealing with a fibre product at every step.

3

The Virtual Surjections Conjecture

In the Introduction we explained that one of the major motivations for studying the n-$(n+1)$-$(n+2)$ Conjecture is that it implies the Virtual Surjections Conjecture, which relates the finiteness properties of subdirect products to the way these groups are embedded in the ambient direct product. The main goal of the present section is to prove this implication. En passant, we will prove the 0-1-2 Lemma, the lowest case of the n-$(n+1)$-$(n+2)$ Conjecture, and observe a curious connection between subdirect products and nilpotent groups.

Let us restate the embedding condition from the Introduction, which we can now phrase using the notation introduced earlier:

Definition 3.1. Let $\Gamma_1, \ldots, \Gamma_n$ be groups and let $P \leq \Gamma_1 \times \cdots \times \Gamma_n$ be a subgroup of their direct product. We say that P virtually surjects to k-tuples of factors if $p_J(P) \leq \Gamma_J$ is a finite index subgroup for every k-element subset of indices $J \subset \{1, \ldots, n\}$.

Recall the conjecture from the Introduction.

Conjecture 3.2 (The Virtual Surjections Conjecture). *Let $k \geq 2$. Let $\Gamma_1, \ldots, \Gamma_n$ be groups of type F_k and $P \leq \Gamma_1 \times \cdots \times \Gamma_n$ a subgroup of their direct product. If P virtually surjects to k-tuples of factors then P is of type F_k.*

We will now embark on the task of showing that the n-$(n+1)$-$(n+2)$ Conjecture implies the Virtual Surjections Conjecture.

3.1 Virtually Nilpotent Quotients

As our first ingredient we need an important consequence of the property of virtually surjecting to k-tuples for some $k \geq 2$. This lemma originated in the "Virtually

Nilpotent Quotients Theorem" from [BM09, Section 4.3], see also [BHMS12, Proposition 3.2]. We will prove a variation of this with a precise bound on the nilpotency class, which we will exploit in Chapter 8.

Lemma 3.3. *Let $\Gamma_1, \dots, \Gamma_n$ be groups and $P \leq \Gamma_1 \times \cdots \times \Gamma_n$ a subdirect product which virtually surjects to k-tuples for some $k \geq 2$. Then for every index i the group Γ_i/N_i is virtually nilpotent of class at most $\lceil \frac{n-1}{k-1} \rceil - 1$.*

Proof. We will show the statement only for $i = 1$. For other choices of i the proof is entirely analogous (or alternatively, can be reduced to the case treated here, using the isomorphism exchanging the first and i^{th} factor of the product).

Let $s := \lceil \frac{n-1}{k-1} \rceil$. Then the set $I := \{2, \dots, n\}$ can be partitioned into s disjoint subsets $I_1, \dots, I_s \subset I$ of size $|I_m| \leq k-1$ for all m. Set $I'_m := \{1\} \cup I_m$. Since P virtually surjects to k-tuples, $p_{I'_m}(P)$ is of finite index in $\Gamma_{I'_m}$ for all m. Therefore $\Gamma'_1 := \bigcap_{m=1}^{s} (p_{I'_m}(P) \cap \Gamma_1)$ is of finite index in Γ_1. Note that Γ'_1 consists of all those $\gamma \in \Gamma_1$ such that for each m with $1 \leq m \leq s$ there exists a $g \in P$ with $p_1(g) = \gamma$ and $p_i(g) = 1$ for all $i \in I_m$. In particular Γ'_1 contains N_1.

Now let $\gamma_1, \dots, \gamma_s \in \Gamma'_1$. We will show that the iterated commutator

$$c := [\gamma_1, [\gamma_2, \dots, [\gamma_{s-1}, \gamma_s] \dots]]$$

is in N_1. Pick elements $g_1, \dots, g_s \in P$ such that $p_1(g_m) = \gamma_m$ for $1 \leq m \leq s$ and $p_i(g_m) = 1$ for $i \in I_m$. Now let

$$c' := [g_1, [g_2, \dots, [g_{s-1}, g_s] \dots]].$$

Then c' is in P and $p_1(c') = c$. Furthermore, for each i with $2 \leq i \leq n$ there is an m with $1 \leq m \leq s$ such that $i \in I_m$, so $p_i(g_m) = 1$. But clearly, if one of the elements in an iterated commutator equals 1, then the commutator itself equals 1, too. So for each i with $2 \leq i \leq n$, we get that $p_i(c') = 1$. But this means that $c = c' \in \Gamma_1 \cap P = N_1$.

We have shown that N_1 contains all the commutators of length s in Γ'_1 and therefore it also contains $\gamma_{s-1}\Gamma'_1$, the $(s-1)^{\text{th}}$ term in the lower central series of Γ'_1. This proves that Γ'_1/N_1 is nilpotent of class at most $s-1$. □

In particular we record the following generalisation of [BM09, Lemma 4.9].

Corollary 3.4. *Let $\Gamma_1, \dots, \Gamma_n$ be groups and let $P \leq \Gamma_1 \times \cdots \times \Gamma_n$ be a subgroup of their direct product which virtually surjects to k-tuples for some $k > \frac{n}{2}$. Then $p_i(P)/N_i$ is virtually abelian for all i.*
Furthermore, P is virtually the kernel of a homomorphism to an abelian group. Precisely: There are finite index subgroups $\Gamma_i^0 \leq \Gamma_i$ for each i and a homomorphism $\phi : \Gamma_1^0 \times \cdots \times \Gamma_n^0 \to A$ to an abelian group A, such that $P \cap (\Gamma_1^0 \times \cdots \times \Gamma_n^0) = \ker \phi$.

Proof. Let $\Gamma'_i := p_i(P) \leq \Gamma_i$. These are finite index subgroups, since P virtually surjects to k-tuples. Then P is a subdirect product in $\Gamma'_1 \times \cdots \times \Gamma'_n$, which still virtually surjects to k-tuples. Now, since $k > \frac{n}{2}$, we have $k \geq \frac{n}{2} + \frac{1}{2}$ and so $\frac{n-1}{k-1} \leq 2$. So Lemma 3.3 says that $\Gamma'_i/(P \cap \Gamma'_i) = \Gamma'_i/N_i$ is virtually abelian for all i.

Hence there are finite index subgroups $\Gamma_i^0 \le \Gamma_i'$, containing N_i, for all i, such that Γ_i^0/N_i is abelian. Then N_i contains the commutator subgroup $[\Gamma_i^0, \Gamma_i^0]$. Therefore $N_1 \times \cdots \times N_n$ contains the commutator subgroup of $\Gamma_1^0 \times \cdots \times \Gamma_n^0$. But then $P \cap (\Gamma_1^0 \times \cdots \times \Gamma_n^0)$ contains this commutator subgroup as well, and in particular it is normal in $\Gamma_1^0 \times \cdots \times \Gamma_n^0$ with abelian quotient group. □

Corollary 3.4 is particularly interesting in light of the following theorem by Kochloukova [Koc10, Theorem 11]:

Theorem 3.5. *If $\Gamma_1, \ldots, \Gamma_n$ are non-abelian free (or limit) groups and $P \le \Gamma_1 \times \cdots \times \Gamma_n$ is a full subdirect product of type F_k then P virtually surjects to k-tuples.* □

We can thus conclude:

Corollary 3.6. *If $\Gamma_1, \ldots, \Gamma_n$ are non-abelian free (or limit) groups and $P \le \Gamma_1 \times \cdots \times \Gamma_n$ is a full subdirect product of type F_k with $k > \frac{n}{2}$ then P is virtually the kernel of a map to an abelian group (in the same sense as in the statement of Corollary 3.4).* □

The subgroups which are virtually kernels of maps to abelian groups are called "of Stallings-Bieri type" in [BM09], alluding to the examples in [Sta63] and [Bie81] of groups of type F_{n-1} with non-finitely generated n^{th} integral homology group, which are kernels of maps from a direct product of n finitely generated free groups to $\mathbb{Z}$. For a while, it was an open question whether every finitely presented full subdirect product of free groups is of this type [BM09, Question 4.3], until examples proving that this is not the case were discovered in [BHMS08, Theorem H] (or see [BHMS12, Theorem H]).

We will have more to say on these kinds of subgroups in Chapter 8, but let us for now get back to the matter at hand.

3.2 The 0-1-2 Lemma and virtual surjection to factors

Next, we will investigate the case $k = 1$ of the Virtual Surjections Conjecture. Of course, we have excluded this case from the Conjecture with good reason: It is not true for $k = 1$; we will obtain an ample source of counterexamples shortly. It does however become true, if we add the extra assumption that the $p_i(P)/N_i$ be virtually nilpotent for all i — as they are for $k \ge 2$ by Lemma 3.3.

To show that this is the case, we will need the 0-1-2 Conjecture. As we will see, this is really no more than an easy lemma. Nevertheless, it is well worth recording the steps carefully here, since we will notice at many points during our later work on the n-$(n+1)$-$(n+2)$ Conjecture that the difficulties frequently lie in formulating and proving higher-dimensional analogues of these very simple steps.

We start with a statements on finite generation and finite presentability in short exact sequences of groups.

Lemma 3.7. *Let $N \hookrightarrow \Gamma \overset{\pi}{\twoheadrightarrow} Q$ be a short exact sequence with Γ finitely generated. If Q is finitely presented then N is finitely normally generated. In case Γ is finitely presented, the converse holds, too, i.e. Q is finitely presented if and only if N is finitely normally generated.*

Proof. Let $\mathcal{X} \subset \Gamma$ be a finite generating set, $F(\mathcal{X})$ the free group on $\mathcal{X}$ and $p : F(\mathcal{X}) \twoheadrightarrow \Gamma$ the canonical epimorphism. Since Q is finitely generated, it has a finite presentation with generating set $\pi(\mathcal{X})$. This means that $K := \ker(\pi \circ p)$ is finitely normally generated in $F(\mathcal{X})$. But then $N = p(K)$ is finitely normally generated in Γ.

Now assume that Γ is finitely presented and N is finitely normally generated in Γ, say by some finite set Υ. Again, take $\mathcal{X}$ to be some finite generating set of Γ and $p : F(\mathcal{X}) \to \Gamma$ the induced epimorphism. Since Γ is finitely presented, $\ker p$ is finitely normally generated, say by a set $\mathcal{Z}$. Now $\ker(\pi \circ p) = p^{-1}(N)$ is easily seen to be finitely normally generated by $\mathcal{Z}$ and a set of lifts of the elements of Υ. □

Lemma 3.8 (The 0-1-2 Lemma). *Let*

$$N_1 \hookrightarrow \Gamma_1 \overset{\pi_1}{\twoheadrightarrow} Q$$
$$N_2 \hookrightarrow \Gamma_2 \overset{\pi_2}{\twoheadrightarrow} Q$$

be two short exact sequences with Γ_1 and Γ_2 finitely generated and Q finitely presented. Then the associated fibre product is finitely generated.

Proof. Let P denote the fibre product. The proof is based on a comparison of the two short exact rows in the following diagram:

$$\begin{tikzcd}
N_1 \arrow[r, hook] \arrow[d, equal] & P \arrow[r, two heads, "p_2"] \arrow[d, two heads, "p_1"] & \Gamma_2 \arrow[d, two heads, "\pi_2"] \\
N_1 \arrow[r, hook] & \Gamma_1 \arrow[r, two heads, "\pi_1"] & Q
\end{tikzcd}$$

First note that, by the previous lemma, N_1 is finitely normally generated in Γ_1, say by some finite set $\mathcal{X} \subset N_1$. Let Υ_1 be a finite generating set for Γ_1 and pick a lift in P, i.e. a finite set $\widetilde{\Upsilon}_1 \subset P$ with $p_1(\widetilde{\Upsilon}_1) = \Upsilon_1$. Let $H_0 := \langle \mathcal{X} \cup \widetilde{\Upsilon}_1 \rangle \leq P$. Then $p_1(H_0) = \Gamma_1$. Now note that the action of any $\gamma = (\gamma_1, \gamma_2) \in P$ by conjugation on $N_1 \leq P$ is the same as the conjugation action of $p_1(\gamma) = \gamma_1$ on $N_1 \leq \Gamma_1$:

$$(\gamma_1, \gamma_2)(n, 1)(\gamma_1, \gamma_2)^{-1} = (\gamma_1 n \gamma_1^{-1}, 1), \quad \text{for all } n \in N_1.$$

In particular if $x \in \mathcal{X}$ and $\gamma_1 \in \Gamma_1$ we can pick a $\gamma \in H_0$ with $p_1(\gamma) = \gamma_1$ and we have

$$\gamma_1 x \gamma_1^{-1} = \gamma x \gamma^{-1} \in H_0.$$

But, by assumption, any element of N_1 is a product of such conjugates. Therefore $N_1 \leq H_0$.

Let Υ_2 be a finite generating set for Γ_2 and $\widetilde{\Upsilon}_2 \subset P$ a lift to P, i.e. $\widetilde{\Upsilon}_2$ is finite and $p_2(\widetilde{\Upsilon}_2) = \Upsilon_2$. Set $H := \langle \mathcal{X} \cup \widetilde{\Upsilon}_1 \cup \widetilde{\Upsilon}_2 \rangle$. Then $p_2(H) = \Gamma_2$. Now let $\gamma \in P$. Then there is a $\gamma' \in H$ with $p_2(\gamma') = p_2(\gamma)$, so $\gamma\gamma'^{-1} \in N_1 \leq H_0 \leq H$ and therefore also $\gamma = \gamma\gamma'^{-1}\gamma' \in H$. Hence we have shown that $P = H$ and so P is finitely generated. □

We now have all we need to establish the "$k = 1$ case" of the Virtual Surjections Conjecture, as announced earlier.

Proposition 3.9. *Let $\Gamma_1, \dots, \Gamma_n$ be finitely generated groups and $P \leq \Gamma_1 \times \cdots \times \Gamma_n$ a subgroup of their direct product. If the image $p_i(P) \leq \Gamma_i$ under the projection to each factor Γ_i is of finite index in Γ_i and $p_i(P)/(P \cap \Gamma_i)$ is virtually nilpotent for all i then P is finitely generated.*

Proof. Note first that all the $\Gamma_i' := p_i(P)$ are finitely generated, since they are of finite index in Γ_i, and that $p_i(P)/(P \cap \Gamma_i') = p_i(P)/(P \cap \Gamma_i)$. Thus, replacing all the Γ_i by Γ_i', we may assume without loss of generality that P is a subdirect product.

We now prove the claim by induction on n. If $n = 1$, there is nothing to prove. So assume $n \geq 2$. Let $T := p_{1,\dots,n-1}(P)$. As described at the end of Chapter 2, P is the fibre product associated to the short exact sequences

$$\begin{aligned} N_{1,\dots,n-1} &\hookrightarrow T \twoheadrightarrow Q \\ N_n &\hookrightarrow \Gamma_n \twoheadrightarrow Q. \end{aligned}$$

Here, T is a subdirect product of $\Gamma_1, \dots, \Gamma_{n-1}$. Note that $T \cap \Gamma_i$ for $1 \leq i \leq n-1$ contains $P \cap \Gamma_i$, therefore $\Gamma_i/(T \cap \Gamma_i)$ is a quotient of the virtually nilpotent group $\Gamma_i/(P \cap \Gamma_i)$, and hence virtually nilpotent itself. Thus we can conclude by induction that T is finitely generated. By assumption, $Q \cong \Gamma_n/N_n$ is finitely generated and virtually nilpotent and therefore finitely presented. Hence we can apply the 0-1-2 Lemma to conclude that P is finitely generated. □

Remark 3.10. Note that the nilpotency condition on the groups Γ_i/N_i is only used in the penultimate sentence of the proof, to establish that certain finitely generated groups are finitely presented. Going through the proof carefully, we can see that the condition that Γ_i/N_i be virtually nilpotent can be replaced with the much weaker condition that $\Gamma_i/(p_J(P) \cap \Gamma_i)$ be finitely presented for all i and all $J \subset \{1, \dots, n\}$ with $i \in J$. In particular, this holds if all quotients of Γ_i/N_i are finitely presented. For finitely presented groups this condition is sometimes called Max-n, as the property of having only finitely presented quotients is equivalent to satisfying the "maximal condition for normal subgroups" (i.e. any ascending chain of normal subgroups is eventually constant). Other examples of finitely presented Max-n groups — apart from the finitely generated virtually nilpotent groups — are virtually polycyclic groups and finitely presented metabelian groups [BS80] (though there are finitely presented soluble groups which are not Max-n [Abe79]).

We can also use Lemma 3.7 to show that the Virtual Surjections Conjecture for $k = 1$ is not true without any extra condition. To do this we prove a sort of converse of the 0-1-2 Lemma (this is [BM09, Corollary 2.4]):

Proposition 3.11. *Let*

$$N_1 \hookrightarrow \Gamma_1 \overset{\pi_1}{\twoheadrightarrow} Q$$
$$N_2 \hookrightarrow \Gamma_2 \overset{\pi_2}{\twoheadrightarrow} Q$$

be two short exact sequences with Γ_1 *and* Γ_2 *finitely presented. Then the associated fibre product is finitely generated if and only if* Q *is finitely presented.*

Proof. The "if-direction" of the proof is simply a special case of the 0-1-2 Lemma.

Conversely, denote the fibre product by P and assume that it is finitely generated. Then consider the short exact sequence

$$N_1 \hookrightarrow P \overset{p_2}{\twoheadrightarrow} \Gamma_2.$$

From Lemma 3.7 we can conclude that N_1 is finitely normally generated in P. But then N_1 is finitely normally generated in Γ_1 as well (by the same observation as in the proof of Lemma 3.8). From the second part of Lemma 3.7 it then follows that Q is finitely presented. □

Let Q be any finitely generated, but not finitely presented group and let $\pi : F \twoheadrightarrow Q$ be an epimorphism from a finitely generated free group to Q. Then the symmetric fibre product associated to π is a subdirect product $P \leq F \times F$, which is not finitely generated, providing a counterexample for the Virtual Surjections Conjecture with $k = 1$.

3.3 The n-$(n+1)$-$(n+2)$ Conjecture implies the Virtual Surjections Conjecture

Before we finally show the main result of this chapter, we need two quick lemmas.

Lemma 3.12. *Let* $P \leq \Gamma_1 \times \dots \times \Gamma_n$. *Let* $I, J \subset \{1, \dots, n\}$ *with* $I \cup J = \{1, \dots, n\}$ *and* $I \cap J \neq \emptyset$. *Then* $p_{I\cap J}(P \cap \Gamma_I) = p_J(P) \cap \Gamma_{I\cap J}$.

Proof. Let $\gamma \in p_{I\cap J}(P \cap \Gamma_I)$. Pick $p \in P \cap \Gamma_I$ with $p_{I\cap J}(p) = \gamma$. Now $p_J(p) \in \Gamma_{I\cap J}$, therefore $\gamma = p_{I\cap J}(p) = p_{I\cap J}(p_J(p)) = p_J(p)$. And so $\gamma \in p_J(P) \cap \Gamma_{I\cap J}$.

Conversely, let $\gamma \in p_J(P) \cap \Gamma_{I\cap J}$. Pick a $p \in P$ with $p_J(p) = \gamma$. Note that $p_{I\cap J}(p) = p_{I\cap J}(p_J(p)) = p_{I\cap J}(\gamma) = \gamma$. But $p_J(p) \in \Gamma_{I\cap J}$, so $p \in \Gamma_I$ and hence $\gamma \in p_{I\cap J}(P \cap \Gamma_I)$. □

Lemma 3.13. *Let* $P \leq \Gamma_1 \times \dots \times \Gamma_n$. *If* P *virtually surjects to* k*-tuples (for some* $k \geq 2$*) then* $N_{1,\dots,n-1} = P \cap (\Gamma_1 \times \dots \times \Gamma_{n-1})$ *virtually surjects to* $(k-1)$*-tuples.*

Proof. Let $J' \subset \{1, \dots, n-1\}$ be some $(k-1)$-element subset. Apply the preceding lemma with $I := \{1, \dots, n-1\}$ and $J := J' \cup \{n\}$ to conclude that $p_{J'}(N_{1,\dots,n-1}) = p_J(P) \cap \Gamma_{J'}$. Note that $p_J(P)$ has finite index in Γ_J, so $p_J(P) \cap \Gamma_{J'}$ has finite index in $\Gamma_{J'}$. □

Theorem 3.14. *Assume that the n-$(n+1)$-$(n+2)$ Conjecture holds whenever the quotient Q is virtually nilpotent.*
Then the Virtual Surjections Conjecture is true.

Proof. We prove the following statement by induction on $k \geq 1$:

Claim. Let $\Gamma_1, \ldots, \Gamma_n$ be groups of type F_k and $P \leq \Gamma_1 \times \cdots \times \Gamma_n$ a subgroup of the direct product. If P virtually surjects to k-tuples and $p_i(P)/(P \cap \Gamma_i)$ is virtually nilpotent for all i, then P is of type F_k.

Of course, the nilpotency assumption is redundant when $k \geq 2$ by Lemma 3.3, so this Claim does imply the Virtual Surjections Conjecture. Furthermore, by the same argument as in the proof of Proposition 3.9 we can always reduce to the case where P is a subdirect product: If P is not subdirect, we replace all the Γ_i by their finite index subgroups $p_i(P)$, which will satisfy the same finiteness properties as the Γ_i (Proposition A.22).

We have already supplied the base case $k = 1$ of the induction in Proposition 3.9. So fix some $k \geq 2$. We now prove the Claim for this k by induction on n.

If $n \leq k$ then there is nothing to prove: P will be of finite index in $\Gamma_1 \times \cdots \times \Gamma_n$ and therefore certainly of type F_k (Propositions A.22 and A.24).

So assume $n > k$. Set $T := p_{1,\ldots,n-1}(P)$, so P is the fibre product associated to the sequences

$$\begin{aligned} N_{1,\ldots,n-1} \hookrightarrow T &\twoheadrightarrow Q \\ N_n \hookrightarrow \Gamma_n &\twoheadrightarrow Q. \end{aligned}$$

Here, $T \leq \Gamma_1 \times \cdots \times \Gamma_{n-1}$ is a subdirect product, which also virtually surjects to k-tuples, since $p_J(T) = p_J(P)$ for all $J \subset \{1, \ldots, n-1\}$. By induction, we can assume that T is of type F_k.

By Lemma 3.13, $N_{1,\ldots,n-1} \leq \Gamma_1 \times \cdots \times \Gamma_{n-1}$ virtually surjects to $(k-1)$-tuples. Furthermore,

$$\Gamma_i/(N_{1,\ldots,n-1} \cap \Gamma_i) = \Gamma_i/P \cap \Gamma_i$$

is virtually nilpotent for all $i \in \{1, \ldots, n-1\}$. Thus, again by induction (on k this time), we can assume that $N_{1,\ldots,n-1}$ is of type F_{k-1}.

By assumption, $Q \cong \Gamma_n/N_n$ is virtually nilpotent and finitely generated (since Γ_n is) and therefore of type F_∞ (Propositions A.27 and A.22). So, assuming the $(k-1)$-k-$(k+1)$ Conjecture, we can conclude that P is of type F_k. □

Two remarks are in order. Firstly, from the proof it is clear that we only need to assume the n-$(n+1)$-$(n+2)$ Conjecture for $n < k$ to conclude that virtual surjection to k-tuples implies type F_k. In particular, the 1-2-3 Theorem implies the Virtual Surjection to Pairs Theorem, i.e. if a subgroup $P \leq \Gamma_1 \times \cdots \times \Gamma_n$ of a direct product of finitely presented groups $\Gamma_1, \ldots, \Gamma_n$ virtually surjects to pairs of factors $\Gamma_i \times \Gamma_j$, then P is finitely presented. This was first proved in [BHMS12, Theorem E].

The second remark concerns the word "virtually" in the statement of Theorem 3.14. In fact, a quick argument shows that it suffices to assume the n-$(n+1)$-$(n+2)$

Conjecture for nilpotent quotients Q, as this implies the conjecture for *virtually* nilpotent quotients.

To see this, assume that the n-$(n+1)$-$(n+2)$ Conjecture holds for some particular quotient group Q and let $\overline{Q}$ be a group that contains Q as a finite index subgroup.

Let

$$N_1 \hookrightarrow \Gamma_1 \overset{\pi_1}{\twoheadrightarrow} \overline{Q}$$
$$N_2 \hookrightarrow \Gamma_2 \overset{\pi_2}{\twoheadrightarrow} \overline{Q}$$

be short exact sequences such that N_1 is of type F_n and Γ_1 and Γ_2 are of type F_{n+1}. We will show that the fibre product of π_1 and π_2, call it P, is of type F_{n+1}. Set $\Gamma_1' := \pi_1^{-1}(Q)$ and $\Gamma_2' := \pi_2^{-1}(Q)$, so we have another pair of short exact sequences

$$N_1 \hookrightarrow \Gamma_1' \overset{\pi_1|}{\twoheadrightarrow} Q$$
$$N_2 \hookrightarrow \Gamma_2' \overset{\pi_2|}{\twoheadrightarrow} Q.$$

Now Γ_1' and Γ_2' have finite index in Γ_1 and Γ_2 respectively, so they are of type F_{n+1} (Proposition A.22). We can therefore apply the assumption that the n-$(n+1)$-$(n+2)$ Conjecture holds for Q to obtain that $P' \leq \Gamma_1' \times \Gamma_2'$, the fibre product associated to the latter pair of short exact sequences, is of type F_{n+1}. But P' is just the intersection of P with the finite index subgroup $\Gamma_1' \times \Gamma_2' \leq \Gamma_1 \times \Gamma_2$, so P' has finite index in P. Thus P is of type F_{n+1} as well (Proposition A.22 again).

The assumptions in the n-$(n+1)$-$(n+2)$ Conjecture Before we embark on proving the first interesting cases of the conjecture in the next chapter, we should make some brief remarks on the necessity of the assumptions in the n-$(n+1)$-$(n+2)$ Conjecture.

There are, for each n, groups that are of type F_n but not of type F_{n+1}. The first such examples to be constructed were given by Stallings [Sta63] in the case $n = 2$ and Bieri [Bie76, Bie81] in the general case. These groups can be defined as follows. Let $F_2 := \langle a, b \mid - \rangle$ be a free group on two generators and $\phi : F_2 \to \mathbb{Z}$ the homomorphism which maps a and b to some fixed generator $t \in \mathbb{Z}$. Then let $\phi^n : (F_2)^n \to \mathbb{Z}$ be the homomorphism from a direct product of n free groups to $\mathbb{Z}$ that restricts to ϕ on each factor. The kernel $\mathrm{SB}_n := \ker \phi^n$ is the n^{th} Stallings-Bieri group and SB_n is of type F_{n-1} but not F_n (see Section A.2 or the references quoted above).

It is easy to see that the conditions on Γ_1 and Γ_2 in the conjecture cannot be weakened to "type F_n". Take Γ_1 to be any group of type F_n but not F_{n+1} (for example $\Gamma_1 \cong \mathrm{SB}_{n+1}$) and pick Γ_2 to be any group of type F_{n+1}. Then let Q be the trivial group, so π_1 and π_2 are the trivial homomorphisms. Then $N_1 = \Gamma_1$ is of type F_n. The fibre product associated to π_1 and π_2 is $\Gamma_1 \times \Gamma_2$, which is of type F_n but not F_{n+1} (Proposition A.24).

More subtly, it will follow from our later results that whenever N_1, Γ_1 and Q are chosen to be of type F_n, F_n and F_{n+1} respectively, the fibre product is of type F_{n+1} *only if* Γ_1 is (Proposition 5.3).

To see that the condition on N_1 in the n-$(n+1)$-$(n+2)$ Conjecture cannot be weakened to F_{n-1}, take $P := \mathrm{SB}_{n+1} \leq (F_2)^{n+1}$. Writing $(F_2)^{n+1} = (F_2)^n \times F_2$, we can consider SB_{n+1} a subgroup of the direct product of $(F_2)^n$ and F_2. It is easy to check that this is in fact a subdirect product and that $\mathrm{SB}_{n+1} \cap (F_2)^n = \mathrm{SB}_n$ and $\mathrm{SB}_{n+1} \cap F_2 = \mathrm{SB}_1$. Therefore SB_{n+1} is the fibre product associated to short exact sequences

$$\mathrm{SB}_n \hookrightarrow (F_2)^n \twoheadrightarrow \mathbb{Z}$$
$$\mathrm{SB}_1 \hookrightarrow F_2 \twoheadrightarrow \mathbb{Z}$$

The middle groups are of type F_∞, as is the quotient and SB_n is of type F_{n-1}. But the fibre product $P = \mathrm{SB}_{n+1}$ is not of type F_{n+1}.

While this shows that the condition on N_1 cannot be weakened while keeping the other conditions fixed, there is a possibility that the condition on N_1 may be weakened if if an accordingly stronger finiteness condition on N_2 is introduced. This in fact happens for the abelian case of the conjecture treated in Chapter 8: Here the condition that N_1 be of type F_{n-1} is replaced by the condition that N_1 be of type F_k and N_2 of type F_{n-1-k} for some k with $0 \leq k \leq n-1$.

The situation is a lot more involved for the F_{n+2}-condition on the quotient group Q, so we will defer this discussion until Section 6.3.

4

The 1-2-3 Theorem

In this section we will turn to the first truly interesting case of the n-$(n+1)$-$(n+2)$ Conjecture: The 1-2-3 Theorem.

This was first established in the symmetrical case by Baumslag, Bridson, Miller and Short in [BBMS00] and then in the general case in [BHMS08, Section 2] (the latter reference contains a sketch of the proof, which is essentially the same as in the symmetrical case; a more detailed account of the same proof is given in [Dis08]).

The original proof is very transparent and there would be little sense in merely copying it here. There are however some very good reasons for including a proof here: The 1-2-3 case is the lowest dimension in which the difficulties of proving the Conjecture become apparent and simultaneously the highest dimension in which a full proof is known. Therefore it is not surprising that many of our methods to prove the Conjecture in higher dimensions are inspired by the 1-2-3 case. The issue then with the proof in [BBMS00], in view of our goals, is that at first glance it does not present any clear opportunities of generalisation to higher dimensions. Our objective here is expressly not to give the easiest or the most succinct proof of the 1-2-3 Theorem, for which we happily refer the reader to [BBMS00] and [BHMS08]. Rather, we aspire to give a proof which readily suggests higher-dimensional generalisations and so can act as a guide to those we present afterwards. As a result, our proof retains the basic idea of [BBMS00], but at many points we will interpret the argument in a new way, present it in a different context, split it up into finer steps or dwell longer on the details.

4.1 Building presentations from short exact sequences

Like the proof of the 0-1-2 Lemma, all the results on the n-$(n+1)$-$(n+2)$ Conjecture rely on a careful analysis of the relationships between the finiteness properties of the groups in a short exact sequence.

So the natural place to start our discussion of the 1-2-3 Theorem is a description of how to construct a presentation of a group from given presentations of a normal subgroup and the corresponding quotient group. The construction is standard, but we describe it in some detail to fix notation for the following sections.

Notation. We will give a presentation of a group Γ in the form $\Gamma \cong \langle \mathcal{X} \mid R, \partial \rangle$. This means that $\mathcal{X}$ and R are sets, $\partial : R \to F(\mathcal{X})$ is a set map and there is some (usually implicit) epimorphism $F(\mathcal{X}) \twoheadrightarrow \Gamma$ with kernel $\langle\langle \partial(R) \rangle\rangle$, so that $\Gamma \cong F(\mathcal{X})/\langle\langle \partial(R) \rangle\rangle$. We will usually not consider $\mathcal{X}$ to be a subset of Γ nor R to be a subset of $F(\mathcal{X})$.[1]

Let
$$N \hookrightarrow \Gamma \overset{\pi}{\twoheadrightarrow} Q$$
be a short exact sequence of groups and let
$$\begin{aligned} N &\cong \langle \mathcal{X} \mid R_N, \partial_N \rangle \\ Q &\cong \langle \mathcal{Z} \mid R_Q, \partial_Q \rangle \end{aligned}$$
be presentations for N and Q. Denote by $p_N : F(\mathcal{X}) \twoheadrightarrow N$ and $p_Q : F(\mathcal{Z}) \twoheadrightarrow Q$ the implied epimorphisms.

Let $\widetilde{\mathcal{Z}}$ be a set isomorphic to $\mathcal{Z}$ with a bijection denoted by $\tilde{z} \leftrightarrow z$. Pick for each $z \in \mathcal{Z}$ an element $\gamma_z \in \Gamma$ with $\pi(\gamma_z) = p_Q(z)$ and define a map $p_\Gamma : \mathcal{X} \cup \widetilde{\mathcal{Z}} \to \Gamma$ by $p_\Gamma(x) = p_N(x)$ for $x \in \mathcal{X}$ and $p_\Gamma(\tilde{z}) = \gamma_z$ for $\tilde{z} \in \widetilde{\mathcal{Z}}$. Then extend p_Γ to a homomorphism $F(\mathcal{X} \cup \widetilde{\mathcal{Z}}) \to \Gamma$ (still denoted by p_Γ). It is clear that this map is an epimorphism. For succinctness we will set
$$F(\mathcal{X}) =: F_N, \quad F(\mathcal{Z}) =: F_Q, \quad F(\mathcal{X} \cup \widetilde{\mathcal{Z}}) =: F_\Gamma.$$

The inclusion map $\mathcal{X} \hookrightarrow \mathcal{X} \cup \widetilde{\Upsilon}$ induces a map $F_N \hookrightarrow F_\Gamma$, which we will take to identify F_N as a subgroup of F_Γ. Furthermore, let $\pi^1 : F_\Gamma \twoheadrightarrow F_Q$ be the epimorphism with $\pi^1(x) = 1$ for $x \in \mathcal{X}$ and $\pi^1(\tilde{z}) = z$ for $\tilde{z} \in \widetilde{\mathcal{Z}}$. We denote the natural splitting

[1] This convention does admittedly make notation more cumbersome at first, but we do not use it for mere misplaced pedantry: Once we introduce higher-dimensional concepts, sloppy notation can quickly make life miserable. More precisely: We will later in this chapter consider the second homotopy group of the 2-dimensional presentation complex associated to some presentation. Naturally, we would like the generators in a presentation to correspond to the edges and the relators to correspond to the 2-cells in this complex. Identifying $x \in \mathcal{X}$ with its image in Γ or $r \in R$ with its image in $F(\mathcal{X})$ amounts to identifying an edge with the corresponding element in the fundamental group or a 2-cell with its boundary, respectively. This is not much of an issue as long as one deals only with elements of the presented group (i.e. the fundamental group of the complex). But it is obviously impractical to use the same notation for a cell and its boundary when trying to investigate more complicated topological properties of a complex. And once we start looking at higher-dimensional skeleta of $B\Gamma$-complexes in later chapters, using careful notation becomes completely unavoidable anyway.

homomorphism by $\sigma : F_Q \to F_\Gamma$, sending $z \in Z$ to $\tilde{z}$.

$$\begin{tikzcd}
F_N \arrow[r, hook] \arrow[d, "p_N"] & F_\Gamma \arrow[r, two heads, "\pi^1"'] \arrow[d, two heads, "p_\Gamma"] & F_Q \arrow[l, bend right, "\sigma"'] \arrow[d, two heads, "p_Q"] \\
N \arrow[r, hook] & \Gamma \arrow[r, two heads, "\pi"] & Q
\end{tikzcd}$$

We will now give a presentation for Γ by specifying three sets of relators R^0_Γ, R^1_Γ and R^2_Γ along with a map $\partial_\Gamma : R^0_\Gamma \cup R^1_\Gamma \cup R^2_\Gamma \to F_\Gamma$ such that $\ker p_\Gamma = \langle\langle \partial_\Gamma(R^0_\Gamma \cup R^1_\Gamma \cup R^2_\Gamma)\rangle\rangle$.

First, set $R^0_\Gamma := R_N$ and $\partial_\Gamma(r) = \partial_N(r)$ for $r \in R^0_\Gamma$.

Second, R^1_Γ contains one element $r_{x,z,\varepsilon}$ for each pair $x \in \mathcal{X}$ and $z^\varepsilon \in \mathcal{Z}^{\pm 1}$. For each such x and z^ε, set

$$d_{x,z,\varepsilon} := \tilde{z}^\varepsilon x \tilde{z}^{-\varepsilon}.$$

Since N is normal, we have $p_\Gamma(d_{x,z,\varepsilon}) \in N$, so we can pick an element $c_{x,z,\varepsilon} \in F(\mathcal{X})$ with $p_\Gamma(d_{x,z,\varepsilon}) = p_N(c_{x,z,\varepsilon})$. Set $\partial_\Gamma(r_{x,z,\varepsilon}) := d_{x,z,\varepsilon}c^{-1}_{x,z,\varepsilon}$.

Finally, R^2_Γ contains one relation r_s for every $s \in R_Q$. For any $s \in R_Q$ let $d_s := \sigma(\partial_Q(s))$. Then $\pi p_\Gamma(d_s) = 1$, so $p_\Gamma(d_s) \in N$. Let $c_s \in F(\mathcal{X})$ be an element with $p_N(c_s) = p_\Gamma(d_s)$ and set $\partial_\Gamma(r_s) := d_s c_s^{-1}$.

Lemma 4.1. *Let $w \in F_\Gamma$ with $p_\Gamma(w) \in N$. Then w is equivalent[2] modulo R^2_Γ to some element $w' \in F_\Gamma$ with $\pi^1(w') = 1$.*

Proof. We have $p_Q\pi^1(w) = \pi p_\Gamma(w) = 1$, so $\pi^1(w) \in \ker p_Q$. Thus we can write

$$\pi^1(w) = \prod_{i=1}^{n} u_i \partial_Q(s_i) u_i^{-1}, \quad \text{for some } u_i \in F_Q, s_i \in R_Q.$$

Let $v := \prod_{i=1}^{n} \sigma(u_i) r_{s_i} \sigma(u_i)^{-1}$. Then $\pi^1(v) = \pi^1(w)$, so $w' := wv^{-1}$ is equivalent to w modulo R^2_Γ and lies in $\ker \pi^1$. □

Lemma 4.2. *Let $w \in F_\Gamma$. Then $w = uw_xw_z$ for some $u \in \langle\langle \partial_\Gamma(R^1_\Gamma)\rangle\rangle$, $w_x \in F(\mathcal{X})$ and $w_z = \sigma\pi^1(w) \in F(\tilde{\mathcal{Z}})$.*

Proof. Let $K := \langle\langle \partial_\Gamma(R^1_\Gamma)\rangle\rangle$ and $p : F_\Gamma \to F_\Gamma/K$ the quotient map. Then

$$p(\tilde{z}^\varepsilon x) = p(c_{x,z,\varepsilon}\tilde{z}^\varepsilon), \quad \text{for all } \tilde{z}^\varepsilon \in \tilde{\mathcal{Z}}^{\pm 1} \text{ and } x \in \mathcal{X}.$$

We can use these relations inductively to swap all the $\tilde{z}$-letters in a word to the right, that is, to find a word w_xw_z with $w_x \in F(\mathcal{X})$ and $w_z \in F(\tilde{\mathcal{Z}})$ and $\pi(w) = \pi(w_xw_z)$. Then $w = uw_xw_z$ for some $u \in K$. Since $\pi^1(R^1_\Gamma) = 1$, we have $\pi^1(w) = \pi^1(w_z)$, and clearly $\sigma\pi^1(w_z) = w_z$, so $\sigma\pi^1(w) = w_z$. □

[2] by which, of course, we mean that $w^{-1}w' \in \langle\langle \partial_\Gamma(R^2_\Gamma)\rangle\rangle$

> LEMMA 4.3. *Let $w \in F_\Gamma$ with $\pi^1(w) = 1$. Then w is equivalent modulo R^1_Γ to some element $w' \in F_N$.*

Proof. Use Lemma 4.2 to write $w = uw_x w_z$ with w_x, w_z and u as in the statement of that lemma. Now we have $1 = \pi^1(w) = \pi^1(w_z)$. But clearly, π^1 is an isomorphism when restricted to $F(\widetilde{\mathcal{Z}})$, so $w_z = 1$ and $w = uw_x$. □

> LEMMA 4.4. *If $w \in F(\mathcal{X}) \leq F_\Gamma$ with $p_\Gamma(w) = 1$, then w is equivalent to 1 modulo R^0_Γ.*

Proof. We can consider w an element of F_N. Then $p_N(w) = p_\Gamma(w) = 1$, so $w \in \langle\!\langle \partial_N(R_N) \rangle\!\rangle_{F(\mathcal{X})}$. By definition, $\partial_N(R_N) = \partial_\Gamma(R^0_\Gamma)$, so certainly $w \in \langle\!\langle \partial_\Gamma(R^0_\Gamma) \rangle\!\rangle_{F_\Gamma}$. □

This shows that $R_\Gamma := R^0_\Gamma \cup R^1_\Gamma \cup R^2_\Gamma$ is a complete set of relators for Γ: If $w \in F_\Gamma$ with $p_\Gamma(w) = 1$, then w is equivalent modulo R^2_Γ to some w' with $\pi^1(w') = 1$ (by Lemma 4.1), which is equivalent modulo R^1_Γ to some $w'' \in F(\mathcal{X})$ (by Lemma 4.3) and this is equivalent to 1 modulo R^0_Γ (by Lemma 4.4).

Hence we have constructed a presentation for Γ from the presentations for N and Q. Note the choices that this construction depended on: First, we chose $p_\Gamma(\tilde{z})$ to be any lift of $p_Q(z)$ along π; then, when defining the relators, we had some freedom in picking the "correction words" $c_{x,z,\varepsilon} \in F(\mathcal{X})$ for each $r_{x,z,\varepsilon} \in R^1_\Gamma$ and $c_s \in F(\mathcal{X})$ for each $r_s \in R^2_\Gamma$.

4.2 Reducing to the "free Γ_2" case

Recall the siutation of the 1-2-3 Theorem:

$$\begin{tikzcd}
 & & N_2 \arrow[d, hook] \\
N_1 \arrow[r, hook] \arrow[d, equal] & P \arrow[r, two heads] \arrow[d, two heads] & \Gamma_2 \arrow[d, two heads] \\
N_1 \arrow[r, hook] & \Gamma_1 \arrow[r, two heads] & Q
\end{tikzcd}$$

We are given two short exact sequences (the bottom row and the right-hand column) with Γ_1 and Γ_2 finitely presented, N_1 finitely generated and Q of type F_3. We want to show that the associated fibre product P is finitely presented.

Transferring finite presentability As in the 0-1-2 Lemma, the proof will be based on comparing the two short exact rows in the diagram. Looking at the middle row, we know that N_1 is finitely generated and Γ_2 is finitely presented.

Using the presentations from the previous section we can prove a useful lemma, which shows that if we have a map between two short exact sequences of this type with isomorphic kernels, then finite presentability can be "transferred" from one to the other. Precisely:

Lemma 4.5. *Let $N \hookrightarrow P \overset{\pi}{\twoheadrightarrow} \Gamma$ be a short exact sequence with N finitely generated and Γ finitely presented. Then P is finitely presented if and only if there is a second short exact sequence $N' \hookrightarrow P' \overset{\pi'}{\twoheadrightarrow} \Gamma'$ with P' finitely presented and a homomorphism $\phi : P' \to P$ which restricts to an isomorphism $\phi|_{N'} : N' \to N$.*

$$\begin{tikzcd} N' \arrow[r, hook] \arrow[d, two heads, tail, "\phi|"] & P' \arrow[r, two heads, "\pi'"] \arrow[d, "\phi"] & \Gamma' \arrow[d] \\ N \arrow[r, hook] & P \arrow[r, two heads, "\pi"] & \Gamma \end{tikzcd}$$

Proof. The "only if" is obvious: ϕ can be chosen as the identity map.

Pick presentations

$$\mathcal{P}_N = \langle \mathcal{X} \mid R_N, \partial_N \rangle \cong N$$
$$\mathcal{P}_\Gamma = \langle \mathcal{Z} \mid R_\Gamma, \partial_\Gamma \rangle \cong \Gamma$$

with $\mathcal{X}$, $\mathcal{Z}$ and R_Γ finite. Let

$$\mathcal{P}_P = \langle \mathcal{X} \cup \widetilde{\mathcal{Z}} \mid R_P^0 \cup R_P^1 \cup R_P^2, \partial_P \rangle$$

be a presentation for P as constructed in the previous section. Then the sets R_P^1 and R_P^2 are finite, but $R_P^0 = R_N$ will typically not be.

Like earlier, let $p_N : F(\mathcal{X}) \twoheadrightarrow N$ denote the epimorphism implied in the presentation $\mathcal{P}_N$. Now pick for each $x \in \mathcal{X}$ an $n'_x \in N'$ with $\phi(n'_x) = p_N(x)$. Then choose some finite set $\Upsilon \subset P'$ so that $\{n'_x \mid x \in \mathcal{X}\} \cup \Upsilon$ generates P' (for example Υ could simply be chosen as a finite generating set of P'). Let $F_{P'} := F(\mathcal{X} \cup \Upsilon)$ and define an epimorphism $p_{P'} : F_{P'} \twoheadrightarrow P'$ by $p'_P(x) = n'_x$ for $x \in \mathcal{X}$ and $p_{P'}(y) = y$ for $y \in \Upsilon$.

Since P' is finitely presented, $\ker p_{P'}$ is finitely normally generated. Let $\partial_{P'} : R_{P'} \to F_{P'}$ be a (set) map from a finite set $R_{P'}$ to $F_{P'}$, such that $\langle\langle \partial_{P'}(R_{P'}) \rangle\rangle = \ker p_{P'}$. In other words,

$$P' \cong \langle \mathcal{X} \cup \Upsilon \mid R_{P'}, \partial_{P'} \rangle.$$

Now define a map $\widetilde{\phi} : F_{P'} \to F_P$ by setting $\widetilde{\phi}(x) = x$ and choosing $\widetilde{\phi}(y)$ so that $p_P(\widetilde{\phi}(y)) = \phi(p_{P'}(y))$. Then the following diagram commutes:

$$\begin{tikzcd} F(\mathcal{X}) \arrow[r, equal] \arrow[d, hook] & F(\mathcal{X}) \arrow[d, hook] \\ F_{P'} \arrow[r, "\widetilde{\phi}"] \arrow[d, two heads, "p_{P'}"] & F_P \arrow[d, two heads, "p_P"] \\ P' \arrow[r, "\phi"] & P \end{tikzcd}$$

We claim that

$$\ker p_P = \langle\langle \widetilde{\phi}(\partial_{P'}(R_{P'})) \cup \partial_P(R_P^1) \cup \partial_P(R_P^2) \rangle\rangle,$$

which would finish the proof that P is finitely presented, since $R_{P'}$, R_P^1 and R_P^2 are finite. Let $w \in \ker p_P$. By Lemmas 4.1 and 4.3, w is equivalent, modulo R_P^2 and R_P^1 to some $w' \in F(\mathcal{X})$. More precisely,

$$ww'^{-1} \in \langle\langle \partial_P(R_P^1) \cup \partial_P(R_P^2) \rangle\rangle.$$

Now $\phi(p_{P'}(w')) = p_P(w') = 1$ and since $p_{P'}(F(\mathcal{X})) = N'$ and ϕ is injective on N', we have $p_{P'}(w') = 1$. Thus $w' \in \langle\langle \partial_{P'}(R_{P'}) \rangle\rangle$ and so

$$w' = \widetilde{\phi}(w') \in \langle\langle \widetilde{\phi}(\partial_{P'}(R_{P'})) \rangle\rangle.$$

Hence

$$w \in \langle\langle \widetilde{\phi}(\partial_{P'}(R_{P'})) \cup \partial_P(R_P^1) \cup \partial_P(R_P^2) \rangle\rangle. \qquad \square$$

The split case and reducing to free Γ_2 How can we apply this criterion in our situation? One obvious case is, when the second short exact sequence $N_2 \hookrightarrow \Gamma_2 \overset{\pi_2}{\twoheadrightarrow} Q$ splits. Let $\sigma_2 : Q \to \Gamma_2$ be a homomorphism with $\pi_2\sigma_2 = \mathrm{id}_Q$. Define $\sigma : \Gamma_1 \to P, \gamma_1 \mapsto (\gamma_1, \sigma_2\pi_1(\gamma_1))$. Then σ maps N_1 isomorphically to $N_1 \times 1 \le P$. So Lemma 4.5 (with $\phi = \sigma$) shows that P is finitely presented. Thus we have proved the 1-2-3 Theorem in the split case.

Of course, a homomorphism $\Gamma_1 \to P$ which maps $N_1 \le \Gamma_1$ isomorphically to $N_1 \times 1 \le P$ exists if *and only if* π_2 splits, so there is no hope of finding such a homomorphism in the general case.

However, we can still put Lemma 4.5 to good use in the general case: We will show that we can reduce to the case in which Γ_2 is finitely generated free. So assume we are in the generic situation of the 1-2-3 Theorem as at the start of the section again. Pick a finitely generated free group F which maps onto Γ_2 via an epimorphism $p : F \to \Gamma_2$ and set $\pi_2' := \pi_2 \circ p$. Let $P' \le \Gamma_1 \times F$ denote the fibre product associated to π_1 and π_2'. Then p induces a homomorphism $\Gamma_1 \times F \to \Gamma_1 \times \Gamma_2$ which maps P' onto P and sends $N_1 \times 1 \le \Gamma_1 \times F$ isomorphically to $N_1 \times 1 \le \Gamma_1 \times \Gamma_2$. If the 1-2-3 Theorem is true whenever the second factor is free then P' is finitely presented and we can apply Lemma 4.5 to the "top" sequences in the following diagram:

$$\begin{tikzcd}
 & N_1 \arrow[rr, hook] \arrow[dd, equal] & & P \arrow[rr, two heads] \arrow[dd] & & \Gamma_2 \arrow[dd] \\
N_1 \arrow[ur, equal] \arrow[rr, hook, crossing over] \arrow[dd, equal] & & P' \arrow[ur, two heads] \arrow[rr, two heads, crossing over] & & F \arrow[ur, two heads] & \\
 & N_1 \arrow[rr, hook] & & \Gamma_1 \arrow[rr, two heads] & & Q \\
N_1 \arrow[ur, equal] \arrow[rr, hook] & & \Gamma_1 \arrow[ur, equal] \arrow[rr, two heads] \arrow[from=uu, crossing over] & & Q \arrow[ur, equal] \arrow[from=uu, crossing over] &
\end{tikzcd}$$

Therefore we obtain

‖ Proposition 4.6. *The* 1-2-3 *Conjecture is true if it is true for free* Γ_2. □

To get some idea of why it might be useful to have Γ_2 be free, we can make a quick obervation: Note that, if $\Gamma_2 = F$ is free, the short exact sequence $N_1 \hookrightarrow P \overset{p_2}{\twoheadrightarrow} F$ splits, so the fibre product P is a semidirect product $N_1 \rtimes F$. We can describe the defining conjugation action of F on N_1 in terms of the original short exact sequences: If $\mathcal{X} \subset F$ is a basis of F, the conjugation action is determined by picking for each $x \in \mathcal{X}$ a lift $\tilde{x} \in P$ with $p_2(\tilde{x}) = x$. Now, as we observed in the proof of the 0-1-2 Lemma, $\tilde{x}$ acts on $N_1 \leq P$ by conjugation the same as $p_1(\tilde{x})$ acts on $N_1 \leq \Gamma_1$ and $p_1(\tilde{x})$ is a lift of $\pi_2(x)$ along π_1. So the conjugation action of F on N_1 can be given by picking for each $x \in \mathcal{X}$ a lift of $\pi_2(x)$ along π_1.

Presenting fibre products with one free factor However, to really appreciate why reducing the original statement of the 1-2-3 Theorem to the case of free Γ_2 simplifies matters, we will now see what the construction from Section 4.1 yields, when applied to the short exact sequences involved in a fibre product with one free factor.

So assume we are given two short exact sequences

$$N_1 \hookrightarrow \Gamma \overset{\pi_1}{\twoheadrightarrow} Q$$
$$N_2 \hookrightarrow F \overset{\pi_2}{\twoheadrightarrow} Q$$

with F finitely generated free, along with a presentation $N_1 \cong \langle \mathcal{X} \mid R_{N_1}, \partial_{N_1} \rangle$, a basis Z for F and a presentation for Q on the generating set $\mathcal{Z}$, say $Q \cong \langle \mathcal{Z} \mid R_Q, \partial_Q \rangle$ (i.e. $\partial_Q : R_Q \to F$). Let P denote the associated fibre product.

Later, of course, we will impose the finiteness conditions from the 1-2-3 Conjecture, namely that N_1 be finitely generated, Γ be finitely presented and Q of type F_3. But for now, we do not need to worry about this.

Our goal here is to apply the construction from Section 4.1 to the rows in the diagram

$$\begin{tikzcd}
N_1 \cong \langle \mathcal{X} \mid R_{N_1} \rangle \arrow[r, hook] \arrow[d, equal] & P \arrow[r, two heads, "p_2"] \arrow[d, "p_1"] & F \cong \langle \mathcal{Z} \mid - \rangle \arrow[d, "\pi_2"] \\
N_1 \cong \langle \mathcal{X} \mid R_{N_1} \rangle \arrow[r, hook] & \Gamma \arrow[r, two heads, "\pi_1"] & Q \cong \langle \mathcal{Z} \mid R_Q \rangle
\end{tikzcd}$$

We will see that by making careful choices during the construction, we can obtain some particularly insightful presentations for the groups Γ and P.

First, let $\tilde{\mathcal{Z}}$ be in bijection with $\mathcal{Z}$, as in Section 4.1. As before, we denote $F_P = F(\mathcal{X} \cup \tilde{\mathcal{Z}}) = F_\Gamma$. Now we have to define epimorphisms $p_P : F_P \to P$ and $p_\Gamma : F_\Gamma \to \Gamma$. Pick lifts $\alpha_z \in P$ for each $z \in Z$ with $p_2(\alpha_z) = z$. Then set $\gamma_z := p_1(\alpha_z)$ and note that $\pi_1(\gamma_z) = \pi_2(z)$. Now define $p_P : \mathcal{X} \cup \tilde{\mathcal{Z}} \to P$ by $p_P(x) = p_N(x)$ for $x \in X$ and

$p_P(\tilde{z}) = \alpha_z$ and extend this to an epimorphism $p_P : F(\mathcal{X} \cup \tilde{\mathcal{Z}}) \to P$. Similarly, define $p_\Gamma : \mathcal{X} \cup \tilde{\mathcal{Z}} \to P$ by $p_\Gamma(x) = p_N(x)$ for $x \in X$ and $p_\Gamma(\tilde{z}) = \gamma_z$, so that $p_\Gamma = p_1 p_P$.

$$\begin{tikzcd}
F(\mathcal{X} \cup \tilde{\mathcal{Z}}) \arrow[r, two heads] \arrow[d, "p_P"] \arrow[dd, bend right, "p_\Gamma"'] & F(\mathcal{Z}) \arrow[d, "\cong"] \\
P \arrow[r, "p_2"] \arrow[d, "p_1"] & F \arrow[d, "\pi_2"] \\
\Gamma \arrow[r, "\pi_1"] & Q
\end{tikzcd}$$

Now we follow the procedure from the Section 4.1 to describe the relators. In both P and Γ we have the relators $R_P^0 = R_\Gamma^0 = R_{N_1}$ coming from N_1 with the boundaries given by $\partial_P(r) = \partial_\Gamma(r) = \partial_{N_1}(r)$ for all $r \in R_{N_1}$.

Now note that for any $n \in N_1$ and $\alpha \in P$ we have $\alpha n \alpha^{-1} = p_1(\alpha) n p_1(\alpha)^{-1}$. Since $p_\Gamma = p_1 \circ p_P$, we have $p_P(\tilde{z}^\varepsilon x \tilde{z}^{-\varepsilon}) = p_\Gamma(\tilde{z}^\varepsilon x \tilde{z}^{-\varepsilon})$ in N_1 for all $x \in \mathcal{X}$ and $z^\varepsilon \in \mathcal{Z}^{\pm 1}$. Thus we can pick the same correction words $c_{x,z,\varepsilon}$ for both presentations. Therefore we have the same relations $w_P(r_{x,z,\varepsilon}) = w_\Gamma(r_{x,z,\varepsilon}) = \tilde{z}^\varepsilon x \tilde{z}^{-\varepsilon} c_{x,z,\varepsilon}$ in both P and Γ.

Finally, the presentation for Γ has additional relations of type R_Γ^2, which are not present in P (there are no relations of type R_P^2). Summarising, we have constructed presentations

$$\begin{aligned} P &\cong \langle \mathcal{X} \cup \tilde{\mathcal{Z}} \mid R_\Gamma^0 \cup R_\Gamma^1, \partial_\Gamma \rangle \\ \Gamma &\cong \langle \mathcal{X} \cup \tilde{\mathcal{Z}} \mid R_\Gamma^0 \cup R_\Gamma^1 \cup R_\Gamma^2, \partial_\Gamma \rangle \end{aligned}$$

Thus Γ is simply the quotient of P by the relations $\partial_\Gamma(R_\Gamma^2)$. Here we can see the fruits of reducing the general 1-2-3 Conjecture to the case of free Γ_2: In this particular situation, we can obtain natural presentations of P and Γ arising from the short exact sequences and the given presentations for N_1 and Q, which differ precisely in the extra relations R_Γ^2, coming from relators in Q. Another way to put this is that the presentation 2-complex for P is a subcomplex of that for Γ.

4.3 The 1-2-3 situation

Let us see what happens when we impose the finiteness restrictions from the 1-2-3 Conjecture. Assume that N_1 is finitely generated, Γ and F are finitely presented and Q is of type F_3. Then we can choose the generating sets $\mathcal{X}$ for N_1 and $\mathcal{Z}$ for F to be finite. Also, we can choose the presentation $Q \cong \langle \mathcal{Z} \mid R_Q, w_Q \rangle$ to be finite (using that a finitely presented group has a finite presentation with respect to any finite generating set).

With these choices, the sets R_Γ^1 and R_Γ^2 in the above presentations for P and Γ will be finite. The set $R_\Gamma^0 = R_N$ of course, will usually be infinite. However, since Γ is finitely presented, finitely many of the relations from R_Γ^0 suffice to present Γ. More precisely, there is a finite subset $R_\Gamma^{0\prime} \subset R_\Gamma^0$ such that $\ker p_\Gamma = \langle\langle \partial_\Gamma(R_\Gamma^{0\prime} \cup R_\Gamma^1 \cup R_\Gamma^2)\rangle\rangle$, i.e. $\Gamma \cong \langle \mathcal{X} \cup \tilde{\mathcal{Z}} \mid R_\Gamma^{0\prime} \cup R_\Gamma^1 \cup R_\Gamma^2, \partial_\Gamma\rangle$. In general, $\langle \mathcal{X} \mid R_\Gamma^{0\prime}, \partial_\Gamma\rangle$ is not a presentation for N, of course. Nor is $\langle \mathcal{X} \cup \tilde{\mathcal{Z}} \mid R_\Gamma^{0\prime} \cup R_\Gamma^1, \partial_\Gamma\rangle$ usually a presentation for P. In other words, not every $a \in \ker p_P$ can be written as a product of conjugates of the $R_\Gamma^{0\prime}$ and R_Γ^1 alone, but it *can* be written as a product of conjugates

$$a = \prod_{i=1}^n w_i \partial_\Gamma(r_i) w_i^{-1}, \quad r_i \in R_\Gamma^{0\prime} \cup R_\Gamma^1 \cup R_\Gamma^2, w_i \in F_\Gamma. \tag{4.1}$$

We want to show that P is finitely presented. This is equivalent to showing that there is some finite subset $R_\Gamma^{0\prime\prime} \subset R_\Gamma^0$ such that $\ker p_P = \langle\langle \partial_\Gamma(R_\Gamma^{0\prime\prime} \cup R_\Gamma^1)\rangle\rangle$. If $a \in \ker p_P$ then it follows from Lemmas 4.1 and 4.3 that a is equivalent modulo R_Γ^1 to some word $a' \in F(\mathcal{X})$. So to show that $\ker p_P = \langle\langle w_\Gamma(R_\Gamma^{0\prime\prime} \cup R_\Gamma^1)\rangle\rangle$ it is enough to show that any word in the x-letters $a \in F(\mathcal{X})$ with $p_P(a) = 1$ is a product of conjugates

$$a = \prod_{i=1}^n w_i \partial_\Gamma(r_i) w_i^{-1}, \quad r_i \in R_\Gamma^{0\prime\prime} \cup R_\Gamma^1, w_i \in F_\Gamma. \tag{4.2}$$

Essentially then, we need to show that we can transform any product (4.1) to one of the form (4.2), getting rid of the terms involving relators from R_Γ^2 while possibly introducing new relators from a finite set $R_\Gamma^{0\prime\prime} \subset R_\Gamma^0$. This is where the F$_3$-condition on Q comes in.

In a moment we will introduce the theory of crossed modules, which formalises the calculus of products of conjugates. But before we do that, let us try first to give a rough geometric sketch of how the F$_3$-condition can help here. We consider the presentation 2-complexes associated to the presentations of N, Γ and Q. For a presentation $\mathcal{P} = \langle \mathcal{X} \mid R, \partial\rangle$, the associated presentation 2-complex has a single 0-cell, 1-cells e_x^1 indexed by $\mathcal{X}$ and 2-cells e_r^2 indexed by R, with e_r^2 attached along the word $\partial(r)$, considered as a loop in the 1-skeleton. In the sequel, we will for convenience sometimes consider $r \in R$ itself as a 2-cell, identifying it with e_r^2.

Apply this construction to the presentations from above

$$\begin{aligned} N &\cong \langle \mathcal{X} \mid R_N, \partial_N\rangle = \mathcal{P}_N \\ Q &\cong \langle \mathcal{Z} \mid R_Q, \partial_Q\rangle = \mathcal{P}_Q \\ \Gamma &\cong \langle \mathcal{X} \cup \tilde{\mathcal{Z}} \mid R_\Gamma^0 \cup R_\Gamma^1 \cup R_\Gamma^2, \partial_\Gamma\rangle = \mathcal{P}_\Gamma \end{aligned}$$

yielding 2-complexes X_N, X_Q and X_Γ respectively. We can treat X_N as a subcomplex of X_Γ, identifying the respective 1-cells associated to the generators $\mathcal{X}$ and the 2-cells associated to the relators $R_N = R_\Gamma^0$. Also, we can define a natural cellular map $\phi : X_\Gamma \to X_Q$, inducing the homomorphism $\pi : \Gamma \to Q$ on fundamental groups: ϕ sends the 1-cell $e_{\tilde{z}}^1 \subset X_\Gamma$ to $e_z^1 \in X_Q$. For the 2-cells, it sends each e_r^2 with $r \in R_\Gamma^0$ to the 0-cell in X_Q, each e_r^2 with $r = r_{x,z,\varepsilon} \in R_\Gamma^1$ to the 1-cell e_z^1 and each e_r^2 with $r = r_s \in R_\Gamma^2$ to e_s^2. Schematically, we may depict the situation as follows:

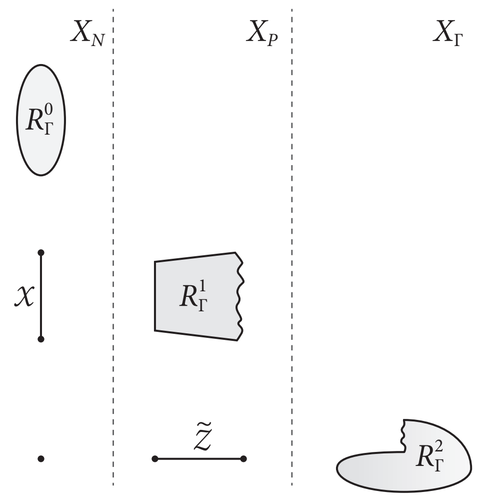

Now X_P, the presentation 2-complex for the presentation $P \cong \langle \mathcal{X} \cup \tilde{\tilde{\mathcal{Z}}} \mid R^0_\Gamma \cup R^1_\Gamma, \partial_\Gamma \rangle$, is just a subcomplex of X_Γ (missing the cells R^2_Γ), and we have $X_N \subset X_P \subset X_\Gamma$. A word $a \in F(\mathcal{X})$ can be represented by a loop $p_a : S^1 \to X_N$. If this word represents the trivial element in N, then we can write it as a product of conjugates of relators, and these products correspond to discs $d : D^2 \to X_\Gamma$ with boundary $d|_{S^1} = p_a$ (we may think of van Kampen diagrams here). Since $R^{0\prime}_\Gamma \cup R^1_\Gamma \cup R^2_\Gamma$ is a complete set of relators for Γ, we can always find such a disc d, which only intersects the 2-cells corresponding to $R^{0\prime}_\Gamma \cup R^1_\Gamma \cup R^2_\Gamma$. The rough idea of the proof of the 1-2-3 Theorem now is this: We will describe a finite set of 3-cells to glue into X_Γ. These can be used to transform discs, by pushing them across the 3-cells. If a disc $d : D^2 \to X_\Gamma$ is pushed across a 3-cell e^3, this yields a new disc $d' : D^2 \to X_\Gamma$, whose image is contained in the union of $\operatorname{im} d$ with the boundary of e^3. We will then give a method to push any disc $d : D^2 \to X_\Gamma$ with boundary contained in X_N across the 3-cells (while keeping the boundary fixed) to yield a new disc d' which does not intersect any of the cells corresponding to R^2_Γ. If d intersected only the 2-cells $R^{0\prime}_\Gamma \cup R^1_\Gamma \cup R^2_\Gamma$, then d' will only intersect the 2-cells $R^{0\prime}_\Gamma \cup R^1_\Gamma$ and those 2-cells from R^0_Γ that lie in the boundary of some 3-cell. Since there are only finitely many 3-cells, each of which can contain only finitely many 2-cells in its boundary, the new disc d' will be contained in $R^{0\prime\prime}_\Gamma \cup R^1_\Gamma$ for some finite subset $R^{0\prime\prime}_\Gamma$. This shows that $R^{0\prime\prime}_\Gamma \cup R^1_\Gamma$ is a complete set of relators for P, establishing that P is finitely presented.

4.4 Crossed modules

One of the main reasons we can actually carry out the idea outlined above, is that we can describe all the geometric notions involved (such as discs in the presentation complex, attaching maps of 3-cells, pushing a disc across a 3-cell) in combinatorial terms, enabling us to do explicit computations.

For general 2-complexes this formulation was essentially introduced by Whitehead [Whi41, Whi46, Whi49], in his program on "combinatorial homotopy". Independently, Reidemeister and Peiffer [Rei49, Pei49] discovered closely related formalisms describing "identities among relators" of a group, which is the special case of Whitehead's work that will interest us here. A good general account of this material is found in [BH82a] (see also [BHS11]).

If $X^{(2)}$ is a 2-complex, then the homotopy classes of discs with boundary in the 1-skeleton $(D^2, S^1) \to (X^{(2)}, X^{(1)})$ form the relative (non-abelian) homotopy group $\pi_2(X^{(2)}, X^{(1)})$. This group fits in an exact sequence (part of the long exact homotopy sequence of the pair $(X^{(2)}, X^{(1)})$)

$$1 \to \pi_2(X^{(2)}) \to \pi_2(X^{(2)}, X^{(1)}) \xrightarrow{\partial} \pi_1(X^{(1)}) \to \pi_1(X^{(2)})$$

Whitehead gave a precise algebraic description of the groups and maps in this sequence. More precisely, he showed that $\pi_2(X^{(2)}, X^{(1)})$ with the boundary map ∂ forms what he called a free crossed $\pi_1(X^{(1)})$-module. This is a deep theorem and we will not recount its proof here, but we will try to give some motivation for the relevant algebraic definitions.

Pre-crossed modules Whitehead's theorem in essence says that discs in the presentation 2-complex associated to $\Gamma \cong \langle \mathcal{X} \mid R, \partial \rangle$ correspond to products of conjugates of the relators[3]

$$\prod_{i=1}^{n} w_i \partial(r_i) w_i^{-1}, \quad w_i \in F(\mathcal{X}), r_i \in R. \tag{4.3}$$

If such a product is trivial in $F(\mathcal{X})$, i.e. corresponds to a sphere in the presentation complex, then it specifies an "identity" between the relators, in the same way as a trivial product of generators in Γ describes a relation between the generators.

To make the notion of a "trivial product of generators" in a group Γ precise, we have to introduce the free group on these generators. Similarly, if we want to speak of a trivial product of conjugates of relators of the form (4.3), we clearly cannot just consider these products as elements of $F(\mathcal{X})$. Instead, we have to introduce a structure whose elements are "formal products of conjugates of relators". This structure is what we will later call a free crossed module. Roughly speaking, this is a group C with a homomorphism $\partial : C \to F$, which maps each "formal product of conjugates" to its value in F. In this way, the situation is not much different from the free group $F(\mathcal{X})$ in a set of generators $\mathcal{X} \subset \Gamma$ with the epimorphism $F(\mathcal{X}) \twoheadrightarrow \Gamma$. But crucially, the crossed module (C, ∂) comes with some extra structure which has to be accounted for, namely a conjugation action of F on C, which is in some way compatible with the homomorphism ∂.

To make all this precise, we start with the following definition:

Definition 4.7. Let F be a group. A pre-crossed F-module is an F-group[4] A along with a homomorphism $\partial : A \to F$ such that

$$\partial(g \cdot a) = g \partial(a) g^{-1}.$$

[3] or rather, equivalence classes of such products, as we will explain shortly

In other words, if we consider F as acting on itself by conjugation, then we ask for ∂ to be a homomorphism of F-groups.

Two easy examples for pre-crossed F-modules are a normal subgroup $N \triangleleft F$ with $\partial : N \to F$ the inclusion map and any $\mathbb{Z}F$-module A with $\partial : A \to F$ the trivial map.

For us, the most important example will be this: Given a presentation $\mathcal{P} = \langle \mathcal{X} \mid R, \partial \rangle$, we define $\Lambda(\mathcal{P})$ to be the free group Λ on the basis $F \times R$ (where $F := F(\mathcal{X})$) together with the homomorphism $\partial : \Lambda \to F$ which sends a basis element $(w, r) \in F \times R$ to wrw^{-1}. Then Λ becomes an F-group by setting $u \cdot (w, r) = (uw, r)$ for each $u \in F$ and $(w, r) \in F \times R$ (and extending the action of u to all of Λ in the unique way that makes $\Lambda \to \Lambda, a \mapsto u \cdot a$ a homomorphism). By definition $\partial(u \cdot (w, r)) = u\partial(w, r)u^{-1}$ and it is easy to show that, in fact, $\partial(u \cdot a) = u\partial(a)u^{-1}$ for all $a \in \Lambda$. Thus (Λ, ∂) is a pre-crossed module. We call $\Lambda(\mathcal{P})$ the *pre-crossed module associated to* $\mathcal{P}$.

We may identify $r \in R$ with $(1, r) \in F \times R$ and thus consider R a subset of Λ. We can then write (w, r) as $w \cdot r$.

More generally, if Γ is any group and $w : R \to \Gamma$ a set map, we can define $\Lambda = F(\Gamma \times R)$ and a map $\partial : \Lambda \to \Gamma$ with $\partial|_R = w|_R$ in the same way as above and Λ is then called the *free pre-crossed* Γ*-module* on w.[5]

Notation. If Γ is a group, A is a Γ-group and $R \subset A$ is some subset, we set

$$\langle\langle R \rangle\rangle_\Gamma := \langle \{ \gamma \cdot r \mid \gamma \in \Gamma, r \in R \} \rangle.$$

If it is unambiguous which group is acting on A, we may also abbreviate this to $\langle\langle R \rangle\rangle = \langle\langle R \rangle\rangle_\Gamma$. Note that this is consistent with the notation for normal subgroups, if we consider a group as acting on itself by conjugation.

Crossed modules If (A, ∂) is any pre-crossed F-module then

$$\partial(aba^{-1}) = \partial(\partial(a) \cdot b), \quad \text{for all } a, b \in A,$$

so $aba^{-1}(\partial(a) \cdot b)^{-1} \in \ker \partial$. We call these elements the *basic Peiffer elements* and the subgroup they generate the *Peiffer subgroup* P. We have:

Lemma 4.8. *Let (A, ∂) be a pre-crossed F-module. Then its Peiffer subgroup P is normal and F-invariant. If $X \subset A$ is a Γ-invariant generating set for A, then P is the normal subgroup generated by the basic Peiffer elements*

$$xyx^{-1}(\partial(x) \cdot y)^{-1}, \quad \text{with } x, y \in X.$$

[4] That is, a group Λ equipped with a homomorphism $\phi : F \to \operatorname{Aut} \Lambda$. We write $w \cdot a := \phi(w)(a)$ for $w \in F$, $a \in \Lambda$.

[5] These are, of course, the free objects in the category of pre-crossed Γ-modules. A morphism between pre-crossed Γ-modules (A, ∂) (A', ∂') is an Γ-group homomorphism $\phi : A \to A'$ with $\phi\partial' = \partial$, and a free pre-crossed module on a set map $\partial : R \to \Gamma$ can be defined by a natural universal mapping property. For details, see [BH82a, Section 3].

For a proof of these basic facts, we refer the reader to [BH82a, Propositions 2 and 3].

In the context of a pre-crossed module $\Lambda(\mathcal{P})$ associated to a presentation $\mathcal{P} = \langle X \mid R, \partial\rangle$ we would like to think of the kernel of ∂ as "identities between relations". The elements of the Peiffer subgroup P then represent trivial identities, those which hold in every presentation. So it is natural to pass to the quotient group Λ/P with the induced map $\bar{\partial} : \Lambda/P \to F$. Since P is Γ-invariant the Γ-action on Λ induces a well-defined Γ-action on Λ/P and clearly, $\bar{\partial}$ is still a homomorphism of F-groups. In other words $(\Lambda/P, \bar{\partial})$ is again a pre-crossed module. However, the Peiffer subgroup of Λ/P is trivial. This is called a *crossed module*.

Definition 4.9. A crossed F-module is a pre-crossed F-module $(A, \bar{\partial})$ with trivial Peiffer subgroup, or in other words,

$$aba^{-1} = \partial(a) \cdot b, \quad \text{for all } a, b \in A.$$

With this terminology, the discussion above can be summarised as saying that to any pre-crossed F-module (A, ∂), there is an induced crossed F-module $(A/P, \bar{\partial})$, where P is the Peiffer subgroup of A and $\bar{\partial} : A/P \to F$ is the induced map.[6]

Definition 4.10. If $\mathcal{P} = \langle \mathcal{X} \mid R, \partial\rangle$ is a presentation with associated pre-crossed module $(\Lambda(\mathcal{P}), \partial)$, then we denote the induced crossed $F(\mathcal{X})$-module by $(C(\mathcal{P}), \bar{\partial})$ and call this the crossed module associated to $\mathcal{P}$.

The kernel of $\bar{\partial}$ is denoted by $\pi(\mathcal{P})$ and is called the module of identities between relators of $\mathcal{P}$ (or the module of identities of $\mathcal{P}$, for short).

More generally, if Γ is a group and (Λ, ∂) the free pre-crossed module on some set map $w : R \to \Gamma$, then we call the induced crossed module $(\Lambda/P, \bar{\partial})$ the free crossed module on the basis w.[7]

From the properties of crossed modules we obtain immediately:

> Lemma 4.11. *Let F be a group and $(A, \bar{\partial})$ a crossed F-module. Then $\pi := \ker \bar{\partial}$ lies in the centraliser of A. In particular, it is abelian and thus an F-module. Furthermore, $N := \operatorname{im} \bar{\partial}$ acts trivially on π, so there is an induced F/N-module structure on π.*

Proof. We have

$$aba^{-1} = \bar{\partial}(a) \cdot b = 1 \cdot b = b, \quad \text{for all } a \in \pi, b \in A.$$

So π lies in the centraliser of A.

[6] One easily checks that this induced crossed F-module satisfies a universal property: If $(B, \bar{\partial}')$ is a crossed F-module and $\phi : A \to B$ a morphism of pre-crossed F-modules then ϕ factors uniquely through the quotient map $A \twoheadrightarrow A/P$ to a morphism of crossed modules $\bar{\phi} : A/P \to B$.

[7] More precisely, we should take $\bar{R}$ to be the image of R under $\Lambda \to \Lambda/P$ and call the induced map $\bar{w} : \bar{R} \to \Gamma$ the basis of the free crossed module. But in fact, one can show that R is mapped injectively under the quotient map $\Lambda \to \Lambda/P$, so we can still regard R as a subset of the induced crossed module. For proofs of this, see [BH82a, Section 3].

But then

$$\partial(b) \cdot a = bab^{-1} = a, \quad \text{for all } a \in \pi, b \in A.$$

So $\operatorname{im} \bar{\partial}$ acts trivially on π. □

Corollary 4.12. *If $\mathcal{P} = \langle \mathcal{X} \mid R, \partial \rangle \cong \Gamma$, then the module of identities $\pi(\mathcal{P})$ is a Γ-module.* □

We can now state the central theorem by Whitehead alluded to at the beginning of the section:

Theorem 4.13. *Let X be a 2-dimensional CW-complex. Let E^2 denote the set of 2-cells and pick for each $e \in E^2$ a characteristic map $c_e : D^2 \to X$. Then $\pi_2(X, X^{(1)})$ with the map $\partial_2 : \pi_2(X, X^{(1)}) \to \pi_1(X^{(1)})$ is a free crossed $\pi_1(X^{(1)})$-module. A basis is given by the set $\{[c_e] \mid e \in E^2\} \subset \pi_2(X, X^{(1)})$ of homotopy classes of the characteristic maps of the 2-cells.* □

For an exposition of Whitehead's original proof, see [Bro80]. A more algebraic proof was given by Ratcliffe [Rat80]).

Consider now the case where $X = X_{\mathcal{P}}$ is a presentation 2-complex associated to $\mathcal{P} = \langle \mathcal{X} \mid R, \partial \rangle \cong \Gamma$. Then $\pi_1(X^{(1)}) \cong F(\mathcal{X})$ and Whitehead's theorem tells us that $\pi_2(X^{(2)}, X^{(1)}) \xrightarrow{\partial_2} F(\mathcal{X})$ is a free crossed $F(\mathcal{X})$-module with a basis given by a set of characteristic maps for the 2-cells. By definition of X, the boundary of the 2-cell associated to some relator $r \in R$ is $\partial(r) \in F(\mathcal{X})$. Therefore, $\pi_2(X^{(2)}, X^{(1)}) \xrightarrow{\partial_2} F(\mathcal{X})$ is isomorphic to the crossed module $C(\mathcal{P}) \xrightarrow{\bar{\partial}} F(\mathcal{X})$ associated to the presentation. Thus we have identified most of the terms in the following excerpt from the long exact homotopy sequence of the pair $(X^{(2)}, X^{(1)})$:

$$\begin{tikzcd}
\pi_2(X^{(1)}) \arrow[r] & \pi_2(X^{(2)}) \arrow[r] & \pi_2(X^{(2)}, X^{(1)}) \arrow[r, "\partial_2"] \arrow[d, "\cong"] & \pi_1(X^{(1)}) \arrow[r] \arrow[d, "\cong"] & \pi_1(X^{(2)}) \arrow[d, "\cong"] \\
& & C(\mathcal{P}) \arrow[r, "\bar{\partial}"] & F(\mathcal{X}) \arrow[r, "p_\Gamma"] & \Gamma
\end{tikzcd}$$

But the higher homotopy groups of any 1-dimensional complex are trivial (since the universal cover of such a complex is contractible and covering maps induce isomorphisms on the higher homotopy groups), so the first group in the sequence vanishes. Therefore, $\pi_2(X^{(2)})$ is the kernel of ∂_2 and can be identified with $\ker \bar{\partial} = \pi(\mathcal{P})$.

Lemma A.20 from the appendix says that a group is of type F_n if and only if some (or, equivalently, every) finite presentation complex has finitely generated second homotopy group. Therefore we can say:

Corollary 4.14. *A group Γ is of type F_3 if for some (or, equivalently, any) finite presentation $\mathcal{P} = \langle \mathcal{X} \mid R, \partial \rangle$, the module $\pi(\mathcal{P})$ of identities of relations of $\mathcal{P}$ is finitely generated as a Γ-module.* □

4.5 Identities among relators in a short exact sequence

Now that we have set up the necessary technology of presentation crossed complexes to analyse identities among relators, we can apply this to the situation of presentations built from short exact sequences.

A description of the generators of the second homotopy module of a such a presentation obtained from a short exact sequence has been given by Baik, Harlander and Pride in [BHP98]. Their description employs the diagrammatical language of "pictures", while we will stick to the purely algebraic description in terms of crossed modules.

Let

$$N \hookrightarrow \Gamma \overset{\pi}{\twoheadrightarrow} Q$$

be a short exact sequence of groups. Assume as before that we are given presentations $\mathcal{P}_N = \langle \mathcal{X} \mid R_N, \partial_N \rangle \cong N$ and $\mathcal{P}_Q = \langle \mathcal{Z} \mid R_Q, \partial_Q \rangle \cong Q$. Let $\mathcal{P}_\Gamma = \langle \mathcal{X} \cup \tilde{\mathcal{Z}} \mid R^0_\Gamma \cup R^1_\Gamma \cup R^2_\Gamma, \partial_\Gamma \rangle$ be the presentation associated to the short exact sequence, as constructed in Section 4.1.

Recall that $R^0_\Gamma = R_N$, and $\partial_\Gamma|_{R^0_\Gamma} = \partial_N$, so we can consider the pre-crossed module $\Lambda(\mathcal{P}_N) = F(F(\mathcal{X}) \times R_N)$ as contained in $\Lambda(\mathcal{P}_\Gamma)$. It is clear that the Peiffer group of $\Lambda(\mathcal{P}_N)$ is contained in that of $\Lambda(\mathcal{P}_\Gamma)$, so this inclusion induces a map $C(\mathcal{P}_N) \to C(\mathcal{P}_\Gamma)$.[8] Furthermore, we can define a natural map $\pi^2 : \Lambda(\mathcal{P}_\Gamma) \to \Lambda(\mathcal{P}_Q)$ by setting

$$\begin{aligned} \pi^2(w \cdot r) &:= 0 && \text{for } w \in F_\Gamma, r \in R^0_\Gamma \cup R^1_\Gamma \\ \pi^2(w \cdot r_s) &:= \pi^1(w) \cdot s && \text{for } w \in F_\Gamma, r_s \in R^2_\Gamma \end{aligned}$$

We have defined the presentation $\mathcal{P}_\Gamma$ in such a way that $\partial_Q \circ \pi^2 = \pi^1 \circ \partial_\Gamma$. This map π^2 thus induces a homomorphism of the free crossed modules $C(\mathcal{P}_\Gamma) \to C(\mathcal{P}_Q)$, which we will denote $\bar{\pi}^2$. In fact, this is a morphism of crossed F_Γ-modules if we regard $C(\mathcal{P}_Q)$ as an F_Γ-group via the map $\pi^1 : F_\Gamma \to F_Q$.

We obtain a commutative diagram with exact columns (but note that the top row, like the middle row, is not exact):

$$\begin{tikzcd}
\pi(\mathcal{P}_N) \arrow[d, hook] & \pi(\mathcal{P}_\Gamma) \arrow[d, hook] & \pi(\mathcal{P}_Q) \arrow[d, hook] \\
C(\mathcal{P}_N) \arrow[r, hook] \arrow[d, "\partial_N"] & C(\mathcal{P}_\Gamma) \arrow[r, two heads, "\bar{\pi}^2"] \arrow[d, "\partial_\Gamma"] & C(\mathcal{P}_Q) \arrow[d, "\partial_Q"] \\
F_N \arrow[r, hook] \arrow[d, two heads, "p_N"] & F_\Gamma \arrow[r, two heads, "\pi^1"'] \arrow[d, two heads, "p_\Gamma"] & F_Q \arrow[l, bend right, "\sigma"'] \arrow[d, two heads, "p_Q"] \\
N \arrow[r, hook] & \Gamma \arrow[r, two heads, "\pi"] & Q
\end{tikzcd}$$

[8] In fact this map is again an inclusion, but we will not need this fact here.

Notation. In the following we will work with elements of the crossed modules associated to the presentations. We will denote the image of $c \in \Lambda(\mathcal{P})$ under the quotient map $\Lambda(\mathcal{P}) \to C(\mathcal{P})$ by $[c] \in C(\mathcal{P})$. We will also denote images of sets $A \subset \Lambda(\mathcal{P})$ under this quotient map by $[A]$.

Now assume that we are also given a generating set of the identities of relators of Q. More precisely, let $I_Q \subset \ker \partial_Q \leq \Lambda(\mathcal{P}_Q)$ be such that $[I_Q] \subset \pi(\mathcal{P}_Q)$ generates $\pi(\mathcal{P}_Q)$ as a Q-module.

We will now describe two sets of identities between the relators of $\mathcal{P}_\Gamma$.

The first set of identities $I^2_\Gamma \subset \Lambda(\mathcal{P}_\Gamma)$ contains one element $i_{x,s}$ for each $x \in \mathcal{X}$ and $s \in R_Q$. Recall that we have for each $s \in R_Q$ a relator $r_s \in R^2_\Gamma$ with

$$\partial_\Gamma(r_s) = d_s c_s^{-1}$$

where

$$d_s = \sigma(\partial_Q(s)) \in F(\widetilde{\mathcal{Z}}) \quad \text{and} \quad c_s \in F(\mathcal{X}) \text{ with } p_N(c_s) = p_\Gamma(d_s).$$

Define

$$d_{x,s} := (x \cdot r_s) r_s^{-1}.$$

Then

$$\pi^1 \partial_\Gamma(d_{x,s}) = \pi^1(x d_s c_s^{-1} x^{-1} (d_s c_s^{-1})^{-1}) = \partial_Q(s) \partial_Q(s)^{-1} = 1.$$

By Lemmas 4.3 and 4.4 there exists

$$c_{x,s} \in \langle\langle R^0_\Gamma \cup R^1_\Gamma \rangle\rangle_{F_\Gamma} \leq \Lambda(\mathcal{P}_\Gamma) \quad \text{with} \quad \partial_\Gamma(c_{x,s}) = \partial_\Gamma(d_{x,s}).$$

So, if we set

$$i_{x,s} := d_{x,s} c_{x,s}^{-1},$$

then $\partial_\Gamma(i_{x,s}) = 1$, so $[i_{x,s}] \in \pi(\mathcal{P}_\Gamma)$.

The second set of identities I^3_Γ contains one element i_j for each $j \in I_Q$. Each $j \in I_Q$ is of the form

$$(w_1 \cdot s_1) \cdots (w_n \cdot s_n), \quad \text{with } w_1, \ldots, w_n \in F_Q \text{ and } s_1, \ldots, s_n \in R_Q.$$

Set

$$d_j := (\sigma(w_1) \cdot r_{s_1}) \cdots (\sigma(w_n) \cdot r_{s_n}),$$

so that $\pi^2(d_j) = j$. We then have

$$\pi^1(\partial_\Gamma(d_j)) = \partial_Q(\pi^2(d_j)) = \partial_Q(j) = 1$$

By Lemmas 4.3 and 4.4 there exists

$$c_j \in \langle\langle R^0_\Gamma \cup R^1_\Gamma \rangle\rangle_{F_\Gamma} \leq \Lambda(\mathcal{P}_\Gamma) \quad \text{with} \quad \partial_\Gamma(c_j) = \partial_\Gamma(d_j).$$

So set

$$i_j := d_j c_j^{-1},$$

and we have $\partial_\Gamma(i_j) = 1$, i.e. $[i_j] \in \pi(\mathcal{P}_\Gamma)$. In particular, since π^2 is trivial on R^0_Γ and R^1_Γ, we have $\pi^2(i_j) = \pi^2(d_j) = j$.

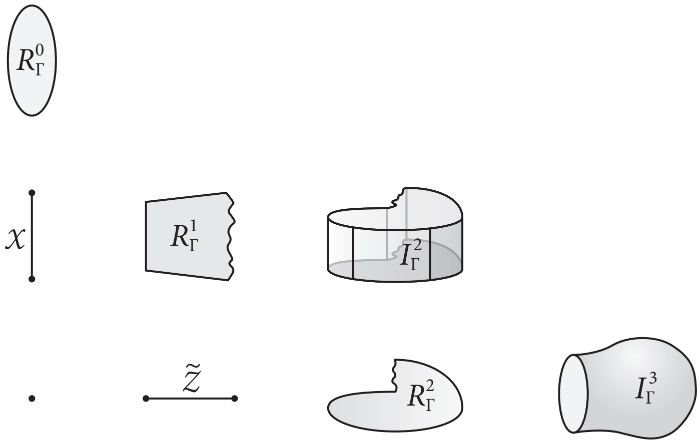

The next lemma tells us how we can use the identities I^3_Γ. Recall that each $[i] \in \pi(\mathcal{P}_\Gamma)$ can be thought of as a 2-sphere in the presentation complex X_Γ of $\mathcal{P}_\Gamma$. In Section 4.3 we saw that the presentation complex X_N for N can be regarded as a subcomplex of X_Γ and there is a map $\phi : X_\Gamma \to X_Q$ which sends X_N to the single vertex of X_Q. So if $d : D^2 \to X_\Gamma$ is a disc in X_Γ with boundary in X_N then $\phi \circ d$ is a 2-sphere in X_Q. If we glue 3-cells into X_Q along the set $[I_Q]$ of generators of $\pi_2(X_Q)$ then every 2-sphere in the resulting complex is null-homotopic.

The lemma can then be interpreted geometrically as saying: If we glue in a 3-cell along each $[i] \in I^3_\Gamma$ then every disc c in X_Γ with boundary in X_N can be homotoped to yield a c' with the same boundary such that the image of c' in X_Q is null-homotopic in X_Q (that is, null-homotopic in the 2-skeleton).

Lemma 4.15. *For any $c \in \Lambda(\mathcal{P}_\Gamma)$ with $\partial_\Gamma(c) \in F(\mathcal{X})$ there is a $c' \in \langle\!\langle I^3_\Gamma \rangle\!\rangle_{F_\Gamma}$ such that $\bar{\pi}^2([cc'^{-1}]) = 1$.*

Proof. We have $\bar{\partial}_Q(\bar{\pi}^2([c])) = \pi^1(\partial_\Gamma(c)) = 1$, so $\bar{\pi}^2([c]) \in \pi(\mathcal{P}_Q)$. Since $[I_Q]$ generates $\pi(\mathcal{P}_Q)$ as an F_Q-module, we can find $w_1, \ldots, w_n \in F_Q$ and $j_1, \ldots, j_n \in I_Q$ such that $\bar{\pi}^2([c]) = (w_1 \cdot j_1) \cdots (w_n \cdot j_n)$. Let

$$c' := (\sigma(w_1) \cdot i_{j_1}) \cdots (\sigma(w_n) \cdot i_{j_n}) \in \langle\!\langle I^3 \rangle\!\rangle_{F_\Gamma},$$

and clearly $\bar{\pi}^2([c]) = \bar{\pi}^2([c'])$. □

Next, we will see what the identities I^2_Γ can be used for.

Lemma 4.16. *Let $s \in R_Q$ and $w \in F_\Gamma$. Then $[w \cdot r_s]$ is equivalent[9] modulo $[I^2_\Gamma]$, to $[c(w_z \cdot r_s)]$ for some $c \in \langle\!\langle R^0_\Gamma \cup R^1_\Gamma \rangle\!\rangle_{F_\Gamma}$ and $w_z = \sigma\pi^1(w) \in F(\tilde{\mathcal{Z}})$.*

Proof. First note that we can write w in the form $w = uw_zw_x$ with $u \in \langle\!\langle \partial_\Gamma(R^1_\Gamma) \rangle\!\rangle$ and $w_x \in F(\mathcal{X})$, $w_z = \sigma\pi^1(w) \in F(\tilde{\mathcal{Z}})$ by Lemma 4.2. Now, note that for any $x \in \mathcal{X}$ we have

$$i^{-1}_{x,s}(x \cdot r_s) = c_{x,s} r_s$$

[9] precisely, we say that some $[c], [c'] \in C(\mathcal{P}_\Gamma)$ are equivalent modulo $[I^2_\Gamma]$ if there is an $i \in \langle\!\langle I^2_\Gamma \rangle\!\rangle_{F_\Gamma}$ with $[c] = [i][c']$

and $c_{x,s} \in \langle\langle R_\Gamma^0 \cup R_\Gamma^1 \rangle\rangle_{F_\Gamma}$. By induction we can then show that $w_x \cdot r_s$ is equivalent modulo I_Γ^2 to some $c_0 r_s$ with $c_0 \in \langle\langle R_\Gamma^0 \cup R_\Gamma^1 \rangle\rangle_{F_\Gamma}$. Thus $w \cdot r_s = uw_Z w_X \cdot r_s$ is equivalent modulo I_Γ^2 to $uw_Z \cdot (c_0 r_s)$. Since $u \in \langle\langle \partial_\Gamma(R_\Gamma^1) \rangle\rangle$, there is a $c_1 \in \langle\langle R_\Gamma^1 \rangle\rangle_{F_\Gamma}$ with $\partial_\Gamma(c_1) = u$. We have

$$\begin{aligned}[uw_z \cdot (c_0 r)] &= [(uw_z \cdot c_0)(uw_z \cdot r)] \\ &= [uw_z \cdot c_0][\partial(c_1) \cdot (w_z \cdot r)] \\ &= [uw_z \cdot c_0][c_1(w_z \cdot r)c_1^{-1}] \\ &= [uw_z \cdot c_0][c_1(\partial(w_z \cdot r) \cdot c_1^{-1}][w_z \cdot r]\end{aligned}$$

and

$$c := (uw_z \cdot c_0)c_1(\partial(w_z \cdot r) \cdot c_1^{-1}) \in \langle\langle R_\Gamma^0 \cup R_\Gamma^1 \rangle\rangle_{F_\Gamma}. \qquad \square$$

Lemma 4.15 yields $c \in \Lambda(\mathcal{P}_\Gamma)$ with $\bar\pi^2([c]) = 1$ (geometrically, discs in X_Γ whose images in X_Q are spheres that can be contracted in the 2-skeleton of X_Q). We will now analyse how we can use the identities I_Γ^2 to transform these elements further. First, we treat a special case.

Lemma 4.17. *Let $c \in \langle\langle R_\Gamma^2 \rangle\rangle_{F(\tilde Z)}$ with $\bar\pi^2([c]) = 1$. Then $[c]$ is equivalent modulo $[I_\Gamma^2]$ to some $[c']$ with $c' \in \langle\langle R_\Gamma^0 \cup R_\Gamma^1 \rangle\rangle_{F_\Gamma}$.*

Proof. We can define a group homomorphism $\sigma^2 : \Lambda(\mathcal{P}_Q) \to \Lambda(\mathcal{P}_\Gamma)$ by defining $\sigma^2(w \cdot s) = \sigma(w) \cdot r_s$ for $w \in F_Q$ and $s \in R_Q$ (recall that, as a group, $\Lambda(\mathcal{P}_Q)$ is free on these $w \cdot s$). By definition, σ^2 is an inverse to the restriction of π^2 to $\langle\langle R_\Gamma^2 \rangle\rangle_{F(\tilde Z)}$.

Now, we have $1 = \bar\pi^2([c]) = [\pi^2(c)]$, so $\pi^2(c)$ is in the kernel of the quotient map $\Lambda(\mathcal{P}_Q) \to C(\mathcal{P}_Q)$. This kernel is the Peiffer subgroup $P \le \Lambda(\mathcal{P}_Q)$. Since $c = \sigma^2(\pi^2(c))$, we have $c \in \sigma(P)$.

By Lemma 4.8, P is generated by the basic Peiffer elements

$$aba^{-1}(\partial_Q(a) \cdot b)^{-1}, \quad \text{with } a, b \in F_Q \times R_Q.$$

Without loss of generality we may assume that c is the image of one of these generators, i.e.

$$c = \sigma^2(aba^{-1}(\partial_Q(a) \cdot b)^{-1})$$

for some $a = u \cdot s$ and $b = v \cdot t$ with $u, v \in F_Q$ and $s, t \in R_Q$. We have

$$\sigma^2(\partial_Q(a) \cdot b) = \sigma^2(u\partial_Q(s)u^{-1}v \cdot t) = \sigma(u\partial_Q(s)u^{-1}v) \cdot r_t.$$

Set

$$w := \sigma(u)\partial_\Gamma(r_s)\sigma(u^{-1}v).$$

Recall that $\partial_\Gamma(r_s) = \sigma(\partial_Q(s))c_s^{-1}$ where $c_s \in \langle\langle \partial_\Gamma(R_\Gamma^0 \cup R_\Gamma^1) \rangle\rangle$. Then $\pi^1(c_s) = 1$, so

$$w_z := \sigma\pi^1(w) = \sigma(u\partial_Q(s)u^{-1}v).$$

Then Lemma 4.16 tells us that there is a $c' \in \langle\langle R^0_\Gamma \cup R^1_\Gamma \rangle\rangle_{F_\Gamma}$ such that $[w \cdot r_t]$ is equivalent modulo $[I^2_\Gamma]$ to $[c'(w_z \cdot r_t)]$. So there is some $i \in \langle\langle I^2_\Gamma \rangle\rangle_{F_\Gamma}$, such that

$$\begin{aligned}[\sigma^2(\partial_Q(a) \cdot b)] &= [w_z \cdot r_t] \\ &= [ic'^{-1}][w \cdot r_t] \\ &= [ic'^{-1}][\sigma(u)\partial_\Gamma(r_s)\sigma(u^{-1}v) \cdot r_t] \\ &= [ic'^{-1}][\partial_\Gamma(\sigma(u) \cdot r_s) \cdot (v \cdot r_t)] \\ &= [ic'^{-1}][\partial_\Gamma(\sigma^2(a)) \cdot \sigma^2(b)] \\ &= [ic'^{-1}][\sigma^2(a)\sigma^2(b)\sigma^2(a)^{-1}]\end{aligned}$$

so

$$[c] = [i][c'].$$

□

More generally we can now show the following, which should be compared to Lemma 4.3:

Lemma 4.18. *Let $[c] \in C(\mathcal{P}_\Gamma)$ with $\partial_\Gamma(c) \in F(\mathcal{X})$ and $\bar{\pi}^2([c]) = 1$. Then $[c]$ is equivalent, modulo $[I^2_\Gamma]$ to some $[c'] \in C(\mathcal{P}_\Gamma)$ with $c' \in \langle\langle R^0_\Gamma \cup R^1_\Gamma \rangle\rangle_{F_\Gamma}$.*

Proof. Let $[c] \in C(\mathcal{P}_\Gamma)$ with $\partial_\Gamma([c]) \in F(\mathcal{X})$ and $\pi^2([c]) = 1$. Then c is of the form

$$c = (w_1 \cdot r_1)\cdots(w_n \cdot r_n), \qquad \text{with } w_1, \ldots, w_n \in F(X \cup \widetilde{Z})$$
$$r_1, \ldots, r_n \in R^0_\Gamma \cup R^1_\Gamma \cup R^2_\Gamma.$$

First, we can use the previous lemma to replace all the terms $w_i \cdot r_i$ with $r_i \in R^2$ by some $c_i w^z_i \cdot r_i$ with $c_i \in \langle\langle R^1 \rangle\rangle_\Gamma$ and $w^z_i \in F(\widetilde{\mathcal{Z}})$. We can thus find some $[c_1] \in C(\mathcal{P}_\Gamma)$ which is equivalent to $[c]$ modulo $[I^2_\Gamma]$ and which is of the form

$$c_1 = (w'_1 \cdot r'_1)\cdots(w'_m \cdot r'_m), \qquad \text{with } w'_1, \ldots, w'_n \in F(\mathcal{X} \cup \widetilde{\mathcal{Z}})$$
$$r_1, \ldots, r_n \in R^0_\Gamma \cup R^1_\Gamma \cup R^2_\Gamma$$
$$\text{such that } w'_i \in F(\widetilde{\mathcal{Z}}) \text{ whenever } r_i \in R^2.$$

Now we can use the crossed module relations

$$[w'_i \cdot r'_i][w'_{i+1} \cdot r'_{i+1}] = [(\partial(w'_i \cdot r'_i) \cdot (w'_{i+1} \cdot r'_{i+1})][w'_i \cdot r'_i]$$

to move all the terms with $r'_i \in R^2_\Gamma$ to the right of the word. Thus we have

$$[c_1] = [c_2][(u_1 \cdot r_{s_1})\cdots(u_k \cdot r_{s_k})]$$

for some $c_2 \in \langle\langle R^0 \cup R^1 \rangle\rangle_{F_\Gamma}$ and $u_1, \ldots, u_k \in F(\widetilde{\mathcal{Z}})$ and $s_1, \ldots, s_k \in R_Q$. Now $[c]$ and $[c_1]$ are equivalent modulo $[I^2_\Gamma]$ and we have $\bar{\pi}^2([I^2_\Gamma]) = 1$, so $\bar{\pi}^2([c_1]) = \bar{\pi}^2([c]) = 1$. Furthermore we have $\bar{\pi}^2([c_2]) = 1$, and thus

$$\bar{\pi}^2([(u_1 \cdot r_{s_1})\cdots(u_k \cdot r_{s_k})]) = 1.$$

By Lemma 4.17 it follows that $[(u_1 \cdot r_{s_1})\cdots(u_k \cdot r_{s_k})]$ is equivalent modulo $[I^2_\Gamma]$ to some $[c_3] \in C(\mathcal{P}_\Gamma)$ with $c_3 \in \langle\langle R^0_\Gamma \cup R^1_\Gamma \rangle\rangle_{F_\Gamma}$. Therefore, modulo $[I^2_\Gamma]$, our $[c]$ is equivalent to $[c_1]$ which is equivalent to $[c_2 c_3]$ and $c_2 c_3 \in \langle\langle R^0_\Gamma \cup R^1_\Gamma \rangle\rangle_{F_\Gamma}$. □

4.6 Proof of the 1-2-3 Theorem

Now we are ready to give a proof of the 1-2-3 Theorem. As we showed in Section 4.2, it is sufficient to consider the following situation:

$$\begin{tikzcd}
N \arrow[r, hook] \arrow[d, equal] & P \arrow[r, two heads] \arrow[d, two heads] & F \arrow[d, two heads, "\pi_2"] \\
N \arrow[r, hook] & \Gamma \arrow[r, two heads, "\pi_1"] & Q
\end{tikzcd}$$

where N is finitely generated, Γ is finitely presented, Q is of type F_3, F is finitely generated free, and P is the fibre product of π_1 and π_2. In order to establish the 1-2-3 Theorem, we need to show that P is finitely presented.

As in Section 4.3 pick a presentation $\mathcal{P}_N = \langle \mathcal{X} \mid R_N, \partial_N \rangle$ for N with $\mathcal{X}$ finite and a presentation $\mathcal{P}_Q = \langle \mathcal{Z} \mid R_Q, \partial_Q \rangle$ for Q, with both $\mathcal{Z}$ and R_Q finite and such that $\mathcal{Z}$ is a basis for F. Furthermore, pick a finite set $I_Q \subset \Lambda(\mathcal{P}_Q)$ such that $[I_Q]$ generates $\pi(\mathcal{P}_Q)$ as a Q-module. Let $\mathcal{P}_\Gamma = \langle \mathcal{X} \cup \tilde{\mathcal{Z}} \mid R^0_\Gamma \cup R^1_\Gamma \cup R^2_\Gamma, \partial_\Gamma \rangle$ be the presentation constructed in Section 4.3. Then let

$$K_\Gamma := \ker(p_\Gamma : F(\mathcal{X} \cup \tilde{\mathcal{Z}}) \twoheadrightarrow \Gamma) = \langle\langle \partial_\Gamma(R^0_\Gamma \cup R^1_\Gamma \cup R^2_\Gamma) \rangle\rangle$$

and

$$K_P := \ker(p_P : F(\mathcal{X} \cup \tilde{\mathcal{Z}}) \twoheadrightarrow P) = \langle\langle \partial_\Gamma(R^0_\Gamma \cup R^1_\Gamma) \rangle\rangle.$$

Since Γ is finitely presented there is a finite subset $R^{0\prime}_\Gamma \subset R^0_\Gamma$ such that

$$K_\Gamma = \langle\langle \partial_\Gamma(R^{0\prime}_\Gamma \cup R^1_\Gamma \cup R^2_\Gamma) \rangle\rangle.$$

Now consider the identities I^2_Γ and I^3_Γ from the previous section. Recall that there is one identity $i_{x,s}$ in I^2_Γ for each generator $x \in \mathcal{X}$ of N and relator $s \in R_Q$ of Q and one identity i_j in I^3_Γ for each identity $j \in I_Q$ of Q. Therefore I^2_Γ and I^3_Γ are *finite* subsets of $\Lambda(\mathcal{P}_\Gamma)$. Thus

$$\langle\langle [I^2_\Gamma] \cup [I^3_\Gamma] \rangle\rangle_{F(\mathcal{X} \cup \tilde{\mathcal{Z}})} \subset \langle\langle [R^{0\prime\prime}_\Gamma] \cup [R^1_\Gamma] \cup [R^2_\Gamma] \rangle\rangle$$

for some finite $R^{0\prime\prime}_\Gamma \subset R^0_\Gamma$. Without loss of generality, we assume that $R^{0\prime}_\Gamma \subset R^{0\prime\prime}_\Gamma$.

Now set

$$K'_P := \langle\langle \partial_\Gamma(R^{0\prime\prime}_\Gamma \cup R^1_\Gamma) \rangle\rangle.$$

We will show that $K_P = K'_P$, which finishes the proof that P is finitely presented.

So let $w \in K_P$. Firstly, by Lemmas 4.1 and 4.3, w is equivalent modulo R^1_Γ to some $w' \in F(\mathcal{X})$. So, without loss of generality, assume $w \in F(\mathcal{X})$.

Now $w \in K_P$, so $w \in K_\Gamma = \langle\langle \partial_\Gamma(R^{0\prime}_\Gamma \cup R^1_\Gamma \cup R^2_\Gamma) \rangle\rangle$. Thus we can pick some $c \in \langle\langle R^{0\prime}_\Gamma \cup R^1_\Gamma \cup R^2_\Gamma \rangle\rangle_{F(\mathcal{X} \cup \tilde{\mathcal{Z}})} \subset C(\mathcal{P}_\Gamma)$ with $\partial_\Gamma(c) = w$. We can now apply Lemmas 4.15 and 4.18 to conclude that there is some $i \in \langle\langle I^2_\Gamma \cup I^3_\Gamma \rangle\rangle_{F(\mathcal{X} \cup \tilde{\mathcal{Z}})}$ with $[c'] := [i][c] \in$

$\langle\langle [R_\Gamma^0] \cup [R_\Gamma^1] \rangle\rangle_{F(\mathcal{X} \cup \tilde{Z})}$. Since i is in the kernel of ∂_Γ, we have $\bar{\partial}_\Gamma([c']) = \bar{\partial}_\Gamma([c]) = w$. Furthermore, both c and i lie in $\langle\langle R_\Gamma^{0''} \cup R_\Gamma^1 \cup R_\Gamma^2 \rangle\rangle$ and thus so does c'. Therefore $c' \in \langle\langle R_\Gamma^{0''} \cup R_\Gamma^1 \rangle\rangle$, and hence $w \in K_P'$. □

From 3.14 (and the remark following the proof) this implies the Virtual Surjections to Pairs Theorem, which was originally obtained in [BHMS12, Theorem E]:

Corollary 4.19 (The Virtual Surjections to Pairs Theorem). *Let $\Gamma_1, \ldots, \Gamma_n$ be finitely presented groups and $P \leq \Gamma_1 \times \cdots \times \Gamma_n$ a subgroup of their direct product. If P virtually surjects to pairs of factors, then P is finitely presented.* □

4.7 Towards higher dimensions

We have seen in this chapter how the combinatorics of crossed modules enable us to prove the 1-2-3 Theorem. Why is this approach not available in higher dimensions? In fact, the central theorem by Whitehead (Theorem 4.13), does have a satisfying generalisation to higher dimensions: If $X^{(n)}$ is an n-dimensional CW-complex for $n > 2$ then the relative homotopy group $\pi_n(X^{(n)}, X^{(n-1)})$ is a free $\pi_1(X)$-module, with a basis given by a set of characteristic maps for the n-cells.

All the relative homotopy groups $\pi_n(X^{(n)}, X^{(n-1)})$ can be assembled into an exact sequence by splicing together parts of the long exact homotopy sequences of each pair $(X^{(n)}, X^{(n-1)})$:

$$\begin{tikzcd}
& \pi_n(X^{(n)}) \arrow[dr] & & & \\
\pi_{n+1}(X^{(n+1)}, X^{(n)}) \arrow[ur] \arrow[rr, "\partial_{n+1}"] & & \pi_n(X^{(n)}, X^{(n-1)}) \arrow[rr, "\partial_n"] \arrow[dr] & & \pi_{n-1}(X^{(n-1)}, X^{(n-2)}) \\
& & & \pi_{n-1}(X^{(n-1)}) \arrow[ur] &
\end{tikzcd}$$

This is a sequence of free $\pi_1(X)$-modules and module homomorphisms except in the lowest dimensions

$$\pi_2(X^{(2)}, X^{(1)}) \xrightarrow{\partial_2} \pi_1(X^{(1)}, X^{(0)})$$

where $\pi_1(X^{(1)}, X^{(0)})$ is a free group and $(\pi_2(X^{(2)}, X^{(1)}), \partial_2)$ is a free crossed module on that group.

Such sequences are called *crossed complexes* and were first studied in [Hue80]. In the special case where X is a $B\Gamma$-complex (say with a single vertex) the crossed complex can be extended by an "augmentation map", as we did at the end of Section 4.4:

$$\begin{tikzcd} \pi_2(X^{(2)}, X^{(1)}) \arrow[r, "\partial_2"] \arrow[d, "\cong"] & \pi_1(X^{(1)}, x_0) \arrow[r, two heads] \arrow[d, "\cong"] & \pi_1(X^{(2)}, x_0) \arrow[d, "\cong"] \\ C(\mathcal{P}) \arrow[r] & F \arrow[r, two heads] & \Gamma \end{tikzcd}$$

This is termed a *crossed resolution* of Γ. For more information on the general theory, see [BH82b, BRS99, BHS11]. The kernels of the boundary maps involved have been studied as "homotopical syzygies" in [Lod00] and a choice of generating set for each of the groups has been called a "generalised group presentation" and studied in [Bro92].

So, to recapture, in a way, the combinatorics in fact become *easier* in higher dimensions: Instead of having to deal with crossed modules, we are facing the much more familiar modules over a group ring.

Of course, there is a catch. In dimension 2 the long exact homotopy sequence associated to the pair $(X^{(2)}, X^{(1)})$ is

$$\cdots \to \pi_2(X^{(1)}) \to \pi_2(X^{(2)}) \to \pi_2(X^{(2)}, X^{(1)}) \to \pi_1(X^{(1)}) \to \cdots$$

Crucially, as we observed earlier, the first term vanishes. It is this fact that allowed us to identify, for a given presentation $\mathcal{P} = \langle \mathcal{X} \mid R, \partial \rangle$, the second homotopy group of the presentation complex $X_{\mathcal{P}}$ with the kernel of the boundary map $\partial : C(\mathcal{P}) \to F(\mathcal{X})$.

Life is not so simple in higher dimensions: The relevant part of the long exact homotopy sequence of the pair $(X^{(n)}, X^{(n-1)})$ is

$$\cdots \to \pi_n(X^{(n-1)}) \to \pi_n(X^{(n)}) \to \pi_n(X^{(n)}, X^{(n-1)}) \overset{\partial_n}{\to} \pi_{n-1}(X^{(n-1)}) \to \cdots.$$

This is now a sequence of $\pi_1(X)$-modules and module homomorphisms, but the module $\pi_n(X^{(n)})$ is no longer the kernel of ∂_n if $n > 2$. In general, we cannot expect the term $\pi_n(X^{(n-1)})$ to be trivial, nor the map $\pi_n(X^{(n)}) \to \pi_n(X^{(n)}, X^{(n-1)})$ to be injective. So the characterisation of F_3 that we achieved in Corollary 4.14 has no analogue in higher dimensions.

Of course, this does not mean that nothing can be gained from replicating the approach presented here in higher dimensions, investigating the relative homotopy groups $\pi_n(X^{(n)}, X^{(n-1)})$ in place of $\pi_2(X^{(2)}, X^{(1)})$. In a way, this is what we do in Chapter 7, though couched in a different language: If X is a $B\Gamma$-complex then in dimensions $n \geq 3$ the groups $\pi_n(X^{(n)}, X^{(n-1)})$ are isomorphic to the cellular chain groups $C_n(\widetilde{X})$ of the universal cover $\widetilde{X}$ of X. In fact, the sequence of relative homotopy groups

$$\cdots \to \pi_n(X^{(n)}, X^{(n-1)}) \to \pi_{n-1}(X^{(n-1)}, X^{(n-2)}) \to \cdots \to \pi_3(X^{(3)}, X^{(2)})$$

is chain isomorphic to the cellular chain complex of the universal covering. So in higher dimensions these relative homotopy groups do not contain more information than a resolution of $\mathbb{Z}$ as a trivial Γ-module. We study these resolutions in detail in Chapter 7 and we will see that many of the results from the current chapter carry over to higher dimensions. The issues raised here help to explain why these generalisations do not easily add up to a proof of the n-$(n+1)$-$(n+2)$ Conjecture, but we will leave it to the end of Chapter 7 to explore more thoroughly what they actually do allow us to say.

5

A topological approach

❦

The statement of the n-$(n+1)$-$(n+2)$ Conjecture is at its core a topological question, as the finiteness properties F_n are concerned with the existence of classifiying complexes of a certain type and implicitly with questions about finite generation of higher homotopy groups (see A.20).

Putting aside for now the troubles about the algebra of higher homotopy groups, which we alluded to at then end of the previous chapter, we will present in this chapter an attack on the n-$(n+1)$-$(n+2)$ Conjecture by mostly topological means. Having seen how the proof of the 1-2-3 Theorem involved a fair amount of calculations with crossed modules, maybe we should not expect to get very far without delving into the combinatorics of classifying complexes. But in fact, we can obtain a number of interesting results: Essentially, many of the results from Sections 4.1–4.3 can be generalised to higher dimensions without appeal to explicit combinatorics. In particular, this method will enable us to replicate the main results from Section 4.2 in higher dimensions: The n-$(n+1)$-$(n+2)$ Conjecture holds when the second short exact sequence splits and the general conjecture can be reduced to the case when Γ_2 is free.

5.1 Introducing stacks

A crucial step in proving the 1-2-3 Theorem was the construction of a presentation of the middle group Γ in a short exact sequence $N \hookrightarrow \Gamma \twoheadrightarrow Q$ from presentations for the kernel N and quotient Q. This amounts to constructing the 2-skeleton of a $B\Gamma$-complex. Now, our task will be to construct an entire useful $B\Gamma$-complex from given BN- and BQ-complexes. "Useful" is a vague term here, of course, but recall from 4.3 some of the properties satisfied by the presentation complex X_Γ constructed there: The presentation complex X_N appeared as a subcomplex, we could define a map $X_\Gamma \to X_Q$ inducing $\pi : \Gamma \to Q$ on the fundamental groups and X_Γ could be pictured as a sort of "twisted product" of X_N and X_Q. We shall hope to replicate these properties in some form here.

Various ways of making these notions precise and constructing the kinds of complexes we look for have been given by different authors (e.g. the "complexes of spaces" in [CS98], the construction in [Sta98] and the more algebraic analogue of "complexes of groups" [Cor92, Hae91]), but the most useful for our purposes is Ross Geoghegan's theory of "stacks" from [Geo08, Chapter 6]. We will start by briefly laying out the parts of the theory pertinent to our discussion.

As a first approximation, given a short exact sequence $N \hookrightarrow \Gamma \overset{\pi}{\twoheadrightarrow} Q$, we can "represent" it topologically via the Borel construction. For this we pick a $B\Gamma$-complex Y_0 and a BQ-complex Z. Then Γ acts naturally on the universal cover $\widetilde{Y}_0$ via deck transformations and also on the universal cover $\widetilde{Z}$ through the homomorphism $\Gamma \overset{\pi}{\to} Q \cong \mathrm{Deck}(\widetilde{Z} \twoheadrightarrow Z)$. So we can let Γ act diagonally on $\widetilde{Y}_0 \times \widetilde{Z}$ and noting that this action is free and properly discontinuous we obtain a new $B\Gamma$-complex $Y := (\widetilde{Y}_0 \times \widetilde{Z})/\Gamma$. The covering map $\widetilde{Z} \to Z$ then gives a map $\phi : Y \to Z$, which induces $\pi : \Gamma \to Q$ on the level of fundamental groups. Furthermore ϕ is a fibre bundle with fibre the BN-complex $X_0 := \widetilde{Y}_0/N$. So, topologically, ϕ has very desirable properties but unfortunately Y is not a very useful $B\Gamma$-complex for our purposes, yet. We are interested in "small" classifying complexes (in the sense as demanded by the finiteness properties F_n), and Y is not small: It will typically have infinitely many cells of any dimension (for example, the fibre X_0 will have infinitely many 0-cells if Q is infinite).

Geoghegan's construction builds from Y a new $B\Gamma$-complex Y' and map $\phi' : Y' \to Z$ by swapping out the fibres X for a different BN-complex X' that can be chosen arbitrarily. As a result, the complex Y' will be reasonably small if the initial complexes Z and X' are. The price to pay is that the map ϕ' is no longer a fibre bundle. However, it is what Geoghegan calls a *stack*, which is a cellular projection map that retains some of the structure of a fibre bundle.

We will now make this notion precise in the following definition from [Geo08, Section 6.1].

Definition 5.1. Let X, Y and Z be CW-complexes. Denote by E_n the set of closed n-cells of Z. A map $\phi : Y \to Z$ is called a stack with fibre X, if there is for each $e \in E_n$ a map $f_e : X \times S^{n-1} \to \phi^{-1}(Z^{(n-1)})$ and denoting the disjoint union of these maps by

$$f_n : \bigsqcup_{e \in E_n} X \times S^{n-1} \to \phi^{-1}(Z^{(n-1)})$$

there are for each $n \geq 0$ homeomorphisms[1]

$$k_n : \phi^{-1}(Z^{(n-1)}) \cup_{f_n} \left(\bigsqcup_{e \in E_n} X \times B^n \right) \to \phi^{-1}(Z^{(n)}),$$

such that

1. k_n agrees with inclusion on $\phi^{-1}(Z^{(n-1)})$.

2. k_n maps each closed cell of Y onto a closed cell of Z.

3. The following diagram commutes:

$$\begin{tikzcd}
\phi^{-1}(Z^{(n-1)}) \uplus (\biguplus_{e\in E_n} X \times B^n) \arrow[d, "\text{quotient}"] \arrow[dr, "u"] & \\
\phi^{-1}(Z^{(n-1)}) \cup_{f_n} (\biguplus_{e\in E_n} X \times B^n) \arrow[d, "k_n"] & Z^{(n)} \\
\phi^{-1}(Z^{(n)}) \arrow[ur, "\phi|"] &
\end{tikzcd}$$

where u agrees with ϕ on $\phi^{-1}(Z^{(n-1)})$ and on each copy $X \times B^n$ corresponding to some cell $e \in E_n$ it is projection to B^n followed by a characteristic map $B^n \to e$.

Roughly speaking, if $\phi : Y \to Z$ is a stack, then Y can be built up step-by-step as a filtration $Y_n := \phi^{-1}(Z^{(n)})$, where going from Y_{n-1} to Y_n involves gluing in a complex $X \times B^n$ over each n-cell $e \in E_n$. Over the interior $\mathring{e}$ of the cell the map $\phi : \phi^{-1}(\mathring{e}) \to \mathring{e}$ has the form of a projection $X \times \mathring{e} \to \mathring{e}$.

Note also what the definition means in dimension 0: By convention, $S^{-1} = \varnothing = Z^{(-1)}$, so f_0 is the map $\varnothing \to \varnothing$ and k_0 is a homeomorphism $\biguplus_{e\in E_0} X \to \phi^{-1}(Z^{(0)})$, that is $\phi^{-1}(Z^{(0)})$ consists of a copy of X for each vertex of Z and ϕ merely maps each copy to the corresponding vertex.

The utility of Definition 5.1 lies in two facts: Firstly, the Borel construction described above in fact yields a stack $(\widetilde{Y} \times \widetilde{Z})/\Gamma \to Z$ with fibre $X := \widetilde{Y}/N$. Secondly, if $\phi : Y \to Z$ is a stack with fibre X and X' is homotopy equivalent to X, we can "rebuild" the stack, by gluing in the fibre X' instead of X at each step. This yields a new stack $\phi' : Y' \to Z$, and a homotopy equivalence $h : Y \to Y'$ making the diagram

$$\begin{tikzcd}
Y \arrow[rr, "h"] \arrow[dr, "\phi"'] & & Y' \arrow[dl, "\phi'"] \\
& Z &
\end{tikzcd}$$

commute up to homotopy. This homotopy can be chosen such that the above diagram commutes up to homotopy "over every cell", meaning that there is a homotopy $H : Y \times I \to Z$ from ϕ to $\phi' \circ h$ with $H(\phi^{-1}(e) \times I) \subset e$ for all closed cells $e \subset Z$. Proofs of these facts can be found in [Geo08, Section 6.1].

[1] We write $\cup_{f_n}$ for "gluing along f_n", i.e. the quotient space of the disjoint union identifying each x in the domain of f_n with $f_n(x)$.

In particular, we obtain the following (this is [Geo08, Theorem 7.1.10]):

Theorem 5.2. *Let $N \hookrightarrow \Gamma \overset{\pi}{\twoheadrightarrow} Q$ be a short exact sequence. Furthermore, let X be a $\mathrm{B}N$-complex and Z a $\mathrm{B}Q$-complex and identify their fundamental groups with N and Q, respectively, by fixing isomorphisms $N \cong \pi_1(X, x_0)$ and $Q \cong \pi_1(Z, z_0)$. Then there is a $\mathrm{B}\Gamma$-complex Y and a cellular map $\phi : Y \to Z$ with the following properties:*

1. *ϕ induces the homomorphism $\pi : \Gamma \to Q$ on the level of fundamental groups.*
2. *ϕ is a stack with fibre X.* □

5.2 Stacks and an F_{n+1}-criterion

To apply this construction to our situation, let us now consider a short exact sequence of groups

$$N \hookrightarrow P \overset{\pi}{\twoheadrightarrow} \Gamma$$

where N is of type F_n and Γ is of type F_{n+1} for some $n \geq 2$.[2] We will be interested in conditions for P to be of type F_{n+1} (later, we will apply these considerations to the short exact sequence $N_1 \hookrightarrow P \twoheadrightarrow \Gamma_2$ associated to a fibre product $P \leq \Gamma_1 \times \Gamma_2$).

Pick a $\mathrm{B}N$-complex X with finite n-skeleton and a $\mathrm{B}\Gamma$-complex Y with finite $(n+1)$-skeleton and let $\phi : W \to Y$ be a stack with fibre X, as given by Theorem 5.2. That is, W is a $\mathrm{B}P$-complex and ϕ induces the homomorphism $\pi : P \twoheadrightarrow \Gamma$. We may furthermore assume that Y has a single vertex $*$ (Proposition A.6 and A.18). By the stack property, $\phi^{-1}(*)$ is a subcomplex isomorphic to X. Identifying these complexes, we will frequently regard X as a subcomplex of W.

Note that in any stack $\phi : W \to Y$ with fibre X the m-cells of W are in bijective correspondence with the set $\bigcup_{k+l=m} E_k^Y \times E_l^X$ where E_k^Y is the set of k-cells of Y and E_l^X is the set of l-cells of X. So in particular, if Y and X have finite n-skeleta then the same will be true of W.

The $(n+1)$-skeleton of W will usually be infinite, since there might be infinitely many $(n+1)$-cells in X. However, there are only finitely many $(n+1)$-cells in W that project down to cells in Y of dimension ≥ 1. In other words, only finitely many $(n+1)$-cells of W are not completely contained in $X = \phi^{-1}(*)$. So P is of type F_{n+1} if and only if we can remove from the $(n+1)$-skeleton $W^{(n+1)}$ all but finitely many of the $(n+1)$-cells in X while making sure that the complex stays n-aspherical (see Lemma A.16).

We seek a generalisation of Lemma 4.5, which was the main ingredient in reducing the 1-2-3 Theorem to the case of the second factor being free.

The proposition we will prove is the following:

[2] Most of what follows can, in principle, be applied to the case $n = 1$ as well. However, we have treated this case before and excluding it here sometimes avoids clutter.

Proposition 5.3. *Let $N \hookrightarrow P \twoheadrightarrow \Gamma$ be a short exact sequence of groups with Γ of type F_{n+1} and N of type F_n for some $n \geq 2$. Assume that there is another short exact sequence $N' \hookrightarrow P' \twoheadrightarrow \Gamma'$ and a homomorphism $\phi : P' \to P$ which restricts to an isomorphism $\phi : N' \to N$ such that P' is of type F_{n+1}.*

$$\begin{tikzcd}
\mathrm{F}_n & \mathrm{F}_{n+1} & \mathrm{F}_{n+1} \\[-20pt]
N' \arrow[r, hook] \arrow[d, "\phi|", hook, two heads] & P' \arrow[r, two heads] \arrow[d, "\phi"] & \Gamma' \arrow[d] \\
N \arrow[r, hook] & P \arrow[r, two heads] & \Gamma \\[-20pt]
\mathrm{F}_n & & \mathrm{F}_{n+1}
\end{tikzcd}$$

Then P is of type F_{n+1} as well.

First we will need an easy lemma which essentially states that removing some $(n+1)$-cells from an n-aspherical complex results in a complex whose n^{th} homotopy group is generated, as a π_1-module, by the attaching maps of the $(n+1)$-cells that have been removed.

Lemma 5.4. *Let $n \geq 2$. Let A be a connected, $(n+1)$-dimensional CW-complex with $\pi_n(A) = 0$ and $B, C \subset A$ subcomplexes such that C is connected, $B^{(n)} = A^{(n)}$ and $B \cup C = A$. Then $\pi_n(B)$ is generated as a $\pi_1(B)$-module by the image of the map $\pi_n(C^{(n)}) \to \pi_n(B)$ induced by inclusion.*

Proof. The pair $(A^{(n+1)}, A^{(n)})$ is an $(n+1)$-cellular extension, so by a theorem of J.H.C. Whitehead ([Whi41, Whi49]; see also [Whi78, Chapter V.1, Theorem 1.1]) the kernel of the map $\pi_n(A^{(n)}) \to \pi_n(A^{(n+1)})$ induced by inclusion is generated by the attaching maps of $(n+1)$-cells. To make this statement precise, we will have to introduce basepoints. For each $(n+1)$-cell e of A, pick a characteristic map $h_e : (D^{n+1}, S^n) \to (A^{(n+1)}, A^{(n)})$. Then pick basepoints $v_0 \in A^{(0)}$ and $x_0 \in S^n$ and for each e a path $p_e : [0,1] \to A$ from v_0 to $h_e(x_0)$. Writing ε_n for some fixed generator of $\pi_n(S^n, x_0)$ and d_e for the restriction $h_e|_{S^n}$ (i.e. d_e is an attaching map for the cell e), we have that d_e represents an element $d_{e*}(\varepsilon_n) \in \pi_n(A^{(n)}, h_e(x_0))$ and p_e induces an isomorphism

$$\tau_{p_e} : \pi_n(A^{(n)}, h_e(x_0)) \to \pi_n(A^{(n)}, v_0).$$

The theorem by Whitehead alluded to above then says that the elements $\tau_{p_e} d_{e*}(\varepsilon_n)$ generate $\pi_n(A^{(n)}, v_0)$ as a $\pi_1(A, v_0)$-module.

Now $(A^{(n)}, B)$ is again an $(n+1)$-cellular extension, so $\iota_* : \pi_n(A^{(n)}, v_0) \to \pi_n(B, v_0)$, induced by inclusion, is a surjection. Since B and A have the same 2-skeleton, we have $\pi_1(A) = \pi_1(B)$. Therefore the elements $\iota_*\big(\tau_{p_e} d_{e*}(\varepsilon_n)\big)$ generate $\pi_n(B, v_0)$ as a $\pi_1(B)$-module. Now $\iota_*\big(\tau_{p_e} d_{e*}(\varepsilon_n)\big) = 0$ for those $(n+1)$-cells e which are contained in B, and all the other $(n+1)$-cells are contained in C, since $B \cup C = A$. But if e is in C then $\iota_*\big(\tau_{p_e} d_{e*}(\varepsilon_n)\big)$ lies in the image of the map $\pi_n(C^{(n)}, v_0) \to \pi_n(B, v_0)$ (we can always choose p_e to lie in C), so this image generates $\pi_n(B, v_0)$ as a $\pi_1(B, v_0)$-module. This proves the assertion. $\square$

Theorem 5.5. *Let $N \hookrightarrow P \overset{\pi}{\twoheadrightarrow} \Gamma$ be an exact sequence of groups, such that Γ is of type F_{n+1} and N is of type F_n (where $n \geq 2$). Furthermore let X be a $\mathrm{B}N$-complex with finite n-skeleton and fix an isomorphism to identify $\pi_1(\mathrm{B}N)$ with N. Then P is of type F_{n+1} if and only if there exists a finite $(n+1)$-dimensional, n-aspherical complex $\overline{X}$, such that*

1. *$X^{(n)}$ can be embedded as a subcomplex in $\overline{X}$ and the inclusion $\iota : X^{(n)} \hookrightarrow \overline{X}$ induces an injective map ι_* on fundamental groups.*
2. *There is a homomorphism $f : \pi_1(\overline{X}) \to P$ making the following diagram commute:*

$$\begin{tikzcd}
X^{(n)} \arrow[r, "\iota"] & \overline{X} \\
\pi_1(X^{(n)}) \arrow[r, hook, "\iota_*"] \arrow[d, equal] & \pi_1(\overline{X}) \arrow[d, "f"] \\
N \arrow[r, hook] & P
\end{tikzcd}$$

Proof. Let Y be a $\mathrm{B}\Gamma$-complex with a single vertex and finite $(n+1)$-skeleton and let $\phi : W \to Y$ be a stack with fibre X and inducing π on fundamental groups, as per Theorem 5.2. Then W has finite n-skeleton. As before we regard $X = \phi^{-1}(Y)$ as a subcomplex of W.

To prove one direction of the claim, suppose P is of type F_{n+1}. In this case some finite P-invariant subcomplex $W'^{(n+1)} \subset W^{(n+1)}$ is n-aspherical (see A.16). Clearly $W'^{(n+1)}$ contains $X^{(n)}$ as a π_1-embedded subcomplex and the inclusion $W'^{(n+1)} \to W^{(n+1)}$ induces a homomorphism f as required in the claim.

The idea of the proof for the other direction is to "replace" the potentially infinitely many $(n+1)$-cells of $X \subset W$ with the finitely many $(n+1)$-cells from $\overline{X}$, glued into W using a topological realisation of f.

Let V denote the subcomplex of $W^{(n+1)}$ consisting of $W^{(n)}$ together with those finitely many $(n+1)$-cells of W that are mapped under ϕ to cells of Z of dimension ≥ 1. Applying Lemma 5.4 with $A = W^{(n+1)}$, $B = V$ and $C = X^{(n+1)}$ we obtain that $\pi_n(V)$ is generated as a $\pi_1(V)$-module by the image of the map $\pi_n(X^{(n)}) \to \pi_n(V)$ induced by inclusion. Roughly speaking, we can push any n-sphere in W into $X^{(n)}$ via a homotopy not intersecting any of the $(n+1)$-cells in X. We now proceed to make n-spheres in X contractible.

We have the following diagram of spaces and inclusions and the induced diagram on fundamental groups:

$$\begin{tikzcd}
X^{(n)} \arrow[r, hook, "\iota"] \arrow[d, equal] & \overline{X}^{(n)} \\
X^{(n)} \arrow[r, hook] & W^{(n)}
\end{tikzcd}
\qquad
\begin{tikzcd}
\pi_1(X^{(n)}) \arrow[r, "\iota_*"] \arrow[d, equal] & \pi_1(\overline{X}^{(n)}) \arrow[d, dashed, "f"] \\
N \arrow[r, hook] & P
\end{tikzcd}$$

Since $W^{(n)}$ is $(n-1)$-aspherical there is no obstruction to constructing a cellular map $\psi : \overline{X}^{(n)} \to W^{(n)}$ making the left square commute and inducing f on the level of fundamental groups. Now let $\widetilde{W} := V \cup_\psi \overline{X}$, i.e. $\widetilde{W}$ is constructed by attaching to V the $(n+1)$-cells of $\overline{X}$ along ψ. Denote by $\bar{\psi} : \overline{X} \to \widetilde{W}$ the natural map (i.e. the composition of the inclusion map $\overline{X} \to V \uplus \overline{X}$ with the quotient map $V \uplus \overline{X} \to V \cup_\psi \overline{X}$).

$$\begin{tikzcd}
X^{(n)} \arrow[r, hook] \arrow[d, equal] & \overline{X}^{(n)} \arrow[rr, hook] \arrow[d, "\psi"] & & \overline{X} \arrow[d, "\bar{\psi}"] \\
X^{(n)} \arrow[r, hook] & W^{(n)} \arrow[r, hook] & V \arrow[r, hook] & \widetilde{W}
\end{tikzcd}$$

So $\widetilde{W}$ is a finite, $(n+1)$-dimensional complex with the same n-skeleton as our original $\mathrm{B}P$-complex W. We will now show that $\widetilde{W}$ is in fact n-aspherical, and therefore a finite $\mathrm{B}P^{(n+1)}$, proving that P is of type F_{n+1}.

We consider the following sequence of homomorphisms induced by inclusions:

$$\pi_n(X^{(n)}) \xrightarrow{\iota_1} \pi_n(V) \xrightarrow{\iota_2} \pi_n(\widetilde{W}).$$

First note that, since $V^{(n)} = \widetilde{W}^{(n)}$, the map ι_2 is a surjection. If $\alpha : S^n \to X^{(n)}$ is any n-sphere in X then, since $\overline{X}$ is n-aspherical, α can be contracted in $\overline{X}$, i.e. there is an extension $\hat{\alpha} : D^{n+1} \to \overline{X}$ with $\hat{\alpha}|_{S^n} = \alpha$. Since ψ was assumed to restrict to inclusion on $X^{(n)}$, so does $\bar{\psi}$. So the composition $\bar{\psi} \circ \hat{\alpha} : D^{n+1} \to \widetilde{W}$ shows that α is null-homotopic in $\widetilde{W}$. In other words, the map $\iota_2 \circ \iota_1 : \pi_n(X^{(n)}) \to \pi_n(\widetilde{W}^{(n+1)})$, induced by inclusion, is the zero map.

But we have shown above that the image of ι_1 generates $\pi_n(V)$ as a π_1-module. Therefore a set of generators of $\pi_n(V)$ is sent to 0 under ι_2, proving that ι_2 is the zero map. Combining this with the fact that ι_2 is a surjection, we get that $\pi_n(\widetilde{W}) = 0$, as claimed. □

Proposition 5.3 now follows as a corollary.

Proof of Proposition 5.3. Let X be a $\mathrm{B}N'$-complex with finite n-skeleton. The quotient $\Gamma' \cong P'/N'$ is of type F_n (actually, even F_{n+1}; see either [Geo08, Theorem 7.2.21] or Lemma 6.16 and Theorem 7.2), so we can pick a $\mathrm{B}\Gamma'$-complex Z' with finite n-skeleton and a single vertex as its 0-skeleton.

By Theorem 5.2 we can pick a $\mathrm{B}P'$-complex W which is a stack over Z' with fibre X. In particular W has finite n-skeleton and contains X as a subcomplex. Since P' is of type F_{n+1}, there is a finite n-aspherical subcomplex $W'^{(n+1)} \subset W^{(n+1)}$ with $W'^{(n)} = W^{(n)}$. Since the fundamental group depends only on the 2-skeleton, we have $\pi_1(W'^{(n+1)}) = \pi_1(W^{(n+1)}) \cong P'$ and $W'^{(n+1)}$ still contains $X^{(n)}$ as a subcomplex with π_1-injective embedding $X^{(n)} \hookrightarrow W'^{(n+1)}$. Thus we can take $W'^{(n+1)}$ to be $\overline{X}$ in Theorem 5.5 and conclude that P is of type F_{n+1}. □

5.3 The Split n-$(n+1)$-$(n+2)$ Conjecture and the reduction to free Γ_2

Applying this to the situation of the n-$(n+1)$-$(n+2)$ Conjecture, consider, as always, two short exact sequences

$$N_1 \hookrightarrow \Gamma_1 \overset{\pi_1}{\twoheadrightarrow} Q$$
$$N_2 \hookrightarrow \Gamma_2 \overset{\pi_2}{\twoheadrightarrow} Q$$

such that N_1 is of type F_n and Γ_1 and Γ_2 are of type F_{n+1} (with $n \geq 2$). As in the split 1-2-3 Theorem we do not need an additional assumption on Q (though of course the ones we have made on Γ_1 and N_1 imply that Q is at least of type F_{n+1}).

Then the associated fibre product P fits into a short exact sequence

$$N_1 \hookrightarrow P \overset{p_2}{\twoheadrightarrow} \Gamma_2,$$

which satisfies the finiteness requirements of Proposition 5.3. So to show that P is of type F_{n+1} we need to find a group P' as in the statement of the Proposition. Note that Γ_1 is a group of type F_{n+1} that contains N_1 as a normal subgroup, as required of P'. However, in general there is no homomorphism $\phi : \Gamma_1 \to P$ which maps $N_1 \leq \Gamma_1$ to $N_1 \times 1 \leq P$.[3] Yet there is one particular special case, where such a map does always exist, namely when the second short exact sequence splits. This allows us to prove the split case of the n-$(n+1)$-$(n+2)$ Conjecture.

Corollary 5.6. *Let*

$$N_1 \hookrightarrow \Gamma_1 \overset{\pi_1}{\twoheadrightarrow} Q$$
$$N_2 \hookrightarrow \Gamma_2 \overset{\pi_2}{\twoheadrightarrow} Q$$

be short exact sequences with Γ_1, Γ_2 *of type* F_{n+1} *and* N_1 *of type* F_n *and assume the second sequence splits. Then the associated fibre product is of type* F_{n+1}.

Proof. Let $\sigma : Q \to \Gamma_2$ be a splitting homomorphism for the second sequence, so $\pi_2 \circ \sigma = \mathrm{id}_Q$. Now consider the exact sequence

$$N_1 \hookrightarrow P \overset{p_2}{\twoheadrightarrow} \Gamma_2.$$

Define

$$\phi : \Gamma_1 \to P, \gamma_1 \mapsto (\gamma_1, \sigma\pi_1(\gamma_1)).$$

Note that since $\pi_1(\gamma_1) = \pi_2\sigma\pi_1(\gamma_1)$ the image of ϕ does indeed lie in P. Obviously, ϕ restricts to the identity on N_1. Therefore, Proposition 5.3 applies to yield the conclusion. □

[3] Note that in the symmetric case, where the two short exact sequences are identical, there always is a natural homomorphism $\Gamma_1 \to P \leq \Gamma_1 \times \Gamma_1$, sending γ to (γ, γ). But this homomorphism does not fit into a commutative diagram, as in the statement of Proposition 5.3, since it does not map $N_1 \leq \Gamma_1$ to $N_1 \times 1 \leq P$.

This result was previously obtained, using different methods, by Martin Bridson and Ken Brown in unpublished work.

As for the 1-2-3 Theorem, we can use Proposition 5.3 not just to prove the split case but also to reduce the general case of the n-$(n+1)$-$(n+2)$ Conjecture to the case when Γ_2 is finitely generated free, by the same method as before.

Consider again the sequences

$$N_1 \hookrightarrow \Gamma_1 \overset{\pi_1}{\twoheadrightarrow} Q$$
$$N_2 \hookrightarrow \Gamma_2 \overset{\pi_2}{\twoheadrightarrow} Q$$

with N_1 of type F_n, Γ_1 and Γ_2 of type F_{n+1} and Q of type F_{n+2} Now let F be a finitely generated free group with an epimorphism $p : F \to \Gamma_2$. Let P' denote the fibre product associated to the short exact sequences

$$N_1 \hookrightarrow \Gamma_1 \overset{\pi_1}{\twoheadrightarrow} Q$$
$$N_2' \hookrightarrow F \overset{\pi_2 p}{\twoheadrightarrow} Q$$

where $N_2' = \ker(\pi_2 p) = p^{-1}(N_2)$. Now the homomorphism

$$\mathrm{id} \times p : \Gamma_1 \times F \to \Gamma_1 \times \Gamma_2, (\gamma_1, g) \mapsto (\gamma_1, p(g))$$

maps P' onto P and clearly sends $N_1 \times 1$ to $N_1 \times 1$. Thus by Proposition 5.3, P is of type F_{n+1} if P' is.

Proposition 5.7. *If the n-$(n+1)$-$(n+2)$ Conjecture holds whenever Γ_2 is finitely generated free then it holds in general.*

5.4 Classifying complexes for fibre products with one free factor

Recall how we used the reduction to the case of free Γ_2 in the proof of the 1-2-3 Theorem: In this special case, we constructed presentations for Γ_1 and P (Section 4.2) of the form:

$$\Gamma_1 \cong \langle \mathcal{X} \cup \tilde{\tilde{\mathcal{Z}}} \mid R^0 \cup R^1 \cup R^2, \partial \rangle$$
$$P \cong \langle \mathcal{X} \cup \tilde{\tilde{\mathcal{Z}}} \mid R^0 \cup R^1, \partial \rangle$$

In other words, the presentation complex for P is a subcomplex of that for Γ_1, missing some 2-cells derived from the relators of Q.

Now that we know that the general case of the n-$(n+1)$-$(n+2)$ Conjecture can likewise be reduced to the case of free Γ_2, we might ask whether we can find an analogue of this statement for entire classifying complexes. In fact, we have a very satisfying equivalent:

Proposition 5.8. *Let*

$$N_1 \hookrightarrow \Gamma \overset{\pi_1}{\twoheadrightarrow} Q$$
$$N_2 \hookrightarrow F \overset{\pi_2}{\twoheadrightarrow} Q$$

be short exact sequences with F free and let P be the associated fibre product. Assume that X is a $\mathrm{B}N_1$-complex and that Z is a $\mathrm{B}Q$-complex with one vertex such that the inclusion $Z^{(1)} \to Z$ induces the map π_2. Then there is a complex Y and a cellular map $\phi : Y \to Z$ such that

1. *Y is a $\mathrm{B}\Gamma$-complex,*
2. *ϕ is a stack with fibre X which induces the homomorphism π_1 on fundamental groups, and*
3. *$W := \phi^{-1}(Z^{(1)})$ is a $\mathrm{B}P$-complex, and the commutative (pullback) square of spaces on the left induces the commutative square of fundamental groups on the right:*

$$\begin{tikzcd} W \arrow[r, "\phi|"] \arrow[d, hook] & Z^{(1)} \arrow[d, hook] \\ Y \arrow[r, two heads, "\phi"] & Z \end{tikzcd} \qquad \begin{tikzcd} P \arrow[r, "p_1"] \arrow[d, "p_2"] & F \arrow[d, two heads, "\pi_2"] \\ \Gamma \arrow[r, two heads, "\pi_1"] & Q \end{tikzcd}$$

Remark 5.9. Of course, there is no difficulty in choosing a $\mathrm{B}Q$-complex Z as in the assumption. Pick a generating set $\mathcal{Z}$ of F. We may regard $\mathcal{Z}$ as a generating set for Q via the epimorphism p_2. So pick a presentation $Q \cong \langle \mathcal{Z} \mid R_Q, \partial \rangle$ (with p_2 being the implied epimorphism $F(\mathcal{Z}) \twoheadrightarrow Q$). Then let $Z^{(2)}$ be the corresponding presentation complex and extend it to a classifying complex Z (Lemma A.5). Clearly, $Z^{(1)}$ is a classifying complex for F (the presentation complex associated to $F \cong \langle \mathcal{Z} \mid - \rangle$), and the inclusion $Z^{(1)} \hookrightarrow Z$ induces p_2.

In particular, if F is finitely generated and Q satisfies some finiteness condition F_n, then there is always a Z as in the statement with finite n-skeleton (use Proposition A.18 before Lemma A.5, when extending the presentation complex to a $\mathrm{B}Q$).

Proof. We follow the Borel construction outlined at the start of the chapter. Let Y_0 be a $\mathrm{B}\Gamma$-complex. Then $\mathrm{B}\Gamma$ acts freely on $\widetilde{Y}_0 \times \widetilde{Z}$, where $\widetilde{Y}_0$ and $\widetilde{Z}$ denote the universal covers of Y_0 and Z respectively. Let $Y' := (\widetilde{Y}_0 \times \widetilde{Z})/\Gamma$. Then the projection map $\widetilde{Y}_0 \times \widetilde{Z} \to \widetilde{Z}$ induces a map on the quotients $\phi' : Y' \to Z$ so the following diagram commutes:

$$\begin{tikzcd} \widetilde{Y}_0 \times \widetilde{Z} \arrow[r, two heads] \arrow[d, two heads] & \widetilde{Z} \arrow[d, two heads] \\ Y' \arrow[r, two heads, "\phi'"] & Z \end{tikzcd}$$

Now ϕ induces the homomorphism $\pi_1 : \Gamma \to Q$ on fundamental groups and, as in the discussion at the start of the chapter, ϕ' is a fibre bundle and a stack with fibre $X' := \widetilde{Y}_0/N_1$, which is a $\mathrm{B}N_1$-complex. We can then "rebuild" ϕ' ([Geo08, Proposition 6.1.4]) to yield a new stack $\phi : Y \to Z$ with fibre X and a homotopy equivalence $h : Y' \to Y$ making the following diagram commute up to homotopy:

$$\begin{tikzcd} Y' \arrow[rr, "h"] \arrow[dr, "\phi'"'] & & Y \arrow[dl, "\phi"] \\ & Z & \end{tikzcd}$$

In fact, as we mentioned earlier, this diagram can be made to commute up to homotopy "over every cell", meaning that the homotopy equivalence h and a homotopy $H : Y' \times I \to Z$ from ϕ' to $\phi \circ h$ can be chosen in such a way that $H(\phi'^{-1}(e) \times I) \subset e$ for each closed cell $e \in Z$. In particular, this implies that for any subcomplex $Z_0 \subset Z$, the homotopy equivalence h restricts to a homotopy equivalence of the preimages of Z_0 under ϕ' and ϕ respectively.

So set $W' := \phi'^{-1}(Z^{(1)})$ and $W := \phi^{-1}(Z^{(1)})$. Then $h|_{W'}$ is a homotopy equivalence, making the diagram

$$\begin{tikzcd} W' \arrow[rr, "h|"] \arrow[dr, "\phi'|"'] & & W \arrow[dl, "\phi|"] \\ & Z^{(1)} & \end{tikzcd}$$

commute up to homotopy.

Furthermore, the restriction $\phi'| : W' \to Z^{(1)}$ is still a fibre bundle, so there is a long exact sequence of homotopy groups:

$$\cdots \to \pi_n(X') \overset{\iota_*}{\to} \pi_n(W') \overset{(\phi'|)_*}{\to} \pi_n(Z^{(1)}) \to \pi_{n-1}(X') \to \cdots$$

For $n \geq 2$ the groups $\pi_n(X')$ and $\pi_n(Z^{(1)})$ are trivial, so we can conclude that the same is true of $\pi_n(W')$. Hence W' is aspherical. This means that W' is a classifying complex for $P' := \pi_1(W')$. We will now show that $P' \cong P$.

In the lowest dimensions, the long exact homotopy sequence of the fibre bundle yields a short exact sequence of groups:

$$\begin{tikzcd} \cdots \arrow[r] & \pi_2(Z^{(1)}) \arrow[r] \arrow[d, equal] & \pi_1(X') \arrow[r, "\iota_*"] \arrow[d, equal] & \pi_1(W') \arrow[r, "(\phi'|)_*"] \arrow[d, equal] & \pi_1(Z^{(1)}) \arrow[r] \arrow[d, equal] & \pi_0(X') \arrow[d, equal] \\ & 1 & N & P' & F & 1 \end{tikzcd}$$

Note that the inclusion map $\varepsilon : W' \hookrightarrow Y'$ induces a map $\varepsilon_* : P' \to \Gamma$. The commutative diagram of spaces on the left then induces a commutative diagram on the level of fundamental groups (as before, we identify X' with the preimage of the single vertex in Z):

$$\begin{tikzcd}
X' \arrow[r, hook] \arrow[d, equal] & W' \arrow[r, "\phi'|"] \arrow[d, hook] & Z^{(1)} \arrow[d, hook] \\
X' \arrow[r, hook] & Y' \arrow[r, two heads, "\phi'"] & Z
\end{tikzcd}
\qquad
\begin{tikzcd}
N_1 \arrow[r, hook] \arrow[d, equal] & P' \arrow[r, "(\phi'|)_*"] \arrow[d, "\varepsilon_*"] & F \arrow[d, two heads, "\pi_2"] \\
N_1 \arrow[r, hook] & \Gamma \arrow[r, two heads, "\pi_1"] & Q
\end{tikzcd}$$

We can define a homomorphism $\eta : P' \to \Gamma \times F, g \mapsto (\varepsilon_*(g), \phi_*(g))$ and commutativity of the right-most square just means that η is in fact a homomorphism into the fibre product $P \leq \Gamma \times F$.

$$\begin{tikzcd}
& N_1 \arrow[rr, hook] \arrow[dl, equal] \arrow[dd, equal] & & P' \arrow[rr, two heads, "(\phi|)_*"] \arrow[dl, "\eta"'] \arrow[dd, "\varepsilon_*"] & & F \arrow[dl, equal] \arrow[dd] \\
N_1 \arrow[rr, hook] \arrow[dr, equal] & & P \arrow[rr, two heads, "p_2"] \arrow[dr, "p_1"'] & & F \arrow[dr, two heads] & \\
& N_1 \arrow[rr, hook] & & \Gamma \arrow[rr, two heads, "\pi_1"] & & Q
\end{tikzcd}$$

We have observed above that the "front" of this diagram commutes. The "bottom" part obviously commutes as well, and the middle triangle commutes by definition of η. Therefore the "top" squares commute and thus η is an isomorphism.

So we have showed that W' is a classifying complex for P. But W is homotopy equivalent to W' and thus a BP as well. It only remains to show the claim about the commutative squares in the statement. Again, we have seen that this claim is true when W and Y are replaced by W' and Y'. So the claim about W and Y follows from commutativity (up to homotopy) of the following diagram:

$$\begin{tikzcd}
& W \arrow[dr, two heads, "\phi|"] \arrow[dd, hook] & \\
W' \arrow[ur, "h|"] \arrow[rr, two heads, "\phi'|"] \arrow[dd, hook] & & Z^{(1)} \arrow[dd, hook] \\
& Y \arrow[dr, two heads, "\phi"] & \\
Y' \arrow[ur, "h"] \arrow[rr, two heads, "\phi'"] & & Z
\end{tikzcd}$$

This finishes the proof. □

The classifying complex for P constructed in this proposition gives a good idea of where the problem lies in proving the n-$(n+1)$-$(n+2)$ Conjecture. To apply it in this situation, assume that N_1 is of type F_n, that Γ_1 and Q are of type F_{n+1} and that F is a finitely generated free group. Then we can pick the BN_1-complex X in the statement of Propisition 5.8 to have finite n-skeleton and the BQ-complex Z to have finite $(n+1)$-skeleton. The proposition gives us a $\mathrm{B}\Gamma$-complex Y and a stack $\phi : Y \to Z$ with fibre X, such that ϕ induces the homomorphism $\pi_1 : \Gamma \to Q$ on fundamental groups and such that the preimage $W := \phi^{-1}(Z^{(1)})$ is a $\mathrm{B}P$-complex. As in the previous sections, Y has finite n-skeleton and only finitely many $(n+1)$-cells outside of X (the embedded fibre over the single vertex of Z). The same is then obviously true for W. Since Γ is of type F_{n+1}, we know that we can remove all but finitely many of the $(n+1)$-cells from X while keeping Y n-aspherical.[4] The question is whether we can also remove all but finitely many $(n+1)$-cells from X while keeping W n-aspherical. Another way to put this: There is a finite set E of $(n+1)$-cells of X such that any n-sphere $f : S^n \to X^{(n)}$ can be extended to $\bar{f} : B^{n+1} \to Y^{(n+1)}$ in such a way that the only $(n+1)$-cells of X that $\operatorname{im} \bar{f}$ intersects are from E. Now we want to do the same thing, only this time, we demand that the extension $\bar{f}$ *lie over the* 1*-skeleton of* Z. Precisely: Is there a finite set E' of $(n+1)$-cells of X such that any n-sphere $f : S^n \to X^{(n)}$ can be extended to $\bar{f} : B^{n+1} \to W^{(n+1)}$ in such a way that the only $(n+1)$-cells of X that $\operatorname{im} \bar{f}$ intersects are from E'? If there is such a set, then P is of type F_{n+1}.

Of course, up to this point of the analysis, we did not use the condition that Q be of type F_{n+2} at all. Based on our experiences from the proof of the 1-2-3 Theorem we might predict that this condition could be used for "pushing disks" from Y into W. We will pick up this thread again in Chapter 7.

Further remarks We might remark that the classifying space W for P constructed in Proposition 5.8 is satisfying from a conceptual viewpoint, since it realises the pullback square of groups characterising the fibre product by a pullback square of classifying complexes. As a curious sidenote, the method of how this complex is constructed is in a certain way formally dual to the way a classifying space for the pushout of groups, i.e. the amalgamated product, is usually constructed (see, e.g., [SW79] or [Geo08, Section 6.2]). Given two monomorphisms $\iota_1 : N \hookrightarrow \Gamma_1$ and $\iota_2 : N \hookrightarrow \Gamma_2$ and classifying spaces X, Y_1 and Y_2 for N, Γ_1 and Γ_2 respectively, we can construct a classifying space for the amalgamated product $\Gamma_1 *_N \Gamma_2$ by gluing the two ends of the "cylinder" $X \times [0,1]$ into Y_1 and Y_2 respectively, by maps realising ι_1 and ι_2. There is another way to phrase this: We may start off picking maps $f_1 : X \to Y_1$ and $f_2 : X \to Y_2$ which induce ι_1 and ι_2 on fundamental groups. The pushout of f_1 and f_2 (in the category of topological spaces) will usually not be a CW-complex, though, so we modify one of the maps first. Specifically, we replace Y_1 by the mapping cylinder of f_1. Call this $Y_1' := Y_1 \cup_{f_1} (X \times [0,1])$ and let $f_1' : X \to Y_1'$ be the natural inclusion. Then the inclusion $Y_1 \hookrightarrow Y_1'$ is a homotopy equivalence (so Y_1' is still a $\mathrm{B}\Gamma_1$) and the

[4] more precisely, we should say, as before, that we can find a finite subcomplex $X'^{(n+1)} \subset X^{(n+1)}$ with $X'^{(n)} = X^{(n)}$ and such that $(Y^{(n+1)} \smallsetminus X^{(n+1)}) \cup X'^{(n+1)}$ is n-aspherical (see Lemma A.16)

square

$$\begin{tikzcd} & X \arrow[dl, "f_1"'] \arrow[dr, "f_1'"] & \\ Y_1 \arrow[rr, hook] & & Y_1' \end{tikzcd}$$

commutes. Now the pushout Z of f_1' and f_2 is a CW-complex and a classifying space for $\Gamma_1 *_N \Gamma_2$.

$$\begin{tikzcd} X \arrow[r, "f_2"] \arrow[d, hook, "f_1'"] & Y_2 \arrow[d] \\ Y_1' \arrow[r] & Z \end{tikzcd} \qquad \begin{tikzcd} C \arrow[r, hook] \arrow[d, hook] & B \arrow[d, hook] \\ A \arrow[r, hook] & A *_C B \end{tikzcd}$$

If we choose the classifying complexes X, Y_1 and Y_2 such that X has finite n-skeleton, and Y_1 and Y_2 have finite $(n+1)$-skeleton, then the resulting complex Z has finite $(n+1)$-skeleton. This proves that an amalgamated product of two groups of type F_{n+1} along a common subgroup of type F_n is of type F_{n+1}.

Now note how this is formally similar to our construction for fibre bundles. This time, we are given two epimorphisms $\pi_1 : \Gamma \twoheadrightarrow Q$ and $\pi_2 : F \to Q$. We pick classifying complexes Y_1 and Z for Γ_1 and Q and a map realising π_1. Then we replace this map by a fibration, namely the fibre bundle $\phi' : Y' \to Z$ coming out of the Borel construction (note that above, we replaced f_1 by a cofibration f_1'). In the proof, we essentially showed that the pullback of this map ϕ' and the inclusion $Z^{(1)} \hookrightarrow Z$ (which realises π_2) is a classifying complex for the pullback of π_1 and π_2. Unlike for amalgamated products, the resulting complex does not automatically have good finiteness properties. Specifically, in order to make use of the extra finiteness condition on the kernel N_1, we had to further replace ϕ' with the stack $\phi : Y \to Z$. Still, the resulting pullback falls short of providing the finiteness conditions demanded by the n-$(n+1)$-$(n+2)$ Conjecture, but it certainly conveys much greater insight into the nature of the problem.

6

Homological finiteness properties of fibre products

In the topological approach presented in the previous section, we endeavoured to prove the n-$(n+1)$-$(n+2)$ Conjecture by constructing a finite $(n+1)$-skeleton of a classifying space of the fibre product directly. We have seen how proving a group to be of type F_{n+1} comes down to proving a certain higher homotopy group to be finitely generated.

An essential feature in the proofs of the 0-1-2 Lemma and the 1-2-3 Theorem is the fact that the low-dimensional homotopy groups of CW-complexes have combinatorial descriptions. The fundamental group of a 1-dimensional complex can be described by edge paths and this enables us to represent the 2-skeleton of a classifying complex by generators and relations. One dimension up we can describe the second homotopy group of a 2-dimensional complex in a similar way, using crossed modules. This combinatorial formalism is crucial to the proof of the 1-2-3 Theorem.

But as we mentioned at the end of Section 4 there is no completely satisfying analogue of this combinatorial formalism in higher dimensions. This is one of the reasons why higher homotopy groups and are so notoriously hard to compute, and it should come as no surprise that the same is usually true of the higher finiteness properties. In the previous section we have attempted to make do with the more conceptual understanding that stacks provide, but without the algebraic tools to do explicit combinatorial computation, we had to content ourselves with introducing some rather crude restrictions, such as the split condition.

Here we will take a different course.

It is in part because of the inherent difficulties with computation of higher homotopy groups, that we often consider homology groups when a more computable invariant of spaces is needed. And so it is natural to turn to homological algebra in search for more computable finiteness properties of groups. As is usual, computability comes at a price: The homological analogues will generally be weaker than the properties F_n.

One of these properties is the "type FP_n" condition, which we will consider in Section 6.3 and in Chapter 7. Here we study the weaker property of having finitely generated integral homology groups up to a certain dimension: If Γ is a group of type F_n then $H_k(\Gamma, \mathbb{Z})$ is finitely generated for all $k \leq n$. This implication is commonly used to show that a group is *not* of type F_n. In the major result of this section we will show that under the assumptions of the n-$(n+1)$-$(n+2)$ Conjecture the fibre product has finitely generated integral homology groups up to dimension $n+1$.

6.1 Type wFP_n

We will consider the following homological finiteness condition.[1]

Definition 6.1. A group Γ is of type wFP_n ("weak FP_n") if $H_k(\Gamma', \mathbb{Z})$ is finitely generated for all $k \leq n$ and all finite index subgroups $\Gamma' \leq \Gamma$.

Recall the definition of homology groups. Choosing any free resolution of $\mathbb{Z}$ by Γ-modules

$$\mathbf{F} \twoheadrightarrow \mathbb{Z} : \ \cdots \to F_2 \to F_1 \to F_0 \twoheadrightarrow \mathbb{Z}$$

the homology groups of Γ with coefficients in some Γ-module M are defined as[2]

$$H_k(\Gamma, M) = H_k(\mathbf{F} \otimes_\Gamma M),$$

and this definition does not depend on the choice of free resolution $\mathbf{F} \twoheadrightarrow \mathbb{Z}$ (more generally, one can also choose projective or flat resolutions of $\mathbb{Z}$).

We will now show that "type[3] FP_n" implies "type wFP_n". We make the following observation:

Lemma 6.2. *Let Λ be a ring, A a finitely generated right Λ-module and B a left Λ-module which is finitely generated as an abelian group. Then $A \otimes_\Lambda B$ is finitely generated as an abelian group.*

Proof. Let $X \subset A$ be a finite generating set of A as a Λ-module and let $Y \subset B$ be a finite generating set of B as an abelian group.

[1] A similar condition has been used before by K. Brown in [Bro74, Bro75] in the context of Euler characteristics of groups. He defines a group to be of finite homological type if it has finite virtual cohomological dimension and every finite index subgroup has finitely generated integral homology. A very similar condition (also called "finite homological type") is used in [Bro94], see also [LS06]. We will discuss the latter condition in more detail at the end of the chapter.

[2] Unless otherwise noted, we employ the usual convention of taking all modules over a group ring to be left modules. Any left Γ-module can be turned into a right Γ-module by defining $a \cdot g := g^{-1}a$ for all $a \in A$ and $g \in \Gamma$. The tensor product $A \otimes_G B$ of two left Γ-modules A and B is understood to employ this induced right-module structure on A.

[3] for more information on the FP_n-properties, see Appendix A.2

Let $a \in A$ and $b \in B$. Then

$$a = \sum_{x \in X} x\lambda_x, \quad \text{for certain } \lambda_x \in \Lambda.$$

And for any $x \in X$ and $y \in Y$ we can write

$$\lambda_x b = \sum_{y \in Y} n_{x,y} y, \quad \text{for certain } n_{x,y} \in \mathbb{Z}.$$

Therefore

$$a \otimes b = \sum_{x \in X} x\lambda_x \otimes b = \sum_{x \in X} x \otimes \sum_{y \in Y} n_{x,y} y = \sum_{\substack{x \in X \\ y \in Y}} n_{x,y} (x \otimes y)$$

Hence $A \otimes_\Lambda B$ is generated as an abelian group by the set $\{x \otimes y \mid x \in X, y \in Y\}$. $\square$

This immediately implies:

Proposition 6.3. *If Γ is of type* FP$_n$ *and M is a Γ-module which is finitely generated as an abelian group then $H_k(\Gamma, M)$ is finitely generated for all $k \leq n$.*

Proof. Let Γ be a group of type FP$_n$ and M a Γ-module which is finitely generated as an abelian group. Pick a free resolution

$$\mathbf{F} \twoheadrightarrow \mathbb{Z} : \quad \cdots \to F_2 \to F_1 \to F_0 \twoheadrightarrow \mathbb{Z}$$

of the trivial Γ-module $\mathbb{Z}$ which is finitely generated up to dimension n. Then the groups $F_k \otimes_\Gamma M$ are finitely generated for $k \leq n$ and hence so is $H_k(\Gamma, M) = H_k(\mathbf{F} \otimes_\Gamma M)$. $\square$

In conjunction with Proposition A.34 this yields:

Corollary 6.4. *Any group of type* FP$_n$ *is of type* wFP$_n$. $\square$

That the converse is not true was shown by Leary and Saadetoğlu in [LS06] and we recount their construction in the Appendix (Proposition A.49).

The LHS spectral sequence To prove an analogue of the n-$(n+1)$-$(n+2)$ Conjecture using the wFP$_n$ condition, we need to study how the homology groups of the groups in a short exact sequence are related to one another. The main tool for doing this is the Lyndon-Hochschild-Serre spectral sequence. For a short exact sequence $N \hookrightarrow \Gamma \twoheadrightarrow Q$ this is a first quadrant spectral sequence of the form

$$E^2_{pq} = H_p(Q, H_q(N, M)) \Rightarrow H_{p+q}(\Gamma, M).$$

where M is any Γ-module. Let us briefly recall how the action of Q on $H_q(N, M)$ arises.

Let G and H be two groups and M a G-module, B an H-module. If $\phi : G \to H$ is a homomorphism then $f : M \to N$ is called *compatible with* ϕ, if $f(gm) = \phi(g)f(m)$ for all $g \in G$ and $m \in M$, i.e. f is a homomorphism of G-modules, if N is considered a G-module via $\phi : G \to H$.

Such a pair (ϕ, f) of compatible morphisms induces a map on homology in the following way: Choose a free resolution

$$\mathbf{F} \twoheadrightarrow \mathbb{Z} : \quad \cdots \to F_2 \to F_1 \to F_0 \twoheadrightarrow \mathbb{Z}$$

of $\mathbb{Z}$ by G-modules and a free resolution

$$\mathbf{F}' \twoheadrightarrow \mathbb{Z} : \quad \cdots \to F_2' \to F_1' \to F_0' \twoheadrightarrow \mathbb{Z}$$

of $\mathbb{Z}$ by H-modules. Then there is a map of chain complexes

$$\begin{tikzcd}
\mathbf{F} \twoheadrightarrow \mathbb{Z} : \arrow[d, "g_*"] & \cdots \arrow[r] & F_2 \arrow[r] \arrow[d, "g_2"] & F_1 \arrow[r] \arrow[d, "g_1"] & F_0 \arrow[r, two heads] \arrow[d, "g_0"] & \mathbb{Z} \arrow[d, equal] \\
\mathbf{F}' \twoheadrightarrow \mathbb{Z} : & \cdots \arrow[r] & F_2' \arrow[r] & F_1' \arrow[r] & F_0' \arrow[r, two heads] & \mathbb{Z}
\end{tikzcd}$$

extending the identity map on $\mathbb{Z}$ and compatible with ϕ (i.e. each f_n is compatible with ϕ in the same sense as before).

Tensoring this map with f gives a chain map $\mathbf{F} \otimes M \to \mathbf{F} \otimes N$:

$$\begin{tikzcd}
\mathbf{F} \otimes M : \arrow[d, "g_* \otimes f"] & \cdots \arrow[r] & F_2 \otimes M \arrow[r] \arrow[d, "g_2 \otimes f"] & F_1 \otimes M \arrow[r] \arrow[d, "g_1 \otimes f"] & F_0 \otimes M \arrow[r] \arrow[d, "g_0 \otimes f"] & 0 \\
\mathbf{F}' \otimes N : & \cdots \arrow[r] & F_2' \otimes N \arrow[r] & F_1' \otimes N \arrow[r] & F_0' \otimes N \arrow[r] & 0
\end{tikzcd}$$

This chain map induces a homomorphism on homology

$$(\phi, f)_* : H_*(G, M) \to H_*(H, N),$$

which, as usual, does not depend on any of the choices made during the construction.

Now let $N \hookrightarrow \Gamma \overset{\pi}{\twoheadrightarrow} Q$ be a short exact sequence and M a Γ-module. We can now consider the special case of the above construction where ϕ is a conjugation homomorphism. For $\gamma \in \Gamma$, let $c^\gamma : N \to N, h \mapsto \gamma h \gamma^{-1}$. Then the module homomorphism $t_\gamma : M \to M, m \mapsto \gamma m$ is compatible with c^γ, since

$$t_\gamma(hm) = ghm = \gamma h \gamma^{-1}(\gamma m) = c^\gamma(t_\gamma(m)), \quad \text{for } h \in N, m \in M.$$

So we can follow the above construction to obtain an induced endomorphism $c_*^\gamma : H_*(N, M) \to H_*(N, M)$.

One way of doing this concretely is by picking a free resolution of $\mathbb{Z}$ by Γ-modules

$$\mathbf{F} \twoheadrightarrow \mathbb{Z}: \quad \cdots \to F_2 \to F_1 \to F_0 \twoheadrightarrow \mathbb{Z}$$

and considering it as a complex of N-modules. We then define a chain map g_*^γ : $\mathbf{F} \to \mathbf{F}$ as

$$g_n^\gamma : F_n \to F_n, a \mapsto \gamma a.$$

This extends the identity map on $\mathbb{Z}$ and is compatible with c^γ. Tensoring with t_γ gives a map $f_*^\gamma \otimes t_\gamma : \mathbf{F} \otimes M \to \mathbf{F} \otimes M$ which induces the desired map c_*^γ on homology. Inspecting this construction it is easy to see that the maps c_*^γ are all automorphisms ($c_*^{\gamma^{-1}}$ being an inverse) and that the map $\Gamma \to \mathrm{Aut}(H_*(N,M)), \gamma \mapsto c_*^\gamma$ is a homomorphism. Thus we have defined an action of Γ on the homology groups $H_n(N,M)$ given by

$$\gamma \cdot a = c_*^\gamma(a), \quad \text{for } \gamma \in \Gamma, a \in H_n(N,M).$$

Furthermore, if $h \in N$ then

$$(f_n^h \otimes t^h)(a \otimes m) = ha \otimes hm = a \otimes m, \quad \text{for all } n \geq 0, a \in F_n, m \in M$$

and so N itself acts trivially on $H_*(N,M)$. Thus the action of Γ on $H_*(N,M)$ factors through $\pi : \Gamma \to Q$, giving an action of Q on $H_*(N,M)$.

6.2 The Weak n-$(n+1)$-$(n+2)$ Theorem

Let $N \hookrightarrow \Gamma \twoheadrightarrow Q$ be a short exact sequence. Using the LHS spectral sequence it is easy to prove that[4] $H_n(\Gamma)$ will be finitely generated if Q is of type FP_n and N is of type wFP_n. Looking a little more carefully, we see that in fact, we can get away with a weaker condition on N:

Lemma 6.5. *Let $N \hookrightarrow \Gamma \twoheadrightarrow Q$ be a short exact sequence with Q of type FP_n and N of type wFP_{n-1}. Also, assume that $H_0(Q, H_n(N))$ is finitely generated. Then $H_k(\Gamma,\mathbb{Z})$ is finitely generated for $k \leq n$.*

Proof. This follows easily from the LHS spectral sequence. Using Proposition 6.3, the assumptions on N and Q imply that $E_{pq}^2 = H_p(Q, H_q(N))$ is finitely generated for $q \leq n-1$ and $p \leq n$. In particular, E_{pq}^2 is finitely generated whenever $p+q \leq n-1$.

[4] as usual, we write $H_n(\Gamma) := H_n(\Gamma,\mathbb{Z})$

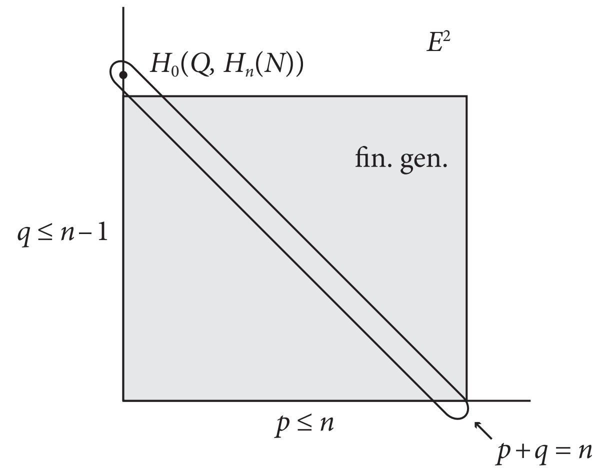

Since $H_0(Q, H_n(N))$ is finitely generated as well, all the entries E^2_{pq} on the diagonal $p + q = n$ are finitely generated. Then the same holds true for the corresponding entries on the E^∞ page. But for any k, the entries E^∞_{pq} on the diagonal $p + q = k$ are the factors of a filtration of $H_k(\Gamma)$, and so $H_k(\Gamma)$ is finitely generated for all $k \leq n$. □

Conversely, if Γ is of type wFP_n, then the stated condition for the kernel N holds, whenever $H_{n+1}(Q, \mathbb{Z})$ is finitely generated.

LEMMA 6.6. *Let $N \hookrightarrow \Gamma \twoheadrightarrow Q$ be a short exact sequence of groups. Assume that N is of type* wFP_{n-1}*, that Γ is of type* wFP_n *and that Q is of type* FP_n*.*
If $H_{n+1}(Q)$ is finitely generated, then so is $H_0(Q, H_n(N))$. If Γ is of type wFP_{n+1} *the converse holds as well.*

Remark 6.7. In the case $n = 1$ this lemma is a "weak" analogue of 3.7.

Proof. Consider the LHS spectral sequence

$$E^2_{pq} = H_p(Q, H_q(N)) \Rightarrow H_{p+q}(\Gamma).$$

The assumptions on N and Q imply that E^2_{pq} is finitely generated for $q \leq n-1$ and $p \leq n$.

Let us first focus on the entry at position $(n+1, 0)$ in the spectral sequence. On the E^{r+1}-page, $E^{r+1}_{n+1,0}$ is the kernel of the map $d_r : E^r_{n+1,0} \to E^r_{n+1-r,r-1}$. The codomain of this map is finitely generated for $2 \leq r \leq n$ and thus $E^{r+1}_{n+1,0} = \ker d_r$ is finitely generated if and only if $E^r_{n+1,0}$ is. Hence, by induction, $E^{n+1}_{n+1,0}$ is finitely generated if and only if $E^2_{n+1,0} = H_{n+1}(Q, \mathbb{Z})$ is.

Similarly, looking at the entry $(0, n)$ we find that $E^{r+1}_{0,n}$ is the cokernel of the map $d_r : E^r_{r,n-r+1} \to E^r_{0,n}$. The groups $E^r_{r,n-r+1}$ are finitely generated for $2 \leq r \leq n$ and thus $E^{r+1}_{0,n} = \operatorname{coker} d_r$ is finitely generated if and only if $E^r_{0,n}$ is. By induction, $E^{n+1}_{0,n}$ is finitely generated if and only if $E^2_{0,n} = H_0(Q, H_n(N))$ is.

On the E^{n+1}-page we have a map

$$d_{n+1} : E^{n+1}_{n+1,0} \to E^{n+1}_{0,n} \tag{6.1}$$

The cokernel of this map is $E^{n+2}_{0,n} = E^{\infty}_{0,n}$, which is finitely generated, since Γ is of type wFP_n. Thus, if $E^{n+1}_{n+1,0}$ is finitely generated, then so is $E^{n+1}_{0,n}$.

If Γ is of type wFP_{n+1}, then the kernel of (6.1), $E^{n+2}_{n+1,0} = E^{\infty}_{n+1,0}$ is finitely generated as well. So in this case, $E^{n+1}_{n+1,0}$ is finitely generated *if and only if* $E^{n+1}_{0,n}$ is. □

Now we can prove the desired "weak" version of the n-$(n+1)$-$(n+2)$ Conjecture.

THEOREM 6.8 (THE WEAK n-$(n+1)$-$(n+2)$ THEOREM). *Let*

$$N_1 \hookrightarrow \Gamma_1 \overset{\pi_1}{\twoheadrightarrow} Q$$
$$N_2 \hookrightarrow \Gamma_2 \overset{\pi_2}{\twoheadrightarrow} Q$$

be short exact sequences with N_1 of type wFP_n, Γ_1 *of type* wFP_{n+1}, Γ_2 *of type* FP_{n+1}, *and Q of type* FP_{n+2}. *Then the associated fibre product is of type* wFP_{n+1}.

Proof. Denote the fibre product by P and the projection homomorphisms by $p_1 : P \to \Gamma_1$ and $p_2 : P \to \Gamma_2$.

We need to show that $H_k(P')$ is finitely generated for all $k \leq n+1$ and all finite index subgroups $P' \leq P$. First we will show that it is in fact enough to prove that $H_k(P)$ is finitely generated for $k \leq n+1$.

Let $P' \leq P$ be a subgroup of finite index and set $\Gamma_1' := p_1(P)$ and $\Gamma_2' := p_2(P)$. Then P' is itself a subdirect product of Γ_1' and Γ_2'. The associated short exact sequences are

$$N_1' \hookrightarrow \Gamma_1' \overset{\pi_1}{\twoheadrightarrow} Q'$$
$$N_2' \hookrightarrow \Gamma_2' \overset{\pi_2}{\twoheadrightarrow} Q'$$

where $N_1' := P' \cap \Gamma_1$, $N_2' := P' \cap \Gamma_2$ and $Q' = P'/(N_1' \times N_2')$. Now the Γ_i' and N_i' are finite index subgroups in Γ_i and N_i respectively, so they share the same finiteness properties. Furthermore $Q' \cong \Gamma_1'/N_1'$ is a finite index subgroup of Γ_1/N_1' which fits in a short exact sequence

$$N_1/N_1' \hookrightarrow \Gamma_1/N_1' \twoheadrightarrow Q.$$

Here N_1/N_1' is finite and Q is of type FP_{n+2}, so Γ_1/N_1' is of type FP_{n+2} as well.

Thus we have shown that any finite index subgroup $P' \leq P$ is itself a fibre product associated to two short exact sequences satisfying the same finiteness conditions as those in the statement of the theorem. Without loss of generality, it therefore suffices to show that $H_k(P, \mathbb{Z})$ is finitely generated for all $k \leq n$.

Consider first the short exact sequence

$$N_1 \hookrightarrow P \overset{p_1}{\twoheadrightarrow} \Gamma_2.$$

Here, Γ_2 is of type FP_{n+1} and N_1 is of type wFP_n. By Lemma 6.5, if we can show that $H_0(\Gamma_2, H_{n+1}(N_1))$ is finitely generated, it follows that $H_k(P)$ is finitely generated for $k \leq n$.

Now consider the short exact sequence

$$N_1 \hookrightarrow \Gamma_1 \overset{\pi_1}{\twoheadrightarrow} Q.$$

Here, we can apply Lemma 6.6 to conclude that $H_0(Q, H_{n+1}(N_1))$ is finitely generated. What remains then is to compare $H_0(\Gamma_2, H_{n+1}(N_1))$ with $H_0(Q, H_{n+1}(N_1))$, and we will show that these two modules are in fact isomorphic, finishing the proof.

Recall that for any group Γ and Γ-module A, the 0^{th} homology group $H_0(\Gamma, A)$ is the module of coinvariants

$$H_0(\Gamma, A) = A_\Gamma = A/\langle\{a - \gamma a \mid a \in A, \gamma \in \Gamma\}\rangle.$$

Let $\gamma_2 \in \Gamma_2$. As we explained above, the action of Γ_2 on $H_{n+1}(N_1)$ is induced by the conjugation action of P on N_1. More precisely, given $\gamma_2 \in \Gamma_2$, pick a lift $\gamma = (\gamma_1, \gamma_2) \in P$. Then the action of γ_2 on $H_{n+1}(N_1)$ is given by the map $c^\gamma_* : H_{n+1}(N_1) \to H_{n+1}(N_1)$ induced by the conjugation homomorphism $c^\gamma : N_1 \to N_1$. As we have remarked before, conjugating an element of $N_1 = N_1 \times 1$ by γ has the same effect as conjugating by γ_1 in Γ_1:

$$c^\gamma(n) = (\gamma_1, \gamma_2)(n, 1)(\gamma_1, \gamma_2)^{-1} = (\gamma_1 n \gamma_1^{-1}, 1).$$

Therefore $c^\gamma = c^{\gamma_1}$. Now set $q := \pi_1(\gamma_1) = \pi_2(\gamma_2)$. By definition of the Q-action on $H_{n+1}(N_1)$ (stemming from the short exact sequence $N_1 \hookrightarrow \Gamma_1 \twoheadrightarrow Q$), we have

$$qa = c^{\gamma_1}_*(a) = c^\gamma_*(a) = \gamma_2 a, \quad \text{for all } a \in H_{n+1}(N_1).$$

Since any $q \in Q$ lifts to some $\gamma_2 \in \Gamma_2$, we obtain that

$$\begin{aligned}(H_{n+1}(N_1))_{\Gamma_2} &= H_{n+1}(N_1)/\langle a - \gamma_2 a \mid \gamma_2 \in \Gamma_2, a \in H_{n+1}(N_1)\rangle \\ &= H_{n+1}(N_1)/\langle a - \pi_2(\gamma_2) a \mid \gamma_2 \in \Gamma_2, a \in H_{n+1}(N_1)\rangle \\ &= H_{n+1}(N_1)/\langle a - qa \mid q \in Q, a \in H_{n+1}(N_1)\rangle \\ &= (H_{n+1}(N_1))_Q\end{aligned}$$

□

Since the n-$(n+1)$-$(n+2)$ Conjecture implies the Virtual Surjections Conjecture, it should come as no surprise that Theorem 6.8 implies a weak variant of the Virtual Surjections Conjecture. The proof follows closely that of Theorem 3.14.

Corollary 6.9 (The Weak Virtual Surjections Theorem). *Let $k \geq 2$. Let $\Gamma_1, \ldots, \Gamma_n$ be groups of type FP_k and $P \leq \Gamma_1 \times \cdots \times \Gamma_n$ a subgroup of their direct product. If P virtually surjects to k-tuples of factors then P is of type wFP_k.*

Proof. We will use the notation established at the end of Chapter 2. As in Theorem 3.14 we prove the following claim by induction on k:

Claim. Let $k \geq 1$. Let $\Gamma_1, \ldots, \Gamma_n$ be groups of type FP_k and $P \leq \Gamma_1 \times \cdots \times \Gamma_n$ a subgroup of their direct product such that $P/(P \cap \Gamma_i)$ is virtually nilpotent for all i with $1 \leq i \leq n$. Then P is of type wFP_k.

For $k \geq 2$ the extra assumption on the $P/(P \cap \Gamma_i)$ is redundant, by 3.3, so the Claim does imply the statement of the Corollary. Furthermore we may assume, by the same argument as used in Chapter 3 (see Proposition 3.9), that P is a subdirect product, replacing each Γ_i by its finite index subgroup $p_i(P)$, if necessary.

In the base case $k = 1$, the groups $\Gamma_1, \ldots, \Gamma_n$ are FP_1 and thus finitely generated. So from Proposition 3.9 we have that P is finitely generated and thus a fortiori of type wFP_1.

Fix a $k \geq 2$. We prove the Claim for this k by induction on n. If $n \leq k$ then P is of finite index in $\Gamma_1 \times \cdots \times \Gamma_n$ and thus of type FP_k.

So let $n > k$. Set $T := p_{1,\ldots,n-1}(P) \leq \Gamma_1 \times \cdots \times \Gamma_{n-1}$. Then P is the fibre product associated to the short exact sequences

$$\begin{aligned} N_{1,\ldots,n-1} \hookrightarrow T &\twoheadrightarrow Q \\ N_n \hookrightarrow \Gamma_n &\twoheadrightarrow Q \end{aligned}$$

Here, $T \leq \Gamma_1 \times \cdots \times \Gamma_{n-1}$ is a subdirect product, which also virtually surjects to k-tuples, since $p_J(T) = p_J(P)$ for all $J \subset \{1, \ldots, n-1\}$. By induction, we can assume that T is of type wFP_k.

By Lemma 3.13, $N_{1,\ldots,n-1} \leq \Gamma_1 \times \cdots \times \Gamma_{n-1}$ virtually surjects to $(k-1)$-tuples. In addition,

$$\Gamma_i/(N_{1,\ldots,n-1} \cap \Gamma_i) = \Gamma_i/P \cap \Gamma_i$$

is virtually nilpotent for all $i \in \{1, \ldots, n-1\}$. Thus, by the inductive assumption on k, we can assume that $N_{1,\ldots,n-1}$ is of type wFP_{k-1}.

Furthermore, by assumption, $Q \cong \Gamma_n/N_n$ is virtually nilpotent and finitely generated and therefore of type F_∞, so certainly of type wFP_2. We can then apply Theorem 6.8 to conclude that P is of type wFP_k, finishing the induction. □

Kochloukova in [Koc10] proved a converse to this theorem in the case where the factors are non-abelian free (or limit) groups, which we quoted earlier (in a weaker form) as Theorem 3.5. Her theorem states [Koc10, Theorem 11]:

Theorem 6.10. *Let $\Gamma_1, \ldots, \Gamma_n$ be non-abelian limit groups and $P \leq \Gamma_1 \times \cdots \times \Gamma_n$ a full subdirect product. If P is of type* wFP_k *(for some $k \geq 2$) then P virtually surjects to k-tuples.* □

So in conjunction with Corollary 6.9 this gives a complete characterisation of the full subdirect products of type wFP_k of free groups:

Corollary 6.11. *Let $\Gamma_1, \ldots, \Gamma_n$ be non-abelian limit groups and $P \leq \Gamma_1 \times \cdots \times \Gamma_n$ a full subdirect product. Then P is of type* wFP_k *(for some $k \geq 2$) if and only if P virtually surjects to k-tuples.* □

Of course, the conjecture (which would follow from the Virtual Surjections Conjecture) is that "wFP_k" in this statement can be replaced by "F_k", which would that these two finiteness properties are equivalent (for $k \geq 2$) among the class of subgroups of direct products of free groups.

As we mentioned, examples of groups which are of type wFP_k but not of type F_k do exist for any k and we give a construction in Proposition A.49 in the Appendix. So we may note that, in principle, the Virtual Surjections Conjecture (and hence the n-$(n+1)$-$(n+2)$ Conjecture) could be disproved by embedding one such example in a direct product of free groups, though the known examples give no indication that this is possible.

Further remarks on Theorem 6.8 Let us investigate how much further we can push the method presented here to improve on the result of Theorem 6.8. First we may notice that we do not actually use the FP-conditions on Γ_2 and Q in the proof beyond the consequences of Proposition 6.3. Let us say that a group Γ is of type FHT_n if $H_k(\Gamma, A)$ is finitely generated for all $k \leq n$ and all Γ-modules A which are finitely generated as abelian groups.[5] Proposition 6.3 shows that property FP_n implies FHT_n and Leary and Saadetoğlu [LS06] showed that FHT_n is in fact strictly weaker than FP_n for all $n \geq 1$.

The proof of Theorem 6.8 only uses the fact that Q is of type FHT_{n+2} and that Γ_2 is of type FHT_{n+1}, so we may substitute these weaker properties for the corresponding FP_k-properties in the statement.

More saliently, the Weak n-$(n+1)$-$(n+2)$ Theorem can be interpreted as saying: If the n-$(n+1)$-$(n+2)$ Conjecture should fail, then this failure cannot be detected in the integral homology of finite index subgroups of the fibre product. This is crucial, since showing that a group has a non-finitely generated integral n^{th} homology group is a very common way of establishing that a group is not of type FP_n.

Having ruled out this way of disproving the Conjecture, how else might we use homology to confirm some suspected counterexample? We could try our luck with other coefficient modules: Finding any P-module A which is finitely generated as an abelian group, such that $H_{n+1}(P, A)$ is not finitely generated, would prove that P is not of type FP_{n+1}. Conversely, can the Weak n-$(n+1)$-$(n+2)$ Theorem be generalised to cover coefficient modules other than $\mathbb{Z}$? Can we show, under the conditions of the n-$(n+2)$-$(n+2)$ Conjecture, that the fibre product is of type FHT_{n+1}?

[5] As we noted before, a similar property was introduced by Brown in [Bro94, Section IX.6]. Brown defines a group Γ to be of finite homological type (or FHT) if it has finite virtual cohomological dimension and $H_k(\Gamma, M)$ is finitely generated for all $k \geq 0$ and any Γ-module M which is finitely generated as an abelian group. Our nomenclature is chosen in such a way that the $\mathrm{FHT}/\mathrm{FHT}_n$ properties are related in roughly the same way as the $\mathrm{FP}/\mathrm{FP}_n$ properties. A group Γ is of type FP if it has finite cohomological dimension and is of type FP_∞.

In fact, two steps of the proof easily generalise to other coefficient modules. Assume that we are given two sequences

$$N_1 \hookrightarrow \Gamma_1 \overset{\pi_1}{\twoheadrightarrow} Q$$
$$N_2 \hookrightarrow \Gamma_2 \overset{\pi_2}{\twoheadrightarrow} Q$$

with N_1 of type FHT_n, both Γ_1 and Γ_2 of type FHT_{n+1} and Q of type FHT_{n+2}. Denote their fibre product by P. By adapting Lemma 6.5, we can show that it is sufficient to show that $H_0(\Gamma_2, H_{n+1}(N_1, A))$ is finitely generated for any P-module A which is finitely generated as an abelian group.

Similarly, by adapting Lemma 6.6, we can show that $H_0(\Gamma_2, H_{n+1}(N_1, A))$ is finitely generated for any Γ_1-module A which is finitely generated as an abelian group.

The problem is that, in this case, we cannot adapt our previous method to show that the second statement implies the first: The first statement is about P-modules, the second about Γ_1-modules. In general, we cannot just consider a P-module as a Γ_1-module in any obvious way. However, this is possible for those P-modules A on which N_2 acts trivially. In this case, the action of P, given by a morphism $P \to \mathrm{Aut}(A)$ factors through $p_1 : P \twoheadrightarrow \Gamma_1$, giving a Γ_1-action on A. Then we can use the same method as in the proof of Theorem 6.8 to show that $H_0(\Gamma_2, H_{n+1}(N_1, A))$ is finitely generated.

To sum up, we have:

PROPOSITION 6.12. *Let*

$$N_1 \hookrightarrow \Gamma_1 \overset{\pi_1}{\twoheadrightarrow} Q$$
$$N_2 \hookrightarrow \Gamma_2 \overset{\pi_2}{\twoheadrightarrow} Q$$

be short exact sequences of groups with N_1 of type FHT_n*, both Γ_1 and Γ_2 of type* FHT_{n+1} *and Q of type* FHT_{n+2}*. Denote their fibre product by P.*
Then $H_{n+1}(P, A)$ is finitely generated for all P-modules A which are finitely generated as abelian groups and on which the subgroup $N_2 \leq P$ acts trivially. □

We do not know at present, whether this statement can be improved to cover arbitrary P-modules A with finitely generated underlying abelian group, i.e. to show that P is of type FHT_n.[6]

We should note that the wFP_n-condition can be rephrased in terms of homology with coefficients as well. If $\Gamma' \leq \Gamma$ is a subgroup then the Eckmann-Shapiro Lemma states that

$$H_k(\Gamma', A) \cong H_k(\Gamma, \Gamma \otimes_{\mathbb{Z}\Gamma'} A), \quad \text{for any } \Gamma'\text{-module } A.$$

In particular, if $\Gamma' \leq \Gamma$ is a finite index subgroup then

$$H_k(\Gamma', \mathbb{Z}) \cong H_k(\Gamma, \Gamma \otimes_{\mathbb{Z}\Gamma'} \mathbb{Z}) \cong H_k(\Gamma, \mathbb{Z}[\Gamma/\Gamma']),$$

[6] In fact, this is the reason why we preferred the wFP_n-condition over the (in some ways more natural) FHT_n-conditions in the statement of the weak n-$(n+1)$-$(n+2)$ Theorem.

where $\mathbb{Z}[\Gamma/\Gamma']$ is the permutation module, which is a free abelian group on the (finite) basis Γ/Γ', with Γ acting on the basis in the obvious way.

It should also be mentioned, though, that we do not know of any examples of groups which are of type wFP_n but not of type FHT_n (see also [LS06, Remark 2]).

6.3 The Bieri-Eckmann Criterion and the FP-n-$(n+1)$-$(n+2)$ Conjecture

When passing from the topological finiteness properties of Chapters 4 and 5 to the homological finiteness properties in the present chapter, we implied that there is a choice to be made: The homology of a group is more readily computable (thanks to tools from homological algebra such as spectral sequences) but it comes with a loss of power, since the stronger F_n and FP_n properties cannot be deduced from finite generation of homology groups alone.

In fact, this is something of a false dichotomy, as we will see in this section and the next chapter. While it is true that the properties wFP_n (or even FHT_n from the end of last section) are strictly weaker than FP_n, there is in fact a way of expressing the FP_n properties in terms of group homology:

The Bieri-Eckmann Criterion. *Let Γ be a group. The following are equivalent:*

1. *Γ is of type FP_n.*
2. *$H_k(\Gamma, -)$ commutes with direct products for $0 \le k < n$.*
3. *Γ is finitely generated and $H_k(\Gamma, \prod \mathbb{Z}\Gamma) = 0$ for $1 \le k < n$, where the product is over some countably infinite index set.*

A criterion of this form first appeared in [BE74, Proposition 3.2], see also [Bie81, Section 2]. We explain the precise meaning of the second condition and give a proof of the Bieri-Eckmann Criterion in the Appendix (Theorem A.39).

Of course, the groups $H_k(\Gamma, \prod \mathbb{Z}\Gamma)$ are much harder to understand than integral homology groups (or even, say, the coefficient modules which are finitely generated as abelian groups that we considered at the end of the previous section). But some of the advantages from before remain. We can still apply spectral sequence techniques, for example. This makes it possible to give some surprisingly simple — but, it has to be said, not necessarily elucidating — proofs of a number of results on the FP_n-properties. For example, we will give a short proof below that an extension of two groups of type FP_n is itself FP_n.

In this section we will show how the Bieri-Eckmann characterisation can be used to give analogues of some of the results from the previous section.

This will allow us to prove a kind of "converse" to the n-$(n+1)$-$(n+2)$ Conjecture (similar to what Proposition 3.11 accomplished vis-à-vis the 0-1-2 Lemma) and specify a concrete Bieri-Eckmann type obstruction to the Conjecture holding in general.

The FP-n-$(n+1)$-$(n+2)$ Conjecture There is an obvious analogue of the n-$(n+1)$-$(n+2)$ Conjecture, replacing all the F_n-properties by corresponding FP_n-properties. We will call this the FP-n-$(n+1)$-$(n+2)$ Conjecture.

THE FP-n-$(n+1)$-$(n+2)$ CONJECTURE. *Let*

$$N_1 \hookrightarrow \Gamma_1 \twoheadrightarrow Q$$
$$N_2 \hookrightarrow \Gamma_2 \twoheadrightarrow Q$$

be short exact sequences such that N_1 is of type FP_n, both Γ_1 and Γ_2 are of type FP_{n+1} and Q is of type FP_{n+2}. Then the associated fibre product is of type FP_{n+1}.

Note that the FP-n-$(n+1)$-$(n+2)$ Conjecture, compared to the original n-$(n+1)$-$(n+2)$ Conjecture has both weaker assumptions and a weaker conclusion, so one does not purely formally imply the other. But it turns out that, in fact, the FP-version of the conjecture is stronger than the F-version. We can show this using the following partial converse to Proposition A.30 (which we will prove in the next chapter):

THEOREM 7.2. *If a group is finitely presented and of type FP_n then it is of type F_n.* □

Now consider the situation in the n-$(n+1)$-$(n+2)$ Conjecture for some $n \geq 2$, i.e.

$$N_1 \hookrightarrow \Gamma_1 \twoheadrightarrow Q$$
$$N_2 \hookrightarrow \Gamma_2 \twoheadrightarrow Q$$

are short exact sequences with N_1 of type F_n, both Γ_1 and Γ_2 of type F_{n+1} and Q of type F_{n+2}. Then the fibre product P is certainly finitely presented, since it fits in an exact sequence

$$N_1 \hookrightarrow P \twoheadrightarrow \Gamma_2$$

with finitely presented kernel and quotient. If the FP-n-$(n+1)$-$(n+2)$ Conjecture held, then P would also be of type FP_{n+1} and thus of type F_{n+1} by Theorem 7.2.

Thus the FP-n-$(n+1)$-$(n+2)$ Conjecture implies the n-$(n+1)$-$(n+2)$ Conjecture.

So we could just try to prove the FP-version of the conjecture instead, and the next section and chapter can be regarded as attempts on this version of the conjecture. One word of warning is in order though:

PROPOSITION 6.13. *The* FP-0-1-2 *Conjecture is false.*

Proof. This follows from the existence of groups which are of type FP_2 but not finitely presented, which was proved by Bestvina and Brady in [BB97]. Let Q be such a group and pick an epimorphism $\pi : F \twoheadrightarrow Q$. Let $P \leq F \times F$ be the symmetric fibre product associated to π. But by Proposition 3.11 (the "converse" to the 0-1-2 Lemma), Q not being finitely presented implies that P is not finitely generated, that is, not of type FP_1. □

FP_n-criteria for short exact sequences We will start by proving an FP_n-analogue of Lemma 6.5. As mentioned above, it is not hard to show, using the Bieri-Eckmann Criterion, that an extension of groups of type FP_n is of type FP_n. We will show something slightly stronger:

Lemma 6.14. *Let $N \hookrightarrow \Gamma \twoheadrightarrow Q$ be a short exact sequence with Q of type FP_n and N of type FP_{n-1} (where $n \geq 2$). Also, assume that $H_0(Q, H_{n-1}(N, \prod \mathbb{Z}\Gamma)) = 0$. Then Γ is of type FP_n.*

Remark 6.15. In particular, the condition on the coinvariants is certainly satisfied if N is of type FP_n, since then $H_{n-1}(N, \prod \mathbb{Z}\Gamma)$ is trivial. The case $n = 1$ has been excluded for convenience. In this case the given condition would have to be replaced with the requirement that the natural map

$$H_0(Q, H_{n-1}(N, \prod \mathbb{Z}\Gamma)) \to \prod H_0(Q, H_0(N, \mathbb{Z}\Gamma)) \cong \prod \mathbb{Z}$$

be onto.

Proof. First note that Γ is certainly finitely generated, since N and Q are. So we need to prove that $H_k(\Gamma, \prod \mathbb{Z}\Gamma) = 0$ for $1 \leq k \leq n-1$ (where the product here and in what follows is always taken over some fixed countable index set).

Consider the LHS spectral sequence

$$E^2_{pq} = H_p(Q, H_q(N, \prod \mathbb{Z}\Gamma)) \Rightarrow H_{p+q}(\Gamma, \prod \mathbb{Z}\Gamma).$$

Since N is of type FP_{n-1} and using that $\mathbb{Z}\Gamma$ is a free N-module, by Lemma A.38 we have

$$H_q(N, \prod \mathbb{Z}\Gamma) \cong \prod H_q(N, \mathbb{Z}\Gamma) = 0, \quad \text{for } 1 \leq q \leq n-2,$$

and

$$H_0(N, \prod \mathbb{Z}\Gamma) \cong \prod H_0(N, \mathbb{Z}\Gamma) \cong \mathbb{Z}Q.$$

In particular $E^2_{pq} = 0$ for $1 \leq q \leq n-2$, and since Q is of type FP_n,

$$E^2_{p,0} = H_p(Q, H_0(N, \prod \mathbb{Z}\Gamma)) \cong H_p(Q, \prod \mathbb{Z}Q) = 0, \quad \text{for } 1 \leq p \leq n-1.$$

Finally, the condition from the statement just says that $E^2_{0,n-1} = 0$. So to sum up, we have $E^2_{pq} = 0$ whenever $1 \leq p+q \leq n-1$. But then $E^\infty_{pq} = 0$ also, for $1 \leq p+q \leq n-1$. This implies that

$$H_k(\Gamma, \prod \mathbb{Z}\Gamma) = 0, \quad \text{for } 1 \leq k \leq n-1,$$

so Γ is of type FP_n. □

The next lemma should be compared to Lemma 6.6.

Lemma 6.16. *Let $N \hookrightarrow \Gamma \twoheadrightarrow Q$ be a short exact sequence with N of type FP_{n-1} and Γ of type FP_n for some $n \geq 2$. Then Q is of type FP_n and there is a surjective homomorphism*

$$H_n(Q, \prod \mathbb{Z}Q) \twoheadrightarrow H_0(Q, H_{n-1}(N, \prod \mathbb{Z}\Gamma))$$

where the products are over some countably infinite index set.
If Γ is of type FP_{n+1} then this map is an isomorphism.

Remark 6.17. In particular, if Q above is of type FP_{n+1}, then $H_n(Q, \prod \mathbb{Z}Q) = 0$, so $H_0(Q, H_{n-1}(N, \prod \mathbb{Z}\Gamma))$ must vanish as well. In this sense the first part can be seen as a kind of converse to Lemma 6.14.

Proof. Consider once again the LHS spectral sequence associated to the given short exact sequence,

$$E^2_{p,q} = H_p(Q, H_q(N, \prod \mathbb{Z}\Gamma)) \Rightarrow H_{p+q}(\Gamma, \prod \mathbb{Z}G).$$

As in the previous proof, the assumption on N implies that $E^2_{p,q} = 0$ for $1 \leq q \leq n-2$ and $E^2_{p,0} \cong H_p(Q, \prod \mathbb{Z}Q)$.

On the E^r-page of the spectral sequence there is a boundary map

$$d_r : E^r_{r,n-r} \to E^r_{0,n-1}$$

and $E^{r+1}_{0,n-1}$ is the cokernel of this map. But for $2 \leq r \leq n-1$ the domain of this map is trivial. It follows that $E^n_{0,n-1} \cong E^2_{0,n-1}$.

Now consider the groups along the bottom edge of the spectral sequence. There is a boundary map on the E^r-page

$$d_r : E^r_{m,0} \to E_{m-r,r-1}$$

and $E^{r+1}_{m,0}$ is the kernel of this map. But for $2 \leq r \leq n-1$ the codomain of this map is trivial. Therefore $E^n_{m,0} \cong E^2_{m,0}$. For $m \leq n-1$, we have $E^n_{m,0} = E^\infty_{m,0}$. But since Γ is of type FP_n, we have $H_m(\Gamma, \prod \mathbb{Z}\Gamma) = 0$ for $1 \leq m \leq n-1$, so all $E^\infty_{p,q}$ with $1 \leq p+q \leq n-1$ must vanish. Thus $H_m(Q, \prod \mathbb{Z}Q) = E^2_{m,0} = 0$ for $1 \leq m \leq n-1$ and so Q is of type FP_n.

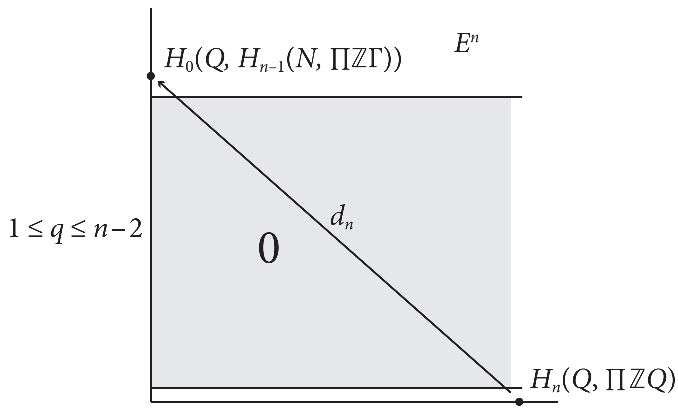

Finally, the boundary map on the E^n-page,

$$d_n : E^n_{n,0} \to E^n_{0,n-1},$$

under the isomorphisms from before becomes a map

$$d_n : H_n(Q, \prod \mathbb{Z}Q) \to H_0(Q, H_{n-1}(N, \prod \mathbb{Z}\Gamma)).$$

Now the cokernel of this map is $E^{n+1}_{0,n-1} = E^{\infty}_{0,n-1}$. But if Γ is of type FP_n, then $H_{n-1}(\Gamma, \prod \mathbb{Z}\Gamma) = 0$, so all the entries $E^{\infty}_{p,q}$ on the diagonal $p + q = n - 1$ must vanish. So in this case the map d_n above has trivial cokernel and hence is surjective.

Similarly, $\ker d_n \cong E^{n+1}_{n,0} = E^{\infty}_{n,0}$ must vanish if Γ is to be of type FP_{n+1}. □

Using these lemmas we can analyse the situation of the n-$(n{+}1)$-$(n{+}2)$ Theorem in much the same way as in the previous section.

Let

$$N_1 \hookrightarrow \Gamma_1 \overset{\pi_1}{\twoheadrightarrow} Q$$
$$N_2 \hookrightarrow \Gamma_2 \overset{\pi_2}{\twoheadrightarrow} Q$$

be short exact sequences and assume that N_1 is of type FP_n, that Γ_1 and Γ_2 are of type FP_{n+1} and that Q is of type FP_{n+2}. Let P denote the associated fibre product and consider the short exact sequence

$$N_1 \hookrightarrow P \overset{p_2}{\twoheadrightarrow} \Gamma_2.$$

To establish the FP-n-$(n+1)$-$(n+2)$ Conjecture we would have to show that P is of type FP_{n+1} and by Lemma 6.14 it would suffice to prove that

$$H_0(\Gamma_2, H_n(N_1, \textstyle\prod \mathbb{Z}P)) = 0. \tag{6.2}$$

Now consider the short exact sequence

$$N_1 \hookrightarrow \Gamma_1 \overset{\pi_1}{\twoheadrightarrow} Q.$$

From Lemma 6.16 it follows that

$$H_0(Q, H_n(N_1, \textstyle\prod \mathbb{Z}\Gamma_1)) = 0. \tag{6.3}$$

So the question becomes how the two homology groups (6.3) and (6.2) are related. We will show that there is a map

$$\phi : H_0(\Gamma_2, H_n(N_1, \textstyle\prod \mathbb{Z}P)) \to H_0(Q, H_n(N_1, \textstyle\prod \mathbb{Z}\Gamma_1)).$$

It is not hard to describe how this map should be defined: The projection $p_1 : P \to \Gamma_1$ induces a map $f : \prod \mathbb{Z}P \to \prod \mathbb{Z}\Gamma_1$; the pair $(\mathrm{id}_{N_1}, f) : (N_1, \prod \mathbb{Z}P) \to (\mathrm{id}_{N_1}, \prod \mathbb{Z}\Gamma_1)$ induces a map g on the n^{th} homology groups and this map factors through to the coinvariants to give ϕ. Of course, we will have to do some work to show that this is a valid definition. Granting this for the moment, the following is now an easy observation:

Proposition 6.18. *In the situation of the* FP-n-$(n+1)$-$(n+2)$ *Conjecture, let* ϕ *be the map defined above. If this map is injective, then* P *is of type* FP_{n+1}.

Proof. We showed that (6.3) holds, and if ϕ is injective this implies that (6.2) holds, which implies that P is of type FP_{n+1}. □

Unfortunately, we do not know whether ϕ is injective. We will say a little more about the sense in which the kernel of ϕ is an obstruction to the n-$(n+1)$-$(n+2)$ Conjecture at the end of the section.

Here, instead, we will study the map ϕ under different conditions to show a "converse" to the FP-n-$(n+1)$-$(n+2)$ Conjecture, a higher-dimensional equivalent to Proposition 3.11. In particular, this shows that the FP_{n+2} condition on the quotient Q is in a certain sense necessary — more on that later.

First of all though, let us show that ϕ can indeed be defined as described. The following construction does not depend on any of the finiteness assumptions made earlier.

Let $f_0 : \mathbb{Z}P \to \mathbb{Z}\Gamma_1$ be the homomorphism of free abelian groups induced by the homomorphism $p_1 : P \to \Gamma_1$. We may consider f_0 as a map of P-modules, with P acting on $\mathbb{Z}\Gamma_1$ via p_1. Now let $f : \prod \mathbb{Z}P \to \prod \mathbb{Z}\Gamma_1$ be the product map $f = \prod f_0$, which maps each component by f_0. Denoting the index set of the product by Λ, we have

$$f(ha)(\lambda) = f_0(ha(\lambda)) = hf_0(a(\lambda)) = hf(a)(\lambda), \quad \text{for } h \in N_1,\ a \in \mathbb{Z}P \text{ and } \lambda \in \Lambda.$$

So f can be considered an N_1-module homomorphism. In other words, the identity homomorphism on N_1 is compatible with f in the sense of Section 6.1. Thus we obtain an induced map on homology

$$\psi_* : H_*(N_1, \prod \mathbb{Z}P) \to H_*(N_1, \prod \mathbb{Z}\Gamma_1).$$

Now $H_*(N_1, \prod \mathbb{Z}P)$ is a Γ_2-module and $H_*(N_1, \prod \mathbb{Z}\Gamma_1)$ is a Q-module and we would like to show that ψ_n induces a map ϕ on the coinvariants.

$$\begin{tikzcd} H_n(N_1, \prod \mathbb{Z}P) \arrow[r, "\psi_n"] \arrow[d, two heads] & H_n(N_1, \prod \mathbb{Z}\Gamma_1) \arrow[d, two heads] \\ H_n(N_1, \prod \mathbb{Z}P)_{\Gamma_2} \arrow[r, "\phi"] & H_n(N_1, \prod \mathbb{Z}\Gamma_1)_Q \end{tikzcd}$$

To achieve this we show that ψ_* is a map of Γ_2-modules, if $H_*(N_1, \prod \mathbb{Z}\Gamma_1)$ is considered a Γ_2-module via the map $\pi_2 : \Gamma_2 \to Q$.

So let us take a look at the respective actions of Γ_2 and Q on these groups. Let

$$\mathbf{F} \twoheadrightarrow \mathbb{Z} : \quad \cdots \to F_2 \to F_1 \to F_0 \twoheadrightarrow \mathbb{Z}$$

be a free resolution of $\mathbb{Z}$ by $\mathbb{Z}P$-modules and let

$$\mathbf{F}' \twoheadrightarrow \mathbb{Z} : \quad \cdots \to F_2' \to F_1' \to F_0' \twoheadrightarrow \mathbb{Z}$$

be a free resolution of $\mathbb{Z}$ by $\mathbb{Z}\Gamma_1$-modules. Via the map $p_1 : P \to \Gamma_1$ we can consider both of these as chain complexes of P-modules. Then there is a chain map $g_* : \mathbf{F} \to \mathbf{F}'$ extending the identity map on $\mathbb{Z}$. This means that $g_k(\gamma a) = p_1(\gamma)g_k(a)$ for all

$\gamma \in P$ and $a \in F_k$ and so, as above, this chain map is compatible with id_{N_1}. Then ψ_* is the map induced on homology by the chain map $g_* \otimes f$:

$$\begin{tikzcd}
\mathbf{F} \otimes_{N_1} \prod \mathbb{Z}P : \arrow[d, "g_*\otimes f"] & \cdots \arrow[r] & F_2 \otimes_{N_1} \prod \mathbb{Z}P \arrow[r] \arrow[d, "g_2\otimes f"] & F_1 \otimes_{N_1} \prod \mathbb{Z}P \arrow[r] \arrow[d, "g_1\otimes f"] & F_0 \otimes_{N_1} \prod \mathbb{Z}P \arrow[d, "g_0\otimes f"] \\
\mathbf{F}' \otimes_{N_1} \prod \mathbb{Z}\Gamma_1 : & \cdots \arrow[r] & F_2' \otimes_{N_1} \prod \mathbb{Z}\Gamma_1 \arrow[r] & F_1' \otimes_{N_1} \prod \mathbb{Z}\Gamma_1 \arrow[r] & F_0' \otimes_{N_1} \prod \mathbb{Z}\Gamma_1
\end{tikzcd}$$

Now recall how Γ_2 acts on $H_*(N_1, \prod \mathbb{Z}P) = H_*(\mathbf{F} \otimes \prod \mathbb{Z}P)$: The group P acts diagonally on the chain groups $F_k \otimes \prod \mathbb{Z}P$ and the induced action on homology factors through $P \to \Gamma_2$. So to describe the action of $\gamma_2 \in \Gamma_2$, first lift it to some $\gamma := (\gamma_1, \gamma_2) \in P$. Then, by definition, for any basic tensor $a \otimes b$ with $a \in F_k$ and $b \in \prod \mathbb{Z}P$ we have[7]

$$\gamma_2[a \otimes b] = [(\gamma a) \otimes (\gamma b)],$$

and so

$$\psi_*(\gamma_2[a \otimes b]) = [(g_k \otimes f)(\gamma a \otimes \gamma b)] = [(\gamma_1 g_k(a)) \otimes (\gamma_1 f(b))].$$

Now let $q := \pi_1(\gamma_1) = \pi_2(\gamma_2)$. By definition Q acts on the homology of $\mathbf{F}' \otimes \prod \mathbb{Z}\Gamma_1$ by lifting to the diagonal action of Γ_1 on the chain complex. This means that

$$q\psi_*([a \otimes b]) = q[g_k(a) \otimes f(b)] = [(\gamma_1 g_k(a)) \otimes (\gamma_1 f(b))].$$

And this proves that ψ_* is compatible with $\pi_2 : \Gamma_2 \to Q$, so that we obtain an induced map[8]

$$\phi : H_0(\Gamma_2, H_n(N_1, \prod \mathbb{Z}P)) \to H_0(Q, H_n(N_1, \prod \mathbb{Z}\Gamma_1)).$$

|| Lemma 6.19. *If N_1 is of type* FP_n *then ϕ is onto.*

Proof. It is clearly sufficient to show that the map

$$\psi_n : H_n(N_1, \prod \mathbb{Z}P) \to H_n(N_1, \prod \mathbb{Z}\Gamma_1)$$

defined earlier is onto, since ϕ is just an induced map on quotient modules.

To show that ψ_n is onto, it will be advantageous to derive ψ_n from a different chain map than earlier. This time we choose a free resolution of $\mathbb{Z}$ by N_1-modules

$$\mathbf{C} \twoheadrightarrow \mathbb{Z} : \quad \cdots \to C_2 \xrightarrow{\partial_2} C_1 \xrightarrow{\partial_1} C_0 \twoheadrightarrow \mathbb{Z},$$

which is finitely generated up to dimension n.

[7] denoting homology classes by brackets

[8] In fact, this proves more generally that there are maps

$$\phi_{p,q} : H_p(\Gamma_2, H_q(N_1, \prod \mathbb{Z}P)) \to H_p(Q, H_q(N_1, \prod \mathbb{Z}\Gamma_1))$$

for any $p, q \geq 0$.

Let $f : \prod \mathbb{Z}P \to \prod \mathbb{Z}\Gamma_1$ be defined as earlier as the natural homomorphism of P-modules induced by the projection $p_1 : P \to \Gamma_1$. Then ψ_* is the map induced on homology by the chain map

$$\mathrm{id} \otimes f : \mathbf{C} \otimes_{N_1} \prod \mathbb{Z}P \to \mathbf{C} \otimes_{N_1} \mathbb{Z}\Gamma_1.$$

To understand the map on n-dimensional homology, we need to understand, first of all, the kernels of the n^{th} boundary map of these two chain complexes.

So let us investigate these a little closer. For $k \leq n$ the natural map

$$C_k \otimes_{N_1} \prod \mathbb{Z}P \to \prod (C_k \otimes_{N_1} \mathbb{Z}P)$$

is an isomorphism by Lemma A.37 since C_k is a finitely generated free N_1-module. So if we consider the chain complex $\mathbf{C} \otimes_{N_1} \prod \mathbb{Z}P$ truncated at dimension n, it is of the form

$$\prod (C_n \otimes_{N_1} \mathbb{Z}P) \to \prod (C_{n-1} \otimes_{N_1} \mathbb{Z}P) \to \cdots$$

The boundary maps in this chain complex are simply products of the maps $\partial_k \otimes \mathrm{id}_{\mathbb{Z}P}$. Now let $Z_n := \ker \partial_n$, so there is an exact sequence

$$0 \to Z_n \to C_n \xrightarrow{\partial_n} C_{n-1}.$$

Note that $\mathbb{Z}P$ is free as an N_1-module and thus flat, so the functor $- \otimes_{N_1} \mathbb{Z}P$ is exact. Therefore

$$0 \to Z_n \otimes_{N_1} \mathbb{Z}P \to C_n \otimes_{N_1} \mathbb{Z}P \xrightarrow{\partial_n \otimes \mathrm{id}} C_{n-1} \otimes_{N_1} \mathbb{Z}P$$

is exact and we have

$$\ker(\partial_n \otimes \mathrm{id}) = Z_n \otimes_{N_1} \mathbb{Z}P,$$

and thus

$$\ker(\partial_n \otimes \mathrm{id}_{\prod \mathbb{Z}P}) = \prod (Z_n \otimes_{N_1} \mathbb{Z}P),$$

by our observation on the boundary maps above. The same argument, with $\mathbb{Z}\Gamma_1$ in place of $\mathbb{Z}P$, shows that

$$\ker(\partial_n \otimes \mathrm{id}_{\prod \mathbb{Z}\Gamma_1}) = \prod (Z_n \otimes_{N_1} \mathbb{Z}\Gamma_1).$$

Now consider the map $\mathrm{id} \otimes f$ on the truncated chain complexes. Under the isomorphisms $C_k \otimes \prod \mathbb{Z}P \cong \prod (C_k \otimes \mathbb{Z}P)$ and $C_k \otimes \prod \mathbb{Z}\Gamma_1 \cong \prod (C_k \otimes \mathbb{Z}\Gamma_1)$ this becomes a map

$$\begin{tikzcd}
\prod (C_n \otimes_{N_1} \mathbb{Z}P) \arrow[r, "\prod \partial_n \otimes \mathrm{id}"] \arrow[d, "\eta_n"] & \prod (C_{n-1} \otimes_{N_1} \mathbb{Z}P) \arrow[r] \arrow[d, "\eta_{n-1}"] & \cdots \\
\prod (C_n \otimes_{N_1} \mathbb{Z}\Gamma_1) \arrow[r, "\prod \partial_n \otimes \mathrm{id}"] & \prod (C_{n-1} \otimes_{N_1} \mathbb{Z}\Gamma_1) \arrow[r] & \cdots
\end{tikzcd}$$

where η_k is just the product of the maps $\mathrm{id}_{C_k} \otimes f_0 : C_k \otimes \mathbb{Z}P \to C_k \otimes \mathbb{Z}\Gamma_1$, with $f_0 : \mathbb{Z}P \to \mathbb{Z}\Gamma_1$ induced by the projection $p_1 : P \to \Gamma_1$ as before. The map ψ_n :

$H_n(N_1, \prod \mathbb{Z}P) \to H_n(N_1, \prod \mathbb{Z}\Gamma_1)$ on homology is induced by the restriction of η_n to the kernels of the left-most boundary maps.

$$\begin{tikzcd}
\ker(\prod \partial_n \otimes \mathrm{id}_{\mathbb{Z}P}) \arrow[r, two heads] \arrow[d, "\eta_n|"] & H_n(N_1, \prod \mathbb{Z}P) \arrow[d, "\psi_n"] \\
\ker(\prod \partial_n \otimes \mathrm{id}_{\mathbb{Z}\Gamma_1}) \arrow[r, two heads] & H_n(N_1, \prod \mathbb{Z}\Gamma_1)
\end{tikzcd}$$

From our analysis of the kernels above, we know that this restriction is just the map

$$\prod(\mathrm{id}_{Z_n} \otimes f_0) : \prod(Z_n \otimes_{N_1} \mathbb{Z}P) \to \prod(Z_n \otimes_{N_1} \mathbb{Z}\Gamma_1).$$

Now f_0 is clearly surjective and then so is $\mathrm{id}_{Z_n} \otimes f_0$, since tensor products are right exact. Hence the product map is surjective as well and it follows that ψ_n is surjective, finishing the proof. □

Now we can easily derive the desired "converse" of the FP-n-$(n+1)$-$(n+2)$ Conjecture.

Theorem 6.20. *Let*

$$N_1 \hookrightarrow \Gamma_1 \overset{\pi_1}{\twoheadrightarrow} Q$$
$$N_2 \hookrightarrow \Gamma_2 \overset{\pi_2}{\twoheadrightarrow} Q$$

be short exact sequences with N_1 of type FP_n and both Γ_1 and Γ_2 of type FP_{n+2} (for some $n \geq 1$). If the associated fibre product is of type FP_{n+1} then Q is of type FP_{n+2}.

Proof. This time, applying Lemma 6.16 to the short exact sequence

$$N_1 \hookrightarrow P \overset{p_2}{\twoheadrightarrow} \Gamma_2,$$

we may conclude that

$$H_0(\Gamma_2, H_n(N_1, \prod \mathbb{Z}P)) = 0.$$

But from Lemma 6.19, we know that the natural map

$$\phi : H_0(\Gamma_2, H_n(N_1, \prod \mathbb{Z}P)) \to H_0(Q, H_n(N_1, \prod \mathbb{Z}\Gamma_1))$$

is surjective, hence

$$H_0(Q, H_n(N_1, \prod \mathbb{Z}\Gamma_1)) = 0.$$

Applying the Lemma 6.16 again, to the short exact sequence $N_1 \hookrightarrow \Gamma_1 \twoheadrightarrow Q$ this time, we obtain that Q is of type FP_{n+1} and that there is an isomorphism

$$H_{n+1}(Q, \prod \mathbb{Z}Q) \cong H_0(Q, H_n(N_1, \prod \mathbb{Z}\Gamma_1)).$$

But then $H_{n+1}(Q, \prod \mathbb{Z}Q) = 0$, so Q is of type FP_{n+2}. □

The necessity of the FP_{n+2} condition It would be very tempting to conclude from the above theorem that the condition in the FP-n-$(n+1)$-$(n+2)$ Conjecture that Q be of type FP_{n+2} cannot be weakened to Q being of type FP_{n+1}. In view of Theorem 6.20, all we need is to find a short exact sequence

$$\overset{\mathrm{FP}_n}{N} \hookrightarrow \overset{\mathrm{FP}_{n+2}}{\Gamma} \xrightarrow{\pi} \overset{\text{not }\mathrm{FP}_{n+2}}{Q} \tag{6.4}$$

with the indicated finiteness properties. Theorem 6.20 allows us to conclude that the symmetric fibre product associated to π is not of type FP_{n+1}.

For $n = 1$ such sequences can be found using a construction due to Rips [Rip82]. Rips' construction supplies to any finitely presented group Q a short exact sequence $N \hookrightarrow \Gamma \twoheadrightarrow Q$ with Γ a $C'(1/6)$ small-cancellation group (and in particular, finitely presented) and N finitely generated. To obtain a sequence of the form (6.4), we only need to pick a group Q which is finitely presented and not of type FP_3. There are many examples of such groups. The first were found by Stallings in [Sta63] and in general the construction by Bestvina and Brady in [BB97] provides an ample source of groups of this type.

We can thus record:

Proposition 6.21. *In the* FP-1-2-3 *Conjecture, the condition that Q be of type* FP_3 *cannot be replaced by the weaker assumption that Q be of type* FP_2. □

Unfortunately, we do not know of any examples for short exact sequences of the type (6.4) for $n > 1$. Let us rephrase this problem as a question:

Question. If Γ is a group of type FP_{n+2} and N a normal subgroup of type FP_n, does the quotient Γ/N have to be of type FP_{n+2}?

Of course, Lemma 6.16 says that Γ/N has to be of type FP_{n+2} if N is of type FP_{n+1}, so the question asks whether we can weaken this assumption to "type FP_n".

If we try to follow the construction for $n = 1$, we would first look for a Q of type FP_{n+1} but not of type FP_{n+2}. There is no shortage of such groups (for example Bieri's generalisation of Stallings' example in [Bie81] or, again, groups built using the Bestvina-Brady construction). However, Rips' construction (or its variants [Wis03, CBW]) will typically not yield a short exact sequence with kernel of type FP_2. The reason is that $C'(1/6)$ small-cancellation groups are of cohomological dimension 2 and Bieri showed that every normal subgroup of type FP_2 in such a group is either free or of finite index ([Bie76], see also [Bie81, Corollary 8.6]).

Conversely, there are also many examples of groups with normal subgroups of type FP_n but not FP_{n+1}, for example using the theory of Σ-invariants (see Chapter 8). But in those cases the quotients will always be "nice" (i.e. of type F_∞). It seems much harder to construct short exact sequences like (6.4), where we have to arrange for both the kernel *and* quotient to have specific finiteness properties.

So, while there is — at least to our knowledge — no good reason to believe that the Question above should have a positive answer, actual counterexamples, for now, remain elusive.

Further remarks The Bieri-Eckmann criterion can be viewed as identifying an obstruction for a group Γ of type FP_{n-1} to be of type FP_n in the form of the homology group $H_{n-1}(\Gamma, \prod \mathbb{Z}\Gamma)$. The group Γ is of type FP_n if and only if this homology group vanishes.

In that sense, we can view Lemmas 6.14 and 6.16 as identifying a precise obstruction for the middle group Γ in a short exact sequence $N \hookrightarrow \Gamma \twoheadrightarrow Q$ to be of type FP_{n+1}, given that N is of type FP_n and Q is of type FP_{n+1}. Namely, Γ is of type FP_{n+1} if and only if the homology group

$$H_0(Q, H_n(N, \prod \mathbb{Z}\Gamma))$$

vanishes (though note that the precise statements of the two lemmas extract some more subtle information from this homology group).

Now consider the map

$$\phi : H_0(\Gamma_2, H_n(N_1, \prod \mathbb{Z}P)) \to H_0(Q, H_n(N_1, \prod \mathbb{Z}\Gamma_1))$$

that we associated to a pair of exact sequences

$$\begin{tikzcd}
N \arrow[r, hook] \arrow[d, equal] & P \arrow[r, two heads, "p_2"] \arrow[d, two heads, "p_1"] & \Gamma_2 \arrow[d, two heads, "\pi_2"] \\
N \arrow[r, hook] & \Gamma_1 \arrow[r, two heads, "\pi_1"] & Q
\end{tikzcd}$$

We may then think of this as a map between the groups of obstructions from above. The kernel and cokernel of ϕ can then be regarded (roughly) as obstructions to transferring finiteness properties from one sequence to the other.

We made this notion precise in two ways earlier: Firstly, we showed that the cokernel of ϕ vanishes if N_1 is of type FP_n and used this in Theorem 6.20 to transfer a finiteness condition on P to one on Q.

Secondly (and maybe even more importantly), we identified the kernel of ϕ as an obstruction to the FP-n-$(n+1)$-$(n+2)$ Conjecture. We should explain what we mean by this exactly. As we have seen, under the finiteness conditions of the FP-n-$(n+1)$-$(n+2)$ Conjecture, the codomain of this map, the homology group $H_0(Q, H_n(N_1, \prod \mathbb{Z}\Gamma_1))$, in fact vanishes. So the kernel is the really the entire homology group $H_0(\Gamma_2, H_n(N_1, \prod \mathbb{Z}P))$. So this group is also an obstruction to P being of type FP_{n+1}, as is, most obviously, the group $H_n(P, \prod \mathbb{Z}P)$. But the point is that ϕ has a chance of being injective under much more general conditions than just those of the FP-n-$(n+1)$-$(n+2)$ Conjecture.

To conclude that P is of type FP_{n+1} from the statement that ϕ is injective, we used all of the finiteness conditions in the Conjecture in an essential way. So it is feasible that we may not actually need them to show that ϕ is injective. Compare this to the situation in Theorem 6.20: We used that ϕ is onto to establish this theorem, but in fact we were able to show that ϕ is onto under much more general conditions than those of Theorem 6.20 (namely whenever N_1 is of type FP_n, with no additional conditions on Γ_1, Γ_2 or Q). In a similar way, the question whether (or when) ϕ is one-to-one can really be studied independently of the rest of the Conjecture. It is in this sense that we call the kernel of ϕ an obstruction to the n-$(n+1)$-$(n+2)$ Conjecture.

Of course, the map and the groups involved are subtle objects to work with, as the proof of Lemma 6.19 may indicate. But the results we obtained demonstrate that improving our understanding of the map ϕ (possibly under additional assumptions or in special cases) would go a long way towards proving or disproving the FP-n-$(n+1)$-$(n+2)$ Conjecture.

7

Resolutions for group extensions

In Chapter 5, we approached the n-$(n+1)$-$(n+2)$ Conjecture from a topological perspective, analysing classifying complexes and their homotopy groups. Then, in Chapter 6, we abandoned homotopy in favour of homology groups. Both approaches have their upsides: Working with homology, we had at our disposal powerful algebraic tools such as the LHS spectral sequence, which allowed us to establish homological finiteness properties for a fibre product under satisfyingly general conditions in the Weak n-$(n+1)$-$(n+2)$ Theorem. On the flipside, the integral homology groups do not contain enough information to establish the stronger finiteness properties from the original n-$(n+1)$-$(n+2)$ Conjecture. The classifying complexes potentially do, but our tools from Chapter 5 seem too blunt to tease out the desired results.

The approach via the Bieri-Eckmann Criterion from Section 6.3 combines some of the advantages from both worlds, as we have seen. Its major disadvantage is its non-constructiveness: We can, in principle, use the criterion to conclude that some particular partial free resolution exists, but the Bieri-Eckmann approach gives little idea what such a resolution might look like. On the one hand this can lead to elegant quick proofs of facts that might be hard to prove in a constructive way. On the other hand, this makes it much harder to apply the geometric intuition gained from the proof of the 1-2-3 Theorem or from the work in Chapter 5. While resolutions relate in a rather straightforward way to geometric notions, the same cannot be said for the homology groups $H_n(\Gamma, \prod \mathbb{Z}\Gamma)$. By "hiding" the low-level structure of chain groups and boundary maps to some degree, the approach from Section 6.3 also makes explicit computation harder.

In this chapter we will actually construct and analyse resolutions. This combines some of the advantages of both of our previous approaches: Resolutions are algebraic objects and as such they retain some of the advantages of computability. As a result we will be able to analyse them in much finer detail than we could the stacks from Chapter 5. At the same time a resolution contains much more data than merely its homology groups. In a certain sense, which will be made precise shortly (Proposition 7.1), a classifying complex can actually be recovered from the cellular chain complex of its universal cover, so in this sense we sacrifice very little information compared to working with the classifying complexes themselves.

As a result, we will be able to derive higher-dimensional analogues of many of the more computational steps involved in the proof we gave of the 1-2-3 Theorem in Chapter 4, especially those of Sections 4.4–4.6. So in a way, this method picks up where the topological approach from Chapter 5 left off.

7.1 Building complexes from resolutions

Resolutions for a group Γ are commonly obtained as the cellular chain complexes of some contractible free Γ-complex, say the universal cover of a $B\Gamma$-complex (see the Appendix A.2). The following proposition, conversely, allows us to convert a resolution back into a complex whose universal cover has this resolution as its cellular chain complex. However, this does not work in the lowest dimensions, so we have to assume the 2-skeleton of the desired complex as given a priori.

The technique used here to construct these complexes is essentially due to Wall [Wal65, Wal66] (who credits Milnor with having invented it). There are several instances in the literature where this technique is used to prove results very similar to the one below (see e.g. [Bro94, Section VIII.7]), but we do not know of a reference for this particular statement, which suits our needs best.

Proposition 7.1. *Let X be a 2-complex with fundamental group Γ. Let*

$$C_2 \xrightarrow{\partial_2} C_1 \xrightarrow{\partial_1} C_0 \twoheadrightarrow \mathbb{Z}$$

be the augmented cellular chain complex of the universal cover of X. Assume that there is a partial free resolution $\mathbf{C}' \twoheadrightarrow \mathbb{Z}$ of dimension n and isomorphisms $\phi_k : C'_k \to C_k$ for $0 \le k \le 2$, so that we have a commutative diagram

$$\begin{tikzcd}
C'_n \arrow[r, "\partial'_n"] & \cdots \arrow[r] & C'_3 \arrow[r, "\partial'_3"] & C'_2 \arrow[r, "\partial'_2"] \arrow[d, "\phi_2"] & C'_1 \arrow[r, "\partial'_1"] \arrow[d, "\phi_1"] & C'_0 \arrow[r, two heads] \arrow[d, "\phi_0"] & \mathbb{Z} \\
 & & & C_2 \arrow[r, "\partial_2"] & C_1 \arrow[r, "\partial_1"] & C_0 \arrow[r, two heads] & \mathbb{Z}
\end{tikzcd}$$

Then there is an n-dimensional complex $\overline{X}$ with $\overline{X}^{(2)} = X$ and an isomorphism ϕ_* of chain complexes

$$\begin{tikzcd}
C'_n \arrow[r, "\partial'_n"] \arrow[d, "\phi_n"] & \cdots \arrow[r] & C'_3 \arrow[r, "\partial'_3"] \arrow[d, "\phi_3"] & C'_2 \arrow[r, "\partial'_2"] \arrow[d, "\phi_2"] & C'_1 \arrow[r, "\partial'_1"] \arrow[d, "\phi_1"] & C'_0 \arrow[d, "\phi_0"] \\
C_n \arrow[r, "\partial_n"] & \cdots \arrow[r] & C_3 \arrow[r, "\partial_3"] & C_2 \arrow[r, "\partial_2"] & C_1 \arrow[r, "\partial_1"] & C_0
\end{tikzcd}$$

where (C_*, ∂_*) is the cellular chain complex of the universal cover of $\overline{X}$, extending the given isomorphism in dimensions ≤ 2.
If for each C'_k with $k \ge 3$ a basis is given, then ϕ_* can be chosen so that it sends each generator to the (singular) homology class of a characteristic map of a cell in the universal cover of $\overline{X}$.

Proof. Assume that $\overline{X}^{(m)}$ has been constructed for some m with $2 \le m < n$ and denote its universal cover by $p : Y^{(m)} \twoheadrightarrow \overline{X}^{(m)}$. Furthermore, assume we are given an isomorphism of chain complexes

$$\begin{tikzcd}
C'_{m+1} \arrow[r, "\partial'_{m+1}"] & C'_m \arrow[r, "\partial'_m"] \arrow[d, "\phi_m"] & \cdots \arrow[r] & C'_3 \arrow[r, "\partial'_3"] \arrow[d, "\phi_3"] & C'_2 \arrow[r, "\partial'_2"] \arrow[d, "\phi_2"] & C'_1 \arrow[r, "\partial'_1"] \arrow[d, "\phi_1"] & C'_0 \arrow[d, "\phi_0"] \\
 & C_m \arrow[r, "\partial_m"] & \cdots \arrow[r] & C_3 \arrow[r, "\partial_3"] & C_2 \arrow[r, "\partial_2"] & C_1 \arrow[r, "\partial_1"] & C_0
\end{tikzcd}$$

where (C_k) is the cellular chain complex of $Y^{(m)}$.

Pick a free generating set $\mathcal{X}$ for C'_{m+1}. For $x \in \mathcal{X}$, we have

$$\phi_m \partial'_{m+1}(x) \in \ker \partial_m = H_m(Y^{(m)}).$$

Now $Y^{(m)}$ is a simply connected complex. Since $\mathbf{C}'$ is exact and (ϕ_*) is an isomorphism up to dimension m, the homology groups $H_k(Y^{(m)}) = H_k(\mathbf{C})$ are trivial for $k \le$

$m-1$. So by Hurewicz' Theorem, the Hurewicz map $h_m : \pi_m(Y^{(m)}) \to H_m(Y^{(m)})$ is an isomorphism. Pick for each $x \in \mathcal{X}$ a map $f_x : S^m \to Y^{(m)}$ with

$$h_m([f_x]) = \phi_m \partial'_{m+1}(x),$$

where $[f_x]$ denotes the pointed homotopy class of f_x in $\pi_m(Y^{(m)})$.

Glue into $\overline{X}^{(m)}$, for each $x \in \mathcal{X}$ one $(m+1)$-cell e_x along $p \circ f_x$ to yield $\overline{X}^{(m+1)}$. Let $c_{e_x} : B^{m+1} \to \overline{X}^{(m+1)}$ be a characteristic map for e_x, such that $c_{e_x}|_{S^m} = pf_x$. We can obtain $Y^{(m+1)}$ from $Y^{(m)}$ by gluing in a Γ-orbit of $(m+1)$-cells along all the Γ-translates of f_x for each $x \in \mathcal{X}$. Denote the cell glued in along f_x by $\tilde{e}_x$ and pick a characteristic map $c_{\tilde{e}_x} : B^{m+1} \to Y^{(m+1)}$ with $c_{\tilde{e}_x}|_{S^m} = f_x$. We can then extend p to $Y^{(m+1)}$, by setting $p(gc_{\tilde{e}_x}(a)) = c_{e_x}(a)$ for all $x \in \mathcal{X}$, $g \in \Gamma$ and $a \in B^{m+1}$.

Recall that, by definition of the cellular homology groups

$$C_{m+1} = H^s_{m+1}(Y^{(m+1)}, Y^{(m)}),$$

where H^s denotes singular homology. Pick a fixed generator

$$\omega \in H^s_{m+1}(B^{m+1}, S^m).$$

By basic homology theory, C_{m+1} is freely generated as an abelian group by

$$\{(c_{\tilde{e}})_*(\omega) \mid \tilde{e} \text{ an } (m+1)\text{-cell of } Y^{(m+1)}\},$$

where $(c_{\tilde{e}})_* : H^s_{m+1}(B^{m+1}, S^m) \to C_{m+1}$ is map induced on homology by the characteristic map $c_{\tilde{e}}$. The $(m+1)$-cells of $Y^{(m+1)}$ are of the form $g\tilde{e}_x$ with $x \in \mathcal{X}$ and $g \in \Gamma$, so as a Γ-module C_{m+1} is freely generated by the $(c_{\tilde{e}_x})_*(\omega)$ with $x \in \mathcal{X}$. We can thus define ϕ_{m+1} as the unique Γ-module homomorphism with $\phi_{m+1}(x) = (c_{\tilde{e}_x})_*(\omega)$. By what we said before, this is an isomorphism.

It remains to show that the square

$$\begin{tikzcd} C'_{m+1} \arrow[r, "\partial'_{m+1}"] \arrow[d, "\phi_{m+1}"] & C'_m \arrow[d, "\phi_m"] \\ C_{m+1} \arrow[r, "\partial_{m+1}"] & C_m \end{tikzcd}$$

commutes. First recall that the Hurewicz map $h_m : \pi_m(Y^{(m)}) \to H_m(Y^{(m)})$ is defined by

$$h_m([f]) = f_*(\xi),$$

where $[f]$ is the (pointed) homotopy class of some $f : S^m \to Y^{(m)}$, $f_* : H_m(S^m) \to H_m(Y^{(m)})$ is the induced map on homology and $\xi \in H_m(S^m)$ is some fixed choice of generator. We may take $\xi := \partial\omega$. Then we have

$$\begin{aligned}\partial_{m+1}(\phi_{m+1}(x)) &= (c_{\tilde{e}_x})_*(\partial\omega)\\ &= (c_{\tilde{e}_x}|_{S^m})_*(\xi)\\ &= (f_x)_*(\xi)\\ &= h_m([f_x])\\ &= \phi_m(\partial'_{m+1}(x))\end{aligned}$$

and that finishes the proof. □

As a direct consequence we obtain:

Theorem 7.2. *A group* Γ *is of type* F_n *with* $n \geq 2$ *if and only if it is finitely presented and of type* FP_n.

Proof. We have already proved that every group of type F_n is of type FP_n.

So let Γ be a finitely presented group of type FP_n. Let X be a finite presentation complex for Γ (i.e. a $\mathrm{B}\Gamma^{(2)}$, in the language of Appendix A.1). Then the cellular chain complex

$$C_2 \to C_1 \to C_0 \to \mathbb{Z}$$

of the universal cover of X is a finitely generated partial free resolution of $\mathbb{Z}$ by Γ-modules. By Lemma A.33 it can be extended to a finitely generated partial free resolution $\mathbf{C}' \twoheadrightarrow \mathbb{Z}$ of dimension n:

$$\begin{tikzcd} C'_n \arrow[r] & \cdots \arrow[r] & C'_3 \arrow[r] & C'_2 \arrow[r] \arrow[d, equal] & C'_1 \arrow[r] \arrow[d, equal] & C'_0 \arrow[r, two heads] \arrow[d, equal] & \mathbb{Z} \\ & & & C_2 \arrow[r] & C_1 \arrow[r] & C_0 \arrow[r, two heads] & \mathbb{Z} \end{tikzcd}$$

By Proposition 7.1 there is then an n-dimensional complex $\overline{X}$ with $\overline{X}^{(2)} = X$ and such that the cellular chain complex of its universal cover Y is isomorphic to $\mathbf{C}$. In particular, since the fundamental group depends only on the 2-skeleton, we have $\pi_1(\overline{X}) \cong \Gamma$. Furthermore, since all the C'_m with $0 \leq m \leq n$ are finitely generated Γ-modules, Y is cocompact with respect to the Γ-action, so $\overline{X}$ is finite. Also, Y is a simply connected space with $H_m(Y) = 0$ for $1 \leq m \leq n-1$. By Hurewicz' Theorem, this means that Y is $(n-1)$-connected. Therefore, $\pi_m(\overline{X}) \cong \pi_m(Y) = 0$ for $2 \leq m \leq n-1$. Thus $\overline{X}$ is a finite $\mathrm{B}\Gamma^{(n)}$, so Γ is of type F_n. □

7.2 Constructing resolutions from short exact sequences

In all of our attempts on the n-$(n+1)$-$(n+2)$ Conjecture the core of the argument is a careful analysis of short exact sequences, and this will be the case here as well. Let $N \hookrightarrow \Gamma \twoheadrightarrow Q$ be a short exact sequence of groups. In Chapter 5 we used Geoghegan's stack construction to describe a classifying complex Y for Γ in terms of given classifying complexes X for N and Z for Q. The stack construction provides a rough conceptual view of what Y looks like: It is a kind of "twisted" product, containing an n-cell for every pair of a k-cell in X and an $(n-k)$-cell in Z. However, it supplies little information about the attaching maps of the cells of Y.

In Chapter 4 we described the lower-dimensional cells of this complex in more detail. The presentation of Γ associated to the short exact sequence, that we described there, may be viewed as an explicit description of the attaching maps of the 2-cells of Y. We also described explicitly (as elements of the presentation crossed module) the attaching maps of certain 3-cells of Y, namely those that project to the 2-cells and the 3-cells of Y.

In this chapter we will start off with free resolutions of $\mathbb{Z}$ as an N- or a Q-module respectively, which should be thought of as the cellular chain complexes of the universal covers of X and Z. We will then build a free resolution of $\mathbb{Z}$ as a Γ-module, which will have the same kind of "twisted product" structure as the stacks from Chapter 5 or the lower-dimensional terms from Chapter 4. The crucial point is that we can not only describe the chain groups but also specify the boundary maps in great detail.

Let

$$N \overset{\iota}{\hookrightarrow} \Gamma \overset{\pi}{\twoheadrightarrow} Q$$

be a short exact sequence of groups. Let

$$\mathbf{A} \twoheadrightarrow \mathbb{Z}: \quad \cdots \to A_2 \overset{\partial_2^A}{\to} A_1 \overset{\partial_1^A}{\to} A_0 \overset{\varepsilon_A}{\twoheadrightarrow} \mathbb{Z}$$

be a free resolution of $\mathbb{Z}$ by N-modules and let

$$\mathbf{C} \twoheadrightarrow \mathbb{Z}: \quad \cdots \to C_2 \overset{\partial_2^C}{\to} C_1 \overset{\partial_1^C}{\to} C_0 \overset{\varepsilon_C}{\twoheadrightarrow} \mathbb{Z}$$

be a free resolution of $\mathbb{Z}$ by Q-modules.

Pick for each n a basis X_n for the free N-module A_n and a basis Z_n and for the free Q-module C_n.

Let $\widetilde{Z}_n = \{\tilde{z} \mid z \in Z_n\}$ be a copy of Z_n and let D_n be the free Γ-module with basis $\widetilde{Z}_n$. Now define

$$B_{k,l} := A_k \otimes_{\mathbb{Z}N} D_l.$$

Since D_l is a direct sum $D_l = \bigoplus_{\tilde{z} \in \widetilde{Z}_l} \mathbb{Z}\Gamma$, it can be given the structure of a Γ-bimodule, considering each $\mathbb{Z}\Gamma$-summand as a bimodule in the natural way and letting Γ act diagonally on the sum D_l. We use the left-module structure on D_l to form the tensor product. The right-module structure can then be used to supply $B_{k,l}$ with the structure of a (right) Γ-module, by setting $(a \otimes d)\gamma = a \otimes (d\gamma)$ for $a \in A_k$, $d \in D_l$ and $\gamma \in \Gamma$. As usual, we will consider $B_{k,l}$ a left Γ-module with $\gamma(a \otimes d) = a \otimes (d\gamma^{-1})$.

Note that

$$\begin{aligned} B_{k,l} &\cong \left(\bigoplus_{x \in X_k} \mathbb{Z}N \right) \otimes_{\mathbb{Z}N} \left(\bigoplus_{\tilde{z} \in \tilde{Z}_l} \mathbb{Z}G \right) \\ &\cong \bigoplus_{\substack{x \in X_k \\ \tilde{z} \in \tilde{Z}_l}} (\mathbb{Z}N \otimes_{\mathbb{Z}N} \mathbb{Z}G) \\ &\cong \bigoplus_{\substack{x \in X_k \\ \tilde{z} \in \tilde{Z}_l}} \mathbb{Z}G, \end{aligned}$$

so $B_{k,l}$ is a free Γ-module on the basis $\{x \otimes \tilde{z} \mid x \in X_k, \tilde{z} \in \tilde{Z}_l\}$.[1]

Fix some $l \geq 0$. Tensoring the resolution $\mathbf{A} \twoheadrightarrow \mathbb{Z}$ with D_l gives a chain complex

$$\cdots \to A_n \otimes_N D_l \xrightarrow{\partial_n^A \otimes \mathrm{id}} A_{n-1} \otimes_N D_l \to \cdots \to A_1 \otimes_N D_l \xrightarrow{\partial_1^A \otimes \mathrm{id}} A_0 \otimes_N D_l \xrightarrow{\varepsilon_A \otimes \mathrm{id}} \mathbb{Z} \otimes_N D_l \tag{7.1}$$

By definition $A_k \otimes_N D_l = B_{k,l}$ and we have

$$\mathbb{Z} \otimes_N D_l \cong \mathbb{Z} \otimes_N \left(\bigoplus_{z \in Z_l} \mathbb{Z}G \right) \cong \bigoplus_{z \in Z_l} \mathbb{Z} \otimes_N \mathbb{Z}G \cong \bigoplus_{z \in Z_l} \mathbb{Z}Q \cong C_l$$

Following all these natural identifications, we find that an explicit isomorphism $\mathbb{Z} \otimes_N D_l \to C_l$ is given by $1 \otimes \tilde{z} \mapsto z$ for $z \in Z_l$. Thus we obtain, for each $l \geq 0$, a chain complex

$$\cdots \to B_{n,l} \xrightarrow{d_n^l} B_{n-1,l} \to \cdots \to B_{1,l} \xrightarrow{d_1^l} B_{0,l} \xrightarrow{\pi_l} C_l \tag{7.2}$$

where d_k^l is given by

$$d_k^l(x \otimes \tilde{z}) = \partial_k^A(x) \otimes \tilde{z}, \quad \text{for } x \in X_k, z \in Z_l$$

and

$$\pi_l(x \otimes \tilde{z}) = \varepsilon_A(x) z, \quad \text{for } x \in X_0, z \in Z_l.$$[2]

Furthermore note that, since D_l is a free $\mathbb{Z}N$-module the functor $- \otimes_N D_l$ is exact, so (7.1) and (7.2) are in fact exact sequences, so (7.2) is a free resolution of C_l as a Γ-module.

[1]In fact, we could have used this as the definition of $B_{k,l}$, but the tensor notation will come in handy in the sequel.

[2]Note that we can always pick the resolution $\mathbf{A} \overset{\varepsilon_A}{\twoheadrightarrow} \mathbb{Z}$ such that $\varepsilon_A(x) = 1$ for all $x \in X_0$. If we do this, we have $\pi_l(x \otimes \tilde{z}) = z$. We could go further and assume that $A_0 \cong \mathbb{Z}N$ (since there is always a resolution where this holds) in which case $B_{0,l} \cong D_l$ and π_l can simply be thought of as the Γ-module homomorphism that sends $\tilde{z} \in D_l$ to $z \in C_l$.

We have now obtained a diagram

$$
\begin{tikzcd}
A_3 \arrow[d, "\partial_3^A"] & B_{3,0} \arrow[d, "d_3^0"] & B_{3,1} \arrow[d, "d_3^1"] & B_{3,2} \arrow[d, "d_3^2"] & B_{3,3} \arrow[d, "d_3^3"] \\
A_2 \arrow[d, "\partial_2^A"] & B_{2,0} \arrow[d, "d_2^0"] & B_{2,1} \arrow[d, "d_2^1"] & B_{2,2} \arrow[d, "d_2^2"] & B_{2,3} \arrow[d, "d_2^3"] \\
A_1 \arrow[d, "\partial_1^A"] & B_{1,0} \arrow[d, "d_1^0"] & B_{1,1} \arrow[d, "d_1^1"] & B_{1,2} \arrow[d, "d_1^2"] & B_{1,3} \arrow[d, "d_1^3"] \\
A_0 \arrow[d, "\varepsilon_A"] & B_{0,0} \arrow[d, two heads, dotted, "\pi_0"] & B_{0,1} \arrow[d, two heads, dotted, "\pi_1"] & B_{0,2} \arrow[d, two heads, dotted, "\pi_2"] & B_{0,3} \arrow[d, two heads, dotted, "\pi_3"] \\
\mathbb{Z} & C_0 & C_1 \arrow[l, "\partial_1^C"] & C_2 \arrow[l, "\partial_2^C"] & C_3 \arrow[l, "\partial_3^C"]
\end{tikzcd}
$$

where the columns are resolutions of the C_l as Γ-modules.

We shall now show how to make a resolution of $\mathbb{Z}$ by Γ-modules from the $B_{k,l}$. The theorem we will prove applies to a more general situation.

Theorem 7.3. *Let Γ be a group. Let*

$$(\mathbf{C}, \partial_*^C): \quad \cdots \to C_2 \overset{\partial_2^C}{\to} C_1 \overset{\partial_1^C}{\to} C_0 \overset{\partial_0^C}{\to} 0$$

be a chain complex of Γ-modules and for each $l \geq 0$ let

$$(B_{*,l}, d_*^l) \overset{\pi_l}{\twoheadrightarrow} C_l: \quad \cdots \to B_{2,l} \overset{d_2^l}{\to} B_{1,l} \overset{d_1^l}{\to} B_{0,l} \overset{\pi_l}{\twoheadrightarrow} C_l$$

be a free resolution of C_l by Γ-modules (so we have a diagram as above). Defining

$$B_n := \bigoplus_{k+l=n} B_{k,l},$$

we can extend the augmentation map π_n to B_n by setting $\pi_n(B_{k,l}) = 0$ for $k \geq 1$. Then we can define boundary maps $\partial_n^B : B_n \to B_{n-1}$ for each $n \geq 1$ to give a chain complex

$$(\mathbf{B}, \partial_*^B): \quad \cdots \to B_2 \overset{\partial_2^B}{\to} B_1 \overset{\partial_1^B}{\to} B_0 \overset{\partial_0^B}{\to} 0$$

such that the maps $\pi_n : B_n \to C_n$ induce an isomorphism on homology.

Remark 7.4. Note that in the situation considered before, $\mathbf{C} \twoheadrightarrow \mathbb{Z}$ was actually a resolution of $\mathbb{Z}$ (which is not required in the statement of the theorem), so the homology of the chain complex (C_*, ∂_*^C) is $\mathbb{Z}$ in dimension 0 and trivial in higher dimensions. The theorem then says that the same is true for (B_*, ∂_*^B), so that

$$\cdots \to B_2 \overset{\partial_2^B}{\to} B_1 \overset{\partial_1^B}{\to} B_0 \to B_0 / \operatorname{im} \partial_1^B = \mathbb{Z}$$

is a free resolution of $\mathbb{Z}$ by Γ-modules.

Theorem 7.3 was proved for the case of extensions of groups by Wall in [Wal61], and in full generality by Ken Brown [Bro82, Lemma 1.5]. The theorem as stated above is sufficient to prove, for example, that any extension of groups of type FP_n is itself of type FP_n (or more generally, that a group is of type FP_n if it acts on an acyclic space such that the stabiliser of every k-cell is of type FP_{n-k}). We will, however, need to analyse extensions of groups in finer detail later, for which we will require a more precise description of the maps ∂_n^B. This is provided by the proposition below, which implies the theorem. Brown in his proof of Theorem 7.3 [Bro82] does not give a description of the boundary maps. Wall does prove something similar to the following proposition in [Wal61], but he treats only the case of group extensions, and his description of the boundary maps is rather different from ours (in particular, the description we give is detailed enough to prove the conclusions of Theorem 7.3 directly, without the need to appeal to a spectral sequence argument). Further variations of Wall's construction have been given in [Sta94] and [EW05]. Finally, a related and very explicit construction for dimension 2 can be found in [Mei97].

Proposition 7.5. *Assume we are in the situation of Theorem 7.3. Set*

$$B_n^{(m)} := \bigoplus_{l=0}^{m} B_{n-l,l} \subset B_n.$$

Then we can define boundary maps

$$d_0^l : B_{0,l} \to B_{0,l-1}, \quad \text{for } l \geq 1$$

making the following diagram commute

$$\begin{tikzcd}
B_{0,0} \arrow[d, dotted, two heads, "\pi_0"] & B_{0,1} \arrow[l, "d_0^1"'] \arrow[d, dotted, two heads, "\pi_1"] & B_{0,2} \arrow[l, "d_0^2"'] \arrow[d, dotted, two heads, "\pi_2"] & B_{0,3} \arrow[l, "d_0^3"'] \arrow[d, dotted, two heads, "\pi_3"] & \cdots \arrow[l] \\
C_0 & C_1 \arrow[l, "\partial_1^C"] & C_2 \arrow[l, "\partial_2^C"] & C_3 \arrow[l, "\partial_3^C"] & \cdots \arrow[l]
\end{tikzcd}$$

as well as "correction maps"

$$c_k^l : B_{k,l} \to B_{n-1}, \quad \text{for } k, l \geq 0$$

and the boundary maps

$$\partial_n^B : B_n \to B_{n-1}, \quad \text{for } n \geq 1$$

such that the following hold:

(i) $\partial_n^B|_{B_{k,l}} = d_k^l + c_k^l$ *for* $n \geq 1$, $k + l = n$.

(ii) $\operatorname{im} c_k^l \subset B_{n-1}^{(l-1)}$ *for* $n \geq 1$, $k + l = n$ *and* $k \geq 1$ *and* $\operatorname{im} c_0^n \subset B_{n-1}^{(n-2)}$.

(iii) $\partial_{n-1}^B \partial_n^B = 0$ *for* $n \geq 1$.

If the maps have been defined up to some n, in a way that (i)–(iii) hold, then additionally we have:

(iv) If $x \in B_{n-1}$ *with* $\pi_{n-1}(x) \in \operatorname{im} \partial_n^C$ *then there is a* $y \in B_{0,n}$ *such that* $x - \partial_n^B(y) \in \ker \pi_{n-1}$.

(v) If $x \in B_{n-1}$ *with* $\pi_{n-1}(x) = 0$ *then there is a* $y \in B_{1,n-1}$ *with* $x - \partial_n^B(y) \in B_{n-1}^{(n-2)}$.

(vi) If $x \in B_{n-1}^{(m)}$ *for some* $m < n-1$ *with* $\partial_{n-1}^B(x) \in B_{n-2}^{(m-1)}$ *then there is a* $y \in B_{n-m,m}$ *with* $x - \partial_n^B(y) \in B_{n-1}^{(m-1)}$.

(vii) If $x \in B_{n-1}^{(m)}$ *for some* $m \leq n-1$ *with* $\partial_{n-1}^B(x) = 0$ *and* $\pi_{n-1}(x) = 0$ *then* $x \in \partial_n^B(B_n^{(m)})$.

(viii) If $x \in B_{n-1}$ *with* $\partial_{n-1}^B(x) = 0$ *and* $\pi_{n-1}(x) \in \operatorname{im} \partial_n^C$ *then* $x \in \operatorname{im} \partial_n^B$.

(ix) π_* *as a chain map from* (B_*, ∂_*^B) *to* (C_*, ∂_*^C) *induces an isomorphism on* H_{n-1}.

Remark 7.6. To gain a better geometric understanding of these conditions, compare this to the setup of stacks from Chapter 5. To a short exact sequence $N \hookrightarrow \Gamma \twoheadrightarrow Q$ and classifying spaces X for N and Z for Q, we associated a stack $X \hookrightarrow Y \overset{\phi}{\twoheadrightarrow} Z$. where ϕ is a map such that for each open cell $\sigma \subset Z$, the preimage $\phi^{-1}(\sigma)$ is homeomorphic to $X \times \sigma$. So an open cell of Y is a product of a cell of X and a cell of Z.

In this setting, $B_{k,l}$ corresponds to the cells of Y which are products of a k-cell in X and an l-cell in Z (cells of "dimension $k \times l$"). The chain map π_* corresponds to ϕ and $B_n^{(m)}$ correponds to $\phi^{-1}(Z^{(m)})$. So $B_n^{(m)}$ should be seen as the n-chains lying "over the m-skeleton of C". In the conditions (iv) to (vi), subtracting $\partial_n^B(y)$ from x, should be read as "pushing the chain x across y". So, for example condition (v) says that any $(n-1)$-chain in the kernel of π_* can be pushed across cells of dimension $1 \times (n-1)$ to lie over the $(n-2)$-skeleton of C. Condition (vi) says that a chain over the m-skeleton of C which bounds over the $(m-1)$-skeleton, can be pushed across cells of dimension $(n-m) \times m$ to lie over the $(m-1)$-skeleton of C.

Proof. Since the modules $B_{k,l}$ are free, we can surely choose maps $d_0^k : B_{0,k} \to B_{0,k-1}$ as in the statement: Simply choose for every $x \in X_0$ and $z \in Z_l$ a $b_{x,\tilde{z}} \in B_{0,k-1}$ with $\pi_{k-1}(b) = \partial_k^C(\pi_k(x \otimes \tilde{z}))$ and define d_0^k by $d_0^k(x \otimes \tilde{z}) = b_{x,\tilde{z}}$.

We set (necessarily, to satisfy (ii)) $c_1^0 = c_0^1 := 0$ and define (to satisfy (i)) ∂_1^B by $\partial_1^B|_{B_{1,0}} = d_1^0$ and $\partial_1^B|_{B_{0,1}} = d_0^1$. Condition (iii) is trivially satisfied.

Next, we will show that, if maps c_k^l and ∂_n^B have been defined for some $n \geq 1$ and all $k, l \geq 0$ with $k + l = n$, satisfying conditions (i)–(iii), then (iv)–(ix) hold as well.

First we will make some useful observations. Recall that, by definition, π_{n-1} is trivial on $B_{k,l}$ for $k \geq 1$. We have

$$B_{n-1}^{(n-2)} = \bigoplus_{k=1}^{n-1} B_{k,n-1-k},$$

so $\pi_{n-1}|_{B_{n-1}^{(n-2)}} = 0$ and by (ii) this implies that $\pi_{n-1} \circ c_0^n = 0$.

Similarly, for $n \geq 1$, $k + l = n$ we have

$$B_{n-1}^{(l-1)} = \bigoplus_{k'=k}^{n-1} B_{k',l},$$

so if $k \geq 1$ this implies $B_{n-1}^{(l-1)} = 0$, and so $\pi_{n-1} \circ c_k^l = 0$ by (ii). To sum up, we have

$$\pi_{n-1} \circ c_k^l = 0, \quad \text{for all } k \text{ and } l \text{ with } k + l = n \tag{$*$}$$

To prove (iv), let $x \in B_{n-1}$ with $\pi_{n-1}(x) \in \operatorname{im} \partial_n^C$. So pick $c \in C_n$ with $\partial_n^C(c) = \pi_{n-1}(x)$ and $y \in B_{0,n}$ with $\pi_n(y) = c$. By (i) and $(*)$ we have

$$\pi_{n-1}\partial_n^B(y) = \pi_{n-1}d_0^n(y) + \pi_{n-1}c_0^n(y) = \pi_{n-1}d_0^n(y).$$

Thus

$$\pi_{n-1}(x - \partial_n^B(y)) = \pi_{n-1}(x) - \pi_{n-1}d_0^n(y) = \partial_n^C(c) - \partial_n^C\pi_n(y) = 0.$$

To prove (v), let $x \in B_{n-1}$ with $\pi_{n-1}(x) = 0$. Note that $B_{n-1} = B_{0,n-1} \oplus B_{n-1}^{(n-2)}$. Write x as

$$x = x' + x'', \quad \text{with } x' \in B_{0,n-1}, x'' \in B_{n-1}^{(n-2)}$$

Then $\pi_{n-1}(x') = \pi_{n-1}(x) = 0$. Since $B_{1,n-1} \xrightarrow{d_1^{n-1}} B_{0,n-1} \xrightarrow{\pi_{n-1}} C_{n-1}$ is an exact sequence, there is a $y \in B_{1,n-1}$ with $d_1^{n-1}(y) = x'$. Now

$$\partial_n^B(y) = x' + c_1^{n-1}(y)$$

and by (ii), $c_1^{n-1}(y) \in B_{n-1}^{(n-2)}$. Thus

$$x - \partial_n^B(y) = x'' - c_1^{n-1}(y) \in B_{n-1}^{(n-2)}.$$

To prove (vi), let $x \in B_{n-1}^{(m)}$ for some $m < n-1$ such that $\partial_{n-1}^B(x) \in B_{n-2}^{(m-1)}$. Note that $B_{n-1}^{(m)} = B_{n-1-m,m} \oplus B_{n-1}^{(m-1)}$, so write

$$x = x' + x'', \quad \text{with } x' \in B_{n-1-m,m}, x'' \in B_{n-1}^{(m-1)}$$

Now

$$\partial_{n-1}^B(x) = d_{n-1-m}^m(x') + c_{n-1-m}^m(x') + \partial_{n-1}^B(x'')$$

Both $c_{n-1-m}^m(x')$ and $\partial_{n-1}^B(x'')$ are in $B_{n-2}^{(m-1)}$, and by assumption, so is $\partial_{n-1}^B(x)$. But $d_{n-1-m}^m(x') \in B_{n-2-m,m}$, and this set has trivial intersection with $B_{n-2}^{(m-1)}$. Hence $d_{n-1-m}^m(x') = 0$. Since $B_{n-m,m} \xrightarrow{d_{n-m}^m} B_{n-1-m,m} \xrightarrow{d_{n-1-m}^m} B_{n-2-m,m}$ is an exact sequence, there is a $y \in B_{n-m,m}$ with $d_{n-m}^m(y) = x'$. Now

$$\partial_n^B(y) = x' + c_{n-m}^m(y)$$

and by (ii), $c_{n-m}^m(y) \in B_{n-1}^{(m-1)}$. Thus

$$x - \partial_n^B(y) = x'' - c_{n-m}^m(y) \in B_{n-1}^{(m-1)}.$$

Property (vii) is a simple inductive application of (vi) and (v). We will do induction on m. First let $m = 0$. If $n = 1$, then (vii) follows directly from (v), noting that $B_{n-1}^{(n-2)} = 0$. If $m = 0$ and $n > 1$ then (vii) follows from (vi), noting that $B_{n-2}^{(m-1)} = 0$ and $B_{n-1}^{(m-1)} = 0$. Now suppose $n > 1$ and that (vii) is true for some $m < n-1$ and let $x \in B_{n-1}^{(m+1)}$. If $m+1 < n-1$, then we can apply (vi) to give a $y \in B_{n-m-1,m+1}$ with $x - \partial_n^B(y) \in B_{n-1}^{(m)}$. Now, by (iii), we have

$$\partial_{n-1}^B(x - \partial_n^B(y)) = 0,$$

and clearly $\pi_{n-1}(x - \partial_n^B(y)) = 0$. Thus, by induction, there is a $y' \in B_n^{(m)}$ with $\partial_n^B(y') = x - \partial_n^B(y)$, and so $x = \partial_n^B(y' + y)$. Note that $y' + y \in B_n^{(m+1)}$ and that finishes the induction for the case $m + 1 < n - 1$.

If $m+1 = n-1$, then (v) says that there is a $y \in B_{1,n-1}$ with $x - \partial_n^B(y) \in B_{n-1}^{(n-2)}$. Now again, $\partial_{n-1}(x - \partial_n^B(y)) = 0$ and $\pi_{n-1}(x - \partial_n^B(y)) = 0$, so by induction there is a $y' \in B_n^{(m)}$ with $\partial_n^B(y') = x - \partial_n^B(y)$. Then $y + y' \in B_n^{(n-1)} = B_n^{(m+1)}$ and $\partial_n^B(y+y') = x$.

To prove (viii) let $x \in B_{n-1}$ with $\partial_{n-1}^B(x) = 0$ and $\pi_{n-1}(x) \in \operatorname{im} \partial_n^C$. By (iv) there is a $y \in B_{0,n}$ such that $x - \partial_n^B(y) \in \ker \pi_{n-1}$. By (iii) we have

$$\partial_{n-1}^B(x - \partial_n^B(y)) = \partial_{n-1}^B(x) - \partial_{n-1}^B \partial_n^B(y) = 0.$$

Thus (vii) can be applied to $x - \partial_n^B(y)$ (with $m = n-1$), to yield a $y' \in B_n^{(m)}$ with $\partial_n^B(y') = x - \partial_n^B(y)$. Then $x = \partial_n^B(y + y')$.

Finally we prove (ix). First we need to prove that π_* as a map from (B_*) to (C_*) is indeed a chain map, i.e. that it commutes with the boundary maps. Let $x \in B_n$. Then we can write

$$x = x_0 + x_1 + x_2, \quad \text{with } x_0 \in B_{0,n}, x_1 \in B_{1,n-1} \text{ and } x_2 \in B_n^{(n-2)}.$$

Note that $\partial_n^B(x_2) \in B_{n-1}^{(n-2)}$, so $\pi_{n-1}\partial_n^B(x_2) = 0$. We also have

$$\pi_{n-1}\partial_{n-1}^B(x_1) = \pi_{n-1} d_{n-1}^1(x_1) = 0.$$

Furthermore, we have

$$\pi_n(x_1 + x_2) = 0.$$

And thus

$$\pi_{n-1}\partial_n^B(x) = \pi_{n-1}\partial_n^B(x_0) = \pi_{n-1} d_n^0(x_0) = \partial_n^C \pi_n(x_0) = \partial_n^C \pi_n(x).$$

This shows that π_* is indeed a chain map up to dimension n, and therefore induces a homomorphism $H_{n-1}(B_*) \to H_{n-1}(C_*)$.

To show that this map is onto, we have to prove that $\ker \partial_{n-1}^C \subset \pi_{n-1}(\ker \partial_{n-1}^B)$. So let $c \in C_{n-1}$ with $\partial_{n-1}^C(c) = 0$. Pick a $b \in B_{0,n-1}$ with $\pi_{n-1}(b) = c$. Now $\partial_{n-2}^B \partial_{n-1}^B(b) = 0$ by (iii) and

$$\pi_{n-2}\partial_{n-1}^B(b) = \pi_{n-2} d_0^{n-1}(b) = \partial_{n-1}^C \pi_{n-1}(b) = 0.$$

So by (vii) there is a $y \in B_{n-1}^{(n-2)}$ with $\partial_{n-1}^B(y) = \partial_{n-1}^B(b)$. Thus $b - y \in \ker \partial_{n-1}^B$. Furthermore $\pi_{n-1}(y) = 0$, so $\pi_{n-1}(b-y) = \pi_{n-1}(b) = c$.

To show that the induced map $H_{n-1}(B_*) \to H_{n-1}(C_*)$ is injective, we have to show that $\operatorname{im} \partial_n^C \subset \pi_{n-1}(\operatorname{im} \partial_n^B)$. So let $c \in \operatorname{im} \partial_{n-1}^C$. Pick $x \in B_{0,n-1}$ with $\pi_{n-1}(x) = c$. By (iv) there is a $y \in B_{0,n}$ such that $x - \partial_n^B(y) \in \ker \pi_{n-1}$. Thus

$$\pi_{n-1}\partial_n^B(y) = \pi_{n-1}(x) = c.$$

This finishes the proof of (ix).

The only thing left to prove now is that, assuming maps ∂_n^B and c_k^l have been defined for some $n \geq 1$ and all k, l with $k + l = n$, such that (i)–(ix) hold, we can define maps ∂_{n+1}^B and c_k^l for $k + l = n+1$ such that (i)–(iii) hold.

Because of (ii), we have to set $c^0_{n+1} = 0$ and $\partial^B_n|_{B_{n+1,0}} = d^0_{n+1}$. Then

$$\partial^B_n \partial^B_{n+1}|_{B_{n+1,0}} = d^0_n d^0_{n+1} = 0,$$

so (iii) holds on $B_{n+1,0}$.

Now we will define ∂^B_n on $B_{k,l}$ with for $1 \leq l \leq n$. For $b \in B_{k,l}$ we have

$$\partial^B_n(d^l_k(b)) = d^l_{k-1} d^l_k(b) + c^l_{k-1} d^l_k(b)$$

Now, if $k > 1$ then $d^l_{k-1} d^l_k = 0$ and $\operatorname{im} c^l_{k-1} \in B^{(l-1)}_{n-1}$. If $k = 1$, then $\operatorname{im} d^l_{k-1} \in B_{k-1,l-1}$ and $\operatorname{im} c_{k-1^l} \in B^{(l-2)}_{n-1}$. So in any case, we have $\partial^B_n(d^l_k(b)) \in B^{(l-1)}_{n-1}$. Furthermore $\partial^B_{n-1}\partial^B_n(d^l_k(b)) = 0$ and $\pi_{n-1}\partial^B_n(d^l_k(b)) = 0$. Thus (vii) applies and so there is a $c_b \in B^{(l-1)}_n$ with $\partial^B_n(c_b) = \partial^B_n(d^l_k(b))$. So we can pick for each $x \in X_k$ and $\tilde{z} \in \tilde{Z}_l$ such a $c_{x,\tilde{z}}$ with $\partial^B_n(c_{x,\tilde{z}}) = \partial^B_n(d^l_k(x \otimes \tilde{z}))$ and define c^l_k by

$$c^l_k(x \otimes \tilde{z}) = -c_{x,\tilde{z}}.$$

Then (ii) is satisfied and setting $\partial^B_{n+1}|_{B_{k,l}} := d^l_k + c^l_k$, we have

$$\partial^B_n \partial^B_{n+1}(x \otimes \tilde{z}) = \partial^B_n(d^l_k(x \otimes \tilde{z})) - \partial^B_n(c_{x,\tilde{z}}) = 0$$

for all $x \in X_n$ and $\tilde{z} \in \tilde{Z}_l$, so (iii) holds on $B_{k,l}$.

Finally, we have to define c^{n+1}_0. For $b \in B_{n+1,0}$ we have

$$\partial^B_n(d^{n+1}_0(b)) = d^n_0 d^{n+1}_0(b) + c^n_0 d^{n+1}_0(b).$$

Now $\partial^B_{n-1}\partial^B_n d^{n+1}_0(b) = 0$ and

$$\pi_{n-1}\partial^B_n d^{n+1}_0(b) = \pi_{n-1} d^n_0 d^{n+1}_0(b) = \partial^C_n \partial^C_{n+1} \pi_{n+1}(b) = 0$$

Thus, by (vii), there is a $c_b \in B^{(n-1)}_n$ with $\partial^B_n(c_b) = \partial^B_n d^{n+1}_0(b)$. Again, pick for each $x \in X_0$ and $\tilde{z} \in \tilde{Z}_{n+1}$ a $c_{x,\tilde{z}}$ with $\partial^B_n(c_{x,\tilde{z}}) = \partial^B_n(d^{n+1}_0(x \otimes \tilde{z}))$ and define c^l_k by

$$c^l_k(x \otimes \tilde{z}) = -c_{x,\tilde{z}}.$$

Then (ii) is satisfied and setting $\partial^B_{n+1}|_{B_{0,n+1}} := d^{n+1}_0 + c^{n+1}_0$, we have

$$\partial^B_n \partial^B_{n+1}(x \otimes \tilde{z}) = \partial^B_n(d^{n+1}_0(x \otimes \tilde{z})) - \partial^B_n(c_{x,\tilde{z}}) = 0$$

for all $x \in X_0$ and $\tilde{z} \in \tilde{Z}_{n+1}$, so (iii) holds on $B_{0,n+1}$. □

7.3 A candidate complex for the n-$(n+1)$-$(n+2)$ Conjecture

With this construction in hand, we can now tackle the sequences arising in the n-$(n+1)$-$(n+2)$-Conjecture.

Proposition 7.7. *Let $N \hookrightarrow \Gamma \twoheadrightarrow Q$ be an extension of groups, where N is of type FP_n, Γ is of type FP_{n+1} and Q is of type FP_{n+2} (with $n \geq 1$). Let (A_*, ∂_*^A) be a free resolution of $\mathbb{Z}$ by N-modules which is finitely generated up to dimension n and (C_*, ∂_*^C) a free resolution of $\mathbb{Z}$ by Q-modules which is finitely generated up to dimension $n+2$. Let (B_*, ∂_*^B) be a resolution of $\mathbb{Z}$ by Γ-modules as given by Proposition 7.5. Then there is a finitely generated submodule $B'_{n+1,0} \leq B_{n+1,0}$, such that any n-cycle in $B_{n,0}$ bounds in $B'_{n+1,0} \oplus B_{n,1}$.*

Proof. First, note that by assumption, $B_{k,l}$ is a finitely generated Γ-module for $k \leq n$ and $l \leq n+2$. Therefore B_m is finitely generated for all $m \leq n$ and hence

$$B_n \to B_{n-1} \to \cdots \to B_1 \to B_0 \twoheadrightarrow \mathbb{Z}$$

is a partial finitely generated free resolution of $\mathbb{Z}$ as a Γ-module. By a basic lemma, $\ker \partial_n^B$ is finitely generated, since Γ is of type FP_n. Thus there is some finitely generated submodule $B''_{n+1} \subset B_{n+1}$ with $\partial_{n+1}(B''_{n+1}) = \ker \partial_n$. Since

$$B_{n+1} = \bigoplus_{k=0}^{n+1} B_{k,n+1-k}$$

and $B_{k,n+1-k}$ is finitely generated for $k \leq n$, we might assume that $B''_{n+1} \supset B_{k,n+1-k}$ for $k \leq n$. Hence B''_{n+1} is of the form

$$B''_{n+1} = B''_{n+1,0} \oplus \bigoplus_{k=0}^{n} B_{k,n+1-k}.$$

where $B''_{n+1,0}$ is a finitely generated submodule of $B_{n+1,0}$.

Furthermore, all the $B_{n+2-l,l}$ are finitely generated for $2 \leq l \leq n+2$. So we can pick a finitely generated $B'_{n+1,0} \leq B_{n+1,0}$ containing $B''_{n+1,0}$ and

$$B'_{n+1} := B'_{n+1,0} \oplus \bigoplus_{k=0}^{n} B_{k,n+1-k}$$

such that $B'_{n+1} \supset \partial_n^B(B_{n+2-l,l})$ for $2 \leq l \leq n+2$.

Let $c \in B_{0,n}$ with $\partial_n^B(c) = 0$. Then there is a $x \in B''_{n+1}$ with $\partial_{n+1}^B(x) = c$. Now

$$\partial_{n+1}^C \pi_{n+1}(x) = \pi_n \partial_{n+1}^B(x) = \pi_n(c) = 0$$

and since (C_*) is acyclic, we have $\pi_{n+1}(x) \in \operatorname{im} \partial_{n+2}^C$. By property (iv) in Proposition 7.5, there is a $y_0 \in B_{0,n+2}$ with $x - \partial_{n+2}^B(y_0) \in \ker \pi_{n+1}$. By (v) there is then a $y_1 \in B_{1,n+1}$ with $x - \partial_{n+2}^B(y_0) - \partial_{n+2}^B(y_1) \in B_{n+1}^{(n)}$. Now assume that elements $y_0, \ldots, y_m$ have been defined for some m with $1 \leq m < n$ such that $y_k \in B_{k,n+2-k}$ and

$$x - \partial_{n+2}^B(y_0 + \cdots + y_m) \in B_{n+1}^{(n+1-m)}.$$

Note that

$$\partial_{n+1}^B(x - \partial_{n+2}^B(y_0 + \cdots + y_m)) = \partial_{n+1}^B(x) = c \in B_n^{(0)}$$

By (vi) there is then a $y_{m+1} \in B_{m+1,n+1-m}$ with

$$x - \partial^B_{n+2}(y_0 + \cdots + y_m + y_{m+1}) \in B^{(n-m)}_{n+1}.$$

By induction we can thus define $y_0, \ldots, y_n$ with $y_k \in B_{k,n+2-k}$ and

$$x' := x - \partial^B_{n+2}(y_0 + \cdots + y_n) \in B^{(1)}_{n+1}.$$

Note that

$$\partial^B_{n+1}(x') = \partial^B_{n+1}(x) = c$$

and by definition $x' \in B'_{n+1}$. Therefore

$$x' \in B'_{n+1} \cap B^{(1)}_{n+1} = B'_{n+1,0} \oplus B_{n,1}. \qquad \square$$

This statement should be compared to the discussion following Proposition 5.8. There we considered a fibre product P associated to two short exact sequences

$$N_1 \hookrightarrow \Gamma_1 \overset{\pi_1}{\twoheadrightarrow} Q$$
$$N_2 \hookrightarrow F \overset{\pi_2}{\twoheadrightarrow} Q$$

with N_1 of type F_n, Γ_1 of type F_{n+1} and F finitely generated free.

Given a $\mathrm{B}N$-complex X with finite n-skeleton and a $\mathrm{B}Q$-complex Z with finite $(n+1)$-skeleton (and a single vertex) we constructed a $\mathrm{B}\Gamma$-complex Y and a stack $\phi : Y \to Z$ realising π_1, with fibre X, such that $W := \phi^{-1}(Z^{(1)})$ was a $\mathrm{B}P$-complex. Then Y and W have finite n-skeleton, and identifying X with the fibre above the single vertex of Z, we considered X a subcomplex of both X and W.

Proving that P is of type FP_{n+1} came down to answering the following question:

Question 7.8. Is there some finite set E of $(n+1)$-cells of X such that every n-sphere $f : S^n \to X^{(n)}$ can be extended to $\bar{f} : B^{n+1} \to W$ such that the only $(n+1)$-cells of X that $\operatorname{im} f$ intersects are in E?

In other words: Is there a finite subcomplex $X'^{(n+1)} \subset X^{(n+1)}$ such that every n-sphere $S^n \to X^{(n)}$ can be contracted in $(W \setminus X) \cup X'^{(n+1)}$?

Since Γ is of type FP_{n+1}, this question has a positive answer when W is replaced by Y. So we speculated whether we can (as in the 1-2-3 case) use the FP_{n+2}-property on Q to "push" a disk $\bar{f} : B^{n+1} \to (Y \setminus X) \cup X'^{(n+1)}$ into W in such a way that the pushing process fixes the boundary and only introduces cells from some finite subcomplex $X'' \subset X$ in the image.

In this chapter we attempted to "algebraise" the notions of Chapter 5, replacing classifying complexes by resolutions. The resolutions from Proposition 7.5 are our "algebraisation" of stacks. Proposition 7.7 now says that Question 7.8 has a positive answer in the algebraic setting.

Let us go through the intuition of the translation from the topological to the algebraic setting in some more detail. As we explained before, the generators of $B_{k,l}$

in the resolution we constructed for Γ correspond to those kl-cells of the stack Y that project to l-cells in Z under ϕ. So the cells of $W = \phi^{-1}(Z^{(1)})$ correspond to the first two columns of the $B_{*,*}$-grid

$$\begin{tikzcd}
A_3 \arrow[d, "\partial_3^A"] & B_{3,0} \arrow[d, "d_3^0"] & B_{3,1} \arrow[l, dotted] \arrow[d, "d_3^1"] & B_{3,2} \arrow[l, dotted] \arrow[d, "d_3^2"] & B_{3,3} \arrow[l, dotted] \arrow[d, "d_3^3"] \\
A_2 \arrow[d, "\partial_2^A"] & B_{2,0} \arrow[d, "d_2^0"] & B_{2,1} \arrow[l, dotted] \arrow[d, "d_2^1"] & B_{2,2} \arrow[l, dotted] \arrow[d, "d_2^2"] \arrow[ull, dotted] & B_{2,3} \arrow[l, dotted] \arrow[d, "d_2^3"] \arrow[ull, dotted] \\
A_1 \arrow[d, "\partial_1^A"] & B_{1,0} \arrow[d, "d_1^0"] & B_{1,1} \arrow[l, dotted] \arrow[d, "d_1^1"] & B_{1,2} \arrow[l, dotted] \arrow[d, "d_1^2"] \arrow[ull, dotted] & B_{1,3} \arrow[l, dotted] \arrow[d, "d_1^3"] \arrow[ull, dotted] \arrow[uulll, dotted] \\
A_0 \arrow[d, "\varepsilon_A"] & B_{0,0} \arrow[d, two heads, "\pi_0"] & B_{0,1} \arrow[l, "d_0^1"] \arrow[d, two heads, "\pi_1"] & B_{0,2} \arrow[l, "d_0^2"] \arrow[d, two heads, "\pi_2"] \arrow[ull, dotted] & B_{0,3} \arrow[l, "d_0^3"] \arrow[d, two heads, "\pi_3"] \arrow[ull, dotted] \arrow[uulll, dotted] \\
\mathbb{Z} & C_0 & C_1 \arrow[l, "\partial_1^C"] & C_2 \arrow[l, "\partial_2^C"] & C_3 \arrow[l, "\partial_3^C"]
\end{tikzcd}$$

The spheres $f : S^n \to X^{(n)}$ correspond to the cycles in $B_{n,0}$. The finitely generated submodule $B'_{n+1,0}$, whose existence Proposition 7.7 asserts, corresponds to a finite set E of $(n+1)$-cells of X. The proposition tell us that any n-cycle in $B_{n,0}$ does bound in $B'_{n+1,0} \oplus B_{n,1}$, i.e. it is the boundary of a chain that uses only the generators corresponding to the finite set E of $(n+1)$-cells of X and the finitely many $(n+1)$-cells of Y projecting onto 1-cells in Z.

However, the correspondence between the topological and algebraic setting is not a purely intuitive notion. As we have seen in Proposition 7.1, resolutions can actually be "translated" into complexes in a precise sense. So how far off is Proposition 7.7 from actually answering Question 7.8 in the topological setting and thus solving the n-$(n+1)$-$(n+2)$ Conjecture?

To address this question, assume that we are in the situation as above,

$$\begin{tikzcd}
 & & N_2 \arrow[d, hook] & \\
N_1 \arrow[r, hook] \arrow[d, equal] & P \arrow[r, two heads] \arrow[d, two heads] & F \arrow[d, two heads, "\pi_2"] & \text{f.g. free} \\
N_1 \arrow[r, hook] & \Gamma_1 \arrow[r, two heads, "\pi_1"] & Q & \\
\mathrm{F}_n & \mathrm{F}_{n+1} & \mathrm{F}_{n+2} &
\end{tikzcd}$$

with $n \geq 2$. Let

$$\mathcal{P}_{N_1} = \langle \mathcal{X} \mid R_{N_1}, \partial_{N_1} \rangle \cong N_1$$
$$\mathcal{P}_Q = \langle \mathcal{Z} \mid R_Q, \partial_Q \rangle \cong Q$$

be finite presentations for N_1 and Q respectively, where $\mathcal{Z}$ is a basis for F. Let $X^{(2)}$ and $Z^{(2)}$ be corresponding presentation 2-complexes and extend these to $\mathrm{B}N_1$- and $\mathrm{B}Q$-complexes X and Z respectively (Lemma A.5). Let $(\mathbf{A}, \partial_*^A)$ and $(\mathbf{C}, \partial_*^C)$ denote the cellular chain complexes of the universal covers of X and Z respectively.

Let

$$\mathcal{P}_\Gamma = \langle \mathcal{X} \cup \widetilde{\mathcal{Z}} \mid R_\Gamma^0 \cup R_\Gamma^1 \cup R_\Gamma^2, \partial_\Gamma \rangle$$

be a presentation for Γ as constructed in Section 4.1. Let $Y^{(2)}$ be the corresponding presentation 2-complex and denote the set of k-cells in Y by E_k. The 1-cells E_1 correspond to the generators, so $E_1 = E_{1,0} \cup E_{0,1}$ with $E_{1,0}$ and $E_{0,1}$ denoting the cells corresponding to generators from $\mathcal{X}$ and $\widetilde{\mathcal{Z}}$ respectively. Similarly, the 2-cells of Y correspond to the relators of $\mathcal{P}_\Gamma$, so $E_2 = E_{2,0} \cup E_{1,1} \cup E_{0,2}$ where $E_{n-k,k}$ denotes the set of 2-cells corresponding to the relators R_Γ^k (see the picture in page 34). For completeness, let $E_{0,0} := E_0$.

Let $p_Y : \widetilde{Y}^{(2)} \to Y^{(2)}$ be the universal cover of $Y^{(2)}$ and

$$B_2 \overset{\partial_2^B}{\to} B_1 \overset{\partial_1^B}{\to} B_0 \to 0$$

its cellular chain complex. The chain group B_k is freely generated, as a Γ-module, by a set of representatives for the Γ-orbits of k-cells and each such Γ-orbit is mapped by p_Y to one k-cell in $Y^{(2)}$.

Thus $B_1 = B_{1,0} \oplus B_{0,1}$ and $B_2 = B_{2,0} \oplus B_{1,1} \oplus B_{0,2}$ where $B_{k,l}$ is the submodule generated by the $(k+l)$-cells that map to cells from the set $E_{k,l}$.

As we explained at the end of Section 4.3, we can also define a natural map $\phi : Y^{(2)} \to Z$ which maps the cells from $E_{k,l}$ to the l-cells of Z. This lifts to a map $\widetilde{\phi} : \widetilde{Y}^{(2)} \to \widetilde{Z}$ on universal covers, which induces a chain map $\phi_* : \mathbf{B} \to \mathbf{C}$.

From our definition of the boundary map ∂_Γ in the presentation $\mathcal{P}_\Gamma$ it is now clear that one can find isomorphisms $B_{k,l} \cong A_k \otimes_{N_1} D_l$ such that the boundary maps ∂_n^B are in fact of the form as in Proposition 7.5, so we have a grid (where the dotted arrows stem from the "correction words" $c_{x,z,\varepsilon}$ and c_s in Section 4.1):

$$\begin{tikzcd}
A_2 \arrow[d, "\partial_2^A"] & B_{2,0} \arrow[d, "d_2^0"] & & \\
A_1 \arrow[d, "\partial_1^A"] & B_{1,0} \arrow[d, "d_1^0"] & B_{1,1} \arrow[l, dotted] \arrow[d, "d_1^1"] & \\
A_0 \arrow[d, "\varepsilon_A"] & B_{0,0} \arrow[d, "\phi_0"] & B_{0,1} \arrow[l, "d_0^1"'] \arrow[d, "\phi_1"] & B_{0,2} \arrow[l, "d_0^2"'] \arrow[llu, dotted] \arrow[d, "\phi_2"] \\
\mathbb{Z} & C_0 & C_1 \arrow[l, "\partial_1^C"'] & C_2 \arrow[l, "\partial_2^C"']
\end{tikzcd}$$

Here, ϕ_* is the map called π_* in the statement of Proposition 7.5.

We can now employ Proposition 7.5 to extend this to a full grid as in the statement of the proposition, yielding in particular, a resolution $(\mathbf{B}, \partial_*^B)$ of $\mathbb{Z}$ by Γ-modules.

But now we are in the situation of Proposition 7.1: We have defined a 2-complex $Y^{(2)}$ and we have extended the cellular chain complex of its universal cover to a full resolution $\mathbf{B} \twoheadrightarrow \mathbb{Z}$. This proposition then yields a complex Y, with $Y^{(2)}$ as its 2-skeleton and the cellular chain complex of its universal cover isomorphic to $\mathbf{B}$. From the last sentence in the statement of the proposition it is also clear that we can choose the complex Y in such a way that its cells correspond precisely to a chosen set of generators of B_n (say, the generators of the form $x \otimes \tilde{z}$ used in Section 7.2). So, let $E_{k,l}$ be the set of cells in Y corresponding to[3] $B_{k,l}$.

It is less clear that we can extend the map $\phi : Y^{(2)} \to Z$ to all of Y in such a way that it maps each cell from $E_{k,l}$ onto an l-cell of Z, but in fact this is possible.[4] From Section 4.2 we know that $\phi^{-1}(Z^{(1)})$ has fundamental group P. Of course, we would like to know that, once we have extended ϕ to Y, the preimage $\phi^{-1}(Z^{(1)})$ is a $\mathrm{B}P$-complex, like in Proposition 5.8. Unfortunately, we do not know how to show this at present.

However, in view of Proposition 5.8, this certainly seems like a reasonable expectation, so let us assume this for the moment:

Assumption. $W := \phi^{-1}(Z^{(1)})$ is indeed aspherical and thus a $\mathrm{B}P$-complex.

Let $\overline{W} := p_Y^{-1}(W) \subset \widetilde{Y}$ be the preimage of W under the universal convering map of Y.

We have a commutative diagram:

$$\begin{tikzcd}
 & \overline{W} \arrow[rr, two heads, "\tilde{\phi}|"] \arrow[dl] \arrow[dd, two heads] & & \widetilde{Z}^{(1)} \arrow[dl] \arrow[dd, two heads] \\
\widetilde{Y} \arrow[rr, two heads, "\tilde{\phi}"] \arrow[dd, two heads] & & \widetilde{Z} \arrow[dd, two heads] & \\
 & W \arrow[rr, two heads, "\phi|"] \arrow[dl] & & Z^{(1)} \arrow[dl] \\
Y \arrow[rr, two heads, "\phi"] & & Z &
\end{tikzcd}$$

Here, all the vertical maps are covering maps and the slanted maps are inclusions.

The n-skeleton of W is finite, so to prove that P is of type F_{n+1}, we would have to show that there is a finite subcomplex $W' \subset W^{(n+1)}$ with $W'^{(n)} = W^{(n)}$ (see

[3] What this means precisely, is that for any n-cell $\tilde{e}$ in the universal cover $\widetilde{Y}$ which maps to a cell $e \in E_{k,l}$ under the covering projection $p_Y : \widetilde{Y} \to Y$, the singular homology class $[c_e] \in H^s(\widetilde{Y}^{(n)}, \widetilde{Y}^{(n-1)})$ of a characteristic map for e is in $B_{k,l}$ (recall that the singular relative homology group $H_n^s(\widetilde{Y}^{(n)}, \widetilde{Y}^{(n-1)})$ is the n^{th} cellular chain group).

[4] Ideally, of course, we would like for ϕ to be a stack with fibre X, but we do not know at the moment whether it is possible to construct Y and ϕ in a way to ensure this.

Lemma A.16) and $\pi_n(W') = 0$. Similarly, $\overline{W}^{(n)}$ is cocompact. It would be enough to show that there is a cocompact, Γ_1-invariant subcomplex $\overline{W}'$ with $\overline{W}'^{(n)} = \overline{W}^{(n)}$ with $\pi_n(\overline{W}') = 0$, since the covering map p_Y induces an isomorphism.

Now apply Proposition 7.7. This tells us that there is a finitely generated submodule $B'_{n+1,0} \leq B_{n+1,0}$ such that any n-cycle in $B_{n,0}$ bounds in $B'_{n+1,0} \oplus B_{n,1}$. Recall that the Γ-orbits of cells in $\overline{W}$ correspond precisely to a fixed set of generators of $B_{n+1,0}$ and $B_{n,1}$. We may assume that $B'_{n+1,0}$ is generated by some finite subset of this fixed set of generators. Let $\overline{W}'$ be the subcomplex of $\overline{W}^{(n+1)}$ with $\overline{W}'^{(n)} = \overline{W}^{(n)}$ and containing exactly those $(n+1)$-cells of Y which correspond to the finitely many generators of $B'_{n+1,0}$ and $B_{n,1}$. Then $\overline{W}'$ is cocompact and another way of expressing Proposition 7.7 is to say that $H_n(\overline{W}') = 0$.

As we said above, it would be sufficient to show that $\pi_n(\overline{W}') = 0$, but even though $\overline{W}'$ has zero homology in dimensions 2 to n, we cannot apply Hurewicz' Theorem here: $p_Y| : \overline{W} \to W$ is *not* the universal cover of W. In fact, since this is just the restriction of p_Y, the deck transformation group of $\overline{W} \to W$ is still Γ_1, so $\pi_1(\overline{W}) = N_2$, the kernel of the map $P \to \Gamma_1$ induced by the inclusion $W \hookrightarrow Y$ (note that also, $\pi_1(\tilde{Z}^{(1)}) = N_2$, since Z is a BQ-complex, so $\tilde{Z}^{(1)}$ is just the fundamental group of the Cayley graph of Q on the generating set $\mathcal{Z}$, and $\ker(F(\mathcal{Z}) \twoheadrightarrow Q) = \ker \pi_2 = N_2$).

Let $\tilde{W}$ denote the universal cover of W (and $\overline{W}$). and let $\tilde{W}' \subset \tilde{W}^{(n+1)}$ be the preimage of $\overline{W}'$ under the covering projection. Also, let $W' := p_Y(\overline{W}') \subset W^{(n+1)}$. So we have a diagram:

$$\begin{tikzcd}
\tilde{W}' \arrow[r, hook] \arrow[d, two heads] & \tilde{W} \arrow[r, two heads] \arrow[d, two heads] & \widetilde{Z^{(1)}} \arrow[d, two heads] \\
\overline{W}' \arrow[r, hook] \arrow[d, two heads] & \overline{W} \arrow[r, two heads] \arrow[d, two heads] & \tilde{Z}^{(1)} \arrow[d, two heads] \\
W' \arrow[r, hook] & W \arrow[r, two heads] & Z^{(1)}
\end{tikzcd}$$

In a way, we can quantify how far $\pi_n(\overline{W}')$ is from being trivial. For that, consider the Cartan-Leray spectral sequence associated to the regular covering $\tilde{W}' \to \tilde{W}'/N_2 = \overline{W}'$ (see [Bro94, Theorem VII.7.9]):

$$E^2_{pq} = H_p(N_2, H_q(\tilde{W}')) \Rightarrow H_{p+q}(\overline{W}')$$

Since $\tilde{W}'^{(n)} = \tilde{W}^{(n)}$ and $\tilde{W}$ is contractible, we have

$$H_q(\tilde{W}') = H_q(\tilde{W}) = 0, \quad \text{for } 1 \leq q \leq n-1, \tag{7.3}$$

and thus $E^r_{p,q} = 0$ for $1 \leq q \leq n-1$ and all $r \geq 2$. Also,

$$E^2_{p,0} = H_p(N_2, H_0(\tilde{W}')) = H_p(N_2) = 0, \quad \text{for } p \geq 2,$$

since N_2 is free (it is a subgroup of the free group F).

Now focus on the term at position $(0,n)$. By what we have just shown, the boundary maps

$$d_r : E^r_{r,n-r+1} \to E^r_{0,n}, \quad \text{for } r \geq 2$$

are all trivial. Therefore

$$E^\infty_{0,n} = E^2_{0,n} = H_0(N_2, H_n(\widetilde{W}')).$$

But this sequence abuts to $H_*(\overline{W}')$ and we have picked $\overline{W}'$ such that $H_n(\overline{W}') = 0$. Hence,

$$H_0(N_2, H_n(\widetilde{W}')) = 0.$$

Now $\widetilde{W}'$ is a simply connected space with trivial homology up to dimension $n-1$ (by (7.3)). So, by Hurewicz' Theorem, $H_n(\widetilde{W}') \cong \pi_n(\widetilde{W}')$. But $\widetilde{W}'$ is the universal cover of W', therefore $\pi_n(\widetilde{W}') \cong \pi_n(W')$. To sum up, this shows that

$$H_0(N_2, \pi_n(W')) = 0.$$

It is not inconceivable that, in fact, $\pi_n(W') = 0$. If this were the case, then W' would be a finite $\mathrm{B}P^{(n+1)}$-complex, proving (modulo the assumption made earlier) that P is of type F_{n+1}.

8

Abelian quotients

We have found in Chapter 3 that there is particular interest in resolving the n-$(n+1)$-$(n+2)$ Conjecture in the case where the quotient Q is nilpotent, as that would suffice to establish the Virtual Surjections Conjecture. One might try then, as a first step towards that goal, to prove the conjecture in the case where the quotient Q is abelian, hoping to then establish the nilpotent case by induction on the nilpotency class. We will prove in this chapter the abelian case, followed by an analysis of the difficulties one comes up against in a potential induction argument.

The case of abelian Q is in fact quite special. Recall the standard setup: We are given two short exact

$$N_1 \hookrightarrow \Gamma_1 \overset{\pi_1}{\twoheadrightarrow} Q$$
$$N_2 \hookrightarrow \Gamma_2 \overset{\pi_2}{\twoheadrightarrow} Q$$

where Γ_1 and Γ_2 are of type F_{n+1}, N_1 is of type F_n and Q now is a finitely generated abelian group (and therefore in particular of type F_∞). The fibre product

$$P = \{(\gamma_1, \gamma_2) \in \Gamma_1 \times \Gamma_2 \mid \pi_1(\gamma_1) = \pi_2(\gamma_2)\}$$

is then the kernel of a homomorphism to an abelian group, namely

$$\Gamma_1 \times \Gamma_2 \to Q, (\gamma_1, \gamma_2) \mapsto \pi_1(\gamma_1) - \pi_2(\gamma_2).$$

This innocuous observation simplifies matters enormously: While there is a rather short supply of generally applicable tools to answer the question which subgroups of a group of type F_n are themselves of this type, there is one large exception. That is the case where the subgroup in question is the kernel of a homomorphism to an abelian group. For here, the theory of Σ-invariants can be brought to bear.

8.1 Σ-invariants, Meinert's Inequality and the abelian n-$(n+1)$-$(n+2)$ Conjecture

We content ourselves here with a brief summary of results from Σ-theory pertinent to our discussion. The invariants were first introduced in the context of metabelian group by Bieri and Strebel in [BS80] and later generalised in [BNS87, BR88, BG03]. A general reference is [BS]. The higher-dimensional topological invariants used here were introduced in [Ren88]. For a survey, see [Bie93, Bie99].

Consider for any group Γ of type F_n the set $\mathrm{Hom}(\Gamma,\mathbb{R}) \smallsetminus \{0\}$ of non-trivial homomorphisms from Γ to the additive reals (the "characters" of Γ). We take two such homomorphisms $\chi_1, \chi_2 : \Gamma \to \mathbb{R}$ to be equivalent if they only differ by a positive real factor, i.e. $\chi_1 = c\chi_2$ for some $c > 0$. The set of equivalence classes is called the character sphere $S(\Gamma)$ (note that $\mathrm{Hom}(\Gamma,\mathbb{R})$ is a finite-dimensional real vector space if Γ is finitely generated, so $S(\Gamma)$ is a topological sphere).

We can then associate to Γ a chain of invariants

$$S(\Gamma) \supset \Sigma^1(\Gamma) \supset \Sigma^2(\Gamma) \supset \cdots \supset \Sigma^n(\Gamma),$$

the Σ- or Bieri-Neumann-Strebel-Renz invariants of Γ.

The Σ-invariants describe certain connectivity properties of $E\Gamma$-complexes and are related to a variety of structural properties of the group, for which we refer the reader to the sources cited above. For us, the utility of these invariants lies in the fact that they contain all the information on the higher finiteness properties of the normal subgroups of Γ with abelian quotient. The following theorem is due to Renz [Ren88, Ren89] (see [BR88] for a proof of the analogous result for the homological invariant or [BG03] for a proof in a much more general setting).

Theorem 8.1. *Let Γ be a group of type F_n and $N \triangleleft \Gamma$ a normal subgroup with Γ/N abelian. Then N is of type F_n if and only if*

$$S(G,N) := \{[\chi] \in S(G) \mid \chi(N) = 0\}$$

is contained in Σ^n. □

To apply this theorem to the problem at hand we also need a way of relating the Σ-invariants of a direct product to those of the factor groups. Fortunately, all our needs are served by Meinert's Inequality. This is due to Meinert (unpublished, based on work in [Mei90]; for a proof see [Geh98, Lemma 9.1] or the sketch in [Bie99]).

Theorem 8.2 (Meinert's Inequality). *Let Γ_1 and Γ_2 be groups of type F_n and let $\chi : \Gamma_1 \times \Gamma_2 \to \mathbb{R}$ be a homomorphism with both $\chi|_{\Gamma_1}$ and $\chi|_{\Gamma_2}$ non-trivial. If $[\chi|_{\Gamma_1}] \in \Sigma^k(\Gamma_1)$ and $[\chi|_{\Gamma_2}] \in \Sigma^l(\Gamma_2)$ for some $k, l \geq 1$ with $k + l < n$ then $[\chi] \in \Sigma^{k+l+1}(\Gamma_1 \times \Gamma_2)$.* □

Now we are equipped to give a proof of the n-$(n+1)$-$(n+2)$ Conjecture in the case where the quotient Q is virtually abelian.

Theorem 8.3. *Let*

$$N_1 \hookrightarrow \Gamma_1 \overset{\pi_1}{\twoheadrightarrow} Q$$
$$N_2 \hookrightarrow \Gamma_2 \overset{\pi_2}{\twoheadrightarrow} Q$$

be short exact sequences of groups with Γ_1 and Γ_2 of type F_{n+1}, Q finitely generated virtually abelian, N_1 of type F_k and N_2 of type F_l for some $k, l \geq 0$ with $k + l \geq n$. Then the fibre product associated to these sequences is of type F_{n+1}.

Remark 8.4. Note that the assumptions on the finiteness properties of the kernels here are more general than in the original statement, where we asked for N_1 to be of type F_n but made no assumptions on N_2. This asymmetry disappears in the abelian case, where we have the more attractive assumption that the finiteness properties of the kernels N_1 and N_2 "add up to n".[1]

Proof. First, we note that we may assume that Q is abelian instead of merely virtually abelian, by the remark after Theorem 3.14.

We have to show that $[\chi] \in \Sigma^{n+1}(\Gamma_1 \times \Gamma_2)$ for all non-trivial homomorphisms $\chi : \Gamma_1 \times \Gamma_2 \to \mathbb{R}$ with $P \subset \ker \chi$.

So let χ be any such a homomorphism. Since we can pick for any $\gamma = (\gamma_1, \gamma_2) \in \Gamma_1 \times \Gamma_2$ a $\gamma' = (\gamma_1', \gamma_2) \in P$, so that $\gamma^{-1}\gamma' \in \Gamma_1$, we have $\Gamma_1 \times \Gamma_2 = \Gamma_1 P$. So $\chi|_{\Gamma_1}$ cannot be trivial, otherwise χ would be, too. Similarly $\chi|_{\Gamma_2} \neq 0$.

Now $\chi|_{\Gamma_1}$ vanishes on $P \cap \Gamma_1 = N_1$. Since N_1 is of type F_k, Theorem 8.1 tells us that $[\chi|_{\Gamma_1}] \in \Sigma^k(\Gamma_1)$. By an analogous argument, $[\chi|_{\Gamma_2}] \in \Sigma^l(\Gamma_2)$. By Meinert's inequality this implies $[\chi] \in \Sigma^{k+l+1}(\Gamma_1 \times \Gamma_2) \subset \Sigma^{n+1}(\Gamma_1 \times \Gamma_2)$. □

Solving the n-$(n+1)$-$(n+2)$ Conjecture in this special case means that we can also establish a corresponding case of the Virtual Surjections Conjecture. This was previously proved by [Koc10].

Corollary 8.5. *Let $\Gamma_1, \dots, \Gamma_n$ be groups of type F_k and $P \leq \Gamma_1 \times \dots \times \Gamma_n$ a subgroup of their direct product. If P virtually surjects to k-tuples of factors and $p_i(P)/(P \cap \Gamma_i)$ is virtually abelian for all i with $1 \leq i \leq n$ then P is of type F_k.*

Proof. This follows directly from Theorem 3.14 (or, more precisely, from the Claim in the proof). □

In fact Lemma 3.3 tells us that if $P \leq \Gamma_1 \times \dots \times \Gamma_n$ is a subdirect product that virtually surjects to k-tuples, where $2k > n$, then the quotients Γ_i/N_i are virtually abelian, so we obtain:

Corollary 8.6. *The Virtual Surjections Conjecture holds if $2k > n$ and in particular for $n \leq 5$.*

[1] Of course it is tempting to speculate whether this alternative assumption might be sufficient in the general case as well, but at this point, the abelian case aside, we have very little evidence to either support or reject this.

Proof. The latter statement follows from the Virtual Surjections to Pairs Theorem (Corollary 4.19). ☐

8.2 The problem of nilpotent quotients

What are the challenges one faces when trying to extend this statement to nilpotent Q? First off, while the Σ-invariants, as presented here, do allow certain generalisations,[2] there simply is no analogue of these invariants giving a similar degree of control over subgroups that contain some term of the lower central series but not the whole commutator subgroup. So there seems to be little hope of directly extending the argument given above to the more general class of nilpotent groups.

Let us try to approach the problem naïvely, by attempting an induction on the nilpotency class. A finitely generated nilpotent group is built up from finitely generated abelian groups by iterated central extensions. So the problem we have to solve is this: Assuming the n-$(n+1)$-$(n+2)$ Conjecture holds for some particular group $\overline{Q}$ (say, of type F_∞) and

$$Z \hookrightarrow Q \overset{p}{\twoheadrightarrow} \overline{Q}$$

is a central extension with Z finitely generated, show that the Conjecture holds for Q.

So suppose the groups Q, $\overline{Q}$ and Z are given as above. Let

$$N_1 \hookrightarrow \Gamma_1 \overset{\pi_1}{\twoheadrightarrow} Q$$
$$N_2 \hookrightarrow \Gamma_2 \overset{\pi_2}{\twoheadrightarrow} Q$$

be short exact sequences with Γ_1 and Γ_2 of type F_{n+1} and N_1 of type F_n and denote the associated fibre product by P. We can then compare this fibre product with the one associated to the short exact sequences

$$\overline{N}_1 \hookrightarrow \Gamma_1 \overset{p\pi_1}{\twoheadrightarrow} \overline{Q}$$
$$\overline{N}_2 \hookrightarrow \Gamma_2 \overset{p\pi_2}{\twoheadrightarrow} \overline{Q}$$

Let us call this fibre product $\overline{P}$. Here, $\overline{N}_1 = \ker(p\pi_1)$ is the preimage $\pi_1^{-1}(\ker p) = \pi_1^{-1}(Z)$. Thus it fits in a short exact sequence

$$N_1 \hookrightarrow \overline{N}_1 \overset{\pi_1}{\twoheadrightarrow} Z.$$

Since N_1 is of type F_n and Z is finitely generated abelian (and hence of type F_∞), we obtain that $\overline{N}_1$, also, is of type F_n.

[2] providing, for example, similar descriptions of the finiteness properties of point stabilisers of action on $\mathrm{CAT}(0)$ spaces

Since we have assumed that the n-$(n+1)$-$(n+2)$ Conjecture holds for $\overline{Q}$, we can apply it here and conclude that $\bar{P}$ is of type F_{n+1}.

How does $\bar{P}$ relate to P? Clearly, P is a subgroup of $\bar{P}$. It is easy to see, in fact, that P is normal in $\bar{P}$: If $\alpha = (\alpha_1, \alpha_2) \in \bar{P}$, then $p\pi_1(\alpha_1) = p\pi_2(\alpha_2)$, and so $\pi_1(\alpha_1) = \pi_2(\alpha_2)z$ for some $z \in Z$. Since Z is central in Q, taking $\gamma = (\gamma_1, \gamma_2) \in P$ we then find that

$$\pi_1(\alpha_1^{-1}\gamma_1\alpha_1) = z^{-1}\pi_2(\alpha_2)^{-1}\pi_2(\gamma_2)\pi_2(\alpha_2)z = \pi_2(\alpha_2\gamma_2\alpha_2)$$

and hence $\alpha^{-1}\gamma\alpha \in P$. It is not much more difficult to identify the quotient $\bar{P}/P$.

Proposition 8.7. *With the above definitions, $\bar{P}/P \cong Z$.*

Proof. Define a map

$$\phi : \bar{P} \to Z, (\gamma_1, \gamma_2) \mapsto \pi_1(\gamma_1)\pi_2(\gamma_2)^{-1}.$$

This map has image in $Z = \ker p$ since, by definition, $p\pi_1(\gamma_1) = p\pi_2(\gamma_2)$ for any $(\gamma_1, \gamma_2) \in \bar{P}$. Moreover, ϕ is a homomorphism, since for $\alpha = (\alpha_1, \alpha_2), \beta = (\beta_1, \beta_2) \in \bar{P}$, we have

$$\begin{aligned}
\phi(\alpha\beta) &= \pi_1(\alpha_1\beta_1)\pi_2(\alpha_2\beta_2)^{-1} \\
&= \pi_1(\alpha_1)\big(\pi_1(\beta_1)\pi_2(\beta_2)^{-1}\big)\pi_2(\alpha_2)^{-1} \\
&= \pi_1(\alpha_1)\pi_2(\alpha_2)^{-1}\big(\pi_1(\beta_1)\pi_2(\beta_2)^{-1}\big) \\
&= \phi(\alpha)\phi(\beta).
\end{aligned}$$

For any $z \in Z$, there is a $\gamma_1 \in \Gamma_1$ with $\pi_1(\gamma_1) = z$ and then $(\gamma_1, 1) \in \bar{P}$ and $\phi((\gamma_1, 1)) = z$. So ϕ is onto. Furthermore, it is clear that the kernel of ϕ is exactly P. □

We have thus shown that $\bar{P}$ is of type F_{n+1} and P is a normal subgroup of $\bar{P}$ with abelian quotient. Potentially then, the problem could again be attacked using Σ-invariants: P is of type F_{n+1} if and only if $[\chi] \in \Sigma(\bar{P})$ for all non-trivial homomorphisms $\chi : \bar{P} \to \mathbb{R}$ with $P \subset \ker\chi$. This raises the question what the Σ-invariants of $\bar{P}$ are.

Our knowledge here, alas, is scant. There is certainly little hope of relating the Σ-invariants of the fibre product $\bar{P}$ to those of the factor groups Γ_1 and Γ_2 in a similarly elegant way as Meinert's Inequality accomplishes for the direct product. As an indication of the difficulties involved, consider the (much more elementary) question of how the character spheres $S(\bar{P})$ and $S(\Gamma_1 \times \Gamma_2)$ are related. The inclusion $\bar{P} \hookrightarrow \Gamma_1 \times \Gamma_2$ induces a restriction map $\mathrm{Hom}(\Gamma_1 \times \Gamma_2, \mathbb{R}) \to \mathrm{Hom}(\bar{P}, \mathbb{R})$, and thus a map $S(\Gamma_1 \times \Gamma_2) \to S(\bar{P}) \cup \{0\}$. But in general, there is no reason to even expect this map to be onto, meaning that every character $\bar{P} \to \mathbb{R}$ could be extended to $\Gamma_1 \times \Gamma_2$. This is not even the case when restricting attention to the homomorphisms whose kernels contain P.

Using the notation from Theorem 8.1 we have observed that

$$S(\bar{P}, P) = \{[\chi] \in S(\bar{P}) \mid P \subset \ker\chi\}$$

is naturally isomorphic to $S(Z)$. Similarly, $S(\Gamma_1 \times \Gamma_2, P) \cong S(Q_{\text{ab}})$. However, the natural map $Z \to Q_{\text{ab}}$, can easily turn out to be trivial (just consider a Q that is nilpotent of class 2 and $Z = [Q, Q]$). In those cases the natural map $S(\Gamma_1 \times \Gamma_2, P) \to S(\bar{P}, P) \cup \{0\}$ will also be trivial. This suggests that knowledge of the Σ-invariants of $\Gamma_1 \times \Gamma_2$ might not be of much help in determining whether $S(\bar{P}, P)$ is contained in $\Sigma^{n+1}(P)$.

A

Finiteness properties

In this appendix we collect proofs for some fundamental results on topological finiteness properties for the convenience of the reader. All of this material is standard, but we compile it here since it often appears somewhat scattered across the literature. Of course, some excellent general references do exist. In particular, [Geo08], [Bro94] and [Bie81] cover many of the results we prove here (and, obviously, many more interesting results).

A.1 Topological finiteness properties

Classifying complexes One of the basic methods of geometric group theory is to associate to a given group some topological space, which serves as a “model” of Γ, and then to investigate geometric properties of this space and finding how they reflect the structure of the group.

There are essentially two ways of “modelling” a group Γ topologically. One is to consider spaces which have Γ as their fundamental group, the other is to consider spaces with a Γ-action. In both approaches the spaces are usually endowed with some additional structure and further conditions are imposed such as connectivity assumptions or restrictions on the stabilisers in the second approach.

Here, the spaces we consider will be CW-complexes which satisfy strong connectivity assumptions, making them unique up to homotopy equivalence, as we will see later. Let us first introduce terminology for our conditions:

Definition A.1. A complex X is n-aspherical if it is path-connected and $\pi_k(X) = 0$ for all k with $2 \leq k \leq n$ or, equivalently, if every map $S^k \to X$ with $2 \leq k \leq n$ can be extended to a map $B^{k+1} \to X$.

A complex is called aspherical if it is n-aspherical for all n.

Remark A.2. Therefore, a complex is 0-aspherical or 1-aspherical exactly if it is path-connected. It is n-connected if and only if it is simply connected and n-aspherical.

By Whitehead's Theorem, a CW-complex is contractible if and only if it is simply connected and aspherical. Since any covering map $p : \widetilde{X} \to X$ induces isomorphisms $p_* : \pi_k(\widetilde{X}, x_0) \to \pi_k(X, p(x_0))$ for all $k \geq 2$ (this is a simple application of the basic lifting theorems) a CW-complex is then aspherical if and only if its universal cover is contractible.

Definition A.3. Let Γ be a group. A classifying complex for Γ, or a $\mathrm{B}\Gamma$-complex, is an aspherical CW-complex with fundamental group isomorphic to Γ.

An $\mathrm{E}\Gamma$-complex is a contractible CW-complex with a free Γ-action[1].

Remark A.4. In the case of $\mathrm{B}\Gamma$-complexes it would be better to take a particular isomorphism between Γ and the fundamental group to be part of the data, but in practice we will usually just tacitly assume that a fixed isomorphism has been chosen.

These two notions are very closely related: If X is a $\mathrm{B}\Gamma$-complex then its universal cover $\widetilde{X}$ will be an $\mathrm{E}\Gamma$-complex. Conversely, if Y is an $\mathrm{E}\Gamma$-complex then the quotient complex Y/G is a $\mathrm{B}\Gamma$-complex. So the two are in some way dual viewpoints and usually to any statement about $\mathrm{B}\Gamma$-complexes there is a "dual" statement about $\mathrm{E}\Gamma$-complexes.

As we shall see, classifying complexes (and $\mathrm{E}\Gamma$-complexes) exist for any group and any two classifying complexes for the same group are homotopy equivalent.

To show that classifying complexes exist, let Γ be a group and pick a presentation $\Gamma \cong \langle \mathcal{X} \mid R, \partial \rangle$. Let $X^{(0)}$ consist of a single 0-cell x_0. For each $x \in \mathcal{X}$ glue a 1-cell e^1_x into $X^{(0)}$ to give a complex $X^{(1)}$ with a single 0-cell and one loop for each generator. The fundamental group of $X^{(1)}$ is isomorphic to $F(\mathcal{X})$, the free group on $\mathcal{X}$, with an isomorphism given by the map $\phi : F(\mathcal{X}) \to \pi_1(X^{(1)}, x_0)$ which maps each $x \in \mathcal{X}$ to the homotopy class of a characteristic map $(B^1, S^0) \to (X^{(1)}, \{x_0\})$ of e^1_x.

For each $r \in R$, pick a path representing $\phi(\partial(r))$ and glue in a 2-cell e^2_r along this path to yield a 2-complex $X^{(2)}$. By the Seifert-Van Kampen Theorem, we have $\pi_1(X^{(2)}, x_0) \cong \Gamma$, with a particular isomorphism induced by the composition of ϕ with the natural map $\pi_1(X^{(1)}, x_0) \to \pi_1(X^{(2)}, x_0)$.

$$\begin{tikzcd}
F(\mathcal{X}) \arrow[r, "\phi"] \arrow[d, two heads] & \pi_1(X^{(1)}, x_0) \arrow[r] & \pi_1(X^{(2)}, x_0) \\
\Gamma \arrow[urr, "\cong"'] & &
\end{tikzcd}$$

The following lemma then says that $X^{(2)}$ is the 2-skeleton of an aspherical complex, which will therefore be a $\mathrm{B}\Gamma$-complex.

[1] All the group actions on CW-complexes that we consider will be so-called *rigid* actions, i.e. Γ acts by cell-permuting homeomorphisms and the (set-wise) stabiliser of each closed cell fixes the cell point-wise. From now on, we will not spell out this assumption explicitly; whenever we talk of Γ-complexes this is understood to mean a CW-complex with a rigid Γ-action.

Lemma A.5. *Let $X^{(n)}$ be an n-dimensional, $(n-1)$-aspherical CW-complex for some $n \geq 2$. Then there is an aspherical complex Y with $X^{(n)}$ as its n-skeleton.*

Proof. We will inductively construct the skeleta for Y, i.e. we will construct a chain of complexes

$$Y_0 \subset Y_1 \subset Y_2 \subset \cdots$$

such that Y_m is m-dimensional and $(m-1)$-aspherical and Y_{m-1} is the $(m-1)$-skeleton of Y_m. Set $Y_m := X^{(m)}$ for $m \leq n$. Assuming that Y_m has been constructed for some $m \geq n$, pick representatives for a set of generators of $\pi_m(X_m)$ and glue in an $(m+1)$-cell along each one to yield an $(m+1)$-dimensional complex Y_{m+1}. Then clearly $\pi_m(Y_{m+1}) = 0$.

We may then take $Y := \bigcup_{m \geq n} Y_m$ and by definition $Y^{(m)} = Y_m$. Since the m^{th} homotopy group depends only on the $(m+1)$-skeleton, we have

$$\pi_m(Y) = \pi_m(Y^{(m+1)}) = \pi_m(Y_{m+1}) = 0, \quad \text{for all } m \geq 2.$$

So Y is aspherical. □

We record:

Proposition A.6. *Every group Γ has a* $\mathrm{B}\Gamma$*-complex and an* $\mathrm{E}\Gamma$*-complex. These can be chosen to have a single vertex or a single orbit of vertices, respectively, as their* 0*-skeleton.* □

It is useful to consider the same construction from the dual viewpoint, where it appears a little cleaner: For Γ a group, let $Y^{(0)}$ be a single Γ-orbit of 0-cells. Then apply the following equivalent of Lemma A.5:

Lemma A.7. *Let Γ be a group and $X^{(n)}$ an n-dimensional, $(n-1)$-connected Γ-complex for some $n \geq 0$. Then there is a contractible Γ-complex Y with $X^{(n)}$ as its n-skeleton. If the action on $X^{(n)}$ is free then Y can be chosen as a free Γ-complex as well.*

Proof. The method is essentially the same as above. We construct a chain of complexes

$$Y_0 \subset Y_1 \subset Y_2 \subset \cdots$$

such that Y_m is an m-dimensional, $(m-1)$-connected Γ-complex (free if $X^{(n)}$ is), with Y_{m-1} as its $(m-1)$-skeleton. Set $Y_m := X^{(m)}$ for $m \leq n$.

If Y_m has been constructed for some $m \geq n$, pick representatives for a set of generators of $\pi_m(Y_m)$ and glue in an orbit of spheres along the Γ-translates of each map to yield Y_{m+1}. Clearly, the Γ-action can be extended to an action on Y_{m+1}, so that Γ acts freely on $Y_{m+1} \smallsetminus Y_m$.

Set $Y := \bigcup_{m \geq n} Y_m$, and by the same argument as earlier, Y is n-connected for all $n \geq 0$, and therefore contractible. □

The complexes in the previous two lemmas will be used very frequently, so we will give them a name:[2]

[2] This notation is not standard.

Definition A.8. An n-dimensional, $(n-1)$-aspherical complex (for some $n \geq 2$) with fundamental group Γ is called a $\mathrm{B}\Gamma^{(n)}$-complex.

An n-dimensional, $(n-1)$-connected[3] free Γ-complex (for some $n \geq 0$) is called an $\mathrm{E}\Gamma^{(n)}$-complex.

Clearly, if X is a $\mathrm{B}\Gamma$-complex then the n-skeleton $X^{(n)}$ is a $\mathrm{B}\Gamma^{(n)}$-complex (for $n \geq 2$) and Lemma A.5 shows that any $\mathrm{B}\Gamma^{(n)}$-complex is the n-skeleton of some $\mathrm{B}\Gamma$-complex. Analogously, by Lemma A.7, $\mathrm{E}\Gamma^{(n)}$-complexes are exactly the n-skeleta of $\mathrm{E}\Gamma$-complexes (for $n \geq 0$).

Having shown existence we will now show that classifying complexes are unique up to homotopy equivalence. The key reason why aspherical or contractible spaces are so useful is that it is easy to construct maps from CW-complexes into them. Precisely, we have:

Lemma A.9. *Let Γ be a group. Let X be a free Γ-complex with a Γ-invariant subcomplex $X' \leq X$. Let Y be another Γ-complex and $\phi : X' \to Y$ a cellular Γ-map. If Y is contractible or every cell of X outside of X' is of dimension at most $n+1$ and Y is n-connected, then ϕ can be extended to a cellular Γ-map $\bar{\phi} : X \to Y$.*

Proof. We set $\bar{\phi}|_{X'} := \phi$ and extend $\bar{\phi}$ inductively by skeleta. If $x_0 \in X$ is any 0-cell not in X', we can just pick $\bar{\phi}(x_0)$ to be some arbitrary 0-cell of Y and then extend equivariantly to the Γ-orbit of x_0 by setting $\bar{\phi}(\gamma x_0) = \gamma\phi(x_0)$ for $\gamma \in \Gamma$. Since we assumed that Γ acts freely on X, this is well-defined. Also note that, since X' is Γ-invariant, any orbit of 0-cells (or any open cells) is either contained in X' or disjoint from it, so this does not interfere with the earlier definition that $\bar{\phi}|_{X'} = \phi$.

Assume $\bar{\phi}$ has already been defined on $X' \cup X^{(k)}$. Let σ be a $(k+1)$-cell of X with characteristic map $c_\sigma : B^{k+1} \to X$. Then $\bar{\phi}$ has been defined on $c_\sigma(S^k)$, so we can define an n-sphere in Y by $f := \phi \circ c_\sigma|_{S^k}$. Since $\pi_k(Y) = 0$, this k-sphere can be extended to a map $F : B^{k+1} \to Y$. Define ϕ on σ by $\phi(c_\sigma(x)) = F(x)$ for $x \in B^{k+1}$. Again, thanks to freeness of the action, we can extend $\bar{\phi}$ equivariantly to the entire Γ-orbit of σ. The same can be done on all the other orbits of $(n+1)$-cells, defining $\bar{\phi}$ on $X^{(n+1)}$. □

If X and Y are Γ-complexes and $\phi_1, \phi_2 : X \to Y$ are Γ-maps, then a Γ-homotopy from ϕ_1 to ϕ_2 is simply a homotopy $H : X \times I \to Y$ from ϕ_1 to ϕ_2, which is Γ-equivariant with Γ acting on $X \times I$ by $\gamma(x,t) = (\gamma x, t)$ for $\gamma \in \Gamma$, $x \in X$ and $t \in I$.

Then we can define the notions of "Γ-homotopic", "Γ-homotopy equivalent" etc., in the obvious way.

Lemma A.10. *Let Γ be a group. Let X be a free Γ-complex and Y a Γ-complex. If either Y is contractible or X is n-dimensional and Y is n-connected then any two cellular Γ-maps $\phi_1, \phi_2 : X \to Y$ are Γ-homotopic.*

Proof. By the above lemma we can extend the map $H : X \times \{0,1\} \to Y$ defined by $H|_{X \times 0} := \phi_1$ and $H|_{X \times 1} := \phi_2$ to $X \times I$ to give a Γ-homotopy from ϕ_1 to ϕ_2. □

[3] "(-1)-connected" means non-empty

Proposition A.11. *Any two* BΓ*-complexes for a group* Γ *are homotopy equivalent.*
Any two EΓ*-complexes are* Γ*-homotopy equivalent.*

Proof. We prove the second claim first. Let Y_1 and Y_2 be EΓ-complexes. By Lemma A.9 there are Γ-maps $\phi : Y_1 \to Y_2$ and $\psi : Y_2 \to Y_1$ (extending the empty map) and by Lemma A.10 there are Γ-homotopies $H_1 : Y_1 \times I \to Y_2$ from id_{Y_1} to $\psi\phi$ and $H_2 : Y_2 \times I \to Y_2$ from id_{Y_2} to $\phi\psi$, proving the second claim.

To prove the first claim, let X_1 and X_2 be classifying complexes for Γ. Then their universal covers Y_1 and Y_2 are EΓ-complexes. So let ϕ, ψ and H_1, H_2 be as above. Then ϕ and ψ descend to maps $\phi/\Gamma : X \to Y$ and $\psi/\Gamma : Y \to X$. Furthermore, $H_1/\Gamma : X \times I \to I$ and $H_2/\Gamma : Y \times I \to Y$ are homotopies from id_X to $\psi\phi$ and id_Y to $\phi\psi$, respectively. □

Type F_n Having proved existence and uniqueness (up to homotopy equivalence) for BΓ-complexes, we can now feel justified in calling these *classifying* complexes. However, the fact that any two classifying complexes are homotopy equivalent does not mean that each one is as useful as any other. The situation is similar to presentations, which correspond, after all, to the 2-skeleta of classifying complexes: It is easy enough to show that every group has a presentation, but for a presentation of a group to be actually useful it has to, at least, be of a manageable size — preferably finite.

There are different ways of imposing finiteness conditions on classifying complexes, but the one we will focus on is the one that naturally generalises the notion of finite presentation: We ask for the classifying complex to have finite n-skeleton for some n. Of course, not every group will have such a classifying complex for any n, so we define:

Definition A.12. A group Γ is of type F_n if there is a classifying complex for Γ with finite n-skeleton.

It is of type F_∞ if it has a classifying complex with finite n-skeleton for every $n \geq 0$.

The following characterisations follow easily:

Proposition A.13. *Let* Γ *be a group. The following are equivalent:*

1. Γ *is of type* F_n.
2. *There is a cocompact* $\mathrm{E}\Gamma^{(n)}$.

And for $n \geq 2$ *both are equivalent to*

3. *There is a finite* $\mathrm{B}\Gamma^{(n)}$.

Proof. If X is a $\mathrm{B}\Gamma$ with finite n-skeleton for some $n \geq 2$ then this n-skeleton will be a finite $\mathrm{B}\Gamma^{(n)}$. If $X^{(n)}$ is a finite $\mathrm{B}\Gamma^{(n)}$ then we can embed $X^{(n)}$ as the n-skeleton of some $\mathrm{B}\Gamma$ by Lemma A.5.

If X is a $\mathrm{B}\Gamma$ with finite n-skeleton for some $n \geq 0$, then the n-skeleton of the universal cover of X is a cocompact $\mathrm{E}\Gamma^{(n)}$. If $Y^{(n)}$ is a cocompact $\mathrm{E}\Gamma^{(n)}$ for some $n \geq 0$ then we can embed $Y^{(n)}$ as the n-skeleton of some $\mathrm{E}\Gamma$-complex Y and Y/Γ will be a $\mathrm{B}\Gamma$ with finite n-skeleton. □

Proposition A.14. *A group is of type* F_1 *if and only if it is finitely generated. It is of type* F_2 *if and only if it is finitely presented.*

Proof. The “if”-statements follow immediately from our construction of $\mathrm{B}\Gamma$-complexes.

If X is a path-connected complex with finite 1-skeleton then the inclusion $X^{(1)} \hookrightarrow X$ induces an epimorphism on fundamental groups and $\pi_1(X^{(1)})$ is the fundamental group of a finite graph and thus finitely generated free.

And if X is a path-connected space with finite 2-skeleton then the Seifert-van Kampen Theorem implies that its fundamental group is the quotient of the finitely generated free group $\pi_1(X^{(1)})$ by the normal subgroup generated by the attaching maps of the 2-cells.[4] □

A classical result says that a finitely presented group is finitely presented with respect to any finite generating set. More precisely, if Γ is finitely presented and $\mathcal{X} \subset \Gamma$ is any finite generating set then the kernel of the canonical epimorphism $F(\mathcal{X}) \twoheadrightarrow \Gamma$ is finitely normally generated, i.e. there is a finite presentation of the form $\langle \mathcal{X} \mid R, \partial \rangle$.

There is a higher-dimensional analogue: If we try to build a finite $\mathrm{B}\Gamma^{(n)}$ skeleton-by-skeleton, we cannot “start off wrong”: If Γ is of type F_m then any finite $\mathrm{B}\Gamma^{(n)}$ for $2 \leq n < m$ can serve as the n-skeleton of some finite $\mathrm{B}\Gamma^{(m)}$.

This will follow easily from the following lemma, which roughly says that if Γ is a group of type F_{n+1} and Y is an $\mathrm{E}\Gamma^{(n+1)}$ with cocompact n-skeleton then all but finitely many Γ-orbits of $(n+1)$-cells can be discarded:

Lemma A.15. *Let* Γ *be a group of type* F_{n+1}. *Let* Y *be an* $\mathrm{E}\Gamma^{(n+1)}$*-complex with cocompact* n*-skeleton (for some* $n \geq 0$*). Then there is a cocompact,* n*-connected* Γ*-invariant subcomplex* $Y' \subset Y$ *(so* Y' *is still an* $\mathrm{E}\Gamma^{(n+1)}$*) with* $Y'^{(n)} = Y^{(n)}$.

Proof. Since Γ is of type F_{n+1} there is a cocompact $\mathrm{E}\Gamma^{(n+1)}$-complex Z.

Let $\phi : Y \to Z$ and $\psi : Z \to Y$ be Γ-maps (which exist by Lemma A.9) and $H : Y^{(n)} \times I \to Y$ a homotopy from the inclusion $Y^{(n)} \hookrightarrow Y$ to $\psi\phi|_{Y^{(n)}}$ (which exists by Lemma A.10). Since $Y^{(n)} \times I$ and Z are cocompact, the images of H and ψ are cocompact Γ-invariants subspaces of Y. Hence they are contained in a cocompact Γ-invariant subcomplex $Y' \subset Y$. Note that Y' contains $Y^{(n)}$ (since $\operatorname{im} H$ does). It

[4] To make this completely precise, one would have to choose basepoints, of course, but we will skip the technical details.

remains to show that Y' is n-connected. Let $f : S^n \to Y^{(n)}$ be an n-sphere. Then $H \circ (f \times \mathrm{id})$ is a homotopy in Y from f to $\psi\phi f$. Since Z is n-connected, ϕf is null-homotopic, say via a homotopy $H' : S^n \times I \to Z$. Then $\psi \circ H'$ is a null-homotopy of $\psi H'|_{S^n \times 0} = \psi\phi f$ with image in Y'. Therefore the concatenation of $H \circ (f \times \mathrm{id})$ and $\psi H'$ is a null-homotopy of f in Y'.

Thus Y' is an n-connected, cocompact Γ-invariant subcomplex. □

The analogous statement for $\mathrm{B}\Gamma$-complexes can, as usual, be proved by passing to universal covers and we omit the proof:

Lemma A.16. *Let Γ be a group of type F_{n+1}. Let X be a $\mathrm{B}\Gamma^{(n+1)}$ with finite n-skeleton (for some $n \geq 2$). Then there is a finite, n-aspherical subcomplex $X' \subset X$ (so X' is still a $\mathrm{B}\Gamma^{(n+1)}$) with $X'^{(n)} = X^{(n)}$.* □

Using Lemma A.15, we can easily prove the following:

Proposition A.17. *If Y is a cocompact $\mathrm{E}\Gamma^{(n)}$ and Γ is of type F_m for some $m \geq n$ then there is a cocompact $\mathrm{E}\Gamma^{(m)}$-complex Y' with $Y'^{(n)} = Y$.*

Proof. It is enough to show the statement for $m = n+1$. The proposition then clearly follows by induction.

But by Lemma A.7, Y can be embedded in an $\mathrm{E}\Gamma$-complex Z with $Z^{(n)} = Y$. Then $Z^{(n+1)}$ is an $\mathrm{E}\Gamma^{(n+1)}$-complex with finite n-skeleton, so Lemma A.15 produces a cocompact subcomplex Z' which is a $\mathrm{E}\Gamma^{(n+1)}$ with $Z'^{(n)} = Y$. □

When we try to show that a group is of type F_n this proposition enables us to proceed inductively by skeleta: We can start off with a single Γ-orbit of 0-cells, then try to glue in finitely many orbits of 1-cells to make the space connected, then try to glue in finitely many orbits of 2-cells to make the resulting space 1-connected and so on. If at any step it is not possible to continue, i.e. to glue in finitely many orbits of $(n+1)$-cells to make the space n-connected, then Γ is not of type F_{n+1}.

The analogous proposition for $\mathrm{B}\Gamma$-complexes follows from Proposition A.17 (or can be proved in the same way, using Lemma A.16 in place of Lemma A.15).

Proposition A.18. *If X is a finite $\mathrm{B}\Gamma^{(n)}$ and Γ is of type F_m for some $m \geq n$ then there is a $\mathrm{B}\Gamma^{(m)}$-complex Y with $Y^{(n)} = X$.* □

In particular, using this lemma inductively, we may conclude:

Proposition A.19. *A group is of type F_∞ if and only if it is of type F_n for all $n \geq 0$.* □

If Y is a simply connected Γ-complex then the higher homotopy groups $\pi_n(Y)$ for $n \geq 2$ are Γ-modules in a natural way. Let us briefly recall how this works. Pick a basepoint $y_0 \in Y$. For each $\gamma \in \Gamma$ the translation map $t^\gamma : Y \to Y, y \mapsto \gamma y$ induces an isomorphism $t^\gamma_* : \pi_n(Y, y_0) \to \pi_n(Y, \gamma y_0)$. The assumption of simply connectedness is needed to make the latter group naturally isomorphic to $\pi_n(Y, y_0)$. If $y, z \in Y$ are any points in the space, pick a path $p : [0,1] \to Y$ from y to z. This path induces a map $\alpha_p : \pi_n(Y,z) \to \pi_n(Y,y)$ by concatenating the path with spheres[5]. Since any two paths from y to z in Y are homotopic, α_p in fact does not depend on p, but only on the endpoints y and z. We write $\alpha_{y,z} := \alpha_p$. Thus we can define an action of Γ on $\pi_n(Y, y_0)$ by letting γ act as $\alpha_{y_0,\gamma y_0} \circ t^\gamma$.

Now assume that $\widetilde{X}$ is the universal cover of a $\mathrm{B}\Gamma^{(n)}$-complex X with $n \geq 2$. Then $\widetilde{X}$ is simply connected, so the Γ-action on $\widetilde{X}$ induces a Γ-action on the higher homotopy groups as above. Of course, the homotopy group $\pi_n(X)$ for $n \geq 2$ are Γ-modules as well: For some loop $p : I \to X$ based at x_0, the element $\gamma = [p] \in \pi_1(X,x_0) = \Gamma$ acts as $\alpha_p : \pi_n(X,x_0) \to \pi_n(X,x_0)$, defined as above. This action is easily seen to be compatible with the action on $\widetilde{X}$ in the sense that the homomorphisms $\pi_n(\tilde{X}) \to \pi_n(\tilde{X})$ induced by the covering map $\widetilde{X} \to X$ are homomorphisms of Γ-modules for $n \geq 2$ (in fact, isomorphisms).

Lemma A.20. *Let X be a finite $\mathrm{B}\Gamma^{(n)}$ for $n \geq 2$. Then Γ is of type F_{n+1} if and only if $\pi_n(X)$ is finitely generated as a Γ-module.*

Proof. If the Γ-module $\pi_n(X)$ has a finite generating set then gluing in one $(n+1)$-cell along a representative for each generator yields a finite $\mathrm{B}\Gamma^{(n+1)}$.

Conversely, if X is a finite $\mathrm{B}\Gamma^{(n)}$ and Γ is of type F_{n+1} then by Proposition A.18, there is a finite $\mathrm{B}\Gamma^{(n+1)}$-complex X' with $X'^{(n)} = X$. Thus (X',X) is an $(n+1)$-cellular extension and a theorem by Whitehead (originally proved in [Whi41, Whi49]; see also [Whi78, Chapter V.1, Theorem 1.1]) says that the relative homotopy group $\pi_{n+1}(X',X)$ is a free Γ-module with basis given by a set of characteristic maps for the $(n+1)$-cells of X'. In particular, $\pi_{n+1}(X',X)$ is finitely generated. But the long exact homotopy sequence

$$\begin{array}{ccccccccc} \cdots \longrightarrow & \pi_{n+1}(X') & \longrightarrow & \pi_{n+1}(X',X) & \longrightarrow & \pi_n(X) & \longrightarrow & \pi_n(X') & \longrightarrow \cdots \\ & & & & & & & \| & \\ & & & & & & & 0 & \end{array}$$

then shows that $\pi_n(X)$ is finitely generated as well. □

With the observation from above that the universal covering projection induces Γ-module isomorphisms between higher homotopy groups, the following analogue is easily deduced:

Lemma A.21. *Let Y be a cocompact $\mathrm{E}\Gamma^{(n)}$ for $n \geq 2$. Then Γ is of type F_{n+1} if and only if $\pi_n(Y)$ is finitely generated as a Γ-module.* □

[5]Precisely: We represent elements of $\pi_n(Y,y)$ by maps $(B^n, S^{n-1}) \to (Y,\{y\})$. Let $C := (S^{n-1} \times I) \cup_g B^n$, where $g : S^{n-1} \to S^{n-1} \times I$, $a \mapsto (a,1)$, and pick a fixed homeomorphism $h : (B^n, S^{n-1}) \to (C, S^{n-1} \times 0)$. If $f : (B^n, S^{n-1}) \to (Y,\{z\})$ is an n-sphere based at z then define $f_p : (C, S^{n-1} \times 0) \to (Y,y)$ by setting $f_p(a,t) = p(t)$ for $(a,t) \in S^{n-1} \times I$ and $f_p(b) = f(b)$ for $b \in B^n$. Finally, define $\alpha_p : \pi_n(Y,z) \to \pi_n(Y,y)$, $[f] \mapsto [f_p \circ h]$. Clearly, if p' is another path homotopic to p then $\alpha_{p'} = \alpha_p$.

Basic properties The finiteness properties F_n behave well with respect to finite index subgroups and finite extensions.

Proposition A.22. *Let Γ be a group and $H \le \Gamma$ a subgroup of finite index. Then Γ is of type F_n if and only if H is.*

Proof. This is classical for $n \le 2$.

Suppose $n \ge 3$. If Γ is of type F_n, then let Y be a cocompact $\mathrm{E}\Gamma^{(n)}$. By restricting the action we can also regard Y as an n-dimensional, $(n-1)$-connected free H-complex. Since H is of finite index the restricted action will still be cocompact.

Conversely, assume H is of type F_n. Inductively, we may assume that Γ is of type F_{n-1}. Let Y be a cocompact $\mathrm{E}\Gamma^{(n-1)}$. Restricting the action as above, we may regard Y as a cocompact $\mathrm{E}H^{(n-1)}$. By Lemma A.21, $\pi_{n-1}(Y)$ is finitely generated as an H-module. But then $\pi_{n-1}(Y)$ is certainly finitely generated as a Γ-module. □

Proposition A.23. *Let Γ be a group of type F_n and H a retract of Γ. Then H is of type F_n as well.*

Proof. This is trivial for $n = 1$, so assume $n \ge 2$.

Let $\pi : \Gamma \twoheadrightarrow H$ be an epimorphism and $\sigma : H \to \Gamma$ a splitting, i.e. $\pi \circ \sigma = \mathrm{id}_H$. Inductively, we may assume that H is of type F_{n-1}. Let Y be a cocompact $\mathrm{E}\Gamma^{(n)}$ and Z an $\mathrm{E}H^{(n)}$ with finite $(n-1)$-skeleton (which exists by A.15).

We may consider Z a Γ-complex by defining $\gamma z = \pi(\gamma)z$ for $\gamma \in \Gamma$ and $z \in Z$. Let $\phi : Y \to Z$ be a Γ-map in this sense (which exists by A.9), so $\phi(\gamma y) = \pi(\gamma)\phi(y)$ for $\gamma \in \Gamma$ and $y \in Y$.

Similarly, we may consider Y an H-complex, via σ. Let $\psi : Z \to Y$ be an H-map, i.e. $\psi(hz) = \sigma(h)\psi(z)$ for $h \in H$ and $z \in Z$.

Then $\phi\psi$ is an H-map, since

$$h \cdot \phi\psi(z) = \pi\sigma(h) \cdot \phi\psi(z) = \phi(\sigma(h)\psi(z)) = \phi\psi(hz), \quad \text{for } h \in H, z \in Z.$$

Therefore, there is an H-homotopy $h : Z^{(n-1)} \times I \to Z$ from id_Z to $\phi\psi$. Also, $\operatorname{im}\phi$ is a cocompact H-invariant subspace of Z. To see this, let $Y' \subset Y$ be a compact subspace with $\Gamma Y' = Y$. Then $\phi(Y')$ is compact and $H\phi(Y') = \pi(H)\phi(Y') = \phi(\Gamma Y') = \phi(Y)$. Hence, $\operatorname{im}\phi \cup \operatorname{im} h$ is a cocompact H-invariant subspace and thus contained in some cocompact H-invariant subcomplex $Z' \subset Z$. We will show that Z' is an $\mathrm{E}H^{(n)}$. Since $Z'^{(n-1)} = Z^{(n-1)}$ it only remains to show that every $(n-1)$-sphere in $Z^{(n-1)}$ can be contracted in Z'.

Let $f : S^{n-1} \to Z^{(n-1)}$. Then ψf is an n-sphere in Y which is null-homotopic since Y is n-aspherical. Let $g : S^{n-1} \times I \to Y$ be a null-homotopy of f. Now f is homotopic in Z' to $\phi\psi f$ via $h \circ (f \times \mathrm{id})$ and $\phi\psi f$ is null-homotopic in Z' via $\phi \circ g$. □

Proposition A.24. *A direct product $\Gamma = \Gamma_1 \times \cdots \times \Gamma_n$ is of type F_n if and only if the factors Γ_i are.*

Proof. The factors are retracts of Γ, so one direction follows directly from Proposition A.23.

The other direction is the simple observation that if $X_1, \ldots, X_n$ are classifying complexes for $\Gamma_1, \ldots \Gamma_n$ respectively then $X_1 \times \cdots \times X_n$ is a classifying complex for Γ. This is because $\pi_n(X \times Y) \cong \pi_n(X) \times \pi_n(Y)$ holds for all (path-connected) spaces X and Y and all $n \geq 1$. □

Corollary A.25. *Every finitely generated abelian group is of type* F_∞.

Proof. Clearly, $\mathbb{Z}$ is of type F_∞, so $\mathbb{Z}^n$ is, by the previous proposition, and every finitely generated abelian group has some $\mathbb{Z}^n$ as a finite index subgroup, so the claim follows from Proposition A.22. □

Proposition A.26. *If* $N \hookrightarrow \Gamma \twoheadrightarrow Q$ *is a short exact sequence with* N *and* Q *of type* F_n *then* Γ *is of type* F_n.

Proof. There are several ways of proving this. For example, it follows immediately from Theorem 5.2, which we quote from [Geo08]. Alternatively, using the FP_n-properties discussed in the next section, it follows from Lemma 6.14 (or Section 7.2) and Theorem 7.2 for which we give full proofs. □

Proposition A.27. *Every finitely generated nilpotent group is of type* F_∞.

Proof. We do an induction on nilpotency class. We proved the statement for finitely generated abelian groups.

Assume that every group which is nilpotent of class at most n is of type F_∞. Let Γ be a group which is nilpotent of class $n+1$. Then Γ is a central extension

$$Z \hookrightarrow \Gamma \twoheadrightarrow Q,$$

where Q is nilpotent of class at most n. Therefore Q is of type F_∞. Furthermore, Z is finitely normally generated by Lemma 3.7. But since Z is central, Γ acts trivially on Z by conjugation, so Z is actually finitely generated. Therefore Z is of type F_∞ by Proposition A.25 and so Γ is of type F_∞ by the preceding proposition. □

A.2 Homological finiteness properties

The topological finiteness properties F_n from the previous section are natural generalisations of finite generation and finite presentability, but the fact that they are characterised by topological properties such as the finite generation of certain homotopy groups (as in Lemma A.20) can make them difficult to work with. "Algebraisations" of these properties were introduced in [Wal66] and [Ser71], which essentially replace the classifying complexes from the last chapter with resolutions.

Resolutions Let Γ be a group and Y an $E\Gamma$-complex. Then consider its (augmented) cellular chain complex:

$$\cdots \to C_n(Y) \xrightarrow{\partial_n} C_{n-1}(Y) \to \cdots \to C_1(Y) \to C_0(Y) \twoheadrightarrow \mathbb{Z}.$$

The action of Γ on Y induces an action on the chain groups $C_n(Y)$.[6] Since Γ acts freely on Y, this gives each $C_n(Y)$ the structure of a free Γ-module. $\mathbb{Z}$ can be considered as a Γ-module with a trivial action and in this way the above chain complex forms what we call a free resolution of $\mathbb{Z}$.

Definition A.28. Let Λ be a ring and M a Λ-module. A free resolution of M is an exact sequence of Λ-modules

$$\cdots \xrightarrow{\partial_3} F_2 \xrightarrow{\partial_2} F_1 \xrightarrow{\partial_1} F_0 \overset{\varepsilon}{\twoheadrightarrow} M$$

where all the F_n are free. We will also denote this resolution by $\mathbf{F} \overset{\varepsilon}{\twoheadrightarrow} M$, meaning that $\mathbf{F}$ (or, more accurately, $(\mathbf{F}, \partial_*)$) is the (unaugmented) chain complex

$$\cdots \xrightarrow{\partial_3} F_2 \xrightarrow{\partial_2} F_1 \xrightarrow{\partial_1} F_0 \to 0$$

A partial free resolution of M (of length n) is an exact sequence

$$F_n \to F_{n-1} \to \cdots \to F_1 \to F_0 \twoheadrightarrow M.$$

We call a resolution or partial resolution finitely generated if all the F_i for $i \geq 0$ are finitely generated (as Λ-modules). We will call it finitely generated up to dimension k if F_i is finitely generated for $i \leq k$.

The partial free resolutions of $\mathbb{Z}$ by Γ-modules are the "algebraisations" of the $E\Gamma^{(n)}$ complexes from Section A.1. Most of the notions and results from that section can now be "algebraised", by replacing maps by chain maps, homotopies by chain homotopies and so on. In particular, the results on extending maps or building homotopies carry over and one can use them to prove that resolutions are unique up to chain homotopy. We will not carry out these standard constructions of homological algebra here, however, and instead refer the reader to [Bro94].

Type FP_n Going back to our earlier example of a resolution, if Γ is a group of type F_n and Y is an $E\Gamma$ with cocompact n-skeleton, then the chain groups $C_k(Y)$ with $0 \leq k \leq n$ are finitely generated Γ-modules. This motivates the next definition:

[6] To see how this action is defined, recall that these chain groups are just defined to be the relative singular homology groups $H_n^s(Y^{(n)}, Y^{(n-1)})$. There is an obvious action of Γ on the set of singular simplices $\Delta^n \to X$ and thus on the singular chain groups. This action respects skeleta, since Γ acts cellularly, and thus induces an action on the relative homology groups. By standard results $C_n(Y)$ is freely generated as an abelian group by the characteristic maps of the n-cells of Y, and since Γ acts by cell-permuting homeomophisms the action of Γ on $C_n(Y)$ just permutes these generators.

Definition A.29. A group Γ is of type FP_n if there is a free resolution

$$\cdots \xrightarrow{\partial_3} F_2 \xrightarrow{\partial_2} F_1 \xrightarrow{\partial_1} F_0 \xrightarrow{\varepsilon} \mathbb{Z}$$

of $\mathbb{Z}$ by Γ-modules which is finitely generated up to dimension n.

So as an immediate consequence we have:

Proposition A.30. *If a group is of type* F_n *then it is of type* FP_n. □

In particular, every finitely generated group is of type FP_1. In fact the converse holds as well (see [Bro94, Section VIII.5]).

For $n \geq 2$ the converse of Proposition A.30 does not hold, as was proved in [BB97]. However, there is a partial converse (essentially due to Wall [Wal65, Wal66]), saying that F_n and FP_n are equivalent among finitely presented groups, which we prove as Theorem 7.2.

Let us now give a homological analogue of Proposition A.18, showing that if Γ is a group of type FP_n, one can build a finitely generated partial free resolution of $\mathbb{Z}$ of length n step-by-step. In other words, any finitely generated partial free resolution of length $k < n$ can be extended to one of length n.

We could simply imitate the proof of Proposition A.18 in the homological setting, but we will give a different proof instead.

Schanuel's Lemma. *If* Λ *is a ring and*

$$K_1 \hookrightarrow F_1 \overset{\pi_1}{\twoheadrightarrow} M$$
$$K_2 \hookrightarrow F_2 \overset{\pi_2}{\twoheadrightarrow} M$$

are short exact sequences of Λ*-modules with* F_1 *and* F_2 *free, then*

$$K_1 \oplus F_2 \cong K_2 \oplus F_1.$$

Proof. Let
$$P := \{(a,b) \in F_1 \oplus F_2 \mid \pi_1(a) = \pi_2(b)\}.$$

Then the projection map $p_1 : P \to F_1$ has kernel $1 \times K_2$. Furthermore it is easy to check that p_1 is onto so we have an exact sequence

$$K_2 \hookrightarrow P \twoheadrightarrow F_1.$$

Since F_1 is free, this sequence splits and therefore $P \cong K_2 \oplus F_1$.

By considering the projection $p_2 : P \to F_2$, we can show that $P \cong K_1 \oplus F_2$. □

Lemma A.31. *Let M be a Λ-module and let*

$$0 \to F_n \xrightarrow{\partial_n} F_{n-1} \xrightarrow{\partial_{n-1}} \cdots \to F_1 \xrightarrow{\partial_1} F_0 \xrightarrow{\partial_0} M$$
$$0 \to F'_n \xrightarrow{\partial'_n} F'_{n-1} \xrightarrow{\partial'_{n-1}} \cdots \to F'_1 \xrightarrow{\partial'_1} F'_0 \xrightarrow{\partial'_0} M$$

be exact sequences with F_k, F'_k free for $0 \leq k \leq n-1$. Then

$$F_0 \oplus F'_1 \oplus F_2 \oplus \cdots \cong F'_0 \oplus F_1 \oplus F'_2 \oplus \cdots.$$

Proof. Let $K_n := \ker \partial_n$ and $K'_n := \ker \partial'_n$. By Scanuel's Lemma, we have

$$K_0 \oplus F'_0 \cong K'_0 \oplus F_0.$$

We also have short exact sequences

$$K_1 \hookrightarrow F_1 \oplus F'_0 \twoheadrightarrow K_0 \oplus F'_0$$

and

$$K'_1 \hookrightarrow F'_1 \oplus F_0 \twoheadrightarrow K'_0 \oplus F_0.$$

Using that $F_1 \oplus F'_0$ and $F'_1 \oplus F_0$ are free and $K'_0 \oplus F'_0 \cong K'_0 \oplus F_0$, Schanuel's Lemma implies that

$$K_1 \oplus F'_1 \oplus F_0 \cong K'_1 \oplus F_1 \oplus F'_0.$$

Continuing inductively in this manner proves the proposition. □

Proposition A.32. *Let Γ be a group of type* FP_n *(for some $n \geq 1$) and*

$$F_n \xrightarrow{\partial_n} F_{n-1} \xrightarrow{\partial_{n-1}} \cdots \xrightarrow{\partial_2} F_1 \xrightarrow{\partial_1} F_0 \xrightarrow{\partial_0} \mathbb{Z}$$

a finitely generated partial free resolution of $\mathbb{Z}$ by Γ-modules. Then Γ is of type FP_{n+1} *if and only if* $\ker \partial_n$ *is finitely generated.*

Proof. If $\ker \partial_n$ is finitely generated then we can pick a finitely generated free Γ-module F_{n+1} and a surjection $\partial_{n+1} : F_{n+1} \twoheadrightarrow \ker \partial_n$ to extend the resolution above to a finitely generated partial free resolution

$$F_{n+1} \xrightarrow{\partial_{n+1}} F_n \xrightarrow{\partial_n} F_{n-1} \to \cdots \to F_1 \to F_0 \to \mathbb{Z}.$$

Thus Γ is of type FP_{n+1}.

Conversely, if Γ is of type FP_{n+1} then there is a finitely generated partial free resolution

$$F'_{n+1} \xrightarrow{\partial'_{n+1}} F'_n \xrightarrow{\partial'_n} F'_{n-1} \to \cdots \to F'_1 \to F'_0 \twoheadrightarrow \mathbb{Z}.$$

We thus have two exact sequences

$$0 \to \ker \partial_n \to F_n \xrightarrow{\partial_n} F_{n-1} \to \cdots \to F_1 \to F_0 \twoheadrightarrow \mathbb{Z}$$
$$0 \to \ker \partial'_n \to F'_n \xrightarrow{\partial'_n} F'_{n-1} \to \cdots \to F'_1 \to F'_0 \twoheadrightarrow \mathbb{Z}$$

which satisfy the assumptions of Lemma A.31. Hence,

$$\ker \partial_n \oplus F'_n \oplus F_{n-1} \oplus \cdots \cong \ker \partial'_n \oplus F_n \oplus F'_{n-1} \oplus \cdots.$$

But $\ker \partial'_n = \operatorname{im} \partial'_{n+1}$ is finitely generated, since F'_{n+1} is, and so are all the F_k and F'_k with $0 \le k \le n$. Thus the right hand side is finitely generated, hence the left hand side is, and thus also the direct factor $\ker \partial_n$. □

As a corollary we get:

Lemma A.33. *Let* Γ *be a group of type* FP_n *and*

$$F_m \to F_{m-1} \to \cdots \to F_1 \to F_0 \to \mathbb{Z}$$

a finitely generated partial free resolution of $\mathbb{Z}$ *as a* Γ*-module of dimension* $m < n$. *Then this resolution can be extended to the left to a finitely generated free resolution of dimension* n*:*

$$F_n \to \cdots \to F_m \to F_{m-1} \to \cdots \to F_1 \to F_0 \to \mathbb{Z}$$

Proof. This is a simple inductive application of the previous proposition. □

Like the topological properties, type FP_n behaves well with respect to passing to finite index subgroups or finite extensions:

Proposition A.34. *Let* Γ *be a group of type* FP_n *and* $H \le \Gamma$ *a finite index subgroup. Then* Γ *is of type* FP_n *if and only if* H *is.*

Proof. First assume that Γ is of type FP_n. Let

$$\mathbf{F} \twoheadrightarrow \mathbb{Z} : \quad \cdots \to F_2 \to F_1 \to F_0 \twoheadrightarrow \mathbb{Z}$$

be a free resolution of $\mathbb{Z}$ by Γ-modules, which is finitely generated up to dimension n. Then we can regard the F_i as H-modules. If F is a free Γ-module on a basis $\mathcal{X} \subset F$, and $T \subset \Gamma$ is a set of representatives of the cosets of H, then it is easy to check that $T\mathcal{X}$ is a basis of F considered as an H-module. In particular any finitely generated free Γ-module is finitely generated free when considered as an H-module. So the resolution $\mathbf{F} \twoheadrightarrow \mathbb{Z}$ can in fact be considered a free resolution of $\mathbb{Z}$ by H-modules, which is finitely generated up to dimension n.

Conversely, let H be of type FP_n. Inductively, we may assume that Γ is of type FP_{n-1}. Let

$$F_{n-1} \overset{\partial_{n-1}}{\to} F_{n-2} \overset{\partial_{n-2}}{\to} \cdots \overset{\partial_2}{\to} F_1 \overset{\partial_1}{\to} F_0 \twoheadrightarrow \mathbb{Z}$$

be a finitely generated partial free resolution of $\mathbb{Z}$ by Γ-modules. Again, we can consider this as a complex of H-modules, in which case it still is a finitely generated partial free resolution. Since H is of type FP_n, the kernel $\ker \partial_{n-1}$ is finitely generated as an H-module and a fortiori, finitely generated as a Γ-module. □

All of the other results that we proved for the topological properties at the end of Section A.1 have analogues for the FP_n-properties, and mostly they can be proved by simply "algebraising" the proofs given there.

The Bieri-Eckmann criterion In this section we will introduce an important criterion, due to Bieri and Eckmann, which comes down to the statement that a group Γ is of type FP_n if and only if the functors $H_k(\Gamma, -)$ commute with arbitrary direct products.

A proof of this statement can be found in [Bie81, Section 2]. Since we use this criterion extensively in Section 6.3, we take the time to explain it carefully and prove it here. The proof we give is different from the one in [Bie81]. In particular, it is essentially completely elementary, not using much homological algebra beyond the definitions.

The computations involved are fairly simple, but sometimes awkward to write down, since they involve following through long sequences of isomorphisms between sets of functions. Therefore, we will use, just in this section, two pieces of uncommon but intuitive notation, which will make many formulae a lot cleaner.

Notation. If $f : X \to Y$ is a function, we write

$$x \overset{f}{\rightsquigarrow} y \quad \text{for} \quad f(x) = y.$$

and

$$x \mapsto f(x) \quad \text{for} \quad f.$$

The first convention allows us to write, for example

$$a \overset{f}{\rightsquigarrow} b \overset{g}{\rightsquigarrow} c \overset{h}{\rightsquigarrow} d$$

for $h(g(f(a))) = d$. The second allows us to describe functions without having to name them, which is useful when working with direct sums and products (see the Lemma below for an example). We will consistently regard the direct product $\prod_{\lambda\in\Lambda} A_\lambda$ or direct sum $\bigoplus_{\lambda\in\Lambda} A_\lambda$ as a set of functions $\Lambda \to \bigcup_{\lambda\in\Lambda} A_\lambda$.

We first note two basic properties of tensor products.

LEMMA A.35. *If R is a ring, B a left R-module, Λ an index set and A_λ a right R-module for each $\lambda \in \Lambda$ then the map*

$$\left(\bigoplus_{\lambda\in\Lambda} A_\lambda\right) \otimes_R B \to \bigoplus_{\lambda\in\Lambda} (A_\lambda \otimes_R B)$$

given by $a \otimes b \mapsto (\lambda \mapsto a(\lambda) \otimes b)$ is an isomorphism. □

LEMMA A.36. *If R is a ring and A a left R-module then the map $R \otimes_R A \to A$ given by $r \otimes a \mapsto ra$ is an isomorphism (with inverse $a \mapsto 1 \otimes a$).* □

As alluded to above, the Bieri-Eckmann criterion is concerned with whether the homology functor $H_k(\Gamma, -)$ commutes with direct products. We will now make this statement precise. Let R be a ring, N a right R-module, Λ some index set and M_λ a left R-module for each $\lambda \in \Lambda$. While we do not have an analogue of Lemma

A.35, replacing direct sums with direct products, we can nevertheless always define a homomorphism

$$\phi : N \otimes_R \prod_{\lambda\in\Lambda} M_\lambda \to \prod_{\lambda\in\Lambda} (N \otimes M_\lambda)$$

with $\phi(n \otimes m) = (\lambda \mapsto n \otimes m(\lambda))$.

This map is natural in N in the following sense: If N_1 and N_2 are left R-modules and $f : N_1 \to N_2$ is a homomorphism, then f on the one hand induces a map

$$f \otimes \mathrm{id}_{\prod M_\lambda} : N_1 \otimes \prod_{\lambda\in\Lambda} M_\lambda \to N_2 \otimes \prod_{\lambda\in\Lambda} M_\lambda$$

and on the other hand the maps $f \otimes \mathrm{id}_{M_\lambda} : N_1 \otimes M_\lambda \to N_2 \otimes M_\lambda$ can be assembled into a map

$$f' : \prod_{\lambda\in\Lambda} (N_1 \otimes M_\lambda) \to \prod_{\lambda\in\Lambda} (N_2 \otimes M_\lambda)$$

given by $f'(p) = (\lambda \mapsto (f \otimes \mathrm{id}_{M_\lambda})(p(\lambda)))$. Thus we obtain a square:

$$\begin{tikzcd}
N_1 \otimes \prod_{\lambda\in\Lambda} M_\lambda \arrow[r, "f\otimes \mathrm{id}"] \arrow[d, "\phi_1"] & N_2 \otimes \prod_{\lambda\in\Lambda} M_\lambda \arrow[d, "\phi_2"] \\
\prod_{\lambda\in\Lambda} N_1 \otimes M_\lambda \arrow[r, "f'"] & \prod_{\lambda\in\Lambda} N_2 \otimes M_\lambda
\end{tikzcd}$$

where the vertical maps are the maps defined above. This square commutes since

$$n \otimes m \overset{f\otimes\mathrm{id}}{\rightsquigarrow} f(n) \otimes m \overset{\phi_2}{\rightsquigarrow} (\lambda \mapsto f(n) \otimes m(\lambda))$$

and

$$\begin{aligned} n \otimes m &\overset{\phi_1}{\rightsquigarrow} (\lambda \mapsto n \otimes m(\lambda)) \\ &\overset{f'}{\rightsquigarrow} (\lambda \mapsto (f \otimes \mathrm{id})(n \otimes m(\lambda))) = (\lambda \mapsto f(n) \otimes m(\lambda)) \end{aligned}$$

Now let Γ be a group and

$$\cdots \overset{\partial_3}{\to} F_2 \overset{\partial_2}{\to} F_1 \overset{\partial_1}{\to} F_0 \twoheadrightarrow \mathbb{Z}$$

a free resolution of $\mathbb{Z}$ by Γ-modules. Applying the above construction yields a chain map (with all the tensor products over $\mathbb{Z}\Gamma$):

$$\begin{tikzcd}
\cdots \arrow[r] & F_n \otimes \prod_{\lambda\in\Lambda} M_\lambda \arrow[r, "\partial_n\otimes\mathrm{id}"] \arrow[d, "\phi_n"] & F_{n-1} \otimes \prod_{\lambda\in\Lambda} M_\lambda \arrow[r, "\partial_{n-1}\otimes\mathrm{id}"] \arrow[d, "\phi_{n-1}"] & \cdots \arrow[r, "\partial_1\otimes\mathrm{id}"] & F_0 \otimes \prod_{\lambda\in\Lambda} M_\lambda \arrow[d, "\phi_0"] \\
\cdots \arrow[r] & \prod_{\lambda\in\Lambda} F_n \otimes M_\lambda \arrow[r, "\partial'_n"] & \prod_{\lambda\in\Lambda} F_{n-1} \otimes M_\lambda \arrow[r, "\partial'_{n-1}"] & \cdots \arrow[r, "\partial'_1"] & \prod_{\lambda\in\Lambda} F_0 \otimes M_\lambda
\end{tikzcd}$$

This induces a map between the homology groups. By definition, the homology groups of the first chain complex are $H_k(\Gamma, \prod_{\lambda\in\Lambda} M_\lambda)$. The second chain complex is simply the dimensionwise product of the chain complex $(F_* \otimes M_\lambda, \partial_* \otimes \mathrm{id})$, so its homology is the product $\prod_{\lambda\in\Lambda} H_k(\Gamma, M_\lambda)$.

So there is a natural map[7]

$$H_n(\Gamma, \prod M_\lambda) \to \prod H_n(\Gamma, M_\lambda).$$

The Bieri-Eckmann criterion asserts that this map is an isomorphism for all $k < n$ if and only if Γ is of type FP_n.

Our proof is based on the following elementary calculation:

Lemma A.37. *Let R be a ring, Λ an index set and M_λ a left R-module for each $\lambda \in \Lambda$. If F is a finitely generated free right R-module then the natural homomorphism*

$$\phi : F \otimes_R \prod_{\lambda\in\Lambda} M_\lambda \to \prod_{\lambda\in\Lambda} F \otimes_R M_\lambda$$

as defined above is an isomorphism.

Proof. Since F is finitely generated free, it is of the form

$$F \cong \bigoplus_{x\in X} R$$

for some finite set X.

There is a sequence of natural isomorphisms:

$$\begin{aligned}\left(\bigoplus_{x\in X} R\right) \otimes_R \prod_{\lambda\in\Lambda} M_\lambda &\overset{\phi_1}{\to} \bigoplus_{x\in X}\left(R \otimes_R \prod_{\lambda\in\Lambda} M_\lambda\right) \overset{\phi_2}{\to} \bigoplus_{x\in X}\prod_{\lambda\in\Lambda} M_\lambda \\ &\overset{\phi_3}{\to} \prod_{\lambda\in\Lambda}\bigoplus_{x\in X} M_\lambda \overset{\phi_4}{\to} \prod_{\lambda\in\Lambda}\bigoplus_{x\in X}(R\otimes_R M_\lambda) \overset{\phi_5}{\to} \prod_{\lambda\in\Lambda}\left(\bigoplus_{x\in X} R\right)\otimes_R M_\lambda\end{aligned}$$

Here, ϕ_1 and ϕ_5 are given by Lemma A.35, ϕ_2 and ϕ_4 are as in Lemma A.36 and ϕ_3 is the "transpose" map $a \mapsto (\lambda \mapsto (x \mapsto a(x)(\lambda)))$. Note that this latter map is an isomorphism since X is finite, so the direct sum over X can be replaced by a direct product, which clearly commutes with the direct product over Λ.

Following the definitions, we get, for $a \in \bigoplus_{x\in X} R$ and $m \in \prod_{\lambda\in\Lambda} M_\lambda$:

$$\begin{aligned} a\otimes m &\overset{\phi_1}{\rightsquigarrow} (x \mapsto a(x)\otimes m) \\ &\overset{\phi_2}{\rightsquigarrow} (x \mapsto a(x)m) \\ &\overset{\phi_3}{\rightsquigarrow} (\lambda \mapsto (x \mapsto a(x)m(\lambda))) \\ &\overset{\phi_4}{\rightsquigarrow} (\lambda \mapsto (x \mapsto 1\otimes a(x)m(\lambda))) = (\lambda \mapsto (x \mapsto a(x)\otimes m(\lambda))) \\ &\overset{\phi_5}{\rightsquigarrow} (\lambda \mapsto a\otimes m(\lambda)) = \phi(a\otimes m)\end{aligned}$$

Thus $\phi = \phi_5 \circ \phi_4 \circ \phi_3 \circ \phi_2 \circ \phi_1$ is an isomorphism. □

[7] Another way to describe this map: For each $\lambda \in \Lambda$ let $f_\lambda : H_n(\Gamma, \prod M_\lambda) \to H_n(\Gamma, M_\lambda)$ be the map induced by the projection $\prod M_\lambda \to M_\lambda$. Then the f_λ can be assembled into a map $H_n(\Gamma, \prod M_\lambda) \to \prod H_n(\Gamma, M_\lambda)$.

Before we give the proof of the Bieri-Eckmann Criterion, let us make one more observation. The construction before Lemma A.37 can be carried out for direct sums in place of direct products to give a map

$$H_n\Big(\Gamma, \bigoplus_{\lambda\in\Lambda} M_\lambda\Big) \to \bigoplus_{\lambda\in\Lambda} H_n(\Gamma, M_\lambda)$$

for any Γ-modules M_λ. But because of Lemma A.35, this map will be an isomorphism.

Lemma A.38. *If Γ is a group and F is a free Γ-module then $H_n(\Gamma, F) = 0$ for $n \geq 1$.*

Proof. Let

$$\mathbf{F} \twoheadrightarrow \mathbb{Z}: \quad \cdots \to F_2 \to F_1 \to F_0 \twoheadrightarrow \mathbb{Z}$$

be a free resolution of $\mathbb{Z}$ by Γ-modules. Since the map from Lemma A.36 is clearly natural, tensoring this resolution with $\mathbb{Z}\Gamma$ will do nothing, i.e. $\mathbf{F} \otimes \mathbb{Z}\Gamma \cong \mathbf{F}$. Since $\mathbf{F}$ has trivial homology in every dimension ≥ 1 we have

$$H_n(\Gamma, \mathbb{Z}\Gamma) = H_n(\mathbf{F} \otimes \mathbb{Z}\Gamma) = H_n(\mathbf{F}) = 0, \quad \text{for } n \geq 1.$$

With the isomorphism above we obtain

$$H_n\Big(\Gamma, \bigoplus_{\lambda\in\Lambda} \mathbb{Z}\Gamma\Big) \cong \bigoplus_{\lambda\in\Lambda} H_n(\Gamma, \mathbb{Z}\Gamma) = 0, \quad \text{for } n \geq 1.$$

for any direct sum of copies of $\mathbb{Z}\Gamma$. Since all free Γ-modules are of this form, this proves the claim. □

Now we are ready for the proof.

Theorem A.39 (The Bieri-Eckmann Criterion). *Let Γ be a group. Then the following are equivalent (for $n \geq 1$):*

1. *Γ is of type FP_n.*
2. *The natural map $H_k(\Gamma, \prod_{\lambda\in\Lambda} M_\lambda) \to \prod_{\lambda\in\Lambda} H_k(\Gamma, M_\lambda)$ is an isomorphism for $k < n$, where Λ is a countable set and M_λ a left Γ-module for each $\lambda \in \Lambda$.*
3. *The natural map $H_k(\Gamma, \prod \mathbb{Z}\Gamma) \to \prod H_k(\Gamma, \mathbb{Z}\Gamma)$ is an isomorphism for $k < n$, where the product is over a countable index set.*
4. *Γ is finitely generated and $H_k(\Gamma, \prod \mathbb{Z}\Gamma) = 0$ for all $k < n$, where the product is over a countable index set.*

Proof. It is clear that 2 implies 3.

We will show first that 1 implies 2. Let Γ be a group of type FP_n and

$$\cdots \to \xrightarrow{\partial_3} F_2 \xrightarrow{\partial_2} F_1 \xrightarrow{\partial_1} F_0 \twoheadrightarrow \mathbb{Z}$$

a free resolution of $\mathbb{Z}$ by Γ-modules. Then, by Lemma A.37, the natural chain map

$$\begin{array}{ccccccc}
\cdots \to F_n \otimes \prod_{\lambda\in\Lambda} M_\lambda & \xrightarrow{\partial_n\otimes\mathrm{id}} & F_{n-1} \otimes \prod_{\lambda\in\Lambda} M_\lambda & \xrightarrow{\partial_{n-1}\otimes\mathrm{id}} \cdots \xrightarrow{\partial_1\otimes\mathrm{id}} & F_0 \otimes \prod_{\lambda\in\Lambda} M_\lambda \\
\downarrow{\scriptstyle\phi_n} & & \downarrow{\scriptstyle\phi_{n-1}} & & \downarrow{\scriptstyle\phi_0} \\
\cdots \to \prod_{\lambda\in\Lambda} F_n \otimes M_\lambda & \xrightarrow{\partial'_n} & \prod_{\lambda\in\Lambda} F_{n-1} \otimes M_\lambda & \xrightarrow{\partial'_{n-1}} \cdots \xrightarrow{\partial'_1} & \prod_{\lambda\in\Lambda} F_0 \otimes M_\lambda
\end{array}$$

consists of isomorphisms up to dimension n. Thus the induced map on homology groups $H_k(\Gamma, \prod M_\lambda) \to \prod H_k(\Gamma, M_\lambda)$ is an isomorphism for $k < n$. This proves that 1 implies 2.

Now we will show that 3 implies 1. We may inductively assume that Γ is of type FP_{n-1}. Let

$$F_n \overset{\partial_n}{\to} F_{n-1} \overset{\partial_{n-1}}{\to} \cdots \overset{\partial_2}{\to} F_1 \overset{\partial_1}{\to} F_0 \twoheadrightarrow \mathbb{Z}$$

be a partial free resolution, finitely generated up to dimension $n-1$. We then have a commutative diagram

$$\begin{array}{ccccc}
F_n \otimes \prod_{\lambda\in\Lambda} \mathbb{Z}\Gamma & \xrightarrow{\partial_n\otimes\mathrm{id}} & F_{n-1} \otimes \prod_{\lambda\in\Lambda} \mathbb{Z}\Gamma & \xrightarrow{\partial_{n-1}\otimes\mathrm{id}} & F_{n-2} \otimes \prod_{\lambda\in\Lambda} \mathbb{Z}\Gamma \\
\downarrow{\scriptstyle\phi_n} & & \downarrow{\scriptstyle\phi_{n-1}} & & \downarrow{\scriptstyle\phi_{n-2}} \\
\prod_{\lambda\in\Lambda} F_n \otimes \mathbb{Z}\Gamma & \xrightarrow{\partial'_n} & \prod_{\lambda\in\Lambda} F_{n-1} \otimes \mathbb{Z}\Gamma & \xrightarrow{\partial'_{n-1}} & \prod_{\lambda\in\Lambda} F_{n-2} \otimes \mathbb{Z}\Gamma \\
\downarrow{\scriptstyle\psi_n} & & \downarrow{\scriptstyle\psi_{n-1}} & & \downarrow{\scriptstyle\psi_{n-2}} \\
\prod_{\lambda\in\Lambda} F_n & \xrightarrow{\partial''_n} & \prod_{\lambda\in\Lambda} F_{n-1} & \xrightarrow{\partial''_{n-1}} & \prod_{\lambda\in\Lambda} F_{n-2}
\end{array}$$

Where the maps ϕ_k are the natural homomorphisms and the maps ψ_k are isomorphisms as in Lemma A.36. Since ϕ_{n-1} and ϕ_{n-2} are isomorphisms by Lemma A.37, $\psi_{n-1}\phi_{n-1}$ maps $\ker(\partial_{n-1}\otimes\mathrm{id})$ onto $\ker\partial''_{n-1}$. Thus ϕ_{n-1} induces an isomorphism on homology only if $\psi_{n-1}\phi_{n-1}$ also maps the image of $\partial_n\otimes\mathrm{id}$ onto $\mathrm{im}\,\partial''_n$.

Suppose that Γ were not of type FP_n. Then $\ker\partial_{n-1}$ is not finitely generated, so we can find a family of elements $(a_\lambda)_{\lambda\in\Lambda}$ which is not contained in any finitely generated submodule of $\ker\partial_{n-1}$ (recall that Λ is simply a countably infinite index set). Since $\ker\partial_{n-1} = \mathrm{im}\,\partial_n$, we can find elements $b_\lambda \in F_n$ with $\partial_n(b_\lambda) = a_\lambda$ for each $\lambda\in\Lambda$.

Set $a := (\lambda \mapsto a_\lambda) \in \prod_{\lambda\in\Lambda} F_{n-1}$ and $b := (\lambda\mapsto b_\lambda) \in \prod_{\lambda\in\Lambda} F_n$. By definition $\partial''_n(b) = a$, so $a \in \mathrm{im}\,\partial''_n$. Therefore,

$$(\partial_{n-1}\otimes\mathrm{id})\phi_{n-1}^{-1}\psi_{n-1}^{-1}(a) = \phi_{n-1}^{-1}\psi_{n-1}^{-1}\partial''_{n-1}(a) = 0.$$

Hence, there is some $x \in F_n \otimes \prod_{\lambda\in\Lambda}\mathbb{Z}\Gamma$ with $\psi_{n-1}\phi_{n-1}(\partial_n\otimes\mathrm{id})(x) = a$. Now x is of the form

$$c_1\otimes m_1 + \cdots + c_k\otimes m_k, \quad \text{for some } c_1,\ldots,c_k \in F_n,\ m_1,\ldots,m_k \in \prod_{\lambda\in\Lambda}\mathbb{Z}\Gamma,$$

and we have

$$\begin{aligned} x &\overset{\partial_n \otimes \mathrm{id}}{\rightsquigarrow} \partial_n(c_1) \otimes m_1 + \cdots + \partial_n(c_k) \otimes m_k \\ &\overset{\phi_{n-1}}{\rightsquigarrow} (\lambda \mapsto \partial_n(c_1) \otimes m_1(\lambda) + \cdots + \partial_n(c_k) \otimes m_k(\lambda)) \\ &\overset{\psi_{n-1}}{\rightsquigarrow} (\lambda \mapsto \partial_n(c_1) m_1(\lambda) + \cdots + \partial_n(c_k) m_k(\lambda)) \end{aligned}$$

Therefore, for each λ,

$$a_\lambda = \partial_n(c_1) m_1(\lambda) + \cdots + \partial_n(c_k) m_k(\lambda),$$

contradicting the assumption that the set of all a_λ is not contained in any finitely generated submodule of $\ker \partial_{n-1} = \operatorname{im} \partial_n$. This shows that Γ has to be of type FP_n after all and thus proves that 3 implies 1.

Therefore the first three statements are equivalent and we will now show that the third is equivalent to the fourth. Assume that 3 holds. Since 3 implies 1, Γ is of type FP_1 and hence finitely generated. Furthermore 3 implies that

$$H_k(\Gamma, \prod \mathbb{Z}\Gamma) \cong H_k(\Gamma, \mathbb{Z}\Gamma) = 0, \quad \text{for } 2 \le k < n,$$

because of Lemma A.38.

Conversely, assume that 4 holds. Since Γ is finitely generated and 1 implies 3, we know that the map

$$H_k(\Gamma, \prod \mathbb{Z}\Gamma) \to \prod H_k(\Gamma, \mathbb{Z}\Gamma)$$

is an isomorphism for $k = 0$. For $2 \le k < n$ we assumed that the domain of this map is trivial. But so is the codomain by Lemma A.38, so the map is an isomorphism. □

Type wFP_n **and examples** In Section 6.1 we introduced the property "weak FP_n":

Definition A.40. A group Γ is of type wFP_n if $H_k(\Gamma', \mathbb{Z})$ is finitely generated for all $k \le n$ and all finite index subgroups $\Gamma' \le \Gamma$.

It is of type wFP_∞ if $H_k(\Gamma', \mathbb{Z})$ is finitely generated for all $k \ge 0$ and all finite index subgroups $\Gamma' \le \Gamma$.

In this section we will construct some examples, that distinguish the wFP_n-properties from each other and from the FP_n- and F_n-properties.

For the topological properties we have the following description of the finiteness properties of amalgamated products and HNN extensions:

Proposition A.41. *Let $\Gamma = \Gamma_1 *_C \Gamma_2$ be an amalgamated product of groups.*

1. *Assume Γ_1 and Γ_2 are finitely presented. Then Γ is finitely presented if and only if C is finitely generated.*
2. *Assume Γ_1 and Γ_2 are of type F_n. Then Γ is of type F_n if C is of type F_{n-1}.*

Let $\Gamma = \Gamma_0 *_C$ be an HNN extension.

1. *Assume Γ_0 is finitely presented. Then Γ is finitely presented if and only if C is finitely generated.*
2. *Assume Γ_0 is of type F_n. Then Γ is of type F_n if C is of type F_{n-1}.*

Proof. The first statement is due to Baumslag [Bau62] for amalgamated products. The analogous statement for HNN extensions can be found in [Mil02].

The second statement follows from the construction of classifying complexes given in [SW79] or [Geo08, Section 6.2] (which we describe in Section 5.4 for amalgamated products). □

To give an analogous statement for the wFP_n-properties, we need the following extremely useful Mayer-Vietoris sequences describing the homology of amalgamated free products and HNN extensions. These are [Bie81, Theorems 2.10 & 2.12]:

Theorem A.42. *Let $\Gamma = \Gamma_1 *_C \Gamma_2$ be an amalgamated product and let M be a Γ-module. Then there is a Mayer-Vietoris sequence:*

$$\cdots \to H_n(C,M) \to H_n(\Gamma_1,M) \oplus H_n(\Gamma_2,M) \to H_n(\Gamma,M) \to H_{n-1}(C,M) \to \cdots$$

Let $\Gamma = \Gamma_0 *_C$ be an HNN extension and let M be a Γ-module. Then there is a Mayer-Vietoris sequence:

$$\cdots \to H_n(C,M) \to H_n(\Gamma_0,M) \to H_n(\Gamma,M) \to H_{n-1}(C,M) \to \cdots$$

□

Using these sequences for $M = \prod \mathbb{Z}\Gamma$ and the Bieri-Eckmann Criterion from the previous section it is not hard to derive the following ([Bie81, Proposition 2.13]; alternatively, we could use Theorem 7.3 for the "if"-direction):

Proposition A.43. *Let $\Gamma = \Gamma_1 *_C \Gamma_2$ be an amalgamated product of groups with Γ_1 and Γ_2 of type FP_n. Then Γ is of type FP_n if and only if C is of type FP_{n-1}.*
Let $\Gamma = \Gamma_0 *_C$ be an HNN extension with Γ_0 of type FP_n. Then Γ is of type FP_n if and only if C is of type FP_{n-1}. □

As we mentioned before, Bestvina and Brady in [BB97] constructed examples of groups which are of type FP_n for all n, but not finitely presented. Their examples are normal subgroups of right-angled Artin groups. Denote one of these Bestvina-Brady examples by C and the ambient right-angled Artin group by Γ. Then the amalgamated product $\Gamma *_C \Gamma$ (or the HNN group $\Gamma *_C$) is finitely presented by 1 from Proposition A.41 and of type FP_n for all n by Proposition A.43. Thus it is also of type F_n for all n by Theorem 7.2. But C, by assumption, is not of type F_n for any $n \geq 3$. This shows that the converses of the statements 2 in Proposition A.41 do not hold.

But let us now try to give a statement analogous to Propositions A.41 and A.43 for the wFP_n-properties. It is not hard to see the following:

Lemma A.44. *If $\Gamma = \Gamma_1 *_C \Gamma_2$ is an amalgamated product with Γ_1 and Γ_2 of type* wFP_n, *then $H_n(\Gamma)$ is finitely generated if and only if $H_{n-1}(C)$ is.*
If $\Gamma = \Gamma_0 *_C$ is an HNN extension with Γ_0 of type wFP_n *then $H_n(\Gamma)$ is finitely generated if and only if $H_{n-1}(C)$ is.*

Proof. This follows directly from the Mayer-Vietoris sequences, in which the boundary maps

$$\cdots \to H_n(\Gamma_1) \oplus H_n(\Gamma_2) \to H_n(\Gamma) \xrightarrow{\partial} H_{n-1}(C) \to H_{n-1}(\Gamma_1) \oplus H_{n-1}(\Gamma_2) \to \cdots$$

and

$$\cdots \to H_n(\Gamma_0) \to H_n(\Gamma) \xrightarrow{\partial} H_{n-1}(C) \to H_{n-1}(\Gamma_0) \to \cdots$$

have finitely generated kernel and cokernel. □

Proposition A.45. *If $\Gamma = \Gamma_1 *_C \Gamma_1$ is an amalgamated product with Γ_1 and Γ_2 of type* wFP_n, *then Γ is of type* wFP_n *if and only if C is of type* wFP_{n-1}.
If $\Gamma = \Gamma_0 *_C$ is an HNN extension with Γ_1 and Γ_2 of type wFP_n, *then Γ is of type* wFP_n *if and only if C is of type* wFP_{n-1}.

Proof. We consider first the statement about amalgamated products.

Let $\Gamma' \leq \Gamma$ be a finite index subgroup. Let T be the Bass-Serre tree associated to the splitting $\Gamma_1 *_C \Gamma_2$. Then Γ' acts on T with vertex stabilisers isomorphic to finite index subgroups of Γ_1 or Γ_2 and edge stabilisers isomorphic to finite index subgroups of C. Therefore the vertex stabilisers are certainly of type wFP_n.

Now consider first the case where C is of type wFP_{n-1}. Then the edge stabilisers of the action of Γ' on T are of type wFP_{n-1}. Since Γ acts cocompactly on T and Γ' is of finite index, Γ' acts cocompactly on T as well, so Γ' is the fundamental group of a finite graph of groups. More precisely, Γ' is inductively built up from groups of type wFP_n (the vertex groups) by a finite sequence of amalgamated products and HNN extensions along subgroups of type wFP_{n-1} (the edge groups). Thus it follows from Lemma A.44 that $H_k(\Gamma')$ is finitely generated for any $k \leq n$. Since this holds for any finite index subgroup Γ', we can conclude that Γ is of type wFP_n.

Conversely, if Γ is not of type wFP_n, then there is a finite index subgroup $\Gamma' \leq \Gamma$ with $H_k(\Gamma')$ not finitely generated for some $k \leq n$. By the same argument as before and using the other implication of Lemma A.44, this shows that there must be some edge stabiliser C' such that $H_{k-1}(C')$ is not finitely generated. Since C' has finite index in C, this implies that C is not of type wFP_{n-1}.

The proof for the statement about HNN extensions is entirely analogous. □

Using a method from [Bri99] (or [BH99, Section III.Γ.5]), the theorems on amalgamated products can be used to "shift up" finiteness properties in dimension. So, for example, if C is a group of type wFP_{n-1} but not of type wFP_n, we could try to embed it in a group Γ of, say, type F_∞; then $\Gamma *_C \Gamma$ will be of type wFP_n but not wFP_{n+1}.

To make this work, we need a way of finding the required embeddings. For that purpose, the following lemma, which is [Bri99, Lemma 2.8], is very useful:

Lemma A.46. *Let Γ be a group and $N \triangleleft \Gamma$ a normal subgroup such that $\Gamma/N \cong \mathbb{Z}$. Then there is a short exact sequence*

$$\Gamma *_N \Gamma \hookrightarrow \Gamma \times (\mathbb{Z} * \mathbb{Z}) \twoheadrightarrow \mathbb{Z} \qquad \square$$

Let us show how this can be used (as in [Bri99]) to construct the Stallings-Bieri examples of groups of type F_n with non-finitely generated $(n+1)^{\mathrm{st}}$ integral homology group (see [Sta63, Bie76]).

Let $\pi : F_2 \to \mathbb{Z}$ be the epimorphism from the free group on two generators to $\mathbb{Z}$ which sends each generator to $1 \in \mathbb{Z}$. Then $\mathrm{SB}_0 := \ker \pi$ is a normal, infinite index subgroup of a free group and therefore free and not finitely generated. Therefore its abelianisation $H_1(\mathrm{SB}_0)$ is a non-finitely generated free abelian group and so SB_0 is not of type wFP_1. From Lemma A.46 it then follows that there is a short exact sequence

$$F_2 *_{\mathrm{SB}_0} F_2 \hookrightarrow F_2 \times F_2 \twoheadrightarrow \mathbb{Z}.$$

Denote the kernel by SB_1. Then SB_1 is finitely generated but not of type wFP_2, by Proposition A.45. Since $\mathrm{SB}_1 \triangleleft F_2 \times F_2$ we can apply Lemma A.46 again. Continuing this inductively, we obtain:

Proposition A.47. *For each n there is a group SB_n, which is a kernel of a map from $(F_2)^n$ to $\mathbb{Z}$ and is F_{n-1} but not wFP_n.* $\square$

With a more sophisticated choice of base group we can use the same argument to show that wFP_n does not imply FP_n for any $n \geq 1$. The construction we will describe is essentially the one given in [LS06].

We need a lemma first.

Lemma A.48. *A group Γ is of type wFP_n if $H_k(N)$ is finitely generated for every* normal *finite index subgroup N and $k \leq n$.*

Proof. Let $\Gamma' \leq \Gamma$ be an arbitrary finite index subgroup of Γ. Then there is a subgroup $N \leq \Gamma'$ of finite index which is normal in Γ. So Γ' fits into a short exact sequence

$$N \hookrightarrow \Gamma' \twoheadrightarrow F$$

where F is finite. Then F is of type FP_n (by Proposition A.34) and it follows by the same argument as in the proof[8] of Lemma 6.5 that $H_k(\Gamma)$ is finitely generated for $k \leq n$. $\square$

[8] in fact, the proof works verbatim if the assumption that N be of type wFP_{n-1} is replaced by the assumption that $H_k(N)$ is finitely generated for all $k \leq n-1$

Proposition A.49. *There is a sequence of groups K_n such that K_n is of type* wFP_∞ *and* F_{n-1} *but not of type* FP_n *and such that K_n fits into a short exact sequence*

$$K_n \hookrightarrow \Gamma_n \twoheadrightarrow \mathbb{Z}$$

with Γ_n of type F_∞.

Proof. Let S be an infinite group of type F_∞ with no finite quotients and $H_k(S) = 0$ for all $k \geq 1$ (such groups are constructed, for example, in [BG04, Chapter 4]). Now take countably many isomorphic copies $S_m \cong S$ with $m \in \mathbb{Z}$ and consider their free product $K_1 = *_{m\in\mathbb{Z}} S_m$. Then K_1 is not finitely generated but still does not have finite quotients (for any homomorphism $K_1 \to \Gamma$ is determined by the compositions $S_m \hookrightarrow K \to \Gamma$, and if Γ is finite, all these compositions must be trivial) and thus no finite index normal subgroups. A Mayer-Vietoris argument shows that $H_k(K_1) = 0$ for all $k \geq 1$, so by Lemma A.48, K_1 is of type wFP_∞. Furthermore, observe that K_1 embeds in $\Gamma_1 := S * \langle t \mid - \rangle$ by sending S_m to $t^{-m} S t^m$. In fact K_1 is the kernel of the map $S * \mathbb{Z} \to \mathbb{Z}$ mapping S trivially and $\mathbb{Z}$ isomorphically. Note that Γ_1 is of type F_∞. Thus the short exact sequence

$$K_1 \hookrightarrow \Gamma_1 \twoheadrightarrow \mathbb{Z}$$

supplies the base case $n = 1$ of the Proposition.

Now suppose groups K_n and Γ_n with the claimed properties have been constructed. Then the double $K_{n+1} := \Gamma_n *_{K_n} \Gamma_n$ is still of type wFP_∞ (by Proposition A.45) and of type F_n but not of type FP_{n+1} (by Proposition A.43). Furthermore, by Lemma A.46, K_{n+1} fits into a short exact sequence

$$K_{n+1} \hookrightarrow \Gamma_n \times (\mathbb{Z} * \mathbb{Z}) \twoheadrightarrow \mathbb{Z},$$

where $\Gamma_n \times (\mathbb{Z} * \mathbb{Z}) =: \Gamma_{n+1}$ is of type F_∞. □

Bibliography

[Abe79] H. Abels. An example of a finitely presented solvable group. In *Homological group theory (Proc. Sympos., Durham, 1977)*, volume 36 of *London Math. Soc. Lecture Note Ser.*, pages 205–211. Cambridge Univ. Press, Cambridge, 1979.

[Bau62] G. Baumslag. A remark on generalized free products. *Proc. Amer. Math. Soc.*, 13:53–54, 1962.

[BB97] M. Bestvina and N. Brady. Morse theory and finiteness properties of groups. *Invent. Math.*, 129(3):445–470, 1997.

[BBMS00] G. Baumslag, M. R. Bridson, C. F. Miller, III, and H. Short. Fibre products, non-positive curvature, and decision problems. *Comment. Math. Helv.*, 75(3):457–477, 2000.

[BE74] R. Bieri and B. Eckmann. Finiteness properties of duality groups. *Comment. Math. Helv.*, 49:74–83, 1974.

[BG95] C. J. B. Brookes and J. R. J. Groves. Modules over nilpotent group rings. *J. London Math. Soc. (2)*, 52(3):467–481, 1995.

[BG99] K.-U. Bux and C. Gonzalez. The Bestvina-Brady construction revisited: Geometric computation of Σ-invariants for right-angled Artin groups. *J. London Math. Soc. (2)*, 60(3):793–801, 1999.

[BG00] C. J. B. Brookes and J. R. J. Groves. Modules over crossed products of a division ring by a free abelian group. I. *J. Algebra*, 229(1):25–54, 2000.

[BG03] R. Bieri and R. Geoghegan. Connectivity properties of group actions on non-positively curved spaces. *Mem. Amer. Math. Soc.*, 161(765):xiv+83, 2003.

[BG04] M. R. Bridson and F. J. Grunewald. Grothendieck's problems concerning profinite completions and representations of groups. *Ann. of Math. (2)*, 160(1):359–373, 2004.

[BH82a] R. Brown and J. Huebschmann. Identities among relations. In *Low-dimensional topology (Bangor, 1979)*, volume 48 of *London Math. Soc. Lecture Note Ser.*, pages 153–202. Cambridge Univ. Press, Cambridge, 1982.

[BH82b] R. Brown and P. J. Higgins. Crossed complexes and nonabelian extensions. In *Category Theory (Gummersbach, 1981)*, volume 962 of *Lecture Notes in Math.*, pages 39–50. Springer, Berlin, 1982.

[BH99] M. R. Bridson and A. Haefliger. *Metric spaces of non-positive curvature*, volume 319 of *Grundlehren der Mathematischen Wissenschaften*. Springer-Verlag, Berlin, 1999.

[BH07a] M. R. Bridson and J. Howie. Subgroups of direct products of elementarily free groups. *Geom. Funct. Anal.*, 17(2):385–403, 2007.

[BH07b] M. R. Bridson and J. Howie. Subgroups of direct products of two limit groups. *Math. Res. Lett.*, 14(4):547–558, 2007.

[BHMS02] M. R. Bridson, J. Howie, C. F. Miller, III, and H. Short. The subgroups of direct products of surface groups. *Geom. Dedicata*, 92:95–103, 2002. Dedicated to John Stallings on the occasion of his 65th birthday.

[BHMS08] M. R. Bridson, J. Howie, C. F. Miller, III, and H. Short. Finitely presented residually free groups. arXiv:0809.3704, 2008.

[BHMS09] M. R. Bridson, J. Howie, C. F. Miller, III, and H. Short. Subgroups of direct products of limit groups. *Ann. of Math. (2)*, 170(3):1447–1467, 2009.

[BHMS12] M. R. Bridson, J. Howie, C. F. Miller, III, and H. Short. On the finite presentation of subdirect products and the nature of residually free groups. To appear in Amer. J. Math, arXiv:1205.6175, 2012.

[BHP98] Y. G. Baik, J. Harlander, and S. J. Pride. The geometry of group extensions. *J. Group Theory*, 1(4):395–416, 1998.

[BHS11] R. Brown, P. J. Higgins, and R. Sivera. *Nonabelian Algebraic Topology*, volume 15 of *EMS Tracts in Mathematics*. European Mathematical Society (EMS), Zürich, 2011. Filtered spaces, crossed complexes, cubical homotopy groupoids, With contributions by Christopher D. Wensley and Sergei V. Soloviev.

[Bie76] R. Bieri. Normal subgroups in duality groups and in groups of cohomological dimension 2. *J. Pure Appl. Algebra*, 7(1):35–51, 1976.

[Bie81] R. Bieri. *Homological Dimension of Discrete Groups*. Queen Mary College Mathematical Notes. Queen Mary College Department of Pure Mathematics, London, second edition, 1981.

[Bie93] R. Bieri. The geometric invariants of a group. A survey with emphasis on the homotopical approach. In *Geometric Group Theory, Vol. 1 (Sussex, 1991)*, volume 181 of *London Math. Soc. Lecture Note Ser.*, pages 24–36. Cambridge Univ. Press, Cambridge, 1993.

[Bie99] R. Bieri. Finiteness length and connectivity length for groups. In *Geometric group theory down under (Canberra, 1996)*, pages 9–22. de Gruyter, Berlin, 1999.

[BM09] M. R. Bridson and C. F. Miller, III. Structure and finiteness properties of subdirect products of groups. *Proc. Lond. Math. Soc. (3)*, 98(3):631–651, 2009.

[BNS87] R. Bieri, W. D. Neumann, and R. Strebel. A geometric invariant of discrete groups. *Invent. Math.*, 90(3):451–477, 1987.

[BR84] G. Baumslag and J. E. Roseblade. Subgroups of direct products of free groups. *J. London Math. Soc. (2)*, 30(1):44–52, 1984.

[BR88] R. Bieri and B. Renz. Valuations on free resolutions and higher geometric invariants of groups. *Comment. Math. Helv.*, 63(3):464–497, 1988.

[Bri99] M. R. Bridson. Doubles, finiteness properties of groups, and quadratic isoperimetric inequalities. *J. Algebra*, 214(2):652–667, 1999.

[Bro74] K. S. Brown. Euler characteristics of discrete groups and G-spaces. *Invent. Math.*, 27:229–264, 1974.

[Bro75] K. S. Brown. Euler characteristics of groups: the p-fractional part. *Invent. Math.*, 29(1):1–5, 1975.

[Bro80] R. Brown. On the second relative homotopy group of an adjunction space: an exposition of a theorem of J. H. C. Whitehead. *J. London Math. Soc. (2)*, 22(1):146–152, 1980.

[Bro82] K. S. Brown. Complete Euler characteristics and fixed-point theory. *J. Pure Appl. Algebra*, 24(2):103–121, 1982.

[Bro92] R. A. Brown. Generalized group presentation and formal deformations of CW complexes. *Trans. Amer. Math. Soc.*, 334(2):519–549, 1992.

[Bro94] K. S. Brown. *Cohomology of Groups*, volume 87 of *Graduate Texts in Mathematics*. Springer-Verlag, New York, 1994. Corrected reprint of the 1982 original.

[BRS99] R. Brown and A. Razak Salleh. Free crossed resolutions of groups and presentations of modules of identities among relations. *LMS J. Comput. Math.*, 2:28–61, 1999.

[BS] R. Bieri and R. Strebel. Geometric Invariants of Discrete Groups. Unpublished monograph.

[BS80] R. Bieri and R. Strebel. Valuations and finitely presented metabelian groups. *Proc. London Math. Soc. (3)*, 41(3):439–464, 1980.

[CBW] O. Cotton-Barratt and H. Wilton. Conjugacy separability of 1-acylindrical graphs of free groups. To appear in Math. Z., arXiv:0906.0101.

[Cor92] J. M. Corson. Complexes of groups. *Proc. London Math. Soc. (3)*, 65(1):199–224, 1992.

[CS98] J. M. Corson and B. S. Spieler. The qsf property and group actions via complexes of spaces. *Math. Z.*, 228(3):589–603, 1998.

[Dis08] W. Dison. *Isoperimetric functions for subdirect products and Bestvina-Brady groups.* PhD thesis, Imperial College London, 2008.

[EW05] G. Ellis and G. Williams. On the cohomology of generalized triangle groups. *Comment. Math. Helv.*, 80(3):571–591, 2005.

[Geh98] R. Gehrke. The higher geometric invariants for groups with sufficient commutativity. *Comm. Algebra*, 26(4):1097–1115, 1998.

[Geo08] R. Geoghegan. *Topological Methods in Group Theory*, volume 243 of *Graduate Texts in Mathematics*. Springer, New York, 2008.

[Hae91] A. Haefliger. Complexes of groups and orbihedra. In *Group theory from a geometrical viewpoint (Trieste, 1990)*, pages 504–540. World Sci. Publ., River Edge, NJ, 1991.

[Hue80] J. Huebschmann. Crossed n-fold extensions of groups and cohomology. *Comment. Math. Helv.*, 55(2):302–313, 1980.

[Koc10] D. H. Kochloukova. On subdirect products of type FP_m of limit groups. *J. Group Theory*, 13(1):1–19, 2010.

[Lod00] J.-L. Loday. Homotopical syzygies. In *Une Dégustation Topologique: Homotopy Theory in the Swiss Alps (Arolla, 1999)*, volume 265 of *Contemp. Math.*, pages 99–127. Amer. Math. Soc., Providence, RI, 2000.

[LS06] I. J. Leary and M. Saadetoğlu. Some groups of finite homological type. *Geom. Dedicata*, 119:113–120, 2006.

[Mei90] H. Meinert. Über die höheren geometrischen Invarianten für endliche Produkte von Gruppen. Diploma thesis, Johann Wolfgang Goethe-Universität, Frankfurt am Main, 1990.

[Mei94] H. Meinert. The geometric invariants of direct products of virtually free groups. *Comment. Math. Helv.*, 69(1):39–48, 1994.

[Mei97] H. Meinert. Actions on 2-complexes and the homotopical invariant Σ^2 of a group. *J. Pure Appl. Algebra*, 119(3):297–317, 1997.

[Mih66] K. A. Mihaĭlova. The occurrence problem for direct products of groups. *Mat. Sb. (N.S.)*, 70 (112):241–251, 1966.

[Mil92] C. F. Miller, III. Decision problems for groups—survey and reflections. In *Algorithms and classification in combinatorial group theory (Berkeley, CA, 1989)*, volume 23 of *Math. Sci. Res. Inst. Publ.*, pages 1–59. Springer, New York, 1992.

[Mil02] C. F. Miller, III. Subgroups of direct products with a free group. *Q. J. Math.*, 53(4):503–506, 2002.

[MMV98] J. Meier, H. Meinert, and L. VanWyk. Higher generation subgroup sets and the Σ-invariants of graph groups. *Comment. Math. Helv.*, 73(1):22–44, 1998.

[Pei49] R. Peiffer. Über Identitäten zwischen Relationen. *Math. Ann.*, 121:67–99, 1949.

[Rat80] J. G. Ratcliffe. Free and projective crossed modules. *J. London Math. Soc. (2)*, 22(1):66–74, 1980.

[Rei49] K. Reidemeister. Über Identitäten von Relationen. *Abh. Math. Sem. Univ. Hamburg*, 16:114–118, 1949.

[Ren88] B. Renz. *Geometrische Invarianten und Endlichkeitseigenschaften von Gruppen*. PhD thesis, Johann Wolfgang Goethe-Universität, Frankfurt am Main, 1988.

[Ren89] B. Renz. Geometric invariants and HNN-extensions. In *Group Theory (Singapore, 1987)*, pages 465–484. de Gruyter, Berlin, 1989.

[Rip82] E. Rips. Subgroups of small cancellation groups. *Bull. London Math. Soc.*, 14(1):45–47, 1982.

[Ser71] J.-P. Serre. Cohomologie des groupes discrets. In *Prospects in mathematics (Proc. Sympos., Princeton Univ., Princeton, N.J., 1970)*, pages 77–169. Ann. of Math. Studies, No. 70. Princeton Univ. Press, Princeton, N.J., 1971.

[Sta63] J. Stallings. A finitely presented group whose 3-dimensional integral homology is not finitely generated. *Amer. J. Math.*, 85:541–543, 1963.

[Sta94] C. W. Stark. Homologically trivial actions on chain complexes and finiteness properties of infinite groups. *J. Pure Appl. Algebra*, 94(2):217–228, 1994.

[Sta98] C. W. Stark. Approximate finiteness properties of infinite groups. *Topology*, 37(5):1073–1088, 1998.

[SW79] P. Scott and T. Wall. Topological methods in group theory. In *Homological group theory (Proc. Sympos., Durham, 1977)*, volume 36 of *London Math. Soc. Lecture Note Ser.*, pages 137–203. Cambridge Univ. Press, Cambridge, 1979.

[Wal61] C. T. C. Wall. Resolutions for extensions of groups. *Proc. Cambridge Philos. Soc.*, 57:251–255, 1961.

[Wal65] C. T. C. Wall. Finiteness conditions for CW-complexes. *Ann. of Math. (2)*, 81:56–69, 1965.

[Wal66] C. T. C. Wall. Finiteness conditions for CW complexes. II. *Proc. Roy. Soc. Ser. A*, 295:129–139, 1966.

[Whi41] J. H. C. Whitehead. On adding relations to homotopy groups. *Ann. of Math. (2)*, 42:409–428, 1941.

[Whi46] J. H. C. Whitehead. Note on a previous paper entitled "On adding relations to homotopy groups.". *Ann. of Math. (2)*, 47:806–810, 1946.

[Whi49] J. H. C. Whitehead. Combinatorial homotopy. II. *Bull. Amer. Math. Soc.*, 55:453–496, 1949.

[Whi78] G. W. Whitehead. *Elements of Homotopy Theory*, volume 61 of *Graduate Texts in Mathematics*. Springer-Verlag, New York, 1978.

[Wis03] D. T. Wise. A residually finite version of Rips's construction. *Bull. London Math. Soc.*, 35(1):23–29, 2003.