

JOURNAL OF THE ASSOCIATION FOR MATHEMATICAL RESEARCH



VOLUME 2 : ISSUE 2

2024

JOURNAL OF THE ASSOCIATION FOR MATHEMATICAL RESEARCH

Managing Editor

Alex Kontorovich, Rutgers University (USA)

Editorial Board

Scott Aaronson, UT Austin (USA)

Karim Adiprasito, CNRS/Jussieu (France)

Nalini Anantharaman, Collège de France (France)

Mikhail Belolipetsky, IMPA (Brazil)

Richard Borcherds, UC Berkeley (USA)

Gunnar Carlsson, Stanford University (USA)

Maria Chudnovsky, Princeton University (USA)

Henri Darmon, McGill University (Canada)

Bill Duke, UCLA (USA)

Yakov Eliashberg, Stanford University (USA)

Anna Felikson, Durham University (UK)

Alireza Salehi Golsefidy, UC San Diego (USA)

Mark Goresky, IAS Princeton (USA)

Federico Rodriguez Hertz, Penn State University (USA)

Svetlana Jitomirskaya, UC Berkeley (USA)

Erez Lapid, Weizmann Institute of Science (Israel)

Misha Lyubich, Stony Brook University (USA)

Eva Miranda, Universitat Politècnica de Catalunya, and CRM (Spain)

Peter Ozsvath, Princeton University (USA)

Leonid Polterovich, Tel Aviv University (Israel)

Alexander Razborov, University of Chicago (USA) and Steklov Institute (Russia)

Regina Rotman, University of Toronto (Canada)

Zeev Rudnick, Tel Aviv University (Israel)

Mikhail Sodin, Tel Aviv University (Israel)

Benny Sudakov, ETH, Zurich (Switzerland)

Gunther Uhlmann, University of Washington (USA)

Tykal Venkataramana, Tata Institute (India)

Claire Voisin, IMJ-PRG (France)

Dani Wise, McGill University (Canada)

Maciej Zworski, UC Berkeley (USA)

PUBLISHED BY

Association for Mathematical Research, amathr.org,
Davis, CA; Jenkintown PA.

Generating the Goeritz group of S^3

Martin Scharlemann

Received 29 Nov 2022; Revised 15 Feb 2024; Accepted 13 Mar 2024

Abstract: In 1980 J. Powell [Po] proposed that five specific elements sufficed to generate the Goeritz group for any genus Heegaard splitting of S^3 . Here we prove that a natural expansion of Powell’s proposed generators, to include all eyeglass twists and all topological conjugates of Powell’s generators, does suffice.

Key words and phrases: Geometric Structures on 3-Manifolds, Goeritz group

1 Introduction

Suppose that M is a closed orientable 3-manifold and $M = A \cup_T B$ is a Heegaard splitting of M . Following [JM] the *Goeritz group* $G(M, T)$ is the group of isotopy classes of diffeomorphisms $(M, T) \rightarrow (M, T)$ for which the induced diffeomorphism $M \rightarrow M$ is isotopic to the identity. An element of the Goeritz group can be viewed as the final result of a (possibly non-unique) loop T_θ , $0 \leq \theta \leq 2\pi$ of embeddings of T in M , that is an element in $\pi_1(\text{Diff}(M)/\text{Diff}(M, T))$ [JM, Theorem 1]. That is the viewpoint we will take.

Little is known about the Goeritz group, even in the case that $M = S^3$ and the Heegaard splittings are fully described [Wa]. In [Po] J. Powell proposed (indeed he believed he had proven) that five specific isotopies generate the Goeritz group $G(S^3, T)$ of the 3-sphere. (In [Sc2] one is found to be redundant, and so the number is reduced to four.) Figures 1 to 4 give important examples of Goeritz elements acting on the standard Heegaard splitting of S^3 .

Powell’s conjecture has been verified for genus ≤ 3 splittings of S^3 [FS1], but even the question of whether $G(S^3, T)$ is finitely generated remains open when $\text{genus}(T) \geq 4$. Here we define a generalization, shown in Figure 4 and called an *eyeglass twist*, of Powell’s proposed generator D_θ . We ultimately show (Corollary 18.4) that if D_θ is replaced by the collection of all eyeglass twists, and we include also all topological conjugates of the other three Powell generators (i. e. conjugates by elements of $G(S^3, T)$), the subgroup $\mathcal{E} \subset G(S^3, T)$ thereby generated is all of $G(S^3, T)$.

In [Sc3] we show that this leads to the following observation: Suppose T is the standard genus $g + 1$ Heegaard surface in S^3 . Any element of $G(S^3, T)$ that acts trivially on a standard genus 1 summand of T is a consequence of the Powell generators acting on T . In other words, the Powell Conjecture is true, stably.

The strategy for the proof of Corollary 18.4 is rather simple: Suppose T is the standard genus g Heegaard splitting of S^3 , and $\tau \in G(S^3, T)$. In [FS1] it is shown that there is a ‘cycle of weak reductions’ capturing τ . That is, given a loop of embeddings $T_\theta, 0 \leq \theta \leq 2\pi$ of T in S^3 that represents τ , there is a natural topological way to extract, for generic θ , a pair of properly embedded disjoint essential disks $a_\theta \subset A, b_\theta \subset B$ that thereby weakly reduce T_θ . Moreover, at the finite number of non-generic points, the chosen weakly reducing pair does not change much. Since the method of choosing the pair (a_θ, b_θ) is topological, it follows relatively easily that $(a_{2\pi}, b_{2\pi}) = (\tau(a_0), \tau(b_0))$. A landmark result of Casson and Gordon [CG] shows that a weakly reducing pair gives rise to a reducing sphere for T ; one might naturally hope that one could track such a reducing sphere K around θ , as was done for the weakly reducing pairs, and thereby be able to meaningfully compare K with $\tau(K)$ and so understand the action of τ .

Sadly, the transition from a weakly reducing pair (a, b) to a reducing sphere involves much choice, so there is no naturally derived reducing sphere K for T as hoped for above. The program here is to find one, via this natural topological method: Let $F \subset S^3$ be the surface obtained from T by weakly reducing along the pair (a, b) . There is a natural and oft-used way to sweep out S^3 by level 2-spheres S_s and, given F , a natural value of s to pick: a value so that the genus of the part of F lying below S_s matches the genus of the part above. Could S_s be turned into a viable, topologically defined reducing sphere for T that could play the role of K above?

It turns out that although S_s itself may not play that role, it does give a recipe for weak reduction that is robust: it doesn’t change abruptly as T moves through S^3 . And as the weak reduction proceeds (as guided by S_s) and F becomes more complicated, its complementary components in S^3 (called the *chambers*) are more likely to contain reducible components. Reducing spheres in these reducible components naturally cut off summands of the original Heegaard surface, and these (inductively) determine an isotopy to the standard picture that is unique, up to action by $\mathcal{E}$. If no complementary component becomes reducible, then at the end of the process we fall back on S_s as the required reducing sphere for T (see the second bullet in the statement of Proposition 10.2).

One can think of the progression Proposition 3.7 $\rightarrow$ Corollary 3.9 $\rightarrow$ Corollary 8.4 $\rightarrow$ Corollary 8.9 $\rightarrow$ Corollary 17.4 $\rightarrow$ Corollary 18.4 as guideposts for the argument.

Although the program is easy to describe, the technical difficulties encountered below are complex since, in effect, they will involve 4 dimensions of sweep-outs. Much of the machinery is new, and then so is the terminology. We call the attention of the reader to the Index at the end of the paper for a guide to this new terminology. A plausible strategy for a proof is presented in Section 3, including a rough overview of the final proof in Subsection 3.3. Some of the argument is not restricted to $M = S^3$, so perhaps some of the methodology can be useful in understanding stabilized Heegaard splittings of other 3-manifolds as well.

General Remarks: All manifolds will be orientable and, unless obviously not, compact. We mostly work in the TOP category (locally flat embeddings, homeomorphisms) to avoid discussion of corner rounding, etc. But at several points (see especially Section 17) we need results from smooth topology. At those points we will work in the DIFF category. In these dimensions (two and three) the difference

in categories is immaterial, see for example [Moi]. Unexplained notation in the statement of a Lemma, Proposition, etc. may well be found in the discussion that immediately precedes it.

2 The eyeglass subgroup

Powell's proposed generators are expressed in terms of *standard genus 1 summands* in a canonical picture of a genus g Heegaard surface in S^3 . In [FS1, Section 2] some consequences of these Powell generators are presented that include the Powell generators themselves, but also include more general moves on the standard genus 1 summands. These more general moves would make sense on any stabilized Heegaard splitting of any closed orientable manifold, if we drop the requirement that the summands are 'standard', which is not meaningful in the context of arbitrary stabilized splittings. Below is a brief description of those generalized moves; more detail can be found in [FS1]. In addition, note that the last move described (an eyeglass twist) is a significant generalization of Powell's move D_θ .

For a Heegaard splitting $M = A \cup_T B$ define a *bubble* for T to be a 3-ball b in M whose boundary intersects T in a single essential circle. Via Waldhausen [Wa] we know that any bubble is the boundary sum of a collection of genus 1 bubbles.

1. A *bubble move* is an isotopy of b along a closed path in $T - b$ that begins and ends at b , returning $(b, b \cap T)$ to itself. See Figure 1. Note that if the closed path is null-homotopic in M , which is automatic if $M = S^3$, the bubble move is isotopic to the identity on M , so it represents an element of $G(M, T)$.




Figure 1: A bubble move

2. Let b be a genus 1 bubble. A *flip* is the homeomorphism $(b, b \cap T) \rightarrow (b, b \cap T)$ that reverses orientation of both the meridian and longitude of the summand, as shown in Figure 2. Powell's [Po] label for a flip on the first standard summand is D_ω . In a genus 1 splitting, where T is a torus, regard the hyperelliptic involution $(S^3, T) \rightarrow (S^3, T)$ as (a degenerate case of) a flip.
3. Let b_1, b_2 be disjoint genus 1 bubbles, and let $v \subset T - (b_1 \cup b_2)$ be an arc connecting them. Let b be the genus 2 bubble obtained by tubing together the genus 1 bubbles along v . A *bubble exchange* exchanges the two genus 1 bubbles within B , as shown in Figure 3. Powell's label for a bubble exchange on the first two standard summands is $D_{\eta_{12}}$.




Figure 2: A flip




Figure 3: A bubble exchange

4. An *eyeglass* is the union of two disks, ℓ_a, ℓ_b (the *lenses*) with an arc v (the *bridge*) connecting their boundaries. Suppose an eyeglass η is embedded in M so that the 1-skeleton of η (called the *frame*) lies in T , one lens is properly embedded in A , and the other lens is properly embedded in B . The embedded η defines a natural automorphism $(M, T) \rightarrow (M, T)$, as illustrated in Figure 4, called an *eyeglass twist*. Powell's generator D_θ is an example, but topologically special because the lenses are primitive disks in each handlebody.




Figure 4: An eyeglass twist

Let $\text{Mod}(M, T)$ be the group of path components of $\text{Diff}(M, T)$; it is the Goeritz group $G(M, T)$ when, as is the case with S^3 , $\text{Mod}(M)$ is trivial [JM].

Definition 2.1. Suppose $M = A \cup_T B$ is a Heegaard splitting. The subgroup of $\text{Mod}(M, T)$ generated by the four types of moves just described is called the eyeglass group $\mathcal{E} \subset \text{Mod}(M, T)$. Any finite composition of these eyeglass generators will be called an eyeglass move.

When $M = S^3$ call the subgroup of $G(S^3, T) = \text{Mod}(S^3, T)$ generated by Powell's proposed generators the Powell group. Any finite composition of these generators will be called a Powell move.

Remarks: It is shown in [FS1, Section 2] that the Powell group can also be described as the subgroup of $G(S^3, T)$ generated by D_θ together with those flips, bubble moves, and bubble exchanges that act on the standard genus 1 bubbles, not on arbitrary genus 1 bubbles.

It follows that any Powell move (indeed any topological conjugate of a Powell move) is an eyeglass move. Note also that whereas the Powell group is not known to be normal in $G(S^3, T)$, the eyeglass group is normal, since any topological conjugate of a generator is a generator.

The Powell Conjecture is that the Powell group is the entire Goeritz group $G(S^3, T)$; we will eventually show the weaker result that the eyeglass group is the entire Goeritz group. If one could show that each eyeglass twist and any topological conjugate of a Powell generator is in the Powell group, the Powell Conjecture would follow. This seems unlikely.

Let $M = A \cup_T B$ be a Heegaard splitting; we briefly review terminology, and make some elementary observations:

Definition 2.2. A sphere S in M is aligned with the Heegaard splitting if $S \cap T$ is at most one circle. Similarly, a properly embedded disk $(D, \partial D) \subset (M, \partial M)$ is aligned with the splitting if $D \cap T$ is at most one circle, and, if it is one circle, the annulus component of $D - T$ is a spanning annulus in the compression body in which it lies. See [Sc1], [FS2].

For example, the boundary of a bubble $\mathfrak{b}$ is an aligned sphere.

Suppose disjoint spheres S and S' are aligned, and α is a properly embedded arc in T so that α intersects S exactly in one end of α and intersects S' only in the other end of α . A thin tubular neighborhood $N \cong D^2 \times I$ of α will intersect $S \cup S'$ in the two disks $D^2 \times \{0, 1\}$. Delete those disks from $S \cup S'$ and glue on the annulus $\partial D^2 \times I$.

Definition 2.3. Call the resulting aligned sphere the tube sum of S and S' along α .

The same construction on S and a disjoint aligned disk D results in another aligned disk, called the tube sum of D and S .

If $\mathfrak{b}$ and $\mathfrak{b}'$ are disjoint bubbles, say of genus p and q then the tube sum of $\partial \mathfrak{b}$ and $\partial \mathfrak{b}'$ naturally bound a genus $p + q$ bubble, called the tube sum of $\mathfrak{b}$ and $\mathfrak{b}'$.

Since a bubble move is an eyeglass move, the tube sum of a sphere or disk with a bubble, and, in particular, the tube sum of two bubbles, is well-defined up to eyeglass moves; it does not depend on the choice of α .

Definition 2.4. For D an aligned disk in M and $\mathfrak{b}$ a disjoint bubble, the tube sum of D and $\mathfrak{b}$ is an aligned disk D' properly isotopic to D in M . Replacing D with D' is called a bubble pass of $\mathfrak{b}$ through D .

Lemma 2.5. Suppose D is an aligned non-separating disk in M , $\mathfrak{b}$ is a disjoint bubble, and D' is the disk obtained from D by a bubble pass of $\mathfrak{b}$ through D . Then there is an eyeglass move which isotopes D to D' in M .

Proof. Let α be the arc between D and $\partial\mathfrak{b}$ along which the bubble pass is made, and let D' be the resulting aligned disk. Since D is assumed non-separating, there is another arc $\beta \subset T$ disjoint from α so that $\partial\beta = \partial\alpha$ but the end of β at D is on the opposite side of D as the end of α at D . Thus the union $\gamma = \alpha \cup \beta$ is a path in $T - \mathfrak{b}$ that passes exactly once through D . Then a bubble move of $\mathfrak{b}$ via γ will isotope D to D' as required. (See Figure 5. The bubble move of $\mathfrak{b}$ to itself that carries D to D' is along the concatenation loop $\beta\alpha$.) $\square$




Figure 5: A bubble move of $\mathfrak{b}$ to itself that carries D to D'

3 The search for a plausible strategy

The proof that $G(S^3, T) = \mathcal{E}$ will be by induction on the genus of T ; it is known for genus ≤ 3 by [FS1]. In this section we lay some groundwork and describe (see Corollary 3.12) a special circumstance that makes the inductive step straightforward. The goal of the rest of the long and technical argument in Sections 4 through 17 is simply to show that conditions sufficiently similar to this special circumstance always arise.

3.1 A technical lemma

The following technical lemma is needed almost immediately for an inductive step:

Let (S^3, T) be a genus $g \geq 1$ Heegaard splitting and B_p be a small ball around a point $p \in T$ that intersects T in a single disk D_p . Let B_- be the 3-ball that remains when $\text{int}(B_p)$ is removed, and $T_- = T \cap B_- = T - \text{int}(D_p)$. T_- is once-punctured genus g surface properly embedded in B_- . Let $\text{Mod}(B_-, T_-)$ be the group of path components of $\text{Diff}(B_-, T_-)$. (We do not require that a diffeomorphism $f : (B_-, T_-) \rightarrow (B_-, T_-) \in \text{Diff}(B_-, T_-)$ be the identity on $\partial B_p = \partial B_-$.)

Echoing the definition of bubble move above, define a B_p -bubble move $\beta_p : (S^3, T) \rightarrow (S^3, T)$ to be the result of an isotopy of (B_p, D_p) along a closed path in T_- that returns (B_p, D_p) to itself. It is understood that D_p remains in T throughout the isotopy, so T is preserved throughout the isotopy. This means that the final map $\beta_p : (S^3, T) \rightarrow (S^3, T)$ still represents the identity in $\text{Mod}(S^3, T)$. However, the restriction $\beta_p|_{B_-} : (B_-, T_-) \rightarrow (B_-, T_-)$ may not be isotopic to the identity and so may represent a non-trivial element of $\text{Mod}(B_-, T_-)$.

Define $\mathcal{E}_- \subset \text{Mod}(B_-, T_-)$ to be the subgroup generated by B_p -bubble moves, together with the flips, bubble moves, bubble exchanges and eyeglass twists that generate $\mathcal{E}$ and whose associated paths, bubbles, lenses and bridges are disjoint from B_p , so they all lie in B_- .

Lemma 3.1. *If $\text{Mod}(S^3, T) = \mathcal{E}$ then $\text{Mod}(B_-, T_-) = \mathcal{E}_-$.*

This is a natural statement, and is relevant here because $G(S^3, T) = \text{Mod}(S^3, T)$.

Proof. Let $h_- : (B_-, T_-) \rightarrow (B_-, T_-)$ represents an element of $\text{Mod}(B_-, T_-)$; after an isotopy we can assume that h_- is the identity near $\partial B_p = \partial B_-$. Extend h_- to a diffeomorphism $h : (S^3, T) \rightarrow (S^3, T)$ representing a class in $\text{Mod}(S^3, T)$ by setting $h|_{B_p}$ to be the identity. Under the assumption $\text{Mod}(S^3, T) = \mathcal{E}$ there is an eyeglass move $h' : (S^3, T) \rightarrow (S^3, T)$ so that $h'h^{-1} : (S^3, T) \rightarrow (S^3, T)$ is isotopic to the identity. We may as well take the generators of h' to lie in B_- , since each of these generators has support near a 1-complex in T and so, by general position, can be chosen to be disjoint from B_p . For example, for an eyeglass twist choose the frame of the eyeglass disjoint from B_p and it becomes a frame in B_- ; in a bubble move, regard the bubble as a thin neighborhood in T of the bubble's 1-skeleton, which we can also take to be disjoint from B_p . Once this is done, the generators give a diffeomorphism $h'_- : (B_-, T_-) \rightarrow (B_-, T_-)$ representing an element of $\mathcal{E}_-$. (There is no claim that h'_- is well-defined. In fact a particular choice of h_- depends on how the supporting 1-complexes of the generators of h' have been isotoped away from B_p .)

By construction, the diffeomorphisms $h'h^{-1}$ and $\text{id}_{(S^3, T)}$, both of which are the identity on B_p , are isotopic, but such an isotopy might well move B_p , so it does not induce an isotopy from $h'_-h^{-1} : (B_-, T_-) \rightarrow (B_-, T_-)$ to the identity. Our intention is to fix this, by suitably altering h' . Denote by θ_t the isotopy from $\theta_0 = h'h^{-1} : (S^3, T) \rightarrow (S^3, T)$ to $\theta_1 = \text{id}_{(S^3, T)}$ and view θ_t as a path in $\text{Diff}(S^3, T)$. The isotopy θ_t determines a p -based loop $\alpha : [0, 1] \rightarrow T$ in T via $\alpha(t) = \theta_t(p)$. Now alter h' to $h'' : (S^3, T) \rightarrow (S^3, T)$ by post-composing with a B_p -bubble move along the loop $\bar{\alpha}$ defined by $\bar{\alpha}(t) = \alpha(1 - t)$. By definition of $\mathcal{E}_-$, the B_p -bubble move $h''_- : (B_-, T_-) \rightarrow (B_-, T_-)$ still lies in $\mathcal{E}_-$. Furthermore, the maps h' and h'' are isotopic as diffeomorphisms on (S^3, T) , but the isotopy θ_t is replaced by an isotopy θ'_t from $h''h^{-1}$ to the identity for which the loop $\theta'_t(p)$ in T becomes the concatenation $\alpha\bar{\alpha}$, and so is nullhomotopic in T .

Consider the evaluation map $e_p : \text{Diff}(S^3, T) \rightarrow T$ given by $e_p(f) = f(p)$. We can regard T as the space of embeddings $\text{Emb}(p, T)$ and deduce from [Pa] that e_p is a fibration. In particular, the null-homotopy of $\alpha\bar{\alpha}$ to $p \in T$ just described lifts to a homotopy rel end points from the arc θ'_t in $\text{Diff}(S^3, T)$ to an arc in $\text{Diff}(S^3, T)$ that lies entirely over p . This arc then defines an isotopy θ''_t from $h''h^{-1} : (S^3, T) \rightarrow (S^3, T)$ to the identity, but now an isotopy in which the point p is fixed throughout. That is, for every $0 \leq t \leq 1$ $\theta''_t(p) = p$.

Finally, we will alter below the isotopy θ'' so that for every $0 \leq t \leq 1$, $\theta''_t(B_p) = B_p$. This will imply that the restriction $\theta''_- = \theta''|_{(B_-, T_-)}$ is an isotopy from $h''_-h^{-1} : (B_-, T_-) \rightarrow (B_-, T_-)$ to the identity, so h''_- and h_- represent the same element of $\text{Mod}(B_-, T_-)$; since $h''_- \in \mathcal{E}_-$ this will conclude the proof.

Here then is a brief sketch of how, using standard tools of differential topology [GP], the isotopy θ'' can be altered so that not just p , but all of B_p is mapped diffeomorphically to itself during the isotopy. The first step is to observe that the derivative of any diffeomorphism $f : (S^3, T, p) \rightarrow (S^3, T, p)$ determines a linear automorphism Df_p of the tangent space of S^3 at p which also sends the tangent space of T at p to itself. The space of such linear automorphisms deformation retracts to the subspace in which the automorphism is orthogonal. (Such a deformation retraction can be easily constructed, for example, from the proof of [Sp, Theorem 4.2].) Via this deformation retraction, continuously alter θ'' so that for any $0 \leq t \leq 1$ the derivative $D(\theta''_t)_p$ is orthogonal. Once this is done, θ'' can be further altered near p so that each map θ''_t becomes itself the orthogonal map $D(\theta''_t)_p$ near p . In particular it takes a small ball around

p diffeomorphically to itself while continuing to preserve the Heegaard surface T . The final step is to deform θ'' so that the ball on which this is true contains the original B_p .

(Aside: when $\text{genus}(T) = 1$, B_p -bubble moves are not needed, per [EE, Theorem 1b].) $\square$

A bubble exchange between genus 1 bubbles b_1 and b_2 , as shown in Figure 3, takes place in the neighborhood of $b_1 \cup b_2 \cup \nu$, where ν is an embedded arc in T that connects the bubbles. A more general operation, which we will call a *generalized bubble exchange*, isotopes b_1 to b_2 along an arc $\nu_1 \subset T$ while simultaneously isotoping b_2 to b_1 along an arc $\nu_2 \subset T$ disjoint from ν_1 . A generalized bubble exchange is a simple bubble exchange if the arcs ν_1, ν_2 are parallel in T .

Lemma 3.2. *Any generalized bubble exchange is an eyeglass move.*

Proof. A generalized exchange of bubbles b_1 and b_2 using arcs ν_1 and ν_2 (first column of Figure 6) can be written (see second column of Figure 6) as a composition of the simple bubble exchange along ν_1 and a bubble move of b_2 along the closed path $\nu_1 \cup \nu_2$. Each of the latter is an eyeglass move. $\square$

Following Lemma 3.2 we will use the term “bubble exchange” to include generalized bubble exchange, unless the distinction is important.




Figure 6: Generalized exchange as an eyeglass move

3.2 Resuming the search for a strategy

We begin by setting up a standard picture for the genus g Heegaard splitting of S^3 that is somewhat different than the one used by Powell. Let $T_g \subset S^3$ be the standard genus g Heegaard surface in S^3 , dividing S^3 into the genus g handlebodies A_g and B_g .

Let $\{c_1, \dots, c_{g-1}\}$ be the disjoint separating circles on T_g shown in Figure 7, with each c_i separating the first i standard summands from the last $g-i$ standard summands. Note that each c_i bounds a disk in both A and B and so defines a reducing sphere S_i for T_g . Let b_i be the genus i Heegaard split 3-ball component of $S^3 - S_i$ containing S_1 . Both b_i and its complement $S^3 - b_i$ are bubbles.





Figure 7: The standard Heegaard surface $T_g \subset S^3$

It is easy to construct a sequence of (generalized) bubble exchanges (and perhaps a flip) that ultimately moves each bubble b_i , $1 \leq i \leq g-1$ to the bubble $S^3 - b_{g-i}$. The sequence is best seen in the elaboration Figure 8 of Figure 7 that we now describe. The unit sphere S in S^3 is shown intersecting the x - y -plane, with the z -axis coming out of the page. Denote by C the unit circle in which S intersects the x - y -plane. For $g = 2n$ or $2n-1$ (depending on the parity of g) consider n planes $Q_1, \dots, Q_n$ parallel to the x - z plane (horizontal planes in the figure) starting at Q_1 the x - z plane itself and ascending from there. Each Q_i , $1 \leq i \leq n-1$ intersects S in a circle d_i ; $d_n = S \cap Q_n$ is either a circle or a point, depending on the parity of g . Place a genus 1 bubble (shown in green in the figure) on S at each point of $d_i \cap C$. The curves c_i that separate the bubbles, as shown in Figure 7 may be taken to be the intersection of vertical planes (that is, planes parallel to the y - z -plane) $P_1, \dots, P_{g-1}$. These planes are distributed symmetrically across the y - z plane and are shown in red in Figure 8. In this picture, simple π -rotation around the y -axis will take each c_i to c_{g-i} and each genus i bubble b_i to the bubble $S^3 - b_{g-i}$, as we seek.

It remains to show that (up to pairwise isotopy of (S^3, T_g)) this π rotation, which we will denote h_ρ , can be accomplished by generalized bubble exchanges and perhaps a flip. This is easy to see; for each $1 \leq i \leq n$ do a generalized bubble exchange between the two bubbles that lie in d_i (or a flip on the single bubble in d_n if g is odd) using the subarc of d_i having z positive (the front face of the sphere S) for one arc of the exchange and the subarc of d_i on the back face of S (z negative) for the other arc. Since each generalized bubble exchange (and the flip) is an eyeglass move, h_ρ is an eyeglass move. (In fact, since the bubbles are standard, h_ρ is a Powell move.)

Suppose $T \subset S^3$ is a genus g Heegaard surface, and $h, h' : (S^3, T) \rightarrow (S^3, T_g)$ are two orientation-


 Figure 8: h_ρ as sequence of bubble exchanges (g odd)

preserving homeomorphisms.

Definition 3.3. h, h' are eyeglass equivalent (written $h \sim h'$) if the composition $h'h^{-1} : (S^3, T_g) \rightarrow (S^3, T_g)$ is isotopic in (S^3, T_g) to an eyeglass move.

For example, the argument above shows that h_ρ is eyeglass equivalent to the identity.

Throughout the remainder of this section we continue to make this inductive assumption:

Assumption 3.4. $G(S^3, T') = \mathcal{E}$ whenever $\text{genus}(T') \leq g - 1$.

Lemma 3.5. Suppose S is a reducing sphere for a genus g Heegaard splitting (S^3, T) . Then there is an orientation preserving homeomorphism $h_S : (S^3, T) \rightarrow (S^3, T_g)$ so that $h_S(S) \in \{S_i, i = 1, \dots, g - 1\}$.

Suppose that S' is a reducing sphere for (S^3, T) that is disjoint from S and $h_{S'} : (S^3, T) \rightarrow (S^3, T_g)$ is similarly defined, so that $h_{S'}(S') \in \{S_i, i = 1, \dots, g - 1\}$. Under Assumption 3.4, the homeomorphisms h_S and $h_{S'}$ are eyeglass equivalent. In particular h_S is well-defined up to eyeglass equivalence.

Proof. The first statement follows almost immediately from [Wa], as we now describe. S divides T into surfaces $T_\pm$ of genus $g_\pm \geq 1$ respectively, with $g_+ + g_- = g$. Choose any orientation-preserving homeomorphism $f : S^3 \rightarrow S^3$ that carries S to S_{g_+} and the complementary ball component of S that contains T_+ to the ball b_{g_+} . Adjust so that the circle $T \cap S$ is sent to $c_{g_+} \subset S_{g_+}$. Then $f(T_+)$ and $T_g \cap b_{g_+}$ each determine a genus g_+ Heegaard splitting of b_{g_+} and [Wa] then implies that the surfaces are isotopic rel c_{g_+} . Similarly isotope $f(T_-)$ to $T_g \cap (S^3 - b_{g_+})$ in $(S^3 - b_{g_+})$. The resulting homeomorphism h_S takes S to S_{g_+} and T to T_g , as required.

Note that this construction subtly depends on a choice: if we had reversed the labels $g_\pm$ then the requirement that T_+ be sent to b_{g_+} would have sent S to S_{g_-} instead of S_{g_+} . But this could also be

accomplished by composing with the homeomorphism h_ρ defined just before this lemma, and we have shown that $h_\rho \sim id_{(S^3, T_g)}$. So the choice of labelling makes no difference in the construction of h_S , up to eyeglass equivalence.

Proceeding then with the second statement, we first claim that h_S is well-defined up to eyeglass equivalence. To that end, suppose another orientation preserving homeomorphism $h' : (S^3, T) \rightarrow (S^3, T_g)$ has $h'(S) \in \{S_i, i = 1, \dots, g-1\}$. Using possibly h_ρ as above, we may as well assume $h_S h'^{-1}(\mathfrak{b}_{g_+}, T_{g_+}) = (\mathfrak{b}_{g_+}, T_{g_+})$. Noting that $g_+ < g$, apply Lemma 3.1 and Assumption 3.4 to the pair $(\mathfrak{b}_{g_+}, T_{g_+})$ with the goal of showing that there is an eyeglass move on (S^3, T_g) which, when composed with $h_S h'^{-1}$, is the identity on $\mathfrak{b}_{g_+}$ and remains $h_S h'^{-1}$ on $S^3 - \mathfrak{b}_{g_+}$. Lemma 3.1 says that, under the inductive assumption, this is true for some move $h_{g_+} : (\mathfrak{b}_{g_+}, T_{g_+}) \rightarrow (\mathfrak{b}_{g_+}, T_{g_+})$ in $\mathcal{E}_-$. By definition, each flip, bubble move, bubble exchange or eyeglass twist used in the construction of h_{g_+} is also an eyeglass move on (S^3, T_g) . On the other hand, a move in $\mathcal{E}_-$ corresponding to a B_p bubble move in the proof of Lemma 3.1 here corresponds to a bubble move on the bubble $S^3 - \mathfrak{b}_{g_+}$ in (S^3, T_g) , and this is also an eyeglass move. So indeed there is an eyeglass move as we seek. Now apply the same argument on the complementary ball $S^3 - \mathfrak{b}_{g_+}$ and deduce that $h_S h'^{-1}$ is an eyeglass move on (S^3, T_g) , so h' is eyeglass equivalent to h_S , as desired.

Now consider the reducing sphere S' . We may as well choose labels $g_\pm$ so that $h_S(S') \subset \mathfrak{b}_{g_+}$. Let $\mathfrak{b}_S = h_S^{-1}(\mathfrak{b}_{g_+})$, a genus g_+ bubble for (S^3, T) . Apply the first statement again, this time to $h_S|_{\mathfrak{b}_S} : (\mathfrak{b}_S, T_+) \rightarrow (\mathfrak{b}_{g_+}, T_{g_+})$, and deduce that $h_S|_{\mathfrak{b}_{g_+}}$ could have been chosen so that $h_S(S') = S_i$ for some $i \leq g_+$, that is $h_S(S') \in \{S_i, i = 1, \dots, g_+\}$. The same is true, by definition, for the given $h_{S'}$. The argument we have just given that h_S is well-defined up to eyeglass equivalence, repeated now for $h_{S'}$, shows that $h_S \sim h_{S'}$ as required.

The last comment, that under the inductive assumption h_S is well-defined up to eyeglass equivalence, follows simply by taking S' to be a parallel copy of S in the previous argument. $\square$

Our goal is to show that for any genus g Heegaard surface T in S^3 , $G(S^3, T) = \mathcal{E}$. Another way of expressing this is that any two orientation preserving homeomorphisms $(S^3, T) \rightarrow (S^3, T_g)$ are eyeglass equivalent.

Suppose $F \subset S^3$ is a possibly disconnected closed surface, dividing S^3 into possibly disconnected 3-manifolds M_A and M_B . Suppose further

1. Each component C of M_A has a Heegaard splitting $C = A_C \cup_{T_C} B_C$ in which A_C is a handlebody. In particular, $\partial C = \partial_- B_C$. Here we follow [FS2] and [Sc1] in allowing sphere components of $\partial_- B_C$, so the compression body B_C may be reducible.
2. The symmetric statement is true for each component of M_B .
3. The splitting of each ball component of M_A or M_B has genus ≥ 1 . Hence any component that has a genus 0 splitting is a punctured 3-sphere (since it has genus 0) with more than one boundary component (since it is not a ball).
4. The Heegaard splitting (S^3, T) that is obtained by amalgamating all of these Heegaard splittings is of genus g . (See [La, Section 3] for a description of Heegaard surface amalgamation.)

The *defining surface* $F \subset S^3$, together with the Heegaard splittings of M_A and M_B as described above, will be called a *Heegaard split chamber complex* that *supports* the Heegaard splitting (S^3, T) . Each

component of M_A or M_B is called a *chamber*. Heegaard split chamber complexes are formally defined and explored more extensively in section 5. (See for example Definition 5.9)

Before proceeding we list some Elementary Facts about surfaces and amalgamation that will be useful:

- EF1: A compact surface is non-planar if and only if it contains two simple closed curves that intersect in a single point
- EF2: Consequently, if F is subsurface of a compact surface F' and F is non-planar, then so is F' .
- EF3: If F is a non-planar surface and F' is obtained by removing a finite number of disjoint disks from the interior of F (that is F' is a *punctured* F), then F' is non-planar.
- EF4: Suppose Heegaard splittings $(M_1, T_1), (M_2, T_2)$ of compact orientable 3-manifolds M_1 and M_2 are amalgamated along a closed surface F that is a boundary component of each. Then a punctured copy of F lies in each of T_1, T_2 and the amalgamated Heegaard surface $T \subset M = M_1 \cup_F M_2$. See [La, Figure 12]. In particular, if F is non-planar, so are T_1, T_2 and T .
- EF5: Under the amalgamation just described, a punctured copy of each $T_i, i = 1, 2$ lies in T . In particular, if either T_i is non-planar, so is T . Again see [La, Figure 12].

Lemma 3.6. *Suppose, in a Heegaard split chamber complex in S^3 as described above, S is a sphere component of F . Then, after amalgamation, a slight push-off of S becomes a reducing sphere for the splitting (S^3, T) , intersecting T in a single circle that is essential in T .*

Proof. We first show that after amalgamation the part of T lying in each of the ball components of $S^3 - S$ must contain a non-planar surface. There are two cases:

If every chamber within a ball B bounded by S is a punctured 3-sphere, then every component of F within B is a sphere. Pick a sphere (possibly S itself) that is innermost among these components. The ball it bounds contains no other component of F and therefore is a ball chamber. By the third property of Heegaard split chamber complexes described above, the splitting surface for this ball chamber has genus ≥ 1 and so is non-planar. It follows that $T \cap B$ is non-planar, using EF5 above.

On the other hand, if there is a chamber in B that is not a punctured 3-sphere then a boundary component of the chamber is non-planar, so again $T \cap B$ is non-planar, using EF4 above.

We now follow the methodology of [CG]: after amalgamation S becomes a punctured sphere lying in T with some of its boundary components bounding disks in A and others bounding disks in B . Choose a circle $c \subset S$ that divides S into two disks: D_A that contains all disks of $S - T$ that lie in A and D_B that contains all disks of $S - T$ that lie in B . Push $\text{int}(D_A)$ slightly into A and $\text{int}(D_B)$ slightly into B . The resulting sphere S' intersects T only in the circle c . As we have just shown, each of the two components of $T - S'$ contains a non-planar surface, namely the part of T lying in a ball component of $S^3 - S$. Thus neither component of $T - S'$ is planar so, in particular, neither component is a disk. Hence S' is a reducing sphere for T and $c = S' \cap T$ is essential in T . $\square$

Proposition 3.7. *Suppose, for the Heegaard split chamber complex above, S is an incompressible sphere in a chamber of M_A or M_B . Then there is an orientation preserving homeomorphism $h_S : (S^3, T) \rightarrow (S^3, T_g)$ so that $h_S(S) \in \{S_i, i = 1, \dots, g - 1\}$.*

Moreover, suppose for $\tau \in G(S^3, T)$ a homeomorphism $h_{\tau(S)}$ is similarly defined, for the Heegaard split chamber complex whose defining surface is $\tau(F)$. Then, under Assumption 3.4, $h_{\tau(S)}\tau \sim h_S$.

Proof. Let C be the chamber in which S lies, say $C \subset M_A$, so $C = A_C \cup_{T_C} B_C$ with A_C a handlebody.

The main theorem of [Sc1] says that the Heegaard surface T_C in C may be isotoped so that it is aligned with S , that is so that it intersects S in at most one circle.

Claim: When S is aligned with T_C in C and the Heegaard surface $T \subset S$ is created by amalgamation of the splittings of the chambers, S becomes a reducing sphere for the splitting (S^3, T) .

There are three cases to consider:

Case 1: T_C cannot be isotoped in C to be disjoint from S .

S is aligned with T_C , and in this case $T_C \cap S$ cannot be empty, so this intersection must be a single circle c . Furthermore c must be essential in T_C : Indeed, if c bounded a disk $D_T \subset T_C$ then the union of D_T and the disk $S \cap A_C$ would be a sphere in the (irreducible) handlebody A_C , and the ball the sphere bounds could be used to isotope D_T through $S \cap A_C$ and so off of S , contradicting the hypothesis of this case.

Since S is separating in S^3 , the curve c is separating in T_C . Since c is essential in T_C , neither component of $T_C - c$ is a disk, so both components are non-planar surfaces. By the argument of EF5 above, each side of c in T is then non-planar, so c is essential in T . Under amalgamation, the disk $S \cap A_C$ becomes a disk in A .

It is easy to arrange that, similarly, the disk $S \cap B_C$ remains a disk in B after the amalgamation: isotope the disk $S_B = S \cap B_C$ so that it intersects the defining B -disks for the amalgamation $\mathcal{B} \subset B_C$ only in arcs. Then S_B intersects $B_C - \eta(\mathcal{B})$, which is a collar of ∂C in B_C , only in disks, namely the complement in S_B of the arcs $S_B \cap \mathcal{B}$. Finally, properly isotope the disks $S_B - \eta(\mathcal{B})$ so that they avoid the A -disks of the amalgamation, so S_B ends up entirely in B . Then S is the union of two disks, $S_A \subset A$ and $S_B \subset B$ along their common boundary c , an essential circle in T . Hence S is a reducing sphere for (S^3, T) .




Figure 9: A spine Σ (in aqua) of B_C intersecting S in a single point.

Case 2: T_C can be isotoped in C to be disjoint from S and $\text{genus}(T_C) \geq 1$.

Since the chamber $C \subset M_A$, A_C is a handlebody and so is irreducible. Hence if S is disjoint from T_C it must lie entirely in the compression body B_C . This implies that B_C is reducible, so $\partial_- B_C$ contains spheres, some of which lie on the other side of S from $T_C = \partial_+ B_C$ in B_C . Let $S_C \subset \partial_- B_C$ be one of them.

Since S separates S_C from T_C in B_C , any spine of B_C intersects S , and there is some spine Σ of B_C that intersects S in a single point (as pictured in Figure 9). Now isotope T_C to be a neighborhood of that spine, so it intersects S in a single circle c . Amalgamation happens well away from S , so S also ends up intersecting T in c . The circle c divides T_C into a disk, on the side of S containing S_C , and a positive genus surface on the other, since by assumption $\text{genus}(T_C) \geq 1$. By Lemma 3.6 the part of T that lies in the ball in S^3 bounded by S_C (not visible in Figure 9) also has genus ≥ 1 . Hence each of the surfaces into which S divides T has genus ≥ 1 , so S is a reducing sphere for T . Moreover the spine Σ has been chosen so that the disk $S \cap B_C$ is one of the defining disks for the amalgamation, so it automatically persists as a disk in B after amalgamation. Thus S becomes a disk intersecting T in the single essential circle c and so a reducing sphere for (S^3, T) .

Case 3: T_C can be isotoped in C to be disjoint from S and $\text{genus}(T_C) = 0$.

By assumption (3) on F above, C is a multiply punctured 3-sphere. Since S is incompressible in T_C , each side of S in T_C contains some of the punctures, i. e. some sphere components of $\partial_- C$. Just as in Case 2, applying Lemma 3.6, T_C can be isotoped to intersect S in a single circle and the part of T lying on each side of S must have positive genus, so S is a reducing sphere.

This establishes the Claim; the proof of the first statement then follows from [Wa].

The proof of the second statement follows from the observation that $h_S \tau^{-1}(\tau(S)) \in \{S_i, i = 1, \dots, g-1\}$ so for $h_{\tau(S)}$ we could take $h_S \tau^{-1}$. So, under the inductive assumption, Lemma 3.5 says that any other choice of $h_{\tau(S)}$ is eyeglass equivalent to $h_S \tau^{-1}$, or $h_{\tau(S)} \tau \sim h_S$. $\square$

Proposition 3.8. *Suppose S and S' are not necessarily disjoint spheres, each in possibly different chambers of M_A or M_B and each is incompressible in the chamber in which it lies. Suppose $h = h_S : (S^3, T) \rightarrow (S^3, T_g)$ and $h' = h_{S'} : (S^3, T) \rightarrow (S^3, T_g)$ are homeomorphisms as described in Proposition 3.7, so, in particular, S and S' are both reducing spheres for (S^3, T) . Under Assumption 3.4, the homeomorphisms h and h' are eyeglass equivalent.*

Proof. The proof is an examination of interlocking special cases.

Case 1: S and S' are isotopic rel T .

In this case S and S' both divide T into components $T_{\pm}$ of genus $g_{\pm} \geq 1$ where $g_+ + g_- = g$. It follows that $h(S) = S_{g_+}$ or S_{g_-} , say the former. After perhaps composing with the homeomorphism $h_{\rho} : (S^3, T_g) \rightarrow (S^3, T_g)$ defined before Lemma 3.5, we may as well assume that also $h'(S') = S_{g_+}$, so $h'h^{-1}$ leaves S_{g_+} invariant. Apply Lemma 3.5 to two copies of S_{g_+} ; the homeomorphism $h'h^{-1}$; and the identity homeomorphism $id_{(S^3, T_g)}$. Deduce that $h'h^{-1} \sim id_{(S^3, T_g)}$. Hence $h \sim h'$.

Case 2: S and S' both lie in the same chamber C of $S^3 - F$ and there is an eyeglass move of T_C in C which carries S to S' .

As discussed in the proof of [FS2, Theorem 5.1], the handle slides that define the eyeglass move in C determine, up to eyeglass twists, parallel handle slides in T after amalgamation. So in this case there is also an eyeglass move of T in S^3 that carries S to a sphere isotopic to S' rel T . The proof of Case 2 then follows from Case 1.

Case 3: S and S' are disjoint.

Apply Lemma 3.5.

Case 4: S and S' both lie in the same chamber C and are isotopic in C .

The main theorem of [FS2] says that there is a sequence of reducing spheres for T_C starting at S and ending at S' so that each successive pair is either isotopic rel T_C , differ by an eyeglass move in (C, T_C) , or are disjoint, because they differ by a bubble pass in T_C . The result then follows from Cases 1, 2 and 3.

Case 5: S and S' are isotopic in $S^3 - F$ to disjoint spheres.

If S and S' lie in different chambers of $S^3 - F$ the result follows from Case 3. Suppose S and S' lie in the same chamber C of $S^3 - F$ and can be isotoped in C , but not necessarily rel T_C , to be disjoint. According to the main theorem of [Sc1] the Heegaard surface T_C can be aligned with the sphere set $S \cup S'$, so that the S and S' become disjoint reducing spheres for T_C . This new alignment of S and S' may change h and h' , but by Case 4 the eyeglass equivalence class of the homeomorphisms will not change. The result then follows from Lemma 3.5.

The general case If S and S' lie in different chambers of $S^3 - F$ the result follows from Case 3. Suppose they lie in the same chamber C . Ignoring T_C for the moment, recall that in classic 3-manifold theory a standard innermost disk argument shows that there is a sequence of incompressible spheres in C beginning with S and ending with S' so that sequential spheres are isotopic in C to disjoint spheres. The result then follows from Case 5. $\square$

The following corollary is then a summary:

Corollary 3.9. *Suppose in a Heegaard split chamber complex (with notation as in conditions (1)-(4) preceding Lemma 3.6), the manifold $S^3 - F$ contains an incompressible sphere, then under Assumption 3.4, F determines a natural eyeglass equivalence class $\tilde{h}_F$ of homeomorphisms $(S^3, T) \rightarrow (S^3, T_g)$.*

Moreover, for $\tau \in G(S^3, T)$, $\tilde{h}_{\tau(F)}\tau = \tilde{h}_F$.

Proof. Choose any incompressible sphere S in $S^3 - F$ and let $h_S : (S^3, T) \rightarrow (S^3, T_g)$ be a homeomorphism as given by Proposition 3.7. Then Proposition 3.8 shows that, up to eyeglass equivalence, h_S is independent of any choice, including the choice of incompressible sphere S . The last sentence follows from the last sentence of Proposition 3.7. $\square$

To simplify (but slightly abuse) notation, we will often use $h_F : (S^3, T) \rightarrow (S^3, T_g)$ to denote any representative of the eyeglass equivalence class $\tilde{h}_F$. With this notation the last sentence of 3.9 would be $h_{\tau(F)}\tau \sim h_F$.

Definition 3.10. *Suppose there are two Heegaard split chamber complexes in S^3 with corresponding associated surfaces $F, F' \subset S^3$ and each Heegaard split chamber complex supports the same Heegaard splitting surface T for S^3 . Suppose each of $S^3 - F$ and $S^3 - F'$ contains an incompressible sphere, so (eyeglass equivalence classes of) homeomorphisms $h_F, h_{F'} : (S^3, T) \rightarrow (S^3, T_g)$ are defined. If $h_F \sim h_{F'}$ then F and F' cocertify.*

(This definition will be expanded later, see comments following Definition 8.2.)

Lemma 3.11. *Suppose, in Definition 3.10 $F \subset F'$ and a component C' of $S^3 - F'$ contains an incompressible sphere S that is also incompressible in the component $C \supset C'$ of $S^3 - F$ in which it lies. Then F and F' cocertify.*

Proof. Via [Sc1] isotope in C the incompressible sphere to a sphere $S \subset C$ that is aligned with T_C . By the Claim in the proof of Proposition 3.7 S becomes a reducing sphere for the splitting (S^3, T) after amalgamation. Similarly isotope S in C' to a sphere $S' \subset C'$ that is aligned with $T_{C'}$ and again conclude that S' upon amalgamation becomes a reducing sphere for (S^3, T) . Then by definition $h_{S'} \in \tilde{h}_{F'}$ and $h_S \in \tilde{h}_F$. Since $C' \subset C$, the isotopies of both S and S' to alignment take place in C so, by Proposition 3.8, $h_{S'}$ also represents $\tilde{h}_F$. Since $\tilde{h}_F$ and $\tilde{h}_{F'}$ have the representative $h_{S'}$ in common, they cocertify. $\square$

Corollary 3.12. *Suppose $F \subset S^3$ is the surface associated with a Heegaard split chamber complex that supports T . Suppose $\tau \in G(S^3, T)$. If the Heegaard split chamber complexes given by F and $\tau(F)$ cocertify, then $\tau \in \mathcal{E}$.*

Proof. Corollary 3.9 shows that $h_F \sim h_{\tau(F)}\tau$. On the other hand, the assumption that F and $\tau(F)$ cocertify implies that $h_{\tau(F)}\tau \sim h_F\tau$. Hence $h_F \sim h_F\tau$ or $id_{(S^3, T)} \sim \tau$, as required. $\square$

3.3 An overview of the argument

Corollary 3.9 suggests a line of attack towards proving $G(S^3, T) = \mathcal{E}$: For each $\tau \in G(S^3, T)$, find a Heegaard split chamber complex in S^3 with associated surface F so that

- the Heegaard split chamber complex supports T
- some chamber contains an incompressible sphere and
- $h_F \sim h_{\tau(F)}$.

There are three conditions required for this example to work that are difficult to realize:

1. The Heegaard splitting of each ball chamber has genus ≥ 1 , as required by the definition of a Heegaard split chamber complex. This property is needed to show, as in Lemma 3.6 and Proposition 3.7 above, that an incompressible sphere in a chamber determines a reducing sphere for T .
2. Some chamber contains an incompressible sphere
3. The Heegaard split chamber complexes derived from F and $\tau(F)$ cocertify.

Section 1 gave a brief description of how the **second requirement** is addressed in the general case: given a position of T in S^3 , a pair of weakly reducing disks is found for T , with weak reduction defining a chamber complex. Next a particularly useful level sphere (a ‘guiding sphere’, see Section 9) S is found, and innermost disks in $S - T$ are used to decompose the chamber complex into a typically more complex one. The process is repeated until all circles of intersection are removed. It will be shown that at some stage in this *disk decomposition* process a chamber will contain an incompressible sphere or a sphere that is equally useful for determining an eyeglass equivalence class of homeomorphisms from $(S^3, T) \rightarrow (S^3, T_g)$. (See Definition 8.2 of a *certificate*.)

This process raises many technical questions: How to find a guiding sphere whose associated disk decomposition sequence must issue a certificate, Section 10. How to ensure, for a given sphere, that the certificates issued during this process cocertify, Section 8. The most difficult problem is showing that a

different choice of such a sphere will result in an equivalent certificate. This takes many steps, occupying Sections 11 through Section 17. The philosophy behind these steps borrows heavily from [FS1] (see next paragraph): the guiding spheres are isotopic in S^3 through appropriate guiding spheres (Section 17) and such an isotopy between them can be broken into small steps, each of which either adds to, subtracts from, or delays a single disk in the associated disk decomposition sequence. We then examine the effect of each such minimal change. The upshot is that, up to cocertification, the certificate issued depends only on the Heegaard split chamber complex itself, and not on the choice of guiding sphere.

To address the **third requirement**, that is, to show that the Heegaard split chamber complexes derived from F and $\tau(F)$ cocertify, we construct a sequence of Heegaard split chamber complexes beginning with F and ending with $\tau(F)$ so that each successive pair cocertify, see Section 18. The heavy lifting for this part of the argument was done in [FS1]. Recall that a Heegaard split chamber complex is derived from a Heegaard splitting (S^3, T) by weak reduction, see Section 5. In [FS1] a series of pairs of weakly reducing disks is found so that, roughly, the resulting Heegaard split chamber complexes are a series beginning with some F and ending with $\tau(F)$, see Section 18. Moreover, successive pairs of weakly reducing disks are related in such a way that Proposition 17.8 guarantees the resulting Heegaard split chamber complexes cocertify. So in the end F and $\tau(F)$ cocertify and the third requirement above is satisfied.

This leaves the **first requirement**, that each chamber in the chamber complex comes with a splitting surface of genus ≥ 1 . Establishing this is particularly vexing, and adds great technical complexity to the argument, for reasons we now describe:

It is necessary to eliminate, whenever they arise, any ball chambers that have genus 0 splittings. In the argument below these will be the ball chambers (called *goneballs*) that are absorbed into the surrounding chamber and so their bounding spheres disappear from F . Goneballs can arise in troublesome ways: a chamber that is a handlebody at one stage can be cut up by a disk decomposition to become a collection of balls; if the handlebody had trivial Heegaard splitting, the resulting balls will be goneballs and their boundaries need to be deleted from F . Since chambers may be deleted in this way, what is to prevent F from disappearing entirely, as pieces are cut up in this manner? We will show that unless the chamber complex is of a particularly simple type, called *tiny*, its defining surface will never completely disappear, see Subsection 5.4.

There is an additional but related complication: it is a classical result that a non-trivial Heegaard splitting of a ball or handlebody chamber is stabilized, that is it will contain bubbles (standard genus one summands). It follows from [FS2] that disk decomposition of a Heegaard split chamber gives a Heegaard splitting of the resulting chambers, but it is defined only up to eyeglass moves and passing bubbles through the decomposing disks. As noted above, it will be important to know if a newly created handlebody chamber could have trivial Heegaard splitting, so we want to flag handlebody chambers that, because of how they arise, are known to have non-trivial splittings. This gives rise to the notion of *flagged chamber complexes*: the exact Heegaard splitting of each chamber is typically unknown, but certain handlebody chambers (including ball chambers) are flagged as necessarily having non-trivial splittings. See Section 7.

Following this informal overview, we begin the more formal argument.

4 Chamber complexes and disk decomposition

Suppose M is a compact orientable 3-manifold. Let $F \subset S^3$ be a (typically disconnected) closed surface in $\text{int}(M)$ that divides M into two typically disconnected 3-manifolds M_A and M_B , with each component of F incident on one side to M_A and on the other to M_B . Call the collection $\mathbb{C} = (F, M_A, M_B)$ a *chamber complex* with defining surface $F = F(\mathbb{C})$. Each component of $M - F$ is called a *chamber*, with those in M_A called *A-chambers* and those in M_B called *B-chambers*.

A Heegaard splitting $M = A \cup_T B$ is a familiar example of a chamber complex in M , with $\mathbb{C} = (T, A, B)$. A Heegaard splitting belongs to a class of chamber complexes, called *tiny*, that, for our purposes, contain too little information to be useful. However, a Heegaard splitting together with a pair of weakly reducing disks (see Section 5 below) does produce a possibly useful chamber complex, through a process called *disk decomposition*, or, more formally, *chamber complex decomposition*, which we now describe.

Phase one of disk decomposition: For $\mathbb{C}$ a chamber complex in M , let $\mathcal{D}$ be a collection of disjoint properly embedded disks, some perhaps inessential, in the chambers of $\mathbb{C}$, with $\partial\mathcal{D}$ disjoint from ∂W , so $\partial\mathcal{D} \subset F$. We will call $\mathcal{D}$ a *disk set* in $\mathbb{C}$. Let $F_{\mathcal{D}} \subset M$ be the surface obtained by doing surgery on $F(\mathbb{C})$ along the disks $\mathcal{D}$, and denote by $\hat{\mathbb{C}}_{\mathcal{D}}$ the associated chamber complex $M - F_{\mathcal{D}}$. (The reason for the circumflex will be given shortly.) In particular, the surgery alters M_A by deleting a bicollar neighborhood of each disk in $\mathcal{D} \cap M_A$ and adding a bicollar neighborhood of each disk in $\mathcal{D} \cap M_B$. The symmetric statement applies to M_B .

Each disk D in $\mathcal{D}$ is the core of a 2-handle used in the surgery as just described. The bicollar neighborhood of ∂D in F will be called the *belt annulus* of the surgery; the two copies of D in $F_{\mathcal{D}}$ that are the result of the surgery are called the *scars* of D . Let C' be a chamber of $\hat{\mathbb{C}}_{\mathcal{D}}$ and $D \in \mathcal{D}$ a disk which leaves a scar or two on $\partial C'$. If the disk D lies outside the new chamber C' then so does the associated belt annulus and we call the scars *external scars* on $\partial C'$. If the disk lies inside C' we call the scars *internal scars* on $\partial C'$. Clearly if D leaves one internal scar on $\partial C'$ it will leave two, since the associated belt annulus lies in C' , but the two internal scars may be on different components of $\partial C'$.

One can think of a chamber $\hat{C}$ of $\hat{\mathbb{C}}_{\mathcal{D}}$ as obtained in two stages: it starts as the complement C_- of a bicollar of $\mathcal{D} \cap C$ in a chamber C of $\mathbb{C}$, a result of surgery on $\mathcal{D} \cap C$. Then 2-handles are added along those disks in $\mathcal{D}$ that are incident to ∂C but lie outside C . The chamber $\hat{C}$ is called a *remnant* of C , even though strictly speaking it doesn't lie entirely inside of C .

Phase two of disk decomposition: Among the components of the surgered surface $F_{\mathcal{D}}$ may be spheres that bound balls in M . In the second phase of disk decomposition a collection of these balls is chosen and, for each ball $G \subset M$ in the collection (henceforth called a *goneball*) all components of $F_{\mathcal{D}}$ that lie in G , including the sphere ∂G , are removed from $F_{\mathcal{D}}$. Whether a ball in M bounded by a sphere in $F_{\mathcal{D}}$ will be chosen to be a goneball of the decomposition is a delicate and important part of the theory. (See, for example, Rules 4.4 and 5.6.) After eliminating all components of $F_{\mathcal{D}}$ that lie in goneballs, call the resulting chamber complex $\mathbb{C}_{\mathcal{D}}$ (note: no circumflex).

Any chamber $C_{\mathcal{D}}$ in the chamber complex $\mathbb{C}_{\mathcal{D}}$ is obtained from a chamber $\hat{C}_{\mathcal{D}}$ in $\hat{\mathbb{C}}_{\mathcal{D}}$ by attaching balls, each corresponding to a goneball for the decomposition. We then say that $C_{\mathcal{D}}$ is a remnant of the chamber C in $\mathbb{C}$ if $\hat{C}_{\mathcal{D}}$ is a remnant of C in $\hat{\mathbb{C}}_{\mathcal{D}}$. Notice that a chamber C in $\mathbb{C}$ may have no remnants in $\mathbb{C}_{\mathcal{D}}$. For example, C might be handlebody for which $\mathcal{D}$ contains a complete collection of meridian disks, so the remnants in $\hat{\mathbb{C}}_{\mathcal{D}}$ are all balls. If these all happen to be declared goneballs, then there are no

remnants of C in $\mathbb{C}_{\mathcal{D}}$.

For $\mathbb{C}, \mathcal{D}, F, F_{\mathcal{D}}, \mathbb{C}_{\mathcal{D}}$ and $\hat{\mathbb{C}}_{\mathcal{D}}$ as above, suppose F' is a component of the surface $F_{\mathcal{D}}$ such that one of the complementary components W of F' in M is a handlebody. (A ball is considered a genus 0 handlebody). W is not necessarily a chamber of $\hat{\mathbb{C}}_{\mathcal{D}}$, but rather the union of those chambers that lie within it.

Definition 4.1. *The handlebody W is diskly if each component of $F \cap \text{int}(W)$ is a disk.*

Lemma 4.2. *If W is diskly then*

1. *No component of F lies entirely in the interior of W .*
2. *Each component of $F_{\mathcal{D}}$ that lies in the interior of W is a sphere.*
3. *Suppose $D \in \mathcal{D}$ lies in the interior of W . Then the corresponding two scars do not lie on the same component of the surface $F_{\mathcal{D}}$.*

Proof. Take the statements in order:

- 1) Each component of F is a closed surface, so none can be a disk.
- 2) Suppose F'' is a component of $F_{\mathcal{D}}$ that lies in the interior of W , and let F''_{-} be the compact surface obtained from F'' by deleting the interior of all scars of the disk surgery that lie in F'' . Then F''_{-} lies in a component of $F \cap \text{int}(W)$ and, by assumption, this component is a disk. Hence the surface F''_{-} has no genus, so neither can F'' . See Figure 10.
- 3) Suppose the two scars left by surgery on a disk $D \in (\mathcal{D} \cap \text{int}(W))$ were on the same component F'' of $F_{\mathcal{D}}$. If $F'' = F'$ then the belt annulus for the surgery is a component of $F \cap \text{int}(W)$ that is not a disk, contradicting our assumption that W is diskly. Similarly, suppose F'' were in the interior of W , and let α be an arc in F'' running between the two scars. Then the union of the belt annulus and a collar of α would be a punctured torus lying in some component of $F \cap \text{int}(W)$, so that component cannot be a disk, again contradicting our assumption that W is diskly. $\square$




Figure 10: Construction of F''_{-} in proof of Lemma 4.2 (2)

Lemma 4.3. *Suppose the handlebody W is diskly, and F'' is a component of $F_{\mathcal{D}}$ that lies in the interior of W . Then F'' bounds a diskly ball in W .*

Proof. By Lemma 4.2 any such component F'' is a sphere and, since the handlebody W is irreducible, it will bound a ball in W . The issue is to show that such a ball is diskly. By assumption, $F \cap \text{int}(W)$ consists of a collection E of disks; if $E = \emptyset$ then there are no components of $F_{\mathcal{D}}$ in the interior of W and there is nothing to prove.

Otherwise, since E lies inside W , the annulus in E adjacent to any circle $c \in \partial E$ is a belt annulus for a surgery disk $D \in \mathcal{D}$ that lies in W , with one of its scars on ∂W . Let $\mathcal{D}_0 \subset \mathcal{D}$ be the subset of disks that lie in W ; we now induct on $n = |\mathcal{D}_0| > 0$.

Let $D_0 \in \mathcal{D}_0$ be a disk whose boundary is innermost in E among the circles $\partial \mathcal{D}_0$; let E_0 be the subdisk of E that it bounds. Since W is irreducible, the disks D_0 and E_0 together bound a ball B_0 in W . The ball is disjoint from F by Lemma 4.2, so we can think of D_0 as just being parallel to E_0 , via the ball B_0 . Inductively apply the lemma to $\mathcal{D}' = \mathcal{D}_0 - D_0$, creating a collection S' of spheres, each of which, by inductive assumption, bounds a diskly ball in W .

Finally, do surgery on the disk D_0 . The result of the surgery is two-fold: The ball B_0 bounded by E_0 and a copy of D_0 becomes a ball chamber that is diskly because it is empty. And the sphere S'_0 in S' that contains E_0 is altered, in effect, by isotoping E_0 across B_0 to the other copy of D_0 . Depending on which side of S'_0 the disk D_0 lies, this change in S'_0 may add the disk $E_0 \subset F$ to the ball that S'_0 bounds in W . But since E_0 is a disk, it does not change the fact that the ball is diskly. See Figure 11. $\square$




Figure 11: $E_0 \subset F$ in the interior of ball bounded by S'_0 . (Figure not incorporating decomposition by $\mathcal{D}' = \mathcal{D}_0 - D_0$.)

The decision process that will be used to decide if sphere components of $F_{\mathcal{D}}$ in a chamber complex decomposition bound goneballs will have certain properties. The ultimate decision process is subtle, but some of those properties can be described already. At this point, view them as rules that disk decompositions follow; later it will be shown that the decision process that will be used satisfies these rules. To that end, declare the rule:

Rule 4.4. *In a chamber complex decomposition, only diskly balls can be goneballs.*

Following Lemma 4.3 we then have

Corollary 4.5. *In the second phase of a chamber complex decomposition, any component that is removed from $F_{\mathcal{D}}$ is a sphere that bounds a diskly ball.*

Continue with $\mathbb{C}, \mathcal{D}, F, F_{\mathcal{D}}, \hat{\mathbb{C}}_{\mathcal{D}}$ and $\mathbb{C}_{\mathcal{D}}$ be as above. Denote the two-stage operation of chamber complex decomposition just described by

$$\mathbb{C} \xrightarrow{\mathcal{D}} \mathbb{C}_{\mathcal{D}}.$$

The next rule limits the amount of disturbance caused by adding or removing a single disk from the chosen disk set.

Rule 4.6. *Suppose $\mathcal{D}$ is a disk set in $\mathbb{C}$, $D \in \mathcal{D}$, and $\mathcal{D}_-$ is the disk set $\mathcal{D} - D$. Let S be a sphere component of $F_{\mathcal{D}}$ that does not contain either scar from D . Then S bounds a goneball in $\mathbb{C}_{\mathcal{D}}$ if and only if it bounds a goneball in $\mathbb{C}_{\mathcal{D}_-}$.*

Corollary 4.7. *Suppose $\mathcal{D}$ consists of a single disk. Then a sphere in $F_{\mathcal{D}}$ that does not contain a scar of D does not bound a goneball.*

Proof. Since $\mathcal{D}_- = \emptyset$ such a sphere would necessarily be a component of $F = F_{\mathcal{D}_-}$ and so does not bound a goneball in $\mathbb{C} = \mathbb{C}_{\mathcal{D}_-}$. By Rule 4.6 it does not bound a goneball in $\mathbb{C}_{\mathcal{D}}$. $\square$

An example of a decision process that satisfies both rules is to simply declare each disky ball to be a goneball. For example, it satisfies Rule 4.6 because a sphere $S \subset F_{\mathcal{D}}$ not containing either scar will also be in $F_{\mathcal{D}_-}$ and a ball it bounds is disky in either $\mathbb{C}_{\mathcal{D}}$ or $\mathbb{C}_{\mathcal{D}_-}$ if and only if F intersects the ball only in disks. Unfortunately declaring each disky ball to be a goneball would be too broad for our purposes; too much information would be lost. But the rules so far do suffice to prove an important property:

Proposition 4.8. *Suppose C is a chamber of $\mathbb{C}$ so that every remnant of C in $\mathbb{C}_{\mathcal{D}}$ is a disky handlebody. Then C is a handlebody.*

Proof. We proceed by induction on $|\mathcal{D}|$, noting that if $\mathcal{D} = \emptyset$ there is nothing to prove.

Case 1: There is a disk in $\mathcal{D}$ whose boundary is inessential in ∂C .

In this case, let D be a disk whose boundary is innermost among all such disks, let E be the disk in $\partial C - \mathcal{D}$ bounded by ∂D . Push $\text{int}(E)$ slightly into the chamber of $\mathbb{C}$ on the opposite side of D . That chamber is either C' if D lies in C (the top row of Figure 12) or C itself, if D lies in a chamber C' adjacent to C (the bottom row of Figure 12). Consider the sphere $S = D \cup E$, lying either in a remnant of C' or a remnant of C , and in either case parallel in that remnant to a sphere component S' of $F_{\mathcal{D}}$.

Claim 1: With the proper choice of D for the construction above, S bounds a ball B_S in M containing S' . Moreover, if D lies in C' , so S is in a remnant of C , then the subball of B_S bounded by S' is a goneball.

Proof of Claim 1: Suppose first that S lies in a remnant C'_r of C' , so D lies inside C , as in the top panel of Figure 12. Then S' is a sphere in the boundary of a chamber C_r of $\hat{\mathbb{C}}_{\mathcal{D}}$ that is a remnant of C . If C_r is a goneball the claim follows. If instead C_r remains as a chamber in $\mathbb{C}_{\mathcal{D}}$, C_r must be a ball, since every remnant of C in $\mathbb{C}_{\mathcal{D}}$ is a handlebody, proving the claim in this case.

Suppose instead that S lies in a remnant C_r of C , so D lies inside C' , as in the bottom panel of Figure 12. Then by hypothesis C_r is a handlebody, and so is irreducible. Hence S bounds a ball B_S in C_r . If B_S lies on the side of S containing S' the claim is shown. Suppose B_S lies on the other side of S , so S' bounds a ball in C_r and that ball contains the punctured handlebody $\partial C - E$. Since each remnant of C is disky, it follows that $\partial C - E$ must be a disk, so C is a ball. In this case use instead of D a disk D' in

$\mathcal{D}$ whose boundary is innermost among those in the disk $\partial C - E$ and repeat the construction. The new sphere corresponding to S lies in B_S so it bounds a ball entirely in C_r , a ball that contains the disk E' in $\partial C - E$ bounded by $\partial D'$, i. e. the one that now corresponds to E . This proves Claim 1, using the disk D' instead of D .




Figure 12: When a disk $D \in \mathcal{D}$ has ∂D inessential in ∂C

Claim 2: The disk D of Claim 1 is inessential in the chamber (either C or C') in which it lies.

Proof of Claim 2: Suppose instead that D is essential in the chamber in which it lies, so there are closed components of F lying between D and E , i. e. in the ball B_S . By Lemma 4.2 the subball of B_S bounded by S' could not then be a disk ball, so by Rule 4.4, S' remains as a sphere component of $F(\mathbb{C}_D)$. Claim 1 then implies that D lies in C so S is in a remnant of C' . Then, as in Claim 1, S' is a sphere in the boundary of a chamber of $\hat{\mathbb{C}}_D$ that is a remnant of C . Since S' remains in $F(\mathbb{C}_D)$ the ball it bounds in B_S must be a ball remnant B' in $\mathbb{C}_D$ since every remnant of C is a handlebody. But even then there is a contradiction, for every remnant of C is assumed to be disk, and here B' contains closed components of F , contradicting Lemma 4.2, and so proving Claim 2.

Let $\mathcal{D}_- = \mathcal{D} - D$. Observe that since D is inessential by Claim 2, $\mathcal{D}_-$ still satisfies the hypotheses of Proposition 4.8: Indeed, from Case 1 and Rule 4.6, the only difference in the remnants of C if we add D back to $\mathcal{D}_-$ before decomposing is:

- possibly adding a disk ball bounded by S' to the set of remnants of C or
- isotoping E to D , if the ball that S' bounds is a goneball. This is the case when $D \in C'$ and may be the case when $D \in C$.

Neither will affect whether each remnant of C is a disk handlebody. So, by inductive assumption, replacing $\mathcal{D}$ with $\mathcal{D}_-$ leads to the conclusion of Proposition 4.8, namely C is a handlebody. This concludes the proof of Proposition 4.8 for Case 1.

Case 2: Any disk incident to ∂C has essential boundary in ∂C . (In particular, any disk incident to ∂C is essential.)

First consider a remnant C_r of C in $\mathbb{C}_{\mathcal{D}}$, which by hypothesis of the Proposition is a disk handlebody. Since C_r is a remnant of C , each scar on ∂C_r is incident to ∂C , so by the hypothesis of this case its boundary is essential in ∂C . But since C_r is disk, this implies it can have no internal scars, i. e. F is disjoint from $\text{int}(C_r)$, since the belt annulus for an internal scar lies on a component of $\partial C \cap \text{int}(C_r)$ and these are all disks. This in turn implies that $\text{int}(C_r)$ contains no goneballs, so C_r is a remnant of C in $\hat{\mathbb{C}}_{\mathcal{D}}$, before goneballs are absorbed. Thus the remnants of C in $\hat{\mathbb{C}}_{\mathcal{D}}$ consist of a union of disk handlebodies, though among these may be balls that become goneballs in $\mathbb{C}_{\mathcal{D}}$. See left side of Figure 13.




Figure 13: When each disk $D \in \mathcal{D}$ has ∂D essential in ∂C

Let D be any disk in $\mathcal{D}$ and, as in Case 1, let $\mathcal{D}_- = \mathcal{D} - D$. Inductively, it suffices to show that $\mathcal{D}_-$ still satisfies the hypotheses of the Proposition, namely that every remnant of C in $\mathbb{C}_{\mathcal{D}_-}$ is a disk handlebody. If D is not incident to ∂C then removing D from $\mathcal{D}$ has no effect on the remnants of C in $\hat{\mathbb{C}}_{\mathcal{D}_-}$, so the remnants are the same as those in $\hat{\mathbb{C}}_{\mathcal{D}}$. By Rule 4.6, the remnants of C in $\mathbb{C}_{\mathcal{D}_-}$, are then also the same as those in $\mathbb{C}_{\mathcal{D}}$, namely a union of disk handlebodies. Thus in this case $\mathcal{D}_-$ satisfies the hypothesis of Proposition 4.8 as required.

What remains is the case that D is incident to ∂C . See right side of Figure 13. Since there are no internal scars on remnants of C , the two scars in $\hat{\mathbb{C}}_{\mathcal{D}}$ left by D are external scars, so D lies in C . Thus $\hat{\mathbb{C}}_{\mathcal{D}_-}$ is obtained from $\hat{\mathbb{C}}_{\mathcal{D}}$ by simply attaching a 1-handle dual to D to a component, or between two components, of $\hat{\mathbb{C}}_{\mathcal{D}}$. Since each such component is a disk handlebody, the result of adding the 1-handle is also a handlebody, and it is disk because its interior is still disjoint from F . Hence $\hat{\mathbb{C}}_{\mathcal{D}_-}$ is a union

of disk handlebodies. $\mathbb{C}_{\mathcal{D}_-}$ is obtained from $\hat{\mathbb{C}}_{\mathcal{D}_-}$ by removing some ball components (the goneballs). Hence $\mathbb{C}_{\mathcal{D}_-}$ is also a union of disk handlebodies, so it satisfies the hypothesis of Proposition 4.8 as required. $\square$

Definition 4.9. A chamber complex $\mathbb{C}$ in M , with defining surface $F = F(\mathbb{C})$, is tiny if either

- $F = \emptyset$ or
- there is a chamber C of $\mathbb{C}$ so that $M - C$ consists of handlebody chambers. These chambers are called the designated handlebody chambers.

A Heegaard surface in which at least one complementary component is a handlebody is a familiar example of a defining surface of a tiny chamber complex. On the other hand, a weak reduction of a Heegaard surface will yield a chamber complex that is not tiny. That fact (see Proposition 5.18) will be the entryway to our study of Heegaard splittings below. The word “tiny” is used because the tree dual to a chamber complex is particularly small for tiny chamber complexes: at its most complicated, it is a star graph with central vertex corresponding to the chamber C .

Proposition 4.10 (Tinyness pulls back). Suppose $\mathbb{C} \xrightarrow{\mathcal{D}} \mathbb{C}_{\mathcal{D}}$ is a chamber complex decomposition in which the chamber complex $\mathbb{C}_{\mathcal{D}}$ is tiny. Suppose further, if the defining surface of $\mathbb{C}_{\mathcal{D}}$ is not empty, that each of the handlebody chambers named in Definition 4.9 is disk. Then $\mathbb{C}$ is tiny.

Proof. We first consider the case in which the defining surface of $\mathbb{C}_{\mathcal{D}}$ is not empty. That is, per Definition 4.9, there is a single chamber C' of $\mathbb{C}_{\mathcal{D}}$ so that $M - C'$ consists of a union W' of disk handlebody chambers. Without loss, assume C' is a B -chamber, so each component of W' is an A -chamber adjacent to C' in $\mathbb{C}_{\mathcal{D}}$. See Figure 14.

Since each component of W' is a disk handlebody, any disk $D \in \mathcal{D}$ that lies in $\text{int}(W')$ lies on a disk in F and so has inessential boundary on F ; let D be one whose boundary is innermost on F . Surgery on D creates a ball in $\text{int}(W')$, which must be a goneball since W' is irreducible. The effect then of surgery on D followed by declaring the ball gone is merely to isotope $F \cap \text{int}(W')$; it has no effect on either the hypothesis of the proposition or the conclusion. Hence we may inductively assume that $\mathcal{D} \cap \text{int}(W') = \emptyset$, so W' has only external scars.

Let $B_{\mathbb{C}}$ denote the collection of B -chambers in $\mathbb{C}$. C' is a B -chamber in $\mathbb{C}_{\mathcal{D}}$, in fact the only B -chamber in $\mathbb{C}_{\mathcal{D}}$. Recall that it is obtained from $B_{\mathbb{C}}$ in three steps:

- $B_{\mathbb{C}}$ is cut along those disks in $\mathcal{D}$ that lie in $B_{\mathbb{C}}$.
- Two-handles are added along those disks in $\mathcal{D}$ that lie in the A -chambers
- Goneballs are eliminated.

Towards understanding the third step, eliminating goneballs, notice that no goneball can lie in $\text{int}(W')$ since W' has no disks of $\mathcal{D}$ in its interior and, in particular, only external scars. Suppose then that U is a maximal goneball in the B -chamber C' . Since U is maximal and lies in the B -chamber C' , the chamber of $\hat{\mathbb{C}}_{\mathcal{D}}$ immediately inside U is an A -chamber. Rule 4.4 says that U is disk. Applying the same argument

as just applied to W' , we may as well then assume that the only scars on ∂U are external scars, so U itself is an A -chamber in $\hat{\mathbb{C}}_{\mathcal{D}}$. With this simplification, so that the only goneballs of the decomposition are A -chambers, none of the three steps in the process above decreases the number of B -chambers (though the first step may increase it). Since, by hypothesis, there is only one B -chamber in $\mathbb{C}_{\mathcal{D}}$ we conclude that $B_{\mathbb{C}}$ is also a single chamber in $\mathbb{C}$.




Figure 14: Tininess when the defining surface of $\mathbb{C}_{\mathcal{D}}$ is not empty

Consider any chamber C of $\mathbb{C}$ other than $B_{\mathbb{C}}$. C is then an A -chamber in $\mathbb{C}$ and so has the property that each of its remnants in $\mathbb{C}_{\mathcal{D}}$ is in W' , so each is a disk handlebody. It then follows from Proposition 4.8 that C itself is a handlebody. Hence $\mathbb{C}$ consists of a union of handlebodies (the A -chambers) whose complement $B_{\mathbb{C}}$ is a single B -chamber. That is, $\mathbb{C}$ is tiny.

The other possibility from Definition 4.9 is that $\mathbb{C}_{\mathcal{D}}$ is tiny because the defining surface for $\mathbb{C}_{\mathcal{D}}$ is empty. This means that in $\mathbb{C}_{\mathcal{D}}$, M is entirely a B -chamber, say. This implies that either F was empty (completing the argument) or surgery on $\mathcal{D}$ resulted only in spheres bounding goneballs. In that case, the argument is basically the same: Rule 4.4 says that each goneball is disk. Then, examining maximal goneballs as above we can assume without loss that each goneball is an A -chamber with only external scars. This implies that every A -chamber in $\mathbb{C}$ can be recovered from a collection of balls (the maximal goneballs) by attaching 1-handles, dual to the disks that leave external scars. The result is visibly a collection of handlebodies. So again $\mathbb{C}$ is tiny. $\square$

5 Heegaard splittings and chamber complexes

5.1 Review of weak reduction and amalgamation

Chamber complexes are naturally relevant to weakly reducible Heegaard splittings, as we now describe. This first subsection is essentially a review of relevant known results; see, for example, [La, Section 3] to which we refer for relevant notation.

Suppose, in a Heegaard splitting $M = A \cup_T B$ of a compact 3-manifold M , $\mathcal{A}, \mathcal{B}$ are disjoint families of disks, properly embedded in A, B respectively, with $\partial\mathcal{A}, \partial\mathcal{B} \subset T$ and at least one member of each disk family an essential disk. Such a pair of families is called a *weakly reducing disk family*, a notion with a long and important history in the study of Heegaard splittings (cf [CG]). Surger T along $\mathcal{A} \cup \mathcal{B}$ and call the resulting surface $\hat{F}$.

$\hat{F}$ defines a chamber complex $\hat{\mathbb{C}}$ in M as follows. One set of chambers, called the A -chambers, are derived from A in two stages: first bicollars of disks in $\mathcal{A}$ are removed (so these disks lie *outside* the resulting A -chambers) and then bicollar neighborhoods of $\mathcal{B}$ are attached (so these are disks that lie *inside* the resulting A -chambers). The remaining chambers, the B -chambers are derived from B in the symmetric fashion. The result is the perhaps counter-intuitive fact that after the surgery, each A -chamber may contain some disks from $\mathcal{B}$, but none from $\mathcal{A}$, and symmetrically for the B -chambers. In any pair of adjacent chambers, one is an A -chamber and the other is a B -chamber, since the surface between them is (except for some disks) a submanifold of T .

Each A -chamber C (and symmetrically for B -chambers) inherits a natural Heegaard splitting $C = A_C \cup_{T_C} B_C$. See Figure 15. This can be demonstrated by considering the two-stage construction just described, see Figure 15: In the first stage, (NE to SW in Figure 15) the surface T is compressed in A along $\mathcal{A}$ and a component C' is chosen. Since A was a compression body, so is C' . Let $T_C \subset C'$ be the end of a collar of $\partial_+ C'$ in C' , so $T_C \subset \text{int}(C')$ (shown in red in the figure). T_C is a trivial Heegaard surface for C' , dividing it into the collar of $\partial_+ C'$ and a copy A_C of C' lying entirely in $\text{int}(C')$. In the second stage of the construction of C (bottom row of Figure 15) 2-handles lying in $\mathcal{B}$ are attached to C' along $\partial_+ C'$, turning the collar of $\partial_+ C'$ into a compression body B_C . Thus T_C remains a Heegaard surface for C . (Note that, as in the outermost surface in the final panel of Figure 15, C may have spherical boundary components, so A_C and B_C would classically be described as punctured compression bodies, as discussed on [Sc1] and [FS2].). The new boundary components of C , that is $\partial C - \partial M = \partial C - \partial_- A$ are, except for the scars of the surgery, subsurfaces of T .




Figure 15: Inherited Heegaard splittings in disk decomposition

Note that in this construction, the result is technically not a Heegaard split chamber complex as defined before Lemma 3.6 because that definition requires that in each A -chamber C , A_C is a handlebody, whereas in the construction above $\partial_- A_C$ may contain components of $\partial_- A$. We will be mostly concerned with a special form of Heegaard splitting in which this is not an issue:

Definition 5.1. A Heegaard splitting $M = A \cup_T B$ is called a *pure Heegaard splitting* if all the components of ∂M lie in ∂B (or all in ∂A). Put another way, it is pure if either A or B is a handlebody.

Obviously if ∂M has at most one boundary component then the splitting is pure. Moreover,

Proposition 5.2. Suppose $\mathbb{C}$ is the chamber complex resulting from weak reduction on a Heegaard splitting of a closed 3-manifold. Then the Heegaard splitting of each chamber in $\mathbb{C}$ is pure, so the result of the weak reduction is a Heegaard split chamber complex.

Proof. Since $M = A \cup_T B$ is closed, both A and B are handlebodies. Suppose C is an A -chamber and consider the construction of its Heegaard structure above. (In Figure 15 take $\partial M = \emptyset$.) In the first stage, A is cut up by $\mathcal{A}$ into handlebodies and a component C' of the result is chosen. Next T_C is defined as the end of a collar of $\partial C'$ in C' . A_C is the complement of the collar, so in particular the surface $\partial C'$ is entirely disjoint from the handlebody A_C . That doesn't change when B_C is created by adding 2-handles to $\partial C'$ to create ∂C . Thus ∂C lies entirely in $\partial_- B_C$, as required. $\square$

There is a natural construction that is inverse to weak reduction: Suppose we are given a chamber complex $\hat{\mathbb{C}}$ with defining surface $\hat{F}$ in which we are able to alternately label the chambers A -chambers and B -chambers, for example when each component of $\hat{F}$ is separating in M . Pick a pure Heegaard splitting $C = A_C \cup_{T_C} B_C$ for each chamber C . These Heegaard splittings induce a Heegaard splitting $M = A \cup_T B$ by amalgamating the Heegaard splittings along $\hat{F}$, and the amalgamated Heegaard splitting is unique up to isotopy, cf [La, Proposition 3.1].

As described in [La], amalgamation in this setting proceeds as follows: for each A -chamber $C = A_C \cup_{T_C} B_C$ choose a complete set of meridian disks $\mathcal{B}_C$ for the compression body B_C . That is, if B_C is cut along $\mathcal{B}_C$ the result is a collar of $\partial_- B_C = \partial C - \partial M$. The disks $\mathcal{B}_C$ define a natural spine γ_C for $B_C \subset C$, namely the union of the surface $\partial C - \partial M$ and a collection of arcs in the interior of C , specifically the arc cocores of the 2-handle neighborhoods of $\mathcal{B}_C$. The arcs are then extended down through the collar using its product structure. Do the dual construction in each B -chamber. Then A can be recovered from $\hat{\mathbb{C}}$ by deleting from each A -chamber C a neighborhood of the graph γ_C and attaching a neighborhood of the graph constructed in each B -chamber. (And, as usual, symmetrically for B .)

An important point for our purposes is that it is possible to choose a *different* spine γ'_C for B_C in each A -chamber C , by originally making a different choice $\mathcal{B}'_C$ of meridian disks. This will result in a different set of arcs and so a different set of 1-handles but will not change the recovered Heegaard splitting (and dually for B -chambers). As argued more formally in [La] one can get from the arcs of γ_C to those of γ'_C by edge slides, which correspond to handle-slides in B that replace $\mathcal{B}_C$ with $\mathcal{B}'_C$ (and dually for B -chambers).

Two cautionary notes on pure Heegaard splittings:

a). The result of amalgamating pure Heegaard splittings may not be pure. For example, suppose F is a closed surface and $F \times I$ is given a pure Heegaard splitting. Take two copies of this pure Heegaard split

$F \times I$ and attach them along one end of each to get a combined copy of $F \times I$. The Heegaard splitting of $F \times I$ obtained by amalgamating the two splittings will not be pure. See top panel of Figure 16.




Figure 16: Amalgamation and weak reduction may destroy purity

b) Weak reduction, even of a pure Heegaard splitting, may not result in pure Heegaard splittings in every chamber. For example, in $M = F \times I$ consider the following natural splitting: begin with four horizontal copies of F in the interior of M labeled in order $F_1, \dots, F_4$, breaking M into 5 copies of $F \times I$, which we will call collars and which we alternately label A and B , so that both the top and bottom collar are labeled A . In each of the three collars between F_i and F_{i+1} , $i = 1, 2, 3$ tube F_i to F_{i+1} by a vertical tube. The result is a pure Heegaard splitting $M = A \cup_T B$ of genus $4 \cdot \text{genus}(F)$ in which each collar labeled A becomes part of A , and symmetrically. Now weakly reduce along the meridians of the 3 tubes, two of them disks in A and one of them in B , creating a chamber complex $\mathbb{C}$ with 5 chambers, alternately A -chambers and B -chambers and each a copy of $F \times I$. The middle 3 chambers all have pure splittings, but the chambers incident to $\partial_- M$ are each copies of $F \times I$ split into two copies by a single copy of F , and so are not pure. See bottom panel of Figure 16.

Here are some useful preliminary results. We use the standard convention that the genus of a surface is the sum of the genera over all components.

Lemma 5.3. *Suppose $M = A \cup_T B$ is a pure Heegaard splitting of a 3-manifold M . Then for any collection of components $F \subset \partial M$*

- $\text{genus}(T) \geq \text{genus}(F)$.
- $\text{genus}(T) = \text{genus}(F)$ only if each component of $\partial M - F$ is a sphere.

Proof. With no loss assume $\partial M = \partial_- B$. A spine of B consists of the union of $\partial_- B = \partial M$ and a graph γ that intersects $\partial_- B$ exactly in some valence 1 vertices. Since B is connected there has to be at least one such end vertex in each component of $\partial_- B$ and, after edge slides, we can assume there is exactly one in each, and that γ is connected. Recall that for η a regular neighborhood of a graph γ in a 3-manifold, $\chi(\partial\eta) = 2\chi(\gamma)$. T is the boundary of a regular neighborhood of the spine, so it follows immediately from this construction that if $\partial M = \partial_- B$ has n boundary components

$$\begin{aligned}\chi(T) &= (\chi(\partial_- B) - n) + (2\chi(\gamma) - n) = (\chi(\partial_- B) - 2n) + 2\chi(\gamma) \\ &= -2\text{genus}(\partial_- B) + 2\chi(\gamma).\end{aligned}$$

Hence

$$2\text{genus}(T) = 2 - \chi(T) = 2 + 2\text{genus}(\partial_- B) - 2\chi(\gamma) \geq 2\text{genus}(\partial_- B)$$

(with equality if and only if $\chi(\gamma) = 1$, i. e. γ is a tree). Hence

$$\text{genus}(T) \geq \text{genus}(\partial_- B) = \text{genus}(\partial M) = \text{genus}(F) + \text{genus}(\partial M - F)$$

so $\text{genus}(T) \geq \text{genus}(F)$ with equality only if $\text{genus}(\partial M - F) = 0$, that is $\partial M - F$ consists of spheres. $\square$

Lemma 5.4. *Suppose a closed surface F divides a 3-manifold M into two components C_1 and C_2 , and each is given a Heegaard splitting $C_i = A_i \cup_{T_i} B_i$. Let T be the Heegaard surface for the splitting of M obtained by amalgamating the two splittings. Then*

$$\chi(T) = \chi(T_1) + \chi(T_2) - \chi(F).$$

In particular, if F is connected, then

$$\text{genus}(T) = \text{genus}(T_1) + \text{genus}(T_2) - \text{genus}(F)$$

Proof. With no loss let A_1 (resp B_2) be the compression body in the splitting by T_1 (resp T_2) that is incident to F . In amalgamation, the construction of the splitting surface T of the amalgamated splitting replaces a p -punctured copy of F in T_1 and a q -punctured copy of F in T_2 with a single $(p+q)$ -punctured copy of F . The p punctures correspond to scars parallel to F of the chosen set of meridian disks in A_1 , after F_1 is compressed along these disks. A similar and symmetric statement is true for the q punctures. See [La, Figure 12].

Computing Euler characteristics from this description we have

$$\chi(T) = \chi(T_1) + \chi(T_2) - (\chi(F) + p) - (\chi(F) + q) + (\chi(F) + p + q)$$

or, more succinctly

$$\chi(T) = \chi(T_1) + \chi(T_2) - \chi(F).$$

Then, if F is connected,

$$\text{genus}(T) + \text{genus}(F) = \text{genus}(T_1) + \text{genus}(T_2)$$

follows since all surfaces are connected. $\square$

Proposition 5.5. *Suppose M is a 3-manifold with connected boundary in which each closed surface separates. Suppose $\mathbb{C}$ is a chamber complex in M with defining surface F and each chamber in $\mathbb{C}$ has a pure Heegaard splitting. Let $C = A_C \cup_{T_C} B_C$ be the chamber whose boundary contains ∂M . Suppose $M = A \cup_T B$ is the Heegaard splitting of M obtained by amalgamating the Heegaard splittings of all the chambers. Then*

- a) $\text{genus}(T_C) \leq \text{genus}(T)$
- b) *If $\text{genus}(T_C) = \text{genus}(T)$ then T_C is isotopic to T and*
- c) *If $\text{genus}(T) = \text{genus}(\partial M)$ then $M = A \cup_T B$ is a trivially split handlebody, and F consists entirely of spheres bounding balls in M . Moreover, amalgamating the Heegaard splittings in each ball bounded by a component of F gives a trivial splitting of the ball.*

Proof. With no loss, assume $\partial M \subset \partial_- B_C$.

It will be useful to consider the natural graph Γ associated to the chamber complex $\mathbb{C}$, in which each vertex corresponds to a chamber in $\mathbb{C}$ and two such vertices are connected by an edge if they are adjacent, that is there is a component of the defining surface $F = F(\mathbb{C})$ that is incident to both. Since every closed surface in M separates, Γ can contain no cycles, so it is a tree; we can take the vertex v_C corresponding to C as the root of the tree.

Case 1: M is a single chamber, so A is a handlebody and Γ is a single vertex.

In this case T_C is T , so a) and b) are automatic. For c), assume $\text{genus}(T) = \text{genus}(\partial M)$. Then no 2-handles could be attached in the construction of B_C from $\partial_+ B_C = T$, for if any attachment circle were non-separating then $\text{genus}(T) > \text{genus}(\partial M)$ and if any were separating on T then each component of the result would be the Heegaard surface for a separate chamber, contradicting the assumption that M has a single chamber. We conclude that B_C is a collar of ∂M . Since A_C is a handlebody, so then is M .

Case 2: M has two chambers, $C \supset \partial M$ and C' .

Since each component of F separates, $F = C \cap C'$ is connected. By Lemma 5.4

$$\text{genus}(T) = \text{genus}(T_C) + \text{genus}(T_{C'}) - \text{genus}(F).$$

By Lemma 5.3 $\text{genus}(T_{C'}) - \text{genus}(F) \geq 0$, so $\text{genus}(T) \geq \text{genus}(T_C)$, establishing a). Moreover, if $\text{genus}(T) = \text{genus}(T_C)$ then $\text{genus}(T_{C'}) = \text{genus}(F)$ and, by Case 1 applied to C' , C' is a trivially split handlebody. But if C' is a trivially split handlebody then amalgamating it with T_C leaves T_C unchanged, so T is isotopic to T_C , establishing b) in this case.

For c) observe that we so far have the inequalities

$$\text{genus}(T) \geq \text{genus}(T_C) \geq \text{genus}(\partial M).$$

Hence if $\text{genus}(T) = \text{genus}(\partial M)$ then both inequalities are equalities. The first equality, by the second claim, shows T_C is isotopic to T . The second equality, by Lemma 5.3 (applied to C and reversing the roles of F and ∂M) says F is a sphere, and we have already established that C' is a trivially split handlebody. Thus C' is a trivially split ball, as claimed. Moreover, Case 1 applied to the amalgamated splitting of M shows M is a trivially split handlebody, completing the proof of c) in this case.

Case 3: Every chamber in M other than C is adjacent to the chamber C . That is, Γ is a star graph: a collection of edges each of which has a single end incident to v_C .

The proof is by induction on the number n of chambers not C . Cases 1 and 2 above apply when $n = 0, 1$, so we can assume $n \geq 2$, that is there are at least 2 chambers other than C . Let C_1, C_2 be two of them, separated from C by surfaces F_1, F_2 respectively. The chamber complex $\mathbb{C}_1$ obtained from $\mathbb{C}$ by amalgamating along F_1 satisfies the hypotheses of Case 3, but with n reduced. So we inductively assume that the Proposition is true for $\mathbb{C}_1$. But further amalgamation of all of $\mathbb{C}_1$ gives Heegaard surface T and does not change ∂M so conclusion c) applies: If $\text{genus}(T) = \text{genus}(\partial M)$ then M is a trivially split handlebody and each $C_i, i \neq 1$ is a trivial split ball. Repeating the argument for the chamber complex $\mathbb{C}_2$ obtained from $\mathbb{C}$ by amalgamating along F_2 shows that $\mathbb{C}_1$ is also a trivially split ball. Thus c) is true in this case.

To prove a) observe that the inductive hypothesis applied to $\mathbb{C}_1$ shows that the Heegaard surface T_1 of the chamber of $\mathbb{C}_1$ containing ∂M has $\text{genus}(T_1) \leq \text{genus}(T)$. But that chamber is obtained from C and C_1 by amalgamation along F_1 so by Case 2) $\text{genus}(T_C) \leq \text{genus}(T_1)$. Hence $\text{genus}(T_C) \leq \text{genus}(T)$ as required to prove a).

Finally, if $\text{genus}(T_C) = \text{genus}(T)$ then we have just shown $\text{genus}(T_1) = \text{genus}(T)$ and $\text{genus}(T_C) = \text{genus}(T_1)$. The first equality shows, by inductive hypothesis, that T_1 and T are isotopic; the second shows, by Case 2), that T_C is isotopic to T_1 . Thus T_C and T are isotopic, proving b) in this case.

Case 4: There is an edge in Γ that is not incident to v_C .

There is then such an edge so that deleting the edge from Γ leaves two components: one containing v_C and the other a star graph, as in Case 3. (For example, in a path from v_C to a most distant vertex in Γ choose the second to last edge traversed.). The component F' of F corresponding to that edge then bounds a submanifold $N \subset M$ containing a chamber complex $\mathbb{C}_N$ that satisfies the hypothesis of Case 3. That is, the chamber adjacent to F' in N is adjacent to every other chamber in $\mathbb{C}_N$. Let T_N be the Heegaard surface obtained by amalgamating the chambers of N . The resulting splitting of N is pure, since $F' = \partial N$ is connected.

We induct on the number of chambers in $\mathbb{C}$. Let $\mathbb{C}'$ be the chamber complex in M obtained by amalgamating the chambers of N . The inductive hypothesis applies then to $\mathbb{C}'$ and, since $F' \neq \partial M$, the chamber C is unchanged by the amalgamation within N . Moreover, the result of amalgamating the Heegaard splittings of all the chambers in $\mathbb{C}'$ is the same as amalgamating all those in $\mathbb{C}$, so by inductive hypothesis $\text{genus}(T) \geq \text{genus}(T_C)$ and equality implies that T is isotopic to T_C . This verifies a) and b).

For c), suppose $\text{genus}(T) = \text{genus}(\partial M)$. Again applying the inductive hypothesis to $\mathbb{C}'$ we deduce that $M = A \cup_T B$ is a trivially split handlebody, and all components of F , other than those in the interior of N , are spheres bounding balls in which the amalgamated Heegaard splitting is trivial. In particular, N is a ball trivially split by T_N . Finally, apply Case 3 to $\mathbb{C}_N$: all components of F in the interior of N are also spheres bounding trivially split balls. $\square$

5.2 From weak reduction to chamber complex decomposition

There is a lot of inefficiency in the chamber complex $\hat{\mathbb{C}}$ obtained by surgery on the weakly reducing family $\mathcal{A}, \mathcal{B}$. For example, if two disks in $\mathcal{A}$ are parallel, then surgery as above turns the collar between them into its own A -chamber, one that is simply a ball with trivial Heegaard splitting. This problem is

easily addressed by viewing surgery on $\mathcal{A}, \mathcal{B}$ as the first step in a disk-decomposition of the chamber complex $A \cup_T B$ along $\mathcal{A}, \mathcal{B}$, and adopting the following rule:

Rule 5.6. *Following the weak reduction of $A \cup_T B$ along $\mathcal{A}, \mathcal{B}$, let W be a ball in M bounded by a component of $\hat{F}$.*

- *Declare W to be a goneball if and only if the Heegaard splitting of W obtained by amalgamating the Heegaard splittings of all the chambers of $\hat{\mathbb{C}}$ in $\text{int}(W)$ gives a trivial Heegaard splitting of W .*
- *For each goneball G , amalgamate the Heegaard splittings of all the chambers of $\hat{\mathbb{C}}$ in G (including amalgamation along ∂G).*

Denote the resulting chamber complex $\mathbb{C}$ (circumflex removed) and say that it is obtained from the splitting $A \cup_T B$ by *weak reduction* along $\mathcal{A} \cup \mathcal{B}$. In Proposition 5.10 below we show that Rules 4.4 and 5.6 do not conflict, so weak reduction can be viewed as disk decomposition of the chamber complex determined by T . Each chamber in $\mathbb{C}$ is Heegaard split, and furthermore no chamber in $\mathbb{C}$ is a trivially split ball.

When a goneball G is absorbed into the adjacent A -chamber C (or symmetrically), the effect on the spine γ_C of B_C is to replace a collar of ∂G in C by a single vertex whose neighborhood is the ball G . As a result, each A chamber C in $\mathbb{C}$ continues to inherit a natural spine γ_C for B_C , but instead of just a collection of arcs attached to $\partial C - \partial M$ the spine is a graph attached to $\partial C - \partial M$, containing vertices as well as arcs. As before, slides of edges in the graph over other edges and over $\partial C - \partial M$ can be used to replace one spine by another and so a chamber-derived alteration of the disks $\mathcal{B}$.

Proposition 5.7. *Following the weak reduction of $A \cup_T B$ along $\mathcal{A}, \mathcal{B}$, let W be a handlebody in M bounded by a component of $\hat{F}$. Suppose the Heegaard splitting $W = A_W \cup_{T_W} B_W$ obtained by amalgamating the splittings of all the chambers in $\text{int}(W)$ is trivial. Then*

- *$T \cap \text{int}(W)$ consists only of disks and*
- *If W is not a ball after the absorption of the goneballs given by Rule 5.6, then it becomes a single chamber in $\mathbb{C}$. If W is a ball, it becomes a goneball in $\mathbb{C}$.*

Proof. With no loss assume that the chamber in W adjacent to ∂W is an A -chamber, so $\partial W = \partial_- B_W$. Since weak reduction is inverse to amalgamation, it is easy to derive this natural description of the splitting surface T_W : Delete from ∂W the p disks corresponding to internal scars and call the resulting surface ∂W_- . Then attach to ∂W_- the surface $T \cap \text{int}(W)$, which has p boundary components, each corresponding to an internal scar in ∂W . Then push the resulting surface into $\text{int}(W)$ via a collar of ∂W .

The compression body B_W in W has a natural spine: the union of a graph $\gamma \subset \text{int}(W)$ with p end-points attached to $\partial_- B_W = \partial W$ at the centers of the internal scars. We then calculate as in the proof of Lemma 5.3:

$$\chi(T_W) = (\chi(\partial W) - p) + (2\chi(\gamma) - p) = \chi(\partial W) + (2\chi(\gamma) - 2p)$$

By assumption T_W and ∂W are homeomorphic, and this implies $\chi(\gamma) = p$. Let n be the number of components of γ . Since T_W , hence B_W , hence the spine is connected, each component of γ must contain one of the end points of γ on ∂W . This implies $n \leq p$. On the other hand, each component of γ has Euler

characteristic ≤ 1 . Hence $p = \chi(\gamma) \leq n$. Together, these inequalities that $p = n$ and each component of γ has Euler characteristic $= 1$. In other words, each component of γ is a tree. This in turn implies that each component of $T \cap \text{int}(W)$ is a disk, verifying the first assertion of the Proposition.

The last assertion c) of Proposition 5.5 further says that each component of $\hat{F}$ lying in the interior of W bounds a goneball, establishing the second claim. $\square$

Corollary 5.8. *The interior of each goneball, as defined in Rule 5.6, intersects T only in disks.*

Proof. Rule 5.6 says that each goneball has trivial Heegaard splitting, so Proposition 5.7 applies. $\square$

The discussion above leads to

Definition 5.9. *A Heegaard split chamber complex $\mathbb{C}$ for a compact manifold M is a chamber complex so that:*

- *Each chamber is labeled as either an A-chamber or B-chamber.*
- *Adjacent chambers are labeled A and B*
- *Each chamber has a pure Heegaard splitting*
- *The Heegaard splitting of each ball chamber is non-trivial.*

Note that the last property guarantees that a Heegaard split chamber complex satisfies the troublesome third requirement of the chamber complexes discussed preceding Lemma 3.6.

Proposition 5.10. *Let $M = A \cup_T B$ be a Heegaard split closed manifold and $\mathcal{A}, \mathcal{B}$ be a weakly reducing family of disks for the splitting. Let $\mathbb{C}_T$ denote the chamber complex in M defined by T . That is, the chambers are A and B . Then surgery on $A \cup \mathcal{B}$ followed by goneball removal, as just described, defines a chamber complex decomposition*

$$\mathbb{C}_T \xrightarrow{A \cup \mathcal{B}} \mathbb{C},$$

and $\mathbb{C}$ is a Heegaard split chamber complex.

Proof. Surgery on $A \cup \mathcal{B}$ is the first step in chamber complex decomposition. The next step is goneball absorption, where Rule 4.4 requires that any ball that is declared a goneball must be disk. That Rule 5.6 implies this follows immediately from Corollary 5.8. The construction above describes how each chamber is endowed with a pure Heegaard splitting.

It remains to show that the Heegaard splitting of each ball chamber in $\mathbb{C}$ is non-trivial. Suppose S is a sphere in $\hat{F}$ that bounds a ball chamber C in $\mathbb{C}$. Since the ball contains no other chambers, the only components of $\hat{F}$ that lie in the interior of C must be goneballs. Rule 5.6 says that these goneballs are absorbed, so their boundary spheres are removed, by amalgamation of Heegaard splittings. In particular, the resulting Heegaard splitting of C is obtained by amalgamation of all the chambers in $\text{int}(C)$. Since C is not itself a goneball, it follows again from Rule 5.6 that this Heegaard splitting is non-trivial, as required. $\square$

5.3 Heegaard split chamber complex decomposition

Suppose $\mathbb{C}$ is a Heegaard split chamber complex in M and $M = A \cup_T B$ is the Heegaard splitting of M obtained by amalgamation of the Heegaard splittings of the chambers of $\mathbb{C}$. Then we say that $\mathbb{C}$ *supports* the splitting $M = A \cup_T B$. So a weak reduction of $A \cup_T B$ as above produces a Heegaard split chamber complex that supports it. (See the proof of Proposition 5.14 below.)

The process of decomposing chamber complexes, as described in Section 4, extends naturally to Heegaard split chamber complexes, as we now describe. Suppose $\mathbb{C}$ is a Heegaard split chamber complex with defining surface F , and $\mathcal{D}$ is a disk set in $\mathbb{C}$ with $\partial\mathcal{D} \subset F$. In each A -chamber $C = A_C \cup_{T_C} B_C$ (and dually for the B -chambers), isotope the splitting surface T_C until it is aligned with the disk set $\mathcal{D}_C = \mathcal{D} \cap C$, as described in [Sc1]. That is, each disk $D \in \mathcal{D}_C$ will either

- intersect T_C in a single circle, and the complementary component $D \cap B_C$ of the circle in D will be a vertical annulus in B_C , or
- lie entirely in B_C , as D_1, D_2 do in Figure 17.

(D cannot lie in A_C since in an A -chamber A_C is disjoint from F .) In either case B_C has a spine Σ (illustrated in Figure 17) that intersects the disks $\mathcal{D}_C$ exactly in $\partial\mathcal{D}_C$. B_C is isotopic to a thin regular neighborhood of the spine Σ , and after the isotopy every disk $D \in \mathcal{D}_C$ intersects T_C in a single circle, with the annulus $D \cap B_C$ vertical in B_C . We say then that $\mathcal{D}$ is *aligned* with $\mathbb{C}$.




Figure 17: Alignment of disks in $\mathcal{D}_C$ that lie entirely in B_C

The first step in chamber complex decomposition along $\mathcal{D}$ is to surger F along $\mathcal{D}$, making the surface $F_{\mathcal{D}}$. When the Heegaard splittings of the chambers are aligned with $\mathcal{D}$ as above, this has two effects on an A -chamber C :

- Surgery on the B -disks incident to ∂C , lying in adjacent B -chambers, adds 2-handles to C and B_C along $\partial_- B_C$; call the results respectively C_+ and B_{C_+} . The manifold B_{C_+} is still a compression body so, in particular, T_C remains a Heegaard splitting surface for C_+ . Note that a B -disk incident to ∂C , lying in an adjacent B -chamber and inessential in that chamber may become part of an essential sphere in C_+ , cutting off a punctured ball from C_+

- Further surgery on the disks $\mathcal{D}_C$ creates a surface $T_{\hat{C}_D}$, each of whose components is a Heegaard surface for one of the chambers $\hat{C}_D$ obtained from C_+ by the surgery on $\mathcal{D}_C$. For essential disks, this observation is familiar in Heegaard theory as ∂ -reduction of a Heegaard surface (see [Sc1]). Thus $T_{\hat{C}_D}$ is a collection of Heegaard surfaces for the remnants of C in $\hat{C}_D$.

The result of the operation just described, done simultaneously on all the chambers of $\mathbb{C}$ defines a pure Heegaard splitting of each chamber in $\hat{\mathbb{C}}_D$.

Definition 5.11. Let $\mathbb{C}_D$ be the Heegaard split chamber complex in M obtained from $\hat{\mathbb{C}}_D$ by the above process, following Rule 5.6 to declare goneballs. Then $\mathbb{C}_D$ is obtained from $\mathbb{C}$ by Heegaard split chamber complex decomposition and we write

$$\mathbb{C} \xrightarrow{\mathcal{D}} \mathbb{C}_D.$$

This choice of language implies that the process is indeed a chamber complex decomposition, that is that following Rule 5.6 to declare goneballs is consistent with Rules 4.4 and 4.6. For example, we need to show that if G is a goneball as declared by Rule 5.6 then each component of $F \cap G$ is a disk, as was shown for the case $F = T$ in Corollary 5.8. That indeed both rules 4.4 and 4.6 are satisfied will follow from the next two more general propositions, see Corollaries 5.13 and 5.15.

Proposition 5.12. Let $\mathbb{C}$ be a Heegaard split chamber complex in a 3-manifold M and $F = F(\mathbb{C})$ be its defining surface. Suppose $\mathcal{D}$ is a disk set in $\mathbb{C}$ and

$$\mathbb{C} \xrightarrow{\mathcal{D}} \mathbb{C}_D$$

is a Heegaard split chamber complex decomposition. Suppose a component of F_D bounds a handlebody W in M . If the Heegaard splitting of W obtained by amalgamating the splittings of all the chambers contained in W is trivial then each component of $F_W = F \cap \text{int}(W)$ is a disk.

Proof. Denote by $M = A \cup_T B$ the Heegaard splitting of M obtained by amalgamating all the splittings of $\mathbb{C}$. Following the argument in Proposition 5.7 we know that $T_W = T \cap \text{int}(W)$ consists of disks.

We now exploit the fact that when Heegaard splittings of adjacent chambers are amalgamated to the defining surface F , F can be viewed as a subsurface of T to which disks are attached. Hence F_W is a subsurface of T_W with possibly disks attached, that is F_W is a planar surface with at most one boundary circle per component. F_W then consists of properly embedded disks and spheres.

An only slightly more complicated argument rules out the existence of spheres. Suppose there were closed components of F_W , necessarily all spheres. In this case, let F_0 be an innermost such sphere, so that it bounds a ball chamber C containing no other components of F_W . Since $\mathbb{C}$ is a Heegaard split chamber complex, C has non-trivial Heegaard splitting, so its Heegaard surface T_C contains a non-separating circle. But when the Heegaard splittings of the chambers are amalgamated, a punctured copy of T_C persists into T , and the non-separating circle can be taken to be disjoint from those punctures. Again this would contradict the fact that T_W consists of disks. We deduce that F_W consists entirely of disks. $\square$

Corollary 5.13. Let $\mathbb{C}$ be a Heegaard split chamber complex in M , $\mathcal{D}$ be a disk set in $\mathbb{C}$ and $\hat{\mathbb{C}}_D$ be the chamber complex obtained by the process described before Definition 5.11. Then following Rule 5.6 to declare goneballs is consistent with Rule 4.4.

Proof. Suppose G is a ball in M that is bounded by a sphere in $F_{\mathcal{D}}$. If G is a goneball under Rule 5.6 then the Heegaard splitting obtained by amalgamating all the splittings of chambers in $\hat{\mathbb{C}}_{\mathcal{D}}$ that lie in G is trivial. But by Proposition 5.12 this implies that $F = F(\mathbb{C})$ intersects $\text{int}(G)$ only in disks, that is G is disk. Thus G is also a goneball under Rule 4.4. $\square$

Proposition 5.14. *Suppose $\mathbb{C}$ is a Heegaard split chamber complex for M that supports the Heegaard splitting $M = A \cup_T B$. Suppose $\mathcal{D}$ is an aligned disk set in $\mathbb{C}$. Let $\hat{\mathbb{C}}_{\mathcal{D}}$ be the chamber complex obtained by the process described before Definition 5.11. Then amalgamating all the induced Heegaard splittings of the chamber complex $\hat{\mathbb{C}}_{\mathcal{D}}$ yields the original Heegaard splitting $M = A \cup_T B$.*

Proof. We are given that amalgamating all the Heegaard splittings of the chambers in $\mathbb{C}$ gives $M = A \cup_T B$. Our assumption is that each disk in $\mathcal{D}$ is aligned with the splitting of the chamber in which the disk lies. Suppose $D \in \mathcal{D}$ lies, without loss, in an A -chamber $C = A_C \cup_{T_C} B_C$ and $C' = A_{C'} \cup_{T_{C'}} B_{C'}$ is the adjacent B -chamber on whose boundary ∂D lies. The process described before Definition 5.11 affects a bicollar neighborhood $D \times [-1, 1]$ of D as follows (see Figure 18):

- The ball $\text{int}(D) \times (-1, 1)$ is moved from $\text{int}(C)$ to $\text{int}(C')$.
- The annulus $\partial D \times [-1, 1]$ is deleted from the spines of $A_{C'}$ and B_C .
- The disks $D \times \{\pm 1\}$ are added to the spines of $A_{C'}$ and B_C .
- The arc $\{0\} \times [-1, 1]$ is added to the spine of $A_{C'}$.

On the other hand, amalgamation of the splittings of $\hat{\mathbb{C}}_{\mathcal{D}}$ does exactly the opposite near $D \times [-1, 1]$: a tube around $\{0\} \times [-1, 1]$, which we can take to be $\partial D \times [-1, 1]$, is added to $F_{\mathcal{D}}$ after deleting the disks $D \times \{\pm 1\}$, undoing each of the items above. Extending this observation now to all the disks in $\mathcal{D}$, the result of amalgamating all the Heegaard splittings after the process described before Definition 5.11 is the same as amalgamating before the process, namely $M = A \cup_T B$. $\square$




Figure 18: Decomposition near $D \in \mathcal{D}$

Corollary 5.15. *Let $\mathbb{C}$ be a Heegaard split chamber complex in M , $\mathcal{D}$ be a disk set in $\mathbb{C}$ and $\hat{\mathbb{C}}_{\mathcal{D}}$ be the chamber complex obtained by the process described before Definition 5.11. Then following Rule 5.6 to declare goneballs is consistent with Rule 4.6.*

Proof. Suppose G is a ball in M that is bounded by a sphere in $F_{\mathcal{D}}$. Suppose further that the disk $D \in \mathcal{D}$ leaves no scar on ∂G , and define $\mathcal{D}_- = \mathcal{D} - D$. Since D leaves no scar on ∂G , the sphere is also a component of $F_{\mathcal{D}_-}$. Moreover, if we let $\mathcal{D}_{\partial} \subset \mathcal{D}_-$ be the set of disks in $\mathcal{D}$ that do leave scars on ∂G , then ∂G is also a component of $F_{\mathcal{D}_{\partial}}$ and, as described before Definition 5.11, G inherits a chamber complex structure $\hat{\mathbb{C}}_{\partial}$, in which each chamber is Heegaard split. Let T_G be the Heegaard splitting of the ball G obtained by amalgamating all these splittings. That is, $\hat{\mathbb{C}}_{\partial}$ supports T_G .

According to Proposition 5.14 the chamber complexes in G obtained from $\hat{\mathbb{C}}_{\partial}$ by decomposition along $\mathcal{D} \cap G$ or along $\mathcal{D}_- \cap G$ also support T_G . In particular, under Rule 5.6, G is a goneball in $\mathbb{C}_{\mathcal{D}}$ if and only if T_G is a trivial splitting of G and this is true if and only if G is also a goneball in $\mathbb{C}_{\mathcal{D}_-}$. Thus Rule 4.6 holds. $\square$

Theorem 5.16. *Suppose $\mathbb{C}$ is a Heegaard split chamber complex for M that supports the Heegaard splitting $M = A \cup_T B$ and $\mathcal{D}$ is a disk set in $\mathbb{C}$. After perhaps a proper isotopy of $\mathcal{D}$ within the chambers, not moving T , the chamber complex decomposition*

$$\mathbb{C} \xrightarrow{\mathcal{D}} \mathbb{C}_{\mathcal{D}}$$

is a Heegaard split chamber complex decomposition and the resulting Heegaard split chamber complex $\mathbb{C}_{\mathcal{D}}$ also supports T .

Proof. This would seem to be a straightforward consequence of Proposition 5.14 and the discussion at the beginning of Subsection 5.3. Namely:

1. In each chamber C isotope the splitting surface T_C until it is aligned with $\mathcal{D}$, as is shown to be possible in [Sc1].
2. Mimic the handle slides used in the isotopies within each chamber by handleslides on T itself, as discussed for example in [La].
3. Now that T and $\mathcal{D}$ are aligned, apply Proposition 5.14.

This (ultimately successful) strategy has an obvious conceptual weakness: The process requires isotoping T by various handleslides, so T will typically end up in a different position in M than it began. We will show that the process above can be reframed so that T is fixed throughout and only the disks $\mathcal{D}$ and the chambers $\mathbb{C}$ are allowed to move. We must further ensure that the disks $\mathcal{D}$ stay disjoint, in particular that the boundaries of disks in $\mathcal{D}$ that lie in adjacent chambers don't end up intersecting, after their proper isotopies, in the component of $F = F(\mathbb{C})$ that lies between them. That is the goal of the argument that follows.

Claim: There is an isotopy of $\mathcal{D}$, fixed on $\partial \mathcal{D} \subset F$, so that afterwards $\mathcal{D}$ is aligned with the splitting surface of each chamber.

The proof of the Claim is straightforward and well-known: Let C be a chamber, $\mathcal{D}_C$ be the set of disks in $\mathcal{D}$ that lie in C , and $T_C \subset \text{int}(C)$ be the Heegaard splitting surface of the chamber. Per [Sc1] there is an isotopy of embeddings $\phi_t : T_C \rightarrow C$ so that ϕ_0 is the original embedding and $\phi_1(T_C)$ is aligned with $\mathcal{D}_C$. By the isotopy extension theorem, ϕ_t can be extended to an isotopy $\theta_t : C \rightarrow C$ in which θ_0 is the identity. Since $T_C \subset \text{int}(C)$, we can also take θ_t to be fixed on ∂C .

Define an isotopy of embeddings $\rho_t : \mathcal{D}_C \rightarrow C$ by $\rho_t = (\theta_t)^{-1}|_{\mathcal{D}_C}$. Observe that ρ_0 is the original embedding (since θ_0 is the identity). Moreover, it is easy to check that $\rho_1(\mathcal{D})$ is aligned with T_C . Namely, observe that

$$\theta_1(\rho_1(\mathcal{D}_C) \cap T_C) = \mathcal{D}_C \cap \theta_1(T_C) = \mathcal{D}_C \cap \phi_1(T_C)$$

and the last term consists, by construction, of at most one circle in each component of $\mathcal{D}_C$. Hence $\rho_1(\mathcal{D}_C) \cap T_C$ consists of at most one circle in each component of $\mathcal{D}_C$, that is the disks $\rho_1(\mathcal{D}_C)$ are aligned with T_C . This proves the claim.

We continue our effort to bring $\mathcal{D}$ into alignment with T , not by handleslides that move T (mimicking handleslides within chambers), but by proper isotopies of $\mathcal{D}$. In doing so, we do not use the Claim *per se*, but rather note that the proof of the claim implies that we only need to show how to keep T fixed, during the alignment, in a collar of the boundary of each chamber (i. e. near the defining surface F). Then a way to keep T also fixed outside the collar, that is away from F , is provided by the proof of the Claim. So we focus on how handle slides of the original splitting $M = A \cup_T B$ that mimic handle slides in the splitting of a chamber $C = A_C \cup_{T_C} B_C$ can be replaced *near* F by proper isotopies of $\mathcal{D}$.

Consider a bicollar neighborhood $F_0 \times [-1, 1]$ of a component F_0 of F that separates an A -chamber C of $\mathbb{C}$ from a B -chamber C' . Here we take $F_0 \times [0, 1] \subset C$ and $F_0 \times [-1, 0] \subset C'$. F_0 itself is a subsurface of T whose boundary is capped off by disk scars coming from surgery on a collection of disks $(\mathcal{B}, \partial\mathcal{B}) \subset (B_C, \partial B_C)$ and disk scars coming from surgery on a collection of disks $(\mathcal{A}, \partial\mathcal{A}) \subset (A_{C'}, \partial A_{C'})$. B intersects $F_0 \times [0, 1]$ in vertical cylinders, the “legs” of 1-handles in B_C , one for each scar of $\mathcal{B}$ on F_0 . We denote the leg corresponding to a scar b by $b \times [0, 1]$. (The symmetric statement is true for $A \cap (F_0 \times [-1, 0])$.)

Let $\mathcal{D}_C = \mathcal{D} \cap C$ and $\mathcal{D}_{C'} = \mathcal{D} \cap C'$. Also spanning the collar $F_0 \times [0, 1]$ are vertical annuli, each of them a collar neighborhood of the boundary of a disk in $\mathcal{D}_C$. The symmetric statement is true in the collar $F_0 \times [-1, 0]$. By general position in the surface F_0 we can take the legs to be disjoint from the annuli, and, since $\mathcal{D}$ is embedded, we have that $\partial\mathcal{D}_C$ and $\partial\mathcal{D}_{C'}$ are disjoint in F_0 .

Consider how a handle slide in B that mimics a handle slide of B_C in C appears in this collar. As it begins, the end of a leg $b_1 \times [0, 1]$ is slid via a path $\gamma \subset F_0$ to the end of a leg $b_2 \times [0, 1]$ and then up and out of view. At the end of the handle slide the same thing occurs elsewhere in the collar. So it is only this initial move that needs to be understood. By general position we may assume that the path γ is disjoint from all scars. Here are possibilities:

1. the path γ is disjoint from both $\partial\mathcal{D}_C$ and $\partial\mathcal{D}_{C'}$, as in the left side of Figure 19.
2. the path γ is disjoint from $\partial\mathcal{D}_C$ but not $\partial\mathcal{D}_{C'}$.
3. the path γ intersects $\partial\mathcal{D}_C$ and possibly also $\partial\mathcal{D}_{C'}$.

In the first case, where γ is disjoint from $\partial\mathcal{D}$, this first stage of the handle slide can be accomplished, leaving T fixed, by just replacing the disk in $\mathcal{B}$ whose scar is b_2 by the disk b' in $\mathcal{B}$ obtained by band-summing b_1 to b_2 along γ . This new disk is disjoint from other scars and, by the assumption of this case, still disjoint from $\partial\mathcal{D}$. See right side of Figure 19.




Figure 19: When γ is disjoint from both $\partial\mathcal{D}_C$ and $\partial\mathcal{D}_{C'}$

In the second case, γ may intersect only $\partial\mathcal{D}_{C'}$. These intersection points can be removed by isotoping $\partial\mathcal{D}_{C'}$ along γ towards and then across the disk $b_1 \subset F_0 \subset \partial\mathcal{C}'$. See Figure 20. (Another description is that we band sum $\mathcal{D}_{C'}$ at the points it crosses γ to copies of b_1 by bands around subsegments of γ). Such a proper isotopy of $\mathcal{D}_{C'}$ is allowed, and reduces this case to the previous case.





Figure 20: When γ is disjoint from $\partial\mathcal{D}_C$ but not $\partial\mathcal{D}_{C'}$

In the third case, the slide of $b_1 \times [0, 1]$ can be broken up into a series of slides along γ , each successive one across a single intersection of γ with $\partial\mathcal{D}_C$. For example, let γ_1 be the segment of γ lying between b_1 and the closest point p of $\gamma \cap \partial\mathcal{D}_C$. Just as in step 2, $\mathcal{D}_{C'}$ can be properly isotoped across b_1 so that it is disjoint from γ_1 . The slide of the leg of $b_1 \times [0, 1]$ along γ_1 creates a disk intersection of the leg with $\mathcal{D}_C$; as the leg is then straightened from below the disk intersection moves upwards out of view. There is an obvious reinterpretation of this step that keeps the leg (hence T) fixed: instead slide $p \in \partial\mathcal{D}_C$ and neighboring points of $\mathcal{D}_C$ along γ_1 and across the leg $b_1 \times [0, 1]$. See Figure 21. Again this creates an intersection disk and again, as the annulus that contains p is made vertical from below, the disk ascends out of view. Note that because we have already cleared all points of $\partial\mathcal{D}_{C'}$ from γ_1 , $\partial\mathcal{D}_{C'}$ and $\partial\mathcal{D}_C$ remain disjoint, as required. Continue the slide of the leg along the rest of γ , proceeding as just described across each successive segment of $\gamma - \partial\mathcal{D}_C$, eventually reducing this case to the second case, which we have already considered. $\square$


 Figure 21: When γ intersects $\partial\mathcal{D}_C$ and possibly also $\partial\mathcal{D}_{C'}$

5.4 Tiny Heegaard split chamber complexes

Proposition 4.10 has an analogue in Heegaard split chamber complex decompositions, as we now describe:

Definition 5.17. A Heegaard split chamber complex $\mathbb{C}$ is tiny if $\mathbb{C}$ is tiny as a chamber complex, and each designated handlebody has a trivial Heegaard splitting.

Note that in a tiny Heegaard split chamber complex no designated handlebody can be a ball, since, by definition, no chamber in a Heegaard split chamber complex is a trivially split ball. Also, since each designated handlebody has only one boundary component, each must be adjacent to the unique chamber C that is not a designated handlebody. Hence if C is a B -chamber, all the designated handlebodies are A -chambers, and symmetrically.

Proposition 5.18. Suppose $\mathbb{C}$ is a Heegaard split chamber complex obtained from a Heegaard splitting $M = A \cup_T B$ by weak reduction. Then $\mathbb{C}$ is not tiny.

Proof. Let $\mathcal{A}, \mathcal{B}$ be the weakly reducing family of disks, and D_A be an essential disk in $\mathcal{A}$. Then after surgery on $\mathcal{A} \cup \mathcal{B}$, D_A lies in a B -chamber C , with ∂D_A on a component of $T \cap \text{int}(C)$. But that component can't be a disk, since ∂D_A is essential on T . Hence C is not a disk handlebody and so, by Proposition 5.7 it is not a trivially split handlebody. To summarize: some B -chamber of $\mathbb{C}$ is not a trivially split handlebody.

The symmetric argument on a disk $D_B \in \mathcal{B}$ shows that there is some A chamber of $\mathbb{C}$ that is not a trivially split handlebody. This immediately implies that the defining surface $F(\mathbb{C}) \neq \emptyset$ and also contradicts the consequence of Definition 5.17 that either all A -chambers or all B -chambers are designated handlebodies, and so must be trivially split. $\square$

Proposition 5.19 (Tinyness pulls back). Let $\mathbb{C} \xrightarrow{\mathcal{D}} \mathbb{C}_{\mathcal{D}}$ be a Heegaard split chamber complex decomposition as given in Definition 5.11. If $\mathbb{C}_{\mathcal{D}}$ is a tiny Heegaard split chamber complex, so is $\mathbb{C}$.

Proof. Proposition 4.10 and most of its proof apply in this situation: We are given that the designated handlebodies of $\mathbb{C}_{\mathcal{D}}$ have trivial Heegaard splitting so, per Proposition 5.12, each such chamber is disk, and Proposition 4.10 applies. That is, $\mathbb{C}$ is tiny as a chamber complex. What remains to be shown is that the designated handlebodies of $\mathbb{C}$ have trivial Heegaard splittings. The proof of Proposition 4.10 is

also valid in this setting; that proof invokes Proposition 4.8 to conclude that all but one chamber of $\mathbb{C}$ is a handlebody, and these become the designated handlebodies of $\mathbb{C}$. Then this strengthened version of Proposition 4.8 implies, in our setting, that each such designated handlebody is trivially split:

Lemma 5.20. *Suppose C is a chamber of $\mathbb{C}$ so that every remnant of C in $\mathbb{C}_{\mathcal{D}}$ is a handlebody with trivial Heegaard splitting. Then C is a handlebody with trivial Heegaard splitting.*

Proof. Proposition 4.8 suffices to conclude that C is a handlebody and the proof of that proposition continues to apply here. According to that proof, all that is required to complete the proof of Lemma 5.20 is this additional claim:

Claim: Suppose C is a handlebody chamber in $\mathbb{C}$ and the disks $\mathcal{D}_C$ in $\mathcal{D}$ that are incident to ∂C all lie within C . Then C is trivially split.

Proof of Claim: Without loss suppose C is an A -chamber and denote by C_r the collection of remnants of C , by hypothesis each a trivially split handlebody A -chamber. Let $F = \partial C$ and let T_C be the chamber's Heegaard surface. Then, by Proposition 5.5c, $\chi(T_C) \leq \chi(F)$, with equality only if the Heegaard splitting of C is trivial.

Let $\hat{T}_{\mathcal{D}}$ be the union of the Heegaard surfaces for the chambers of $\hat{\mathbb{C}}_{\mathcal{D}}$ that are remnants of C (some of which may be goneballs). When aligned, as described at the beginning of Subsection 5.3, each disk in $\mathcal{D}_C$ intersects F in a single circle (its boundary) and intersects T_C either in a single circle or not at all. In the latter case, the disk lies entirely in B_C and, since $\partial_- B_C$ is incompressible in the compression body B_C , the boundary of the disk is inessential in $\partial_- C$. Hence surgery on such a disk creates a new sphere component of $F_{\mathcal{D}}$ bounding a trivially split ball, so such a surgery adds a sphere component both to $\hat{T}_{\mathcal{D}}$ and to $F_{\mathcal{D}}$. Similarly, surgery on a disk in $\mathcal{D}$ that intersects T in a single circle raises the Euler characteristic of both $\hat{T}_{\mathcal{D}}$ and F by 2. So in the end, surgery on $\mathcal{D}$ raises the Euler characteristic of both $\hat{T}_{\mathcal{D}}$ and F by $2|\mathcal{D}_C|$. (That is, $\chi(F_{\mathcal{D}}) - \chi(F) = 2|\mathcal{D}_C|$.) Consequently, $\chi(\hat{T}_{\mathcal{D}}) \leq \chi(F_{\mathcal{D}})$, again with equality only if C is trivially split.

Consider next what happens when the bounding sphere of a goneball is removed from $F_{\mathcal{D}}$, and $\hat{T}_{\mathcal{D}}$ is altered by amalgamation along that sphere. Since goneballs are exactly those balls that are trivially split, the result is to remove exactly a sphere from both $F_{\mathcal{D}}$ and $\hat{T}_{\mathcal{D}}$, lowering the Euler characteristic of both by 2. Once all spheres bounding goneballs are removed, we then have $\chi(T_{\mathcal{D}}) \leq \chi(\partial C_r)$, again with equality only if $\mathbb{C}$ is trivially split. (Here $T_{\mathcal{D}}$ is the union of all the Heegaard surfaces of the chambers of C_r .) By assumption, each component of $\mathbb{C}_r$ is trivially split. This means that $\chi(T_{\mathcal{D}}) = \chi(\partial C_r)$, so $\chi(T_C) = \chi(F)$ and indeed C is trivially split. This proves the Claim, hence the Lemma and so the Proposition. $\square$

$\square$

6 Sequences of aligned chamber complex decompositions in S^3

There is an important caveat about the structure of a Heegaard split chamber complex decomposition: Typically the Heegaard splitting of a specific chamber of $\mathbb{C}_{\mathcal{D}}$ will not be well-defined. Indeed, even the genus of the chamber may be ambiguous, since, in the process described before Definition 5.11 there is a choice of how to align $\mathcal{D}_C$ with T_C in a chamber C of $\mathbb{C}$. One aspect of this choice (see [FS2]) is that a bubble may be moved across a disk in $\mathcal{D}_C$; if, for example, the disk is separating, this changes the genus

of the chambers on each side of the disk. One result of this ambiguity is that $\mathbb{C}_{\mathcal{D}}$ may not be well-defined even as a chamber complex: depending on alignment, a ball chamber in $\hat{\mathbb{C}}_{\mathcal{D}}$ may or may not be trivially split, and therefore may or may not appear as a chamber in $\mathbb{C}_{\mathcal{D}}$.

Soon (see Section 7) this ambiguity will be directly addressed, by describing a *preferred* way of aligning disks, a way that will eventually suffice to certify unambiguously a chamber complex obtained by weak reduction of a Heegaard splitting of S^3 . In this section we avoid much of this ambiguity by oversimplifying the full theory. The section is motivational and transitional: We present an application of the theory of disk decompositions developed in Sections 4 and 5, one that extends Proposition 3.8 and so connects to our previous discussions. It also is a model for later arguments we will need in a more complicated setting. We begin with a definition that will only be used in this section, as a way of illustrating and simplifying the full argument that will eventually follow.

Definition 6.1. *In the setting of Definition 5.11, suppose $\mathbb{C}_{\mathcal{D}}$ is a Heegaard split chamber complex obtained from $\hat{\mathbb{C}}_{\mathcal{D}}$ by declaring as goneballs all disk balls. Then*

$$\mathbb{C} \xrightarrow{\mathcal{D}} \mathbb{C}_{\mathcal{D}}$$

is an aligned chamber complex decomposition.

Definition 6.2. *Suppose $\mathbb{C}$ is a Heegaard split chamber complex, that supports a genus g Heegaard splitting $M = A \cup_T B$; C is a chamber in $\mathbb{C}$; and $\mathfrak{b}$ is a genus $g' < g$ bubble in the Heegaard splitting $C = A_C \cup_{T_C} B_C$. Let $\mathbb{C}/\mathfrak{b}$ be the chamber complex, with Heegaard splittings of each chamber, obtained by destabilizing $C = A_C \cup_{T_C} B_C$ along $\mathfrak{b}$, replacing $\mathfrak{b}$ with a neighborhood of a point $*$ in the splitting surface.*

Amalgamating all Heegaard splittings of $\mathbb{C}/\mathfrak{b}$ gives a genus $g - g'$ Heegaard splitting of M obtained by destabilizing T along $\mathfrak{b}$.

Although every chamber in $\mathbb{C}/\mathfrak{b}$ is Heegaard split, $\mathbb{C}/\mathfrak{b}$ may fail to be a Heegaard split chamber complex for a somewhat technical reason: when C is a ball and $\mathfrak{b}$ is a maximal bubble in C , the chamber C in $\mathbb{C}/\mathfrak{b}$ becomes a trivially split ball chamber and this is not allowed in a Heegaard split chamber complex. Nonetheless, any disk set $\mathcal{D}$ in $\mathbb{C}/\mathfrak{b}$ can be aligned with the Heegaard splittings of the chambers, the construction described in Subsection 5.3 carried out, and all disk balls declared goneballs. Denote this aligned chamber complex decomposition

$$\mathbb{C}/\mathfrak{b} \xrightarrow{\mathcal{D}} (\mathbb{C}/\mathfrak{b})_{\mathcal{D}}.$$

Lemma 6.3. *Suppose, in the setting of Definition 6.2, $\mathcal{D}$ is a disk set in $\mathbb{C}$. Then $\mathcal{D}$ can be aligned in $\mathbb{C}$ and in $\mathbb{C}/\mathfrak{b}$ so that $\mathfrak{b}$ remains a bubble in a chamber of $\mathbb{C}_{\mathcal{D}}$ and $\mathbb{C}_{\mathcal{D}}/\mathfrak{b} = (\mathbb{C}/\mathfrak{b})_{\mathcal{D}}$.*

Proof. This is immediate from the definitions: Align the disks $\mathcal{D}$ in $\mathbb{C}/\mathfrak{b}$ so that, by general position, they do not contain the point $*$. This defines also an alignment of $\mathcal{D}$ in $\mathbb{C}$. Following surgery on $\mathcal{D}$ absorb goneballs in both $\hat{\mathbb{C}}$ and $\mathbb{C}/\mathfrak{b}$. They are the same goneballs because in both cases the goneballs are exactly the disk balls. These are defined at the level of chamber complexes, and do not depend on the Heegaard splittings of the chambers. $\square$

Call an alignment as in Lemma 6.3 a *b-wise alignment*.

Suppose a Heegaard split chamber complex $\mathbb{C}_0$ supports T . For example $\mathbb{C}_0$ could be the Heegaard splitting $A \cup_T B$ itself. Let

$$\vec{\mathbb{C}}: \quad \mathbb{C}_0 \xrightarrow{\mathcal{D}_0} \mathbb{C}_1 \xrightarrow{\mathcal{D}_1} \mathbb{C}_2 \xrightarrow{\mathcal{D}_2} \dots \xrightarrow{\mathcal{D}_{n-1}} \mathbb{C}_n$$

be a sequence of chamber complex decompositions.

Proposition 6.4. *Suppose chambers $C_0 \in \mathbb{C}_0$ and $C_n \in \mathbb{C}_n$ contain, respectively, incompressible spheres S and S' . Then, for iteratively $0 \leq i \leq n-1$ each disk set $\mathcal{D}_i$ can be aligned in $\mathbb{C}_i$, so that $\mathbb{C}_0$ and the resulting Heegaard split chamber complex structure on $\mathbb{C}_n$ cocertify.*

Proof. We induct on $n \geq 1$.

A standard innermost disk argument shows that there is an incompressible sphere S_0 in the chamber C_0 that is disjoint from the disks $\mathcal{D}_0$. This implies that S_0 will lie in a single chamber C_1 of the chamber complex $\mathbb{C}_1$. Align the disks $\mathcal{D}_0$ as in Theorem 5.16 so that

$$\mathbb{C}_0 \xrightarrow{\mathcal{D}_0} \mathbb{C}_1$$

is an aligned disk decomposition. More specifically, in chamber C_0 choose an alignment, per [Sc1], so that the incompressible sphere S_0 is also aligned with the Heegaard splitting surface of the chamber and hence, after amalgamation, also aligned with T , per Proposition 3.7. Following Proposition 3.8 we may as well substitute S_0 for S .

If S_0 is incompressible in C_1 , then, according to Propositions 3.7 and Corollary 3.9, $h_{S_0}: (S^3, T) \rightarrow (S^3, T_g)$ defines both $h_{F(\mathbb{C}_0)}$ and $h_{F(\mathbb{C}_1)}$. Thus $h_{F(\mathbb{C}_0)} \sim h_{F(\mathbb{C}_1)}$. By inductive assumption $h_{F(\mathbb{C}_1)} \sim h_{F(\mathbb{C}_n)}$, completing the proof in this case.

However, it is possible that the sphere S_0 is compressible in the chamber C_1 . (For example, the interior of a ball W that S_0 bounds in C_1 could have contained a single handlebody chamber of $\mathbb{C}_0$, one that disappears in $\mathbb{C}_1$ because it is decomposed by $\mathcal{D}_0$ into a ball that is a goneball.) If S_0 is compressible in C_1 then it bounds a non-trivial bubble b in (S^3, T) , by [Wa], one that already appears as a bubble in the induced Heegaard splitting of $\mathbb{C}_1$. (The bubble is non-trivial because the presence in $\text{int}(W)$ of other chambers means that W is not a trivially split ball, per Proposition 5.12.) In this case, we argue as follows:

For the rest of the decompositions

$$\mathbb{C}_1 \xrightarrow{\mathcal{D}_1} \mathbb{C}_2 \xrightarrow{\mathcal{D}_2} \dots \xrightarrow{\mathcal{D}_{n-1}} \mathbb{C}_n$$

in the sequence, choose *b-wise* alignments. Then the bubble b (bounded still by S_0) arrives intact in a single chamber of $\mathbb{C}_n$. In $\mathbb{C}_n/b$, S' can be taken to be disjoint from $*$, simply by general position. Then in $\mathbb{C}_n$, b is disjoint from S' . Then Lemma 3.5 shows that h_{S_0} and $h_{S'}$ are eyeglass equivalent. Hence $h_{F(\mathbb{C}_0)} \sim h_{S_0} \sim h_{S'} \sim h_{F(\mathbb{C}_n)}$ as required. $\square$

Having shown there exists a choice of alignment for which $\mathbb{C}_0$ and $\mathbb{C}_n$ cocertify, we next show that any choice of alignments will do. The critical point is that we understand how different alignments of disk-sphere sets in a Heegaard split 3-manifold are related: By [FS2] they differ by a sequence of eyeglass moves and passing bubbles through the disk-sphere set. We begin with a lemma and proposition that are not specific to Heegaard splittings of S^3 .

Lemma 6.5. Suppose $\mathbb{C}_0$ is a Heegaard split chamber complex in M that supports the splitting $M = A \cup_T B$, and

$$\mathbb{C}_0 \xrightarrow{\mathcal{D}_0} \mathbb{C}_1$$

is a chamber complex decomposition. Suppose $E \subset \mathbb{C}_1$ is a disk-sphere set.

Let $\mathcal{D}_0^x$ and $\mathcal{D}_0^y$ be possibly different alignments of the disk set $\mathcal{D}_0$ in $\mathbb{C}_0$, with resulting Heegaard split chamber complexes $\mathbb{C}_1^x$ and $\mathbb{C}_1^y$ respectively, and similarly let E^x, E^y be alignments of the disk-sphere set E in $\mathbb{C}_1^x$ and $\mathbb{C}_1^y$ respectively. Then a sequence of bubble passes and eyeglass moves will change the alignment $\mathcal{D}_0^x$ to $\mathcal{D}_0^y$ and E^x to E^y .

Proof. By [FS2] there is a series of bubble passes and eyeglass moves that change the alignment $\mathcal{D}_0^x$ to $\mathcal{D}_0^y$. The proof of the Lemma is by induction on p , the number of bubble passes required. If $p = 0$ then, after an eyeglass move, we can take $\mathcal{D}_0^x = \mathcal{D}_0^y$, so $\mathbb{C}_1^x = \mathbb{C}_1^y$ and then apply [FS2] to the disk sets E^x and E^y in $\mathbb{C}_1^x = \mathbb{C}_1^y$.

So suppose $p \geq 1$ and assume the Lemma is true whenever the number of bubble passes needed to change the alignment $\mathcal{D}_0^x$ to $\mathcal{D}_0^y$ is less than p . Let b be the first ball that is passed, through a disk $D \in \mathcal{D}_0^x$, in changing the alignment $\mathcal{D}_0^x$ to $\mathcal{D}_0^y$. Say b is passed from a chamber $C \in \hat{\mathbb{C}}_1^x$ to a chamber $C' \in \hat{\mathbb{C}}_1^x$, with possibly $C = C'$. (Enthusiasts will note that if $C = C'$ then the bubble pass is also an eyeglass move, by Lemma 2.5, so p can be reduced and we are done.)

Consider the subset E^C of the aligned disk set E^x that lies in C . As aligned, E^C may well intersect the interior of b . However there is *some* alignment E^b of E^C in C that is disjoint from the bubble b . This follows from general position and [Sc1] applied to E^C in $\hat{\mathbb{C}}_1/b$. Then by [FS2] the alignment E^b can be obtained from E^C by eyeglass moves and bubble passes through E in C . Thus we may as well assume that $E^C = E^b$.

By further bubble passes of b through E^C we can move b adjacent to the disk $D \in \mathcal{D}_0^x$ and complete the pass of b through D . Then only $p - 1$ further bubble passes are needed to move $\mathcal{D}_0^x$ to $\mathcal{D}_0^y$, completing the inductive step. $\square$

This argument generalizes, though the statement and argument become more complex. Let $\mathbb{C}_0$ be a Heegaard split chamber complex in M that supports the Heegaard splitting $M = A \cup_T B$. Suppose

$$\vec{\mathbb{C}}: \quad \mathbb{C}_0 \xrightarrow{\mathcal{D}_0} \mathbb{C}_1 \xrightarrow{\mathcal{D}_1} \mathbb{C}_2 \xrightarrow{\mathcal{D}_2} \dots \xrightarrow{\mathcal{D}_{n-1}} \mathbb{C}_n$$

is a sequence of chamber complex decompositions and E is a disk-sphere set in $\mathbb{C}_n$.

Suppose that the disks $\mathcal{D}_i$ are aligned at each successive stage of the decomposition, that

$$\vec{\mathbb{C}}^x: \quad \mathbb{C}_0 \xrightarrow{\mathcal{D}_0^x} \mathbb{C}_1^x \xrightarrow{\mathcal{D}_1^x} \mathbb{C}_2^x \xrightarrow{\mathcal{D}_2^x} \dots \xrightarrow{\mathcal{D}_{n-1}^x} \mathbb{C}_n^x$$

is the resulting aligned chamber complex decomposition sequence, and that E is aligned with the resulting Heegaard splitting of $\mathbb{C}_n^x$. Suppose further that for some $0 \leq i \leq n$ there is a bubble b in one of the chambers C of $\mathbb{C}_i$ and the bubble b remains disjoint from each set of disks $\mathcal{D}_j, j \geq i$ and from $E \subset \mathbb{C}_n$. That is, the alignment of the disks in the remainder of the sequence is b -wise. Then say that b is an *intact bubble* for the sequence of aligned chamber complex decompositions.

Suppose next that the alignment of some disk set $\mathcal{D}_i$ is altered simply by a bubble pass of b to another chamber C' of $\mathbb{C}_{i+1}$, and each successive alignment of the $\mathcal{D}_j, j \geq i+1$ and the alignment of E all are b -wise, so that b remains an intact bubble for the resulting aligned chamber complex decomposition

$$\vec{\mathbb{C}}^y: \quad \mathbb{C}_0 \xrightarrow{\mathcal{D}_0^y} \mathbb{C}_1^y \xrightarrow{\mathcal{D}_1^y} \mathbb{C}_2^y \xrightarrow{\mathcal{D}_2^y} \dots \xrightarrow{\mathcal{D}_{n-1}^y} \mathbb{C}_n^y$$

and alignment of E . (By construction $\mathbb{C}_k^x = \mathbb{C}_k^y$ and $\mathcal{D}_{k-1}^x = \mathcal{D}_{k-1}^y$ for $1 \leq k \leq i$.) Denote the two alignments of E by E^x and E^y respectively.

Definition 6.6. *In the construction above, the aligned chamber complex decomposition sequence $\vec{\mathbb{C}}^y$ and the alignment E^y of E are said to be obtained from $\vec{\mathbb{C}}^x$ and E^x by an ancestral bubble pass. This includes the degenerate case, in which $\vec{\mathbb{C}}^x = \vec{\mathbb{C}}^y$ but the alignments E^x, E^y differ by a bubble pass in $\mathbb{C}_n$.*

Proposition 6.7. *Suppose aligned chamber complex decomposition sequences $\vec{\mathbb{C}}^x$ and $\vec{\mathbb{C}}^y$ and alignments E^x, E^y are chosen for the chamber complex decomposition sequence $\vec{\mathbb{C}}$ and the disk-sphere set E in $\mathbb{C}_n$. Then there is a sequence of ancestral bubble passes and eyeglass moves that changes the pair $(\vec{\mathbb{C}}^x, E^x)$ to $(\vec{\mathbb{C}}^y, E^y)$.*

Proof. The proof begins much as the proof of Lemma 6.5. By [FS2] there is a series of bubble passes and eyeglass moves that change the alignment $\mathcal{D}_0^x$ to $\mathcal{D}_0^y$. The proof is by induction on the pair (n, p) , lexicographically ordered, where p is the minimal number of bubble passes in such a series. Lemma 6.5 covers the case $n = 1$, so assume $n \geq 2$. In case $p = 0$, we can take $\mathcal{D}_0^x = \mathcal{D}_0^y$, so $\mathbb{C}_1^x = \mathbb{C}_1^y$ and then apply the inductive assumption to the shorter sequences

$$\mathbb{C}_1^x \xrightarrow{\mathcal{D}_1^x} \mathbb{C}_2^x \xrightarrow{\mathcal{D}_2^x} \dots \xrightarrow{\mathcal{D}_{n-1}^x} \mathbb{C}_n^x \supset E^x$$

and

$$\mathbb{C}_1^y \xrightarrow{\mathcal{D}_1^y} \mathbb{C}_2^y \xrightarrow{\mathcal{D}_2^y} \dots \xrightarrow{\mathcal{D}_{n-1}^y} \mathbb{C}_n^y \supset E^y.$$

So we may suppose $p \geq 1$ and, as in the proof of Lemma 6.5, let b be the first bubble that is passed, through a disk $D \in \mathcal{D}_0^x$, from a chamber $C \in \hat{\mathbb{C}}_1^x$ to a chamber $C' \in \hat{\mathbb{C}}_1^x$. Let $\mathcal{D}_0^w$ be the disk set $\mathcal{D}_0$ as realigned by this bubble pass. Let

$$\mathbb{C}_1^x \xrightarrow{\mathcal{D}_1^z} \mathbb{C}_2^z \xrightarrow{\mathcal{D}_2^z} \dots \xrightarrow{\mathcal{D}_{n-1}^z} \mathbb{C}_n^z \supset E^z$$

and

$$\mathbb{C}_1^w \xrightarrow{\mathcal{D}_1^w} \mathbb{C}_2^w \xrightarrow{\mathcal{D}_2^w} \dots \xrightarrow{\mathcal{D}_{n-1}^w} \mathbb{C}_n^w \supset E^w$$

be b -wise alignments of the aligned chamber complex sequences, with $b \subset C$ in $\mathbb{C}_1^x$ (as is given) and $b \subset C'$ in $\mathbb{C}_1^w$. (Here the significant difference between the given $\mathcal{D}_i^x$ and the new $\mathcal{D}_i^z, i \geq 1$, is that each $\mathcal{D}_i^z$ is b -wise aligned.) By definition, the two sequences differ by an ancestral bubble pass and by inductive assumption (on n) the former sequence differs from

$$\mathbb{C}_1^x \xrightarrow{\mathcal{D}_1^x} \mathbb{C}_2^x \xrightarrow{\mathcal{D}_2^x} \dots \xrightarrow{\mathcal{D}_{n-1}^x} \mathbb{C}_n^x \supset E^x$$

by a sequence of ancestral bubble passes and eyeglass moves. Extending to the left we then have sequences

$$\begin{aligned}\vec{\mathbb{C}}^x: \quad \mathbb{C}_0 &\xrightarrow{\mathcal{D}_0^x} \mathbb{C}_1^x \xrightarrow{\mathcal{D}_1^x} \mathbb{C}_2^x \xrightarrow{\mathcal{D}_2^x} \dots \xrightarrow{\mathcal{D}_{n-1}^x} \mathbb{C}_n^x \supset E^x, \\ \vec{\mathbb{C}}^z: \quad \mathbb{C}_0 &\xrightarrow{\mathcal{D}_0^z} \mathbb{C}_1^z \xrightarrow{\mathcal{D}_1^z} \mathbb{C}_2^z \xrightarrow{\mathcal{D}_2^z} \dots \xrightarrow{\mathcal{D}_{n-1}^z} \mathbb{C}_n^z \supset E^z, \\ \vec{\mathbb{C}}^w: \quad \mathbb{C}_0 &\xrightarrow{\mathcal{D}_0^w} \mathbb{C}_1^w \xrightarrow{\mathcal{D}_1^w} \mathbb{C}_2^w \xrightarrow{\mathcal{D}_2^w} \dots \xrightarrow{\mathcal{D}_{n-1}^w} \mathbb{C}_n^w \supset E^w \\ \vec{\mathbb{C}}^y: \quad \mathbb{C}_0 &\xrightarrow{\mathcal{D}_0^y} \mathbb{C}_1^y \xrightarrow{\mathcal{D}_1^y} \mathbb{C}_2^y \xrightarrow{\mathcal{D}_2^y} \dots \xrightarrow{\mathcal{D}_{n-1}^y} \mathbb{C}_n^y \supset E^y\end{aligned}$$

and we have shown that any pair of the first three ($\vec{\mathbb{C}}^x, \vec{\mathbb{C}}^z$, and $\vec{\mathbb{C}}^w$) differ by a sequence of ancestral bubble passes and eyeglass moves. Now notice that, by construction, the two aligned chamber complex decompositions $\mathbb{C}_0 \xrightarrow{\mathcal{D}_0^w} \mathbb{C}_1^w$ and $\mathbb{C}_0 \xrightarrow{\mathcal{D}_0^y} \mathbb{C}_1^y$ differ by a sequence of eyeglass moves and at most $p-1$ bubble passes. Hence, by inductive assumption (on p), the last two sequences (hence also the first and the last) differ by a sequence of ancestral bubble passes and eyeglass moves, as required. $\square$

We now turn to S^3 . Throughout the remaining part of this section, make the inductive Assumption 3.4. That is, the Goeritz group of S^3 is the eyeglass group on splittings of genus $< g$.

Suppose a Heegaard split chamber complex $\mathbb{C}_0$ supports the genus g Heegaard splitting $S^3 = A \cup_T B$ and

$$\vec{\mathbb{C}}: \quad \mathbb{C}_0 \xrightarrow{\mathcal{D}_0} \mathbb{C}_1 \xrightarrow{\mathcal{D}_1} \mathbb{C}_2 \xrightarrow{\mathcal{D}_2} \dots \xrightarrow{\mathcal{D}_{n-1}} \mathbb{C}_n$$

is a sequence of chamber complex decompositions. Suppose further that $\mathbb{C}_n$ contains an incompressible sphere S .

Lemma 6.8. *Suppose*

$$\vec{\mathbb{C}}^x: \quad \mathbb{C}_0 \xrightarrow{\mathcal{D}_0^x} \mathbb{C}_1^x \xrightarrow{\mathcal{D}_1^x} \mathbb{C}_2^x \xrightarrow{\mathcal{D}_2^x} \dots \xrightarrow{\mathcal{D}_{n-1}^x} \mathbb{C}_n^x \supset S^x$$

and

$$\vec{\mathbb{C}}^y: \quad \mathbb{C}_0 \xrightarrow{\mathcal{D}_0^y} \mathbb{C}_1^y \xrightarrow{\mathcal{D}_1^y} \mathbb{C}_2^y \xrightarrow{\mathcal{D}_2^y} \dots \xrightarrow{\mathcal{D}_{n-1}^y} \mathbb{C}_n^y \supset S^y$$

are sequences of aligned chamber complex decompositions resulting from possibly different choices of alignment at each stage. Suppose further that the sequences differ by an ancestral bubble pass, of the bubble $\mathfrak{b}$. (In particular S is aligned with $\mathbb{C}_n^x$ and $\mathbb{C}_n^y$ and is disjoint from $\mathfrak{b}$.) Denote S as so aligned S^x and S^y respectively. Then $\mathbb{C}_n^x$ and $\mathbb{C}_n^y$ cocertify.

Proof. We sketch the proof; more detail appears in the earlier proof of Proposition 3.8.

Let $S_{\mathfrak{b}}$ be the sphere $\partial\mathfrak{b}$, defining, after amalgamation, a reducing sphere for the original Heegaard splitting $A \cup_T B$. By Lemma 3.5 $h_{S^x} \sim h_{S_{\mathfrak{b}}} \sim h_{S^y} : (S^3, T) \rightarrow (S^3, T_g)$ as required. $\square$

Now return to the setting of Proposition 6.4: Suppose $S^3 = A \cup_T B$ is a genus g Heegaard splitting of S^3 , a Heegaard split chamber complex $\mathbb{C}_0$ supports T and

$$\vec{\mathbb{C}}: \quad \mathbb{C}_0 \xrightarrow{\mathcal{D}_0} \mathbb{C}_1 \xrightarrow{\mathcal{D}_1} \mathbb{C}_2 \xrightarrow{\mathcal{D}_2} \dots \xrightarrow{\mathcal{D}_{n-1}} \mathbb{C}_n$$

is a sequence of chamber complex decompositions.

Proposition 6.9. *Suppose chambers $C_0 \in \mathbb{C}_0$ and $C_n \in \mathbb{C}_n$ each contain incompressible spheres. Then, for any alignment of each disk set $\mathcal{D}_i$ in $\mathbb{C}_i$, the Heegaard split chamber complexes $\mathbb{C}_0$ and $\mathbb{C}_n$ cocertify.*

Proof. By Proposition 6.4 there is some alignment $\vec{\mathbb{C}}^x$ for the sequence so that $\mathbb{C}_0$ and $\mathbb{C}_n^x$ cocertify. By Proposition 6.7 any other alignment $\vec{\mathbb{C}}^y$ can be obtained from $\vec{\mathbb{C}}^x$ by a sequence of ancestral bubble passes and eyeglass moves. Lemma 6.8 shows that then $\mathbb{C}_n^x$ and $\mathbb{C}_n^y$ cocertify. Hence $\mathbb{C}_0$ and $\mathbb{C}_n^y$ cocertify. $\square$

Corollary 6.10. *Suppose (S^3, T) is a genus g Heegaard splitting of S^3 and*

$$\mathbb{C}_0 \xrightarrow{\mathcal{D}_0} \mathbb{C}_1 \xrightarrow{\mathcal{D}_1} \mathbb{C}_2 \xrightarrow{\mathcal{D}_2} \dots \xrightarrow{\mathcal{D}_{n-1}} \mathbb{C}_n$$

is a sequence of aligned chamber complex decompositions of Heegaard split chamber complexes supporting T . If S and S' are each incompressible spheres in possibly different chambers of the sequence, then those Heegaard split chamber complexes cocertify. $\square$

Unfortunately, declaring every disk ball to be a goneball, as is done throughout this Section, erases too much information. A critical example is this: consider a sequence of two aligned chamber complex decompositions of Heegaard split chamber complexes

$$\mathbb{C}_0 \xrightarrow{\mathcal{D}_0} \mathbb{C}_1 \xrightarrow{\mathcal{D}_1} \mathbb{C}_2$$

and suppose that a chamber C_1 of $\mathbb{C}_1$ is a handlebody that is not disk in the first decomposition. We know from Proposition 5.12 that the Heegaard splitting of C_1 is non-trivial. Now suppose that the disks of $\mathcal{D}_1$ that are incident to the handlebody C_1 are a complete set of meridians for C_1 . The result is a ball chamber of $\mathbb{C}_2$ that has non-trivial Heegaard splitting. As shown in the proof of Proposition 3.7, this is (under the inductive Assumption 3.4) enough information to determine an eyeglass equivalence class $(S^3, T) \rightarrow (S^3, T_g)$. But if we take all disk balls to be goneballs, this ball would be absorbed because it is disk under the decomposition $\mathbb{C}_1 \xrightarrow{\mathcal{D}_1} \mathbb{C}_2$. Thus we lose the information it contains.

Motivated by this example, we will use a more complicated rule to declare that, under certain conditions, a disk ball should *not* be declared a goneball. The rule will focus on how handlebodies are treated in a disk decomposition. The critical step is to be more restrictive in how we allow disks to be aligned.

7 Preferred alignment and flagged chamber complexes

Return now to the general case, which we briefly review: $M = A \cup_T B$ is a Heegaard splitting of a compact 3-manifold; $\mathbb{C}$ is a Heegaard split chamber complex in M that supports T ; $\mathcal{D}$ is an aligned disk set in $\mathbb{C}$; $\hat{\mathbb{C}}_{\mathcal{D}}$ is the chamber complex obtained by surgery on $\mathcal{D}$; and $\mathbb{C}_{\mathcal{D}}$ is the Heegaard split chamber complex obtained by using Rule 5.6 to declare goneballs. Thus

$$\mathbb{C} \xrightarrow{\mathcal{D}} \mathbb{C}_{\mathcal{D}}$$

is a Heegaard split chamber complex decomposition.

Because the disks are aligned, each chamber inherits a Heegaard splitting, as described before Definition 5.11. We consider 3-balls to be (genus 0) handlebodies and introduce the terminology:

Definition 7.1. A Heegaard split handlebody chamber in $\mathbb{C}$ is empty if the Heegaard splitting is trivial. If the Heegaard splitting is non-trivial, the handlebody chamber is called occupied. A flagged chamber complex is a Heegaard split chamber complex in which each handlebody chamber is labelled either empty or occupied.

Two flagged chamber complexes are the same as flagged chamber complexes if they are the same as chamber complexes and the flagging of each handlebody chamber is the same. The Heegaard splittings of any given chamber are not necessarily isotopic, nor even of the same genus.

Definition 7.2. The alignment of a disk set $\mathcal{D}$ in $\mathbb{C}$ is a preferred alignment if, in each chamber C of $\mathbb{C}$, it has these properties:

1. If any remnant of C is not a disk handlebody, then each disk handlebody remnant is empty.
2. If every remnant of C is a disk handlebody, then, per Proposition 4.8, C is a handlebody. In this case:
 - (a) If C is empty, so is every remnant
 - (b) If C is occupied, then exactly one remnant of C is occupied.

Corollary 7.3. Suppose $\mathcal{D}$ is a disk set that is in preferred alignment in a flagged chamber complex $\mathbb{C}$. Suppose $\mathbb{C}$ has no occupied handlebody chambers. Then, under the decomposition

$$\mathbb{C} \xrightarrow{\mathcal{D}} \mathbb{C}_{\mathcal{D}}$$

each handlebody chamber in $\mathbb{C}_{\mathcal{D}}$ is empty if and only if it is disk.

Proof. If a handlebody chamber in $\mathbb{C}_{\mathcal{D}}$ is not disk, then it is occupied per Proposition 5.12, so the interest is in the other direction. Suppose C' is a disk handlebody chamber in $\mathbb{C}_{\mathcal{D}}$, a remnant of a chamber C in $\mathbb{C}$. By hypothesis, C is not an occupied handlebody so, per Definition 7.2(2a), if every remnant of C is a disk handlebody then every remnant, including C' , is empty. On the other hand, if not every remnant of C is a disk handlebody then per 7.2(1) C' is empty. Thus in every case C' is empty, as required. $\square$

We will first show that any disk set has a preferred alignment. A critical fact from Heegaard theory is that any Heegaard splitting of a handlebody is standard, that is it is a stabilization of the trivial splitting. (See, for example, the discussion at the end of Section 6 or, in more detail, the proof of Proposition 8.1 below.) In particular, if $H = A_H \cup_{T_H} B_H$ is a Heegaard split handlebody with non-trivial splitting then there is a bubble $\mathfrak{b}$ for the splitting so that if the bubble is destabilized (for example by replacing $\mathfrak{b}$ with a ball containing just a properly embedded equatorial disk), the splitting becomes trivial. That is, the splitting surface becomes parallel to ∂H . Such a bubble will be called a *maximal bubble* in H .

Lemma 7.4. Suppose

$$\mathbb{C} \xrightarrow{\mathcal{D}} \mathbb{C}_{\mathcal{D}}$$

is a Heegaard split chamber complex decomposition and $H = A_H \cup_{T_H} B_H$ is an occupied disk handlebody remnant of a chamber $C \in \mathbb{C}$, with $C = A_C \cup_{T_C} B_C$. Then there is an isotopy of T_H in H so that afterwards




Figure 22: Case 1 of Lemma 7.4

- There is a disk $E \subset T_H$ so that $T_H - E \subset T_C$ and
- There is a maximal bubble $\mathfrak{b}$, disjoint from E , in the splitting of H that is also a bubble in the splitting of C .

Proof. With no loss assume C and hence H are A -chambers. We briefly set terminology and review the results of Heegaard split chamber complex decomposition from Subsection 5.3. Let Σ_C be a spine for B_C consisting of the union of ∂C and a graph $\gamma_C \subset C$ incident to ∂C only in its ends. We may as well focus on those disks in $\mathcal{D}$ that are incident to C and retreat to denoting these as $\mathcal{D}$. Since $\mathcal{D}$ is aligned, each disk $D \in \mathcal{D}$ intersects the spine Σ_C only in ∂D ; the interior of D lies either in C or in a chamber adjacent to C . In particular, $\mathcal{D}$ is disjoint from the graph γ_C . Except for scars left on ∂H by $\mathcal{D}$, $\partial H \subset \partial C$. If $D \in \mathcal{D}$ leaves an external scar on ∂H then D lies in C ; if it leaves an internal scar then $\text{int}(D)$ lies in an adjacent chamber. There is a spine Σ_H for B_H which consists of the union of ∂H with two graphs in H (whose union we denote γ_H): the graph $\gamma_C \cap H$ and a graph $\gamma_I \subset H$ determined by the internal scars in ∂H and goneballs in the interior of H . Each edge in γ_I is dual to a disk in $\mathcal{D}$ whose interior lies outside C . More specifically, each end of an edge of γ_I lies either on an internal scar in ∂H or on a vertex in $\gamma_I \cap \text{int}(H)$ that corresponds to a goneball. Because the compression body B_H is connected, each component of γ_H has at least one end on ∂H .

We consider increasingly complicated cases:

Case 1: H is a ball and has no internal scars or goneballs. That is $\gamma_I = \emptyset$.

Connect the (external) scars of $\mathcal{D}$ on ∂H by a collection of arcs in ∂H which, by general position, can be taken to be disjoint from the ends of γ_H in ∂H . Viewing these arcs as edges and the scars as vertices of a graph in ∂H , choose the arcs so that the resulting graph is a tree. Let $E_\partial \subset \partial H$ be a (disk) regular neighborhood in ∂H of the union of the scars and arcs. Since that union is disjoint from the ends of γ_H we can take E_∂ also to be disjoint from the ends of γ_H .

The compression body B_C is a regular neighborhood of the spine Σ_C . Since $\mathcal{D}$ is aligned and so disjoint from γ_C , $\mathcal{D}$ intersects B_C only in a collar of ∂B_C . Hence B_C intersects each scar $\mathfrak{s}$ on ∂H in a collar of $\partial \mathfrak{s}$ in ∂H . A regular neighborhood of Σ_H , to which we can isotope B_H , can then be constructed

from B_C in an obvious way: take the union of $B_C \cap H$ and a collar in H of each scar on ∂H . See Figure 22. With this construction, T_H lies in T_C except for the collection of disks in T_H that are parallel to the scars. If we then take E to be the disk in T_H parallel to E_∂ , we have $T_H - E \subset T_C$, verifying the first assertion of the lemma in this case.

Since E is disjoint from γ_H , the interior of E can be pushed, rel ∂E , into A_H . Call the result E_A . Since H is a ball, the complement of E_∂ in ∂H is also a disk; let E_B be a copy of the disk $\partial H - E_\partial$ pushed up into the compression body B_H until its boundary lies on $\partial E = \partial E_A$ in T_H . The union of E_A and E_B is a sphere bounding a ball in H that contains all of $T_H - E$. That ball is then a maximal bubble for H that is also a bubble for the splitting of C , verifying the lemma in this case.

Case 2: H is a ball that may contain internal scars and goneballs.

Suppose s is an internal scar on ∂H , a scar left by a disk in $\mathcal{D}$ that lies in an adjacent chamber C' . See Figure 23. Since by assumption H is a *disky* remnant, ∂s bounds a disk component E_s of $\partial C \cap \text{int}(H)$. Since H is irreducible, the disks s and E_s are isotopic rel ∂ in H , and all ball components of $\hat{\mathbb{C}}_{\mathcal{D}}$ must be goneballs and therefore trivially split. In particular the Heegaard splitting surface of the adjacent chamber C' does not contribute any genus to T_H .

Let ∂H_s be the copy of ∂H obtained by replacing each interior scar s with the corresponding disk $E_s \subset \partial C$. Now apply the argument of Case 1 to the ball in H bounded by ∂H_s instead of ∂H , with the arcs in ∂H between external scars in that construction chosen to avoid also internal scars. This completes the proof in this case. (Note that with this construction, T_H is disjoint from the ball in H between the scar s and the disk E_s .)




Figure 23: Case 2 of Lemma 7.4

Case 3: The general case.

Define the disk $E \subset T_H$ parallel to a disk $E_\partial \subset \partial H$ as in Case 1, with arcs chosen in its construction to also be disjoint from internal scars. Essentially the same argument as in Cases 1 and 2 shows that we can ensure that $T_H - E \subset T_C$. So we proceed to the second assertion of the Lemma: there is a maximal bubble in H that is disjoint from E and is also a bubble for C .

Let $\mathcal{D}_+$ be a minimal complete collection of meridian disks in H . That is, $\mathcal{D}_+$ is a properly embedded collection of disks in H so that the closed complement of a collar of $\mathcal{D}_+$ in H is a single ball W . By general position, we may take $\partial \mathcal{D}_+ \subset \partial H$ to be disjoint from E and all internal scars, and per [Sc1] we

can take $\mathcal{D}_+$ to be aligned with the Heegaard surface T_H . This means in particular that the graph $\gamma_H \subset \Sigma_H$ can be taken to be disjoint from the meridian disks $\mathcal{D}_+$, so it lies entirely in the ball W .

Let E_{∂_+} be the disk in ∂W obtained by band summing E_{∂} to each copy of $\mathcal{D}_+$ in ∂W . (For each disk $D \in \mathcal{D}_+$ there are two copies $D_{\pm}$ in ∂W .) See Figure 24. The complement in the sphere ∂W of E_{∂_+} is a disk that lies entirely in ∂H . In parallel fashion expand the disk E_A (per Case 1 a copy of E with interior pushed into A) to a disk E_{A+} properly embedded in A by band summing E_A to each disk $D_{\pm} \cap A, D \in \mathcal{D}_+$. Since $\partial W - E_+$ is a disk in ∂H , it follows as in Case 1 that $\partial E_{A+} \subset T_H$ bounds a disk in B_H . The union of the two disks is then a sphere parallel to ∂W and the ball $\mathfrak{b}$ it bounds in W is disjoint from E and contains all but the ends of γ_H . So $\mathfrak{b}$ is a maximal bubble, as required. $\square$




Figure 24: Case 3 of Lemma 7.4

Proposition 7.5. Suppose $\mathbb{C}$ is a Heegaard split chamber complex in $M = A \cup_T B$ that supports T , and $\mathcal{D}$ is a disk set in $\mathbb{C}$. Then $\mathcal{D}$ has a preferred alignment.

Proof. As discussed in Subsection 5.3, $\mathcal{D}$ can be aligned. The strategy will be to begin with any alignment and alter it to a preferred alignment.

We first describe how to realign $\mathcal{D}$ so as to achieve Definition 7.2(1). Suppose some remnant $C' \in \mathbb{C}_{\mathcal{D}}$ of a chamber $C \in \mathbb{C}$ is not a disky handlebody, and let $H = \{H_1, \dots, H_k\}$ be the collection of occupied disky handlebody remnants of C . If $H = \emptyset$ there is nothing to prove. Otherwise, let $\gamma = \{\gamma_1, \dots, \gamma_k\}$ be a collection of arcs in C with the following properties:

- Each γ_i is a path in the Heegaard surface T_C that is transverse to $\mathcal{D}$ and runs from a point in $\text{int}(C')$ to a point in $\text{int}(H_i)$.
- Among all such families of arcs, $|\gamma \cap \mathcal{D}| \geq k$ is minimized.

Let $D \in \mathcal{D}$ be the disk containing the closest intersection point on γ_1 to its endpoint in H_1 . Let C'' be the remnant that is on the other side of D . Since $|\gamma \cap \mathcal{D}|$ is minimized, γ_1 is otherwise disjoint from D and $C'' \neq H_1$. H_1 is assumed to be occupied; let $\mathfrak{b}$ be a maximal bubble for the splitting $H_1 = A_1 \cup_{T_1} B_1$ satisfying the conclusion of Lemma 7.4. That is, $\mathfrak{b}$ is also a bubble for the Heegaard splitting $C = A_C \cup_{T_C} B_C$.

Take the endpoint of γ_1 in H_1 to be a point in the circle $c_b = \partial b \cap T_C$, so the subpath $\gamma_H = \gamma_1 \cap H_1$ is an arc in T_C from D to the sphere ∂b . Tube sum D to ∂b along the path γ_H , creating a new disk D' . Observe that D' intersects T_C in a single circle, namely the band sum along γ_H of the circles $D \cap T_C$ and $\partial b \cap T_C$. Thus D' is an aligned disk, and is isotopic to D in C because b is a ball. Moreover, the bubble b lies on the same side of D' as C'' , not H . Thus D' is a realignment of D that makes H an empty handlebody. If C'' is not a handlebody, or is a handlebody that was already occupied, then k is reduced by 1. If C'' was an empty handlebody then after the realignment C'' becomes an occupied handlebody, leaving k unchanged. But the path γ_1 now has its endpoint in C'' and no longer intersects the realigned D' , so $|\gamma_1 \cap \mathcal{D}|$ (with D realigned) is reduced by 1. Continue the realignment of disks until $k = |\gamma \cap \mathcal{D}| = 0$ as required.

To establish Definition 7.2(2a) note that if any remnant of C is a diskly occupied handlebody then Lemma 7.4 shows C contains a bubble and so cannot have trivial splitting; it is occupied. To establish 7.2(2b), use the same argument as for 7.2(1) above, but make an arbitrary choice of any remnant for C' . $\square$

Definition 7.6. Suppose $\mathbb{C}$ is a flagged chamber complex, $\mathcal{D}$ is a disk set in $\mathbb{C}$ with a preferred alignment, and $\hat{\mathbb{C}}_{\mathcal{D}}$ is the chamber complex obtained by surgery on $\mathcal{D}$. Let $\mathbb{C}_{\mathcal{D}}$ be the flagged chamber complex obtained by using Rule 5.6 to declare goneballs. Then the Heegaard split chamber complex decomposition

$$\mathbb{C} \xrightarrow{\mathcal{D}} \mathbb{C}_{\mathcal{D}}$$

is called a flagged chamber complex decomposition.

For much of our argument, simply keeping track of the flagging of the chamber complexes (that is, whether a handlebody chamber is empty or occupied) will suffice.

Proposition 7.7. Suppose, for $\mathcal{D}$ a disk set in $\mathbb{C}$,

$$\mathbb{C} \xrightarrow{\mathcal{D}} \mathbb{C}_{\mathcal{D}}$$

is a flagged chamber complex decomposition.

1. Any ball chamber in $\mathbb{C}_{\mathcal{D}}$ is occupied. (That is, $\mathbb{C}_{\mathcal{D}}$ is a Heegaard split chamber complex.)
2. If any remnant of a chamber C in $\mathbb{C}$ is an occupied diskly handlebody then C is an occupied handlebody and every remnant of C in $\mathbb{C}_{\mathcal{D}}$ is a diskly handlebody.
3. Suppose $\mathcal{D}$ is given a different preferred alignment and $\mathbb{C}'_{\mathcal{D}}$ is the resulting flagged chamber complex structure. If $\mathbb{C}'_{\mathcal{D}}$ and $\mathbb{C}_{\mathcal{D}}$ differ as flagged chamber complexes, then $\mathbb{C}$ contains an occupied handlebody chamber whose remnants in $\mathbb{C}_{\mathcal{D}}$ and $\mathbb{C}'_{\mathcal{D}}$ are all diskly handlebodies.

Proof. (1) If a ball chamber were not occupied, it would be empty so, by Definition 7.1, it would have trivial Heegaard splitting. But then, by Rule 5.6, it would have been a goneball.

(2) The proof is essentially the same as that of Corollary 7.3: If a remnant of C is an occupied diskly handlebody then from the contrapositive of Definition 7.2(1) every remnant of C is a diskly handlebody. Then the contrapositive of Definition 7.2(2a) says that C must be an occupied handlebody.

(3) In a flagged chamber complex decomposition, the flagging of the chambers after surgery, as described in Definition 7.2, depends only on the disks $\mathcal{D}$ and not on how they are aligned with the Heegaard surfaces in the chamber, except in satisfying Definition 7.2(2b): If C is an occupied handlebody chamber of $\mathbb{C}$ and each remnant of C is a disk handlebody, then a choice of alignment is made to ensure that exactly one remnant (of possibly several) is occupied. Hence if $\mathbb{C}'_{\mathcal{D}}$ and $\mathbb{C}_{\mathcal{D}}$ differ, so $\hat{\mathbb{C}}_{\mathcal{D}}$ and $\hat{\mathbb{C}}'_{\mathcal{D}}$ differ, it must be because of a different such choice. That is, in $\hat{\mathbb{C}}_{\mathcal{D}}$ one remnant is an occupied handlebody, and in $\hat{\mathbb{C}}'_{\mathcal{D}}$ a different remnant is. Thus, per statement (2), C is the required occupied handlebody chamber of $\mathbb{C}$. $\square$

Suppose, as described in the proof of Proposition 7.7(3), two flagged chamber complex decompositions $\mathbb{C} \xrightarrow{\mathcal{D}} \mathbb{C}_{\mathcal{D}}$ and $\mathbb{C} \xrightarrow{\mathcal{D}} \mathbb{C}'_{\mathcal{D}}$ differ only because, for some occupied handlebody chambers in $\mathbb{C}$, different alignments of the disks $\mathcal{D}$ in the chambers result in different remnants being occupied.

Definition 7.8. *The decompositions $\mathbb{C} \xrightarrow{\mathcal{D}} \mathbb{C}_{\mathcal{D}}$ and $\mathbb{C} \xrightarrow{\mathcal{D}} \mathbb{C}'_{\mathcal{D}}$ are sibling decompositions, and the occupied handlebody chambers in $\mathbb{C}$ are called parent chambers in the sibling decompositions.*

Proposition 7.5. is an existence statement; we now move towards a parallel uniqueness statement.

Lemma 7.9. *Suppose $H = A \cup_T B$ is a Heegaard split handlebody. Suppose $\mathfrak{b}^p$ and $\mathfrak{b}^q$ are not necessarily disjoint bubbles for T of genus $p \leq q$ respectively. Then*

- *There is a genus $(q - p)$ bubble $\mathfrak{b}'$ for T that is disjoint from $\mathfrak{b}^p$, and a homeomorphism $h : (H, T) \rightarrow (H, T)$ so that h is the identity on ∂H , and $\mathfrak{b}^q$ is the tube sum of $h(\mathfrak{b}^p)$ and $h(\mathfrak{b}')$.*
- *Under Assumption 3.4, if $\text{genus}(T) - \text{genus}(H) \leq g - 1$ then we may take h to be an eyeglass move.*

Proof. Pick an aligned disk set $\mathcal{D}$ in H so that $H - \eta(\mathcal{D})$ is a single ball. This implies $|\mathcal{D}| = \text{genus}(H)$ and each disk is non-separating. Then surgery on $\mathcal{D}$ gives a genus $g' = \text{genus}(T) - \text{genus}(H)$ Heegaard splitting of the ball which, by [Wa], is unique. More explicitly, echoing the notation surrounding Figure 7, there is a nested sequence of g' bubbles $\mathfrak{b}_1 \subset \dots \subset \mathfrak{b}_{g'}$ for T so that each sphere $S_i = \partial \mathfrak{b}_i$ intersects T in a single circle c_i , and $T/\mathfrak{b}_{g'}$ is parallel to ∂H .

More specifically, align $\mathcal{D}$ first with $T/\mathfrak{b}^q$ so that in the associated alignment of $\mathcal{D}$ with T , $\mathfrak{b}^q$ is disjoint from $\mathcal{D}$. Then with no loss of generality we may assume that $\mathfrak{b}^q = \mathfrak{b}_q$. For the first statement above, it then suffices to find a homeomorphism $h(H, T) \rightarrow (H, T)$ so that $h(\mathfrak{b}^p) = \mathfrak{b}_p$. Denote this alignment of $\mathcal{D}$ by $\mathcal{D}^q$.

Similarly, let $\mathcal{D}^p$ be an alignment of $\mathcal{D}$ found by first aligning with $T/\mathfrak{b}^p$ and then observing that the associated alignment of $\mathcal{D}^p$ with T is disjoint from $\mathfrak{b}^p$. According to [FS2] there is a sequence of bubble passes and eyeglass moves that will take $\mathcal{D}^p$ to $\mathcal{D}^q$ and, since each disk in $\mathcal{D}$ is non-separating disk, it follows from Lemma 2.5 that even the bubble passes are eyeglass moves. (To pass a bubble through a disk $D \in \mathcal{D}$ choose the arc β in the proof to be disjoint from all other disks in $\mathcal{D}$.) Thus there is an eyeglass move that takes $\mathcal{D}^p$ to $\mathcal{D}^q$. Put another way, there is an eyeglass move $h_1 : (H, T) \rightarrow (H, T)$ so that $h_1(\mathfrak{b}^p)$ lies in the ball $H - \eta(\mathcal{D}^q)$. By [Wa] there is then a homeomorphism $h_2 : (H, T) \rightarrow (H, T)$ so that $h_2(h_1(\mathfrak{b}^p)) = \mathfrak{b}_p$. This completes the proof of the first statement, with $h = h_2 h_1$.

To prove the second statement, just note that T induces a genus g' splitting on $H - \eta(\mathcal{D}^q)$, so if $g' \leq g - 1$ then we may take h_2 to be an eyeglass move, so $h_2 h_1$ is also an eyeglass move. $\square$

Throughout the remainder of this section $\mathbb{C}$ is a flagged chamber complex in M supporting a genus g Heegaard splitting, with $\mathcal{D}$ a disk set in $\mathbb{C}$. We further continue with Assumption 3.4. Let $\mathcal{D}_\pm$ denote two preferred alignments of $\mathcal{D}$ that result in the same flagged chamber complex decompositions $\mathbb{C} \xrightarrow{\mathcal{D}} \mathbb{C}_{\mathcal{D}_\pm}$. That is, to repeat from Definition 7.1, $\mathbb{C}_{\mathcal{D}_\pm}$ are identical chamber complexes with identical flagging (handlebody chambers either empty or occupied), though the underlying Heegaard splittings of each chamber may differ.

Proposition 7.10. *There is a sequence of preferred alignments of $\mathcal{D}$ in $\mathbb{C}$*

$$\mathcal{D}_- = \mathcal{D}_0, \mathcal{D}_1, \mathcal{D}_2, \mathcal{D}_3, \dots, \mathcal{D}_n = \mathcal{D}_+$$

so that

- a) for each $1 \leq i \leq n$, $\mathcal{D}_i$ is obtained from $\mathcal{D}_{i-1}$ by eyeglass moves and a single bubble pass, possibly through multiple disks of $\mathcal{D}_{i-1}$.
- b) The result $\mathbb{C}_{\mathcal{D}_i}$ of each flagged chamber complex decomposition

$$\mathbb{C} \xrightarrow{\mathcal{D}_i} \mathbb{C}_{\mathcal{D}_i}$$

is the same flagged chamber complex as $\mathcal{D}_\pm$.

Proof. The central theorem of [FS2] is that there is such a sequence of alignments, but without the condition that each $\mathcal{D}_i$ is in *preferred* alignment and that each resulting Heegaard split chamber complex $\mathbb{C}_{\mathcal{D}_i}$ has the same flagging. We begin with such a sequence and alter it to achieve these conditions.

Claim: It suffices to find a sequence of (not necessarily preferred) alignments satisfying a) and b) so that each disk handlebody chamber of $\mathbb{C}_{\mathcal{D}}$ that is empty in $\mathbb{C}_{\mathcal{D}_-}$ is also empty in each $\mathbb{C}_{\mathcal{D}_i}$.

Proof of claim: We will assume we have found such a sequence of alignments and show how the proof of the proposition follows.

First observe that by Proposition 5.12 any handlebody chamber of $\mathbb{C}_{\mathcal{D}}$ that is not disk is occupied regardless of alignments, so the only difference in possible flaggings of the $\mathbb{C}_{\mathcal{D}_i}$ is in the flagging of disk handlebodies. Since $\mathbb{C}_{\mathcal{D}_-}$ has a preferred alignment, it satisfies Definition 7.2(1), namely that if any remnant of C is not a disk handlebody, then each disk handlebody remnant is empty. But by the assumption, each disk handlebody remnant is then empty in $\mathbb{C}_{\mathcal{D}_i}$, so each $\mathbb{C}_{\mathcal{D}_i}$ also satisfies Definition 7.2(1).

Similarly, if each remnant of a chamber C is a disk handlebody (so C is a handlebody, per Proposition 4.8), then Definition 7.2(2a) says that if C is empty, every remnant is empty in $\mathbb{C}_{\mathcal{D}_-}$ and so, by assumption, also in each $\mathbb{C}_{\mathcal{D}_i}$. Thus each $\mathbb{C}_{\mathcal{D}_i}$ satisfies Definition 7.2(2a). On the other hand, if C is occupied, Definition 7.2(2b) says that exactly one remnant, say $C' \in \mathbb{C}_{\mathcal{D}_-}$ is occupied, so the others are all empty. By assumption, all remnants of C in each $\mathbb{C}_{\mathcal{D}_i}$, except possibly C' , are then empty. But since C is occupied it has non-trivial splitting, that is the genus of its Heegaard splitting surface is greater than that of its boundary. So decomposition cannot leave only trivially split handlebodies. Thus some remnant of C must be empty in each $\mathbb{C}_{\mathcal{D}_i}$, and C' is the only possibility. Thus exactly one remnant of C (namely C') in each

$\mathbb{C}_{\mathcal{D}_i}$ is occupied, and so each $\mathbb{C}_{\mathcal{D}_i}$ also satisfies Definition 7.2(2b), with the same flagging as $\mathbb{C}_{\mathcal{D}_-}$. Thus each $\mathcal{D}_i$ is in preferred alignment, as required.

Having established the claim, we proceed to show that in fact there is a sequence of alignments satisfying a) and b) so that each disk handlebody chamber of $\mathbb{C}_{\mathcal{D}}$ that is empty in $\mathbb{C}_{\mathcal{D}_-}$ is also empty in each $\mathbb{C}_{\mathcal{D}_i}$. This follows from a min-max argument that we now describe.

Let H be the collection of disk handlebodies in $\mathbb{C}_{\mathcal{D}}$ that are designated as empty in $\mathbb{C}_{\mathcal{D}_-}$ and so by assumption, also in $\mathbb{C}_{\mathcal{D}_+}$. Since these are flagged chamber complexes, no chamber in H is a ball.

Consider any sequence of (not necessarily preferred) alignments, given by [FS2], that begins with $\mathcal{D}_-$ and ends with $\mathcal{D}_+$ so that each alignment in the sequence differs from the previous alignment by eyeglass moves before and after a single bubble pass, possibly through multiple disks. For each $0 \leq i \leq n$, let $\sigma_i \geq 0$ be the difference between the sum of all the genera of the Heegaard surfaces for H in $\mathbb{C}_{\mathcal{D}_i}$ and the sum of the genera of ∂H . If we can find a sequence so that each $\sigma_i = 0$, then each Heegaard splitting would be trivial, so each chamber in H would continue to be empty in each $\mathbb{C}_{\mathcal{D}_i}$. Following the claim above, this would conclude the proof of the lemma.

In search of a sequence of alignments for which $\sigma_i = 0$, let

$$\mathcal{D}_- = \mathcal{D}_0, \mathcal{D}_1, \mathcal{D}_2, \mathcal{D}_3, \dots, \mathcal{D}_n = \mathcal{D}_+$$

be one that minimizes, in lexicographic order, the pair (m, x) , where m is the maximum of σ_i throughout the sequence, and x is the number of integers i for which $\sigma_i = m$. The aim is to show that $m = 0$. Towards a proof by contradiction, suppose $m \geq 1$ and let i be the smallest integer for which $\sigma_i = m$. We know that $1 \leq i \leq n-1$ since $\sigma_0 = \sigma_n = 0$.

So we have

1. $\sigma_{i-1} < m$
2. $\sigma_i = m$
3. $\sigma_{i+1} \leq m$.

The first statement implies that the realignment of $\mathcal{D}_{i-1}$ to $\mathcal{D}_i$ moves a bubble from a remnant $R_1 \notin H$ to a remnant $R_2 \in H$. The third statement means that in the realignment of $\mathcal{D}_i$ to $\mathcal{D}_{i+1}$, which moves a bubble from remnant R_3 to remnant R_4 , we cannot have both $R_3 \notin H$ and $R_4 \in H$.

Suppose that $R_2 \neq R_3$. In this case, the bubble in R_3 could be passed to R_4 before the bubble passes from R_1 to R_2 . This does not affect σ_{i-1} or σ_{i+1} but may change σ_i , say to σ'_i . Since we cannot have both $R_3 \notin H$ and $R_4 \in H$, $\sigma'_i \leq \sigma_{i-1}$. (1) then implies $\sigma'_i < m$. Thus we have reduced x (or m if always $j > i \implies \sigma_j < m$).

Next suppose $R_2 = R_3$, and just call it R . We have already seen that this remnant must be in H , and so a disk handlebody. Since bubbles can only be passed between remnants of the same chamber, we deduce that R_1, R, R_4 are all remnants of the same chamber of $\mathbb{C}$. In particular, it is possible to directly pass a bubble from R_1 to R_4 . Let g_1, g_4 be respectively the genera of the bubble b_1 that is moved into R from R_1 and the bubble b_4 that is moved from R into R_4 . There are two cases to consider:

Case 1: $g_1 \leq g_4$

Let $D_1 \subset \partial R$ be the disk through which b_1 is passed into R from R_1 as it lies in $\mathcal{D}_{i-1}$. We proceed as though D_1 is the only disk through which b_1 is passed, so it is passed from one remnant of C to an adjacent one; if it is passed through multiple chambers the argument only needs to be altered by taking multiple parallel copies of D_1 , corresponding to how the thin tube from R to b_1 passes through other disks in $\mathcal{D}_{i-1}$ in the bubble pass. Similarly let $D_4 \subset \partial R$ be the disk (or parallel copies of the disk) in $\mathcal{D}_{i-1}$ through which b_4 is passed into R_4 . (Possibly $R_1 = R_4$ and $D_1 = D_4$.)

Let $(D_s, \partial D_s) \subset (R_1, \partial D_1)$ be a disk parallel to D_1 obtained by tube-summing b_1 to D_1 . Similarly, let H_+ be the handlebody obtained from R by the same tube-summing, replacing D_1 with D_s . Then H_+ contains both b_1 and b_4 . According to Lemma 7.9 there is an eyeglass move $h : H_+ \rightarrow H_+$ and a bubble b' of genus $g_4 - g_1$ so that b_4 is the tube sum of b' and $h(b_1)$. Replace by an eyeglass move $D_1 \in \mathcal{D}_{i-1}$ with $D'_1 = h(D_1)$. See the highly schematic Figure 25.




Figure 25: Case 1 when $R_2 = R_3 = R$

Do the transition from the alignment of $\mathcal{D}_{i-1}$ to that of $\mathcal{D}_{i+1}$ with these bubble passes instead of the given two: Tube sum D'_1 to b' , after which b_4 lies entirely in R_1 . This changes σ to $\sigma_i - g_4$, since $R_1 \notin H$. Then pass the bubble b_4 directly from R_1 into R_4 . This realignment leaves D_4 aligned so that R_4 contains b_4 . The upshot of this replacement of D_1 with D'_1 in $\mathcal{D}_{i-1}$ is that the first occurrence of $\sigma_i = m$ in the sequence is replaced with $m - g_4$. Thus x is reduced by one (or m is reduced, if always $j > i \implies \sigma_j < m$).

Case 2: $g_4 \leq g_1$

The argument is much the same, using H_+ as defined above, with the following alteration: According to Lemma 7.9 there is a homeomorphism $h : H_+ \rightarrow H_+$ and a bubble b' of genus $g_1 - g_4$ so that $h(b_1)$ is

the tube sum of b' and b_4 .

Do the transition from the alignment of $\mathcal{D}_i$ to that of $\mathcal{D}_{i+1}$ with these bubble passes instead of the given two: First replace D_1 with $h(D_1)$. After this realignment $b_4 \subset h(b_1)$ lies entirely in R_1 . Next move b_4 directly to R_4 . This move does not raise $\sigma = \sigma_{i-1}$ and may lower it by g_4 if $R_4 \notin H$. Next move b' into R , returning the alignment to $\mathcal{D}_{i+1}$. This transition replaces the first occurrence of $\sigma_i = m$ in the sequence with a term no larger than σ_{i-1} . Thus again x is reduced by one (or m is reduced, if always $j > i \implies \sigma_j < m$).

In any case we have contradicted the assumption that (m, x) is minimal. We conclude that $m = 0$, as required. $\square$

Definition 7.11. A flagged chamber complex is tiny if it is a tiny chamber complex, as in Definition 4.9, and, in the case that $F(\mathbb{C}) \neq \emptyset$, each of the designated handlebodies is empty.

In particular, if $F(\mathbb{C})$ is a single component dividing M into two handlebodies, so either of them can play the role of designated handlebody, the flagged chamber complex is tiny unless *both* handlebodies are occupied.

Proposition 7.12 (Tinyness pulls back). Suppose

$$\mathbb{C} \xrightarrow{\mathcal{D}} \mathbb{C}_{\mathcal{D}}$$

is a flagged chamber complex decomposition. If $\mathbb{C}_{\mathcal{D}}$ is tiny, so was $\mathbb{C}$.

Proof. This is essentially a restatement of Proposition 5.19. $\square$

Corollary 7.13. Suppose

$$\mathbb{C}_0 \xrightarrow{\mathcal{D}_0} \mathbb{C}_1 \xrightarrow{\mathcal{D}_1} \mathbb{C}_2 \xrightarrow{\mathcal{D}_2} \dots \xrightarrow{\mathcal{D}_{n-1}} \mathbb{C}_n$$

is a sequence of flagged chamber complex decompositions. Then if $\mathbb{C}_0$ is not tiny, neither is $\mathbb{C}_n$. In particular, $F(\mathbb{C}_n) \neq \emptyset$.

8 Occupied handlebodies and certification

In this section we demonstrate that the principal results from Section 6 continue to apply in the context of flagged chamber complexes, and with a broader version of certification. Continue with Assumption 3.4, that $G(S^3, T') = \mathcal{E}$ whenever $\text{genus}(T') \leq g - 1$.

Proposition 8.1. Let $\mathbb{C}$ be a flagged chamber complex for S^3 that supports the genus g Heegaard splitting (S^3, T) . Suppose $C = A_C \cup_{T_C} B_C$ is an occupied handlebody chamber of $\mathbb{C}$.

1. There is a reducing sphere S for the Heegaard surface T_C and a corresponding homeomorphism $h_C : (S^3; T) \rightarrow (S^3; T_g)$ so that $h_C(S) \in \{S_i, i = 1, \dots, g - 1\}$. Moreover, the eyeglass equivalence class of h_C is independent of S .
2. If C' is another occupied handlebody chamber of $\mathbb{C}$ then h_C and $h_{C'}$ are eyeglass equivalent.

3. If S is an incompressible sphere in a chamber of $\mathbb{C}$ then h_C is eyeglass equivalent to the homeomorphism h_S given in Proposition 3.7.
4. Suppose $\mathcal{D}$ is a disk set in $\mathbb{C}$ in preferred alignment, and $\mathbb{C} \xrightarrow{\mathcal{D}} \mathbb{C}_{\mathcal{D}}$ is the corresponding flagged chamber complex decomposition. Suppose all remnants of C are disk handlebodies and C' is the chamber in $\mathbb{C}_{\mathcal{D}}$ that, per Definition 7.2, is the unique remnant of C that is occupied. Then h_C and $h_{C'}$ are eyeglass equivalent.
5. For any $\tau \in G(S^3, T)$ and $h_{\tau(C)}$ similarly defined, $h_{\tau(C)}\tau \sim h_C$.

Proof. We prove each statement in turn:

1. Since C is an occupied handlebody, T_C is a non-trivial Heegaard splitting. That is, $b = \text{genus}(T_C) - \text{genus}(\partial C) \geq 1$. Following Lemma 7.9 there is a genus b bubble $\mathfrak{b}$ in T_C and, so long as $b < g$, $\mathfrak{b}$ is unique up to eyeglass moves.

We have

$$g = \text{genus}(T) \geq \text{genus}(T_C) = b + \text{genus}(\partial C)$$

so indeed $b < g$ unless C is a ball and T_C is genus g . But this is impossible, for the complement of C in S^3 would then be a trivially split ball, and so a goneball, so ∂C could not be a component of F . We conclude that $\mathfrak{b}$ is unique up to eyeglass moves and take $\partial\mathfrak{b}$ for the required reducing sphere.

2. This follows from Lemma 3.5.
3. This again follows from Lemma 3.5, since the handlebody chamber C cannot be the same as the chamber with incompressible sphere.
4. For the non-trivial Heegaard splitting of handlebody C' , let $\mathfrak{b}'$ be a bubble whose boundary defines $h_{C'}$, as we have just described. According to Lemma 7.4, $\mathfrak{b}'$ can be placed so that it is also a bubble for T_C . According to Lemma 7.9 an eyeglass move will put $\mathfrak{b}'$ inside the bubble $\mathfrak{b}$ whose boundary defines h_C . The result then again follows from Lemma 3.5.
5. This follows from the second statement in Proposition 3.7.

□

Motivated by Proposition 8.1(1-3), we have

Definition 8.2. A flagged chamber complex $\mathbb{C}$ in S^3 certifies if some chamber contains an incompressible sphere or some chamber is an occupied handlebody. An incompressible sphere in such a chamber, or the boundary of a bubble in an occupied handlebody, is called a certificate issued by $\mathbb{C}$.

Corollary 8.3. Let $\mathbb{C}$ be a flagged chamber complex for S^3 that supports the genus g Heegaard splitting (S^3, T) . If $\mathbb{C}$ certifies then all homeomorphisms $(S^3, T) \rightarrow (S^3, T_g)$ determined by certificates in $\mathbb{C}$ are eyeglass equivalent.

Let $h_{\mathbb{C}} : (S^3, T) \rightarrow (S^3, T_g)$ denote any homeomorphism given by such a certificate; by Corollary 8.3 $h_{\mathbb{C}}$ is well defined up to eyeglass equivalence. If $\mathbb{C}'$ is another flagged chamber complex that certifies and the homeomorphisms $h_{\mathbb{C}}$ and $h_{\mathbb{C}'}$ are eyeglass equivalent, then we say that $\mathbb{C}$ and $\mathbb{C}'$ *cocertify* and write $\mathbb{C} \sim \mathbb{C}'$.

Corollary 8.4. *If $\mathbb{C}$ certifies and $\tau \in G(S^3, T)$ then $h_{\tau(\mathbb{C})}\tau \sim h_{\mathbb{C}}$.*

Proof. This follows from 8.1(5). □

Suppose $S^3 = A \cup_T B$ is a genus g Heegaard splitting of S^3 , the Heegaard split chamber complex $\mathbb{C}_0$ supports T , and we are given a sequence of flagged chamber complex decompositions

$$\vec{\mathbb{C}}: \quad \mathbb{C}_0 \xrightarrow{\mathcal{D}_0} \mathbb{C}_1 \xrightarrow{\mathcal{D}_1} \mathbb{C}_2 \xrightarrow{\mathcal{D}_2} \dots \xrightarrow{\mathcal{D}_{n-1}} \mathbb{C}_n$$

Suppose further that both $\mathbb{C}_0$ and $\mathbb{C}_n$ certify.

In analogy to Proposition 6.4 we have.

Proposition 8.5. *For iteratively $0 \leq i \leq n$ there is a preferred alignment of each $\mathcal{D}_i$ so that $\mathbb{C}_0$ and $\mathbb{C}_n$ cocertify.*

Proof. The proof is by induction on n .

Initial Case: $n = 1$

Following Proposition 7.5, take any preferred alignment of $\mathcal{D}_0$ in $\mathbb{C}_0$.

One possibility is that in the flagged chamber complex decomposition $\mathbb{C}_0 \xrightarrow{\mathcal{D}_0} \mathbb{C}_1$ there is an occupied handlebody chamber in $\mathbb{C}_0$ for which every remnant in $\mathbb{C}_1$ is a disk handlebody. Then the result follows from Proposition 8.1(4). So assume that this is not the case. Then by Proposition 7.7(2) a handlebody in $\mathbb{C}_1$ is empty if and only if it is disk. This implies that different choices of preferred alignment of $\mathcal{D}_0$ have no effect on the flagging of $\mathbb{C}_1$. In particular, every disk ball in $\hat{\mathbb{C}}_1$ is a goneball.

Consider a chamber $C = A_C \cup_{T_C} B_C$ in $\mathbb{C}_0$ that certifies. There is a (reducing sphere) certificate $S \subset C$ for T_C that is either incompressible in C or, if C is an occupied handlebody, bounds a bubble b for T_C . In either case there is a realignment of $\mathcal{D}_0$ in $\mathbb{C}$ that is disjoint from S , in the former case by the proof of Proposition 6.4 and in the latter case by considering the Heegaard splitting $\mathbb{C}_0/b$. We first need to show that this realignment of $\mathcal{D}_0 \cap C$ can be done in a way that is still a preferred alignment and then show that the homeomorphisms h_S and $h_{S'}$ determined by S and some certificate S' for $\mathbb{C}_1$ are eyeglass equivalent.

Consider the remnant C' of C in $\hat{\mathbb{C}}_1$ that contains S .

Subcase 1: C' is not a disk handlebody. (In particular, C' is not a goneball, and so C' remains a chamber in $\mathbb{C}_1$.)

For example, this is always the case if S is incompressible in C and so is either incompressible in C' or bounds a ball containing components of F .

The disk set $\mathcal{D}_0 \cap C$ can be given a preferred alignment by emptying any remnant of C that is a disk handlebody by passing bubbles to C' , see Definition 7.2. If S is incompressible in C' then it is a certificate for $\mathbb{C}_1$ and we are done. If there is a certificate S' for $\mathbb{C}_1$ in a chamber of $\mathbb{C}_1$ other than C' , then S and S'

are disjoint, so h_S and $h_{S'}$ are eyeglass equivalent by Lemma 3.5. This implies $h_{\mathbb{C}_1} \sim h_{S'} \sim h_S \sim h_{\mathbb{C}_0}$ and again we are done.

So we are reduced to the case where C' is the only certifying chamber in $\mathbb{C}_1$ and S bounds a bubble b in C' . If C' contains an incompressible sphere, then aligning it with the Heegaard surface for C' in $\mathbb{C}_1/b$ gives a certificate S' for $\mathbb{C}_1$ that is disjoint from b and we are done as above. If C' does not contain an incompressible sphere then, since it contains a certificate for $\mathbb{C}_1$, it must be an occupied handlebody in which S' bounds a bubble. By Lemma 7.9 the spheres S and S' can be made disjoint by an eyeglass move. By Lemma 3.5 this again implies $h_{\mathbb{C}_1} \sim h_{S'} \sim h_S \sim h_{\mathbb{C}_0}$. This concludes the argument in this subcase.

Subcase 2: C' is a disk handlebody.

In this case, to obtain a preferred alignment of $\mathcal{D}_0$ in $\mathbb{C}$ we may need to empty C' , see Definition 7.2. Recall that this is done by finding a maximal bubble in C' , one that is also a bubble for T_C , and tube summing its boundary to a possible series of disks in $\mathcal{D}_0$. Lemma 7.9 shows that such a maximal bubble can be found that contains b within it, so b ends up in a remnant of C that is not a disk handlebody. Then Subcase 1 applies. This concludes the argument for Subcase 2 and hence the Initial Case $n = 1$.

The inductive step, $n \geq 2$:

Inductively suppose the result is known for any sequence of length $\leq n - 1$. Then we may as well assume that $\mathbb{C}_i$ does not certify for $1 \leq i \leq n - 1$, else we could break the sequence into two sequences of smaller length. In particular, in none of the decompositions does an occupied handlebody have all remnants disk handlebodies, for when that occurs in a flagged chamber complex decomposition both the occupied handlebody and one of the remnants certify, see Definitions 7.6 and 7.2. So, just as in the case $n = 1$, the alignment of each $\mathcal{D}_i$ is preferred if and only if it has this property: each handlebody remnant is empty if and only if it is disk. As a result, different choices of preferred alignment throughout the decomposition sequence have no effect on the flagging of the flagged chamber complexes.

Since $\mathbb{C}_1$ does not certify, no chamber in $\mathbb{C}_1$ contains an incompressible sphere. Also, no occupied handlebody in $\mathbb{C}_0$ has only handlebody remnants in $\mathbb{C}_1$, since at least one such remnant would certify. There is a certificate for $\mathbb{C}_0$ in some chamber C of $\mathbb{C}_0$. As in the case $n = 1$ we can find such a certificate S and realign $\mathcal{D}_0 \cap C$ so that the sphere S is disjoint from $\mathcal{D}_0$. Then, regardless of whether S is an incompressible sphere or bounds a bubble in C , S must compress in the chamber of $\mathbb{C}_1$ that contains it, so it bounds a bubble b in that chamber.

Following further the case $n = 1$, the disks $\mathcal{D}_0 \cap C$ can be further realigned so that all disk handlebody remnants of C are emptied, by bubble passes to a remnant C' of C that is not a disk handlebody, and this can be done in a way that leaves the bubble b intact. The result is that the first decomposition again has a preferred alignment, but now with b lying intact in the remnant $C' \in \mathbb{C}_1$. Finally note that C' cannot be a handlebody at all: by construction it is not a disk handlebody and if it were not disk it would be occupied and so it would certify.

Consider the flagged chamber complex $\mathbb{C}_1/b$. Since b does not lie in a handlebody, there is no change in flagging. Via Proposition 7.5 there is a preferred alignment for the disks $\mathcal{D}_1$ so that

$$\mathbb{C}_1/b \xrightarrow{\mathcal{D}_1} \mathbb{C}_2/b$$

is a flagged chamber complex decomposition. Since C' is not a handlebody, not every remnant of C' in $\mathbb{C}_2$ is a disk handlebody and, unless $n = 2$, $\mathbb{C}_2$ does not certify, so in fact not every remnant is a handlebody

at all. By isotoping $*$ in $\mathbb{C}_1/\mathfrak{b}$ (or, equivalently, passing the bubble $\mathfrak{b}$ in $\mathbb{C}_1$) we can ensure that $*$ lies in a remnant of C' in $\mathbb{C}_2/\mathfrak{b}$ that is not a handlebody. Continue in this manner until $*$ reaches $\mathbb{C}_n/\mathfrak{b}$ and lies in a chamber C'' that is not a disk handlebody (but may be an occupied handlebody). The proof now concludes as it did for $n = 1$, Subcase 1. $\square$

Just as Proposition 8.5 for flagged decompositions parallels Proposition 6.4, we now work towards finding a parallel of Proposition 6.9 for flagged decompositions. The argument is mostly just a reprise of the previous argument. We begin with an analogue of Lemma 6.5.

Lemma 8.6. *Suppose $\mathbb{C}_0$ is a flagged chamber complex in S^3 that supports the splitting $S^3 = A \cup_T B$, and*

$$\mathbb{C}_0 \xrightarrow{\mathcal{D}_0} \mathbb{C}_1$$

is a flagged chamber complex decomposition.

Let $\mathcal{D}_0^x$ and $\mathcal{D}_0^y$ be possibly different preferred alignments of the disk set $\mathcal{D}_0$ in $\mathbb{C}_0$, with resulting flagged chamber complexes $\mathbb{C}_1^x$ and $\mathbb{C}_1^y$ respectively. Then if $\mathbb{C}_0$ and $\mathbb{C}_1^x$ cocertify, so do $\mathbb{C}_0$ and $\mathbb{C}_1^y$.

Proof. If any occupied handlebody chamber C in $\mathbb{C}_0$ is decomposed entirely into disk handlebody chambers in $\mathbb{C}_1$, then from Proposition 8.1(4) $\mathbb{C}_0$ and $\mathbb{C}_1$ cocertify in any preferred alignment and we are done. So we will assume there is no such decomposition of a chamber. Then by Proposition 7.7(2) a handlebody in $\mathbb{C}_1$ is empty if and only if it is disk. This implies that $\mathbb{C}_1^x$ and $\mathbb{C}_1^y$ have the same flagging.

In the case that $\mathbb{C}_1^x$ certifies because a chamber contains an incompressible sphere, the proof is the same as that of Lemma 6.5, with this augmentation: Following Proposition 7.10, the various bubble passes that are needed can be done in such a way that at every stage the disk alignments are preferred and give the same flagging.

In the case that $\mathbb{C}_1^x$ certifies because a chamber is an occupied handlebody, the proof is similar. From [FS2] there is a sequence of bubble passes and eyeglass moves that change the alignment $\mathcal{D}_0^x$ to $\mathcal{D}_0^y$. From Proposition 7.10 this can be done so that in the sequence the alignments are always preferred and the flagging never changes. So it suffices to consider the case of a single bubble pass. With no loss of generality then assume that $\mathbb{C}_1^x$ and $\mathbb{C}_1^y$ differ by a single bubble pass. If the bubble pass does not involve an occupied handlebody, there is nothing to prove. So assume that $\mathbb{C}_1^y$ is obtained from $\mathbb{C}_1^x$ by passing a bubble $\mathfrak{b}$ out of an occupied handlebody chamber C' of $\mathbb{C}_1$.

Since the flagging remains the same after the bubble passes out of C' , $\mathfrak{b}$ is not a maximal bubble for C' in $\mathbb{C}_1^x$. Following Lemma 7.9 there is a bubble $\mathfrak{b}^y$ in C' that is disjoint from $\mathfrak{b}$ and the tube sum of the two bubbles $\mathfrak{b}^x$ is a maximal bubble for C' in $\mathbb{C}_1^x$. In particular, $\mathfrak{b}^y$ is a maximal bubble for C' in $\mathbb{C}_1^y$. Since the spheres $S^x = \partial\mathfrak{b}^x$ and $S^y = \partial\mathfrak{b}^y$ are disjoint, it follows from Lemma 3.5 that the homeomorphisms $h_{S^x}, h_{S^y} : (S^3, T) \rightarrow (S^3, T_g)$ are eyeglass equivalent, so $\mathbb{C}_1^x$ and $\mathbb{C}_1^y$ cocertify, as required. $\square$

Continue with the hypotheses of Proposition 8.5:

Proposition 8.7. *No matter which preferred alignment occurs at each step in $\vec{\mathbb{C}}$, the flagged chamber complexes $\mathbb{C}_0$ and $\mathbb{C}_n$ cocertify.*

Proof. The case $n = 1$ is Lemma 8.6, so assume $n \geq 2$. We can further inductively assume that none of the $\mathbb{C}_i$, $1 \leq i \leq n - 1$ certify, so each handlebody chamber is empty. In this case, the argument proceeds

essentially as in Proposition 6.9: Proposition 6.7 and Lemma 6.8 remain true for flagged chamber complex decompositions, using Proposition 7.10 to ensure that throughout the argument all decomposing disks are in preferred alignment. Then the proof of Proposition 6.9 suffices if $\mathbb{C}_n$ contains an incompressible sphere. If instead $\mathbb{C}_n$ certifies because it contains an occupied handlebody, the proof concludes as does the proof of Lemma 8.6. $\square$

Following Proposition 8.7 it is possible to make the following definition:

Definition 8.8. Suppose (S^3, T) is a genus g Heegaard splitting of S^3 , and

$$\vec{\mathbb{C}}: \quad \mathbb{C}_0 \xrightarrow{\mathcal{D}_0} \mathbb{C}_1 \xrightarrow{\mathcal{D}_1} \mathbb{C}_2 \xrightarrow{\mathcal{D}_2} \dots \xrightarrow{\mathcal{D}_{n-1}} \mathbb{C}_n$$

is a sequence of flagged chamber complex decompositions supporting T . If any chamber complex $\mathbb{C}_i$ certifies we say that the sequence $\vec{\mathbb{C}}$ certifies and let $h_{\vec{\mathbb{C}}}: (S^3, T) \rightarrow (S^3, T_g)$ be (the eyeglass equivalence class of) any homeomorphism given by a certifying flagged chamber complex in the sequence.

Two flagged chamber complex sequences $\vec{\mathbb{C}}$ and $\vec{\mathbb{C}}'$ cocertify (written $\vec{\mathbb{C}} \sim \vec{\mathbb{C}}'$) if both sequences certify and $h_{\vec{\mathbb{C}}} \sim h_{\vec{\mathbb{C}}'}$.

For example, suppose $\tau \in G(S^3, T)$ and $\vec{\mathbb{C}}$ is a flagged chamber complex decomposition sequence supporting T as above. There is a natural way to define a similar flagged chamber complex decomposition sequence $\tau(\vec{\mathbb{C}})$, namely

$$\tau(\vec{\mathbb{C}}): \quad \tau(\mathbb{C}_0) \xrightarrow{\tau(\mathcal{D}_0)} \tau(\mathbb{C}_1) \xrightarrow{\tau(\mathcal{D}_1)} \tau(\mathbb{C}_2) \xrightarrow{\tau(\mathcal{D}_2)} \dots \xrightarrow{\tau(\mathcal{D}_{n-1})} \tau(\mathbb{C}_n)$$


Corollary 8.9. If $\vec{\mathbb{C}}$ certifies and $\tau \in G(S^3, T)$ then $h_{\tau(\vec{\mathbb{C}})} \tau \sim h_{\vec{\mathbb{C}}}$.

Proof. This follows immediately from Corollary 8.4. $\square$

9 Guiding spheres

Suppose $\mathbb{C}$ is a flagged chamber complex in S^3 with defining surface $F = F(\mathbb{C})$, and $S \subset S^3$ is an embedded sphere transverse to F . In this and following sections we intend to associate to each such sphere S a sequence $\overrightarrow{(\mathbb{C}, S)}$ of flagged chamber complex decompositions so that:

1. For $\mathbb{C}$ not tiny and $S_t, 0 \leq t \leq 1$ a sweep-out of S^3 by spheres, there is a non-empty interval $(a, b) \subset I$ so that, for each $t \in (a, b)$, $\vec{\mathbb{C}}_{S_t}$ certifies.
2. For any $t, t' \in (a, b)$ the sequences $\vec{\mathbb{C}}_{S_t}$ and $\vec{\mathbb{C}}_{S_{t'}}$ cocertify.
3. If, at any point $\mathbb{C}_i \xrightarrow{\mathcal{D}_i} \mathbb{C}_{i+1}$ in the sequence $\overrightarrow{(\mathbb{C}, S)}$, a disk E , disjoint from S , is added to $\mathcal{D}_i$, the resulting sequence of flagged chamber complex decompositions and the sequence $\vec{\mathbb{C}}_S$ cocertify.



 Figure 26: $S \cap F$ and its associated tree

Let $F = F(\mathbb{C})$ be the defining surface of $\mathbb{C}$ and, in the standard way, associate to S and F a tree Y , in which each vertex corresponds to a component of $S - F$ and each edge corresponds to a (circle) component of $S \cap F$. An edge associated to a circle $c \subset S - F$ connects the two vertices that correspond to the components of $S - F$ that are incident to c . Then the leaves of Y correspond to the disk components of $S - F$, and the outermost edges in Y correspond to innermost circles in S of $F \cap S$. See Figure 26.

Definition 9.1. For any edge $e \in Y$, with end vertices $v_{\pm}$, $Y - e$ consists of two trees $Y_{\pm}$, with $v_{\pm} \in Y_{\pm}$. Let

$$m_{\pm} = \max\{d(v_{\pm}, v_{\ell}) \mid v_{\ell} \text{ a leaf in } Y \text{ lying in } Y_{\pm}\}$$

and $\rho_Y(e) = \min\{m_+, m_-\}$. See Figure 27.


 Figure 27: Edge e and its trees Y_{pm} ; $m_+ = 4, m = m_- = 2$

For example, for e an outermost edge, one of its incident vertices, v_- say, is a leaf, so $Y_- = v_-$ and $\rho_Y(e) = m_- = 0$.

Lemma 9.2. Suppose e is an edge in a tree Y .

1. Suppose, with no loss of generality, that $\rho_Y(e) = m_- \leq m_+$ and e' is an edge in Y_- . Then $\rho_Y(e') < \rho_Y(e)$.

2. Let Y' be the tree obtained from Y by trimming: that is, removing all leaves and outermost edges from Y . For $e \in Y'$, $\rho_{Y'}(e) = \rho_Y(e) - 1$.
3. $0 \leq \rho_Y(e) \leq \lfloor \frac{\text{diam}(Y)-1}{2} \rfloor$

Proof. We prove the items in turn:

1. Let Y'_- be the component of $Y - e'$ that is contained in Y_- and v'_- the end of e' in Y'_- . For any leaf v'_ℓ of Y'_- that lies in Y , the path from v_- to v'_ℓ in Y_- passes through e' , so

$$\rho_Y(e) = m_- \geq d(v_-, v'_\ell) \geq d(v'_-, v'_\ell) + 1.$$

Since this is true for each leaf v'_ℓ of Y'_- ,

$$\rho_Y(e) \geq \max\{d(v'_-, v'_\ell) | v'_\ell \text{ a leaf in } Y'_- \text{ lying in } Y\} + 1 \geq \rho_{Y'}(e') + 1$$

2. Note that any leaf v'_ℓ in Y'_+ that lies in Y' is necessarily incident to some outermost edges of Y , so $d(v_+, v'_\ell) \leq m_+ - 1$. This implies that

$$m'_+ = \max\{d(v_+, v'_\ell) | v'_\ell \text{ a leaf in } Y'_+ \text{ lying in } Y'\} \leq m_+ - 1.$$

On the other hand, if v_ℓ is the most distant leaf from v_+ in Y_+ that lies in Y then the outermost edge incident to v_ℓ is incident to a leaf v'_ℓ in Y'_+ , so $m'_+ \geq m_+ - 1$. Together these imply $m'_+ = m_+ - 1$. Similarly $m'_- = m_- - 1$. Hence

$$\rho_{Y'}(e) = \min\{m'_+, m'_-\} = \min\{m_+, m_-\} - 1 = \rho_Y(e) - 1,$$

as required.

3. Repeat the process of trimming $\lfloor \frac{\text{diam}(Y)-1}{2} \rfloor$ times. Each step reduces the diameter of the tree by 2 and, by conclusion (2), reduces the maximum of ρ on any remaining edge by 1. Thus it suffices to verify the case in which $\text{diam}(Y) = 1$ or 2, i. e. Y is a single edge or a star graph, depending on the parity of $\text{diam}(Y)$. In both these cases, $\rho(e) = 0$ as required.

□

Partition the collection of circles $S \cap F$ as follows: For $0 \leq i \leq \lfloor \frac{\text{diam}(Y)-1}{2} \rfloor$ let c_i be the set of all circles for which the corresponding edge e in Y has $\rho_Y(e) = i$.

Lemma 9.3. *Suppose $c \in c_i$. Then c bounds a disk D in S such that each circle $c' \subset \text{int}(D) \cap F$ lies in some c_j , $j < i$.*

Proof. Let Y be the tree corresponding to $S \cap F$ and apply Lemma 9.2(1) to the edge in Y corresponding to c . The trees $Y_\pm$ are trees describing the circles $S \cap F$ lying in the complementary disks $D_\pm$ of c in S , with $v_\pm$ corresponding to the components of $D_\pm - F$ that are adjacent to c . With no loss of generality suppose D_- corresponds to Y_- with $i = \rho_Y(e) = m_-$. Then Lemma 9.2(1) says that any component c' of $\text{int}(D_-) \cap F$ has corresponding edge e' in Y' with $\rho_Y(e') < \rho_Y(e) = i$. Setting $j = \rho_Y(e')$ we have $c' \in c_j$. □

Each circle in c_0 corresponds to a leaf in Y , so each $c \in c_0$ bounds a disk in $S - F$. Let $\mathcal{D}_0$ be the collection of disks, viewed as a disk set in $\mathbb{C}$. Choose a preferred alignment for $\mathcal{D}_0$ and consider the flagged chamber complex decomposition $\mathbb{C} \xrightarrow{\mathcal{D}_0} \mathbb{C}_1$. Before the decomposition each circle in c_1 bounded a disk in S that intersected F only in components of c_0 , by Lemma 9.3. So after the disk decomposition along $\mathcal{D}_0$, which surgers away all components of c_0 , each circle in c_1 bounds a disk in S that is disjoint from the defining surface $F_1 = F(\mathbb{C}_1)$. More correctly, this is true of each *remaining* circle in c_1 , for some circles in c_1 may lie on the boundary of goneballs of the decomposition by $\mathcal{D}_0$, and so not still appear in $F_1 \cap S$. In any case, the collection of disks in $S - F_1$ bounded by (remaining) circles in c_1 will be denoted $\mathcal{D}_1$ and a choice of preferred alignment defines a flagged chamber complex decomposition $\mathbb{C}_1 \xrightarrow{\mathcal{D}_1} \mathbb{C}_2$. Continue in this manner through $\mathbb{C}_{n-1} \xrightarrow{\mathcal{D}_{n-1}} \mathbb{C}_n$, where $n - 1 = \lfloor \frac{\text{diam}(Y)-1}{2} \rfloor$ or $n = \lfloor \frac{\text{diam}(Y)+1}{2} \rfloor$. In the end, $F_n = F(\mathbb{C}_n)$ is then disjoint from S .

There are potentially two sources of ambiguity in this construction:

One ambiguity is this: in the last decomposition in the sequence, $\mathbb{C}_{n-1} \xrightarrow{\mathcal{D}_{n-1}} \mathbb{C}_n$, it's possible that the defining surface $F_{n-1} = F(\mathbb{C}_{n-1})$ intersects S in a single circle, so $|c_{n-1}| = 1$. This circle divides S into two disks, one lying in the incident A -chamber of $\mathbb{C}_{n-1}$ and the other lying in the incident B -chamber. The description above would allow either of these disks to be used as the disk set $\mathcal{D}_{n-1}$ in the last decomposition. Will it make a difference which one we use? This issue will be addressed later - see Lemma 10.4.

The second potential source of ambiguity is that *prima facie* the construction above depends at each stage on the choice of preferred alignment of the disks $\mathcal{D}_i$. So the process is more accurately described as defining a *family* of decomposition sequences, which may differ at several points in the choice of preferred alignment. We will now show that if any sequence in the family certifies, then they all cocertify.

Definition 9.4. Given $\mathbb{C}$ a flagged chamber complex in S^3 , S a sphere transverse to $F(\mathbb{C})$, and Y the tree associated to the circles $F \cap S$ in S , a sequence of flagged chamber complex decompositions

$$\vec{\mathbb{C}}: \quad \mathbb{C} = \mathbb{C}_0 \xrightarrow{\mathcal{D}_0} \mathbb{C}_1 \xrightarrow{\mathcal{D}_1} \mathbb{C}_2 \xrightarrow{\mathcal{D}_2} \dots \xrightarrow{\mathcal{D}_{k-1}} \mathbb{C}_k$$

as constructed above is said to be guided by the sphere S . If $k = n = \lfloor \frac{\text{diam}(Y)+1}{2} \rfloor$ the sequence is complete.

The collection of all complete sequences will be denoted $(\overrightarrow{\mathbb{C}}, S)$.

Continue with Assumption 3.4, that $G(S^3, T') = \mathcal{E}$ whenever $\text{genus}(T') \leq g - 1$.

Proposition 9.5. If $(\overrightarrow{\mathbb{C}}, S)$ contains more than one complete sequence of flagged chamber complex decompositions then each sequence in $(\overrightarrow{\mathbb{C}}, S)$ certifies and all sequences cocertify.

Proof. Suppose

$$\vec{\mathbb{C}}: \mathbb{C}_0 \xrightarrow{\mathcal{D}_0} \mathbb{C}_1 \xrightarrow{\mathcal{D}_1} \dots \xrightarrow{\mathcal{D}_{n-1}} \mathbb{C}_n$$

and

$$\vec{\mathbb{C}}': \mathbb{C}_0 \xrightarrow{\mathcal{D}_0} \mathbb{C}'_1 \xrightarrow{\mathcal{D}'_1} \dots \xrightarrow{\mathcal{D}'_{n-1}} \mathbb{C}'_n$$

are any two sequences in $\overrightarrow{(\mathbb{C}, S)}$. and $i \geq 1$ is the smallest value for which $\mathbb{C}_i \neq \mathbb{C}'_i$. Then $\mathbb{C}_{i-1} = \mathbb{C}'_{i-1}$ and $\mathcal{D}_{i-1} = \mathcal{D}'_{i-1}$ are the same as disk sets but are given different preferred alignments. In particular, by Proposition 7.7(3), the decompositions $\mathbb{C}_{i-1} \xrightarrow{\mathcal{D}_{i-1}} \mathbb{C}_i$ and $\mathbb{C}_{i-1} \xrightarrow{\mathcal{D}'_{i-1}} \mathbb{C}'_i$ contain sibling decompositions. That is, there is at least one parent chamber C in $\mathbb{C}_{i-1}$ so that C is an occupied handlebody and $\mathcal{D}_{i-1}$ is given a different preferred alignment in C for the two decompositions. Then C certifies for both $\vec{\mathbb{C}}$ and $\vec{\mathbb{C}}'$. $\square$

If $\vec{\mathbb{C}} \in \overrightarrow{(\mathbb{C}, S)}$ certifies, denote a homeomorphism $h_{\vec{\mathbb{C}}} : (S^3, T) \rightarrow (S^3, T_g)$ (see Definition 8.8) by $h_{(\mathbb{C}, S)}$. By Proposition 9.5 $h_{(\mathbb{C}, S)}$ is well-defined up to eyeglass equivalence; that is, it does not depend on the choice of decomposition sequence $\vec{\mathbb{C}} \in \overrightarrow{(\mathbb{C}, S)}$.

Corollary 9.6. *If $\vec{\mathbb{C}} \in \overrightarrow{(\mathbb{C}, S)}$ certifies and $\tau \in G(S^3, T)$ then $h_{(\tau(\mathbb{C}), \tau(S))} \tau \sim h_{(\mathbb{C}, S)}$.*

Proof. This follows immediately from Corollary 8.9. $\square$

10 Balanced and almost balanced spheres

Definition 10.1. *Suppose $S \subset S^3$ is a sphere transverse to $F = F(\mathbb{C})$ in a flagged chamber complex $\mathbb{C} \subset S^3$, with X, Y the complementary components of S .*

Then S is balanced for $\mathbb{C}$ if the compact surfaces $F \cap X$ and $F \cap Y$ are either both planar (planar balanced) or both non-planar (non-planar balanced).

S is almost balanced if $\text{genus}(F \cap X) = 0$ and $\text{genus}(F \cap Y) = 1$ or vice versa.

Proposition 10.2. . *Suppose*

$$\vec{\mathbb{C}} : \quad \mathbb{C} = \mathbb{C}_0 \xrightarrow{\mathcal{D}_0} \mathbb{C}_1 \xrightarrow{\mathcal{D}_1} \dots \xrightarrow{\mathcal{D}_{n-1}} \mathbb{C}_n$$

is a complete flagged chamber complex decomposition sequence guided by $S \subset S^3$. That is, $\vec{\mathbb{C}} \in \overrightarrow{(\mathbb{C}, S)}$. Then

- If S is planar (resp non-planar) balanced in any chamber complex in the sequence, then it is in every chamber complex in the sequence. In this case, call S planar (resp non-planar) balanced for the sequence.*
- If S is balanced for the sequence and $\mathbb{C}$ is not tiny, then $\mathbb{C}_n$ certifies, so $\vec{\mathbb{C}}$ certifies.*

Proof. For the first statement, note that the effect on $F \cap X$, say, of decomposing along $\mathcal{D}_i$ is 2-fold:

1. Cap off some of boundary components of $F \cap X$ with disks.
2. Delete some spheres from F (the boundaries of the goneballs).

Neither of these steps affects the genus of $F \cap X$ or $F \cap Y$.

For the second statement, note that since $\mathbb{C}$ is not tiny, $\mathbb{C}_n$ is not empty (Corollary 7.13). Let $F = F(\mathbb{C}_n) \neq \emptyset$ be the defining surface for $\mathbb{C}_n$. When the sequence is non-planar balanced, there are (non-planar) closed components of F in both X and Y . This implies that S is a reducing sphere for the chamber of $\mathbb{C}_n$ in which it lies, so S itself is a certificate. If the sequence is planar balanced, then each component of F is a closed planar surface, i. e. a sphere. Since there are no empty balls in a flagged chamber complex, an innermost sphere in F bounds an occupied ball, which certifies (see Definition 8.2). $\square$

Similarly we have:

Proposition 10.3. *Suppose*

$$\vec{\mathbb{C}}: \quad \mathbb{C} = \mathbb{C}_0 \xrightarrow{\mathcal{D}_0} \mathbb{C}_1 \xrightarrow{\mathcal{D}_1} \dots \xrightarrow{\mathcal{D}_{n-1}} \mathbb{C}_n$$

is a complete flagged chamber complex decomposition sequence guided by $S \subset S^3$. That is, $\vec{\mathbb{C}} \in \overline{(\mathbb{C}, S)}$. Then

- *If S is almost balanced in any chamber complex in the sequence, then it is in every chamber complex in the sequence. In this case, call S almost balanced for the sequence.*
- *If S is almost balanced for the sequence and $\mathbb{C}$ is not tiny, then $\mathbb{C}_n$ certifies, so $\vec{\mathbb{C}}$ certifies.*

Proof. The proof of the first statement is unchanged.

For the proof of the second, observe that since the sequence is complete, $F_n = F(\mathbb{C}_n)$ is disjoint from S , so both $F_n \cap X$ and $F_n \cap Y$ are closed surfaces. Since F_n is almost balanced, $F_n \cap X$ has genus 0 and $F_n \cap Y$ has genus 1, or vice versa. In any case F_n is the union of a torus and some spheres. If there are spheres, at least one must bound an occupied ball and so certifies. If F is just a torus it bounds a solid torus Y in S^3 , or perhaps complementary solid tori $Y_{\pm}$. Since $\mathbb{C}$ is not tiny, $\mathbb{C}_n$ is not tiny (Corollary 7.13), so Y cannot be empty nor, in the second case, can both solid tori $Y_{\pm}$ be empty. Hence one of these is an occupied solid torus and so is a certificate for $\mathbb{C}_n$. $\square$

We now briefly return to what was described, preceding Definition 9.4, as the first source of ambiguity in the construction of a sequence of flagged chamber complex decompositions guided by S . Suppose $\mathbb{C}$ is a flagged chamber complex supporting the genus g splitting $S^3 = A \cup_T B$ and S is a sphere in S^3 intersecting $\mathbb{C}$ in a single circle c and dividing S^3 into two 3-balls, X and Y . The circle c divides S into two disks D_A lying in an A -chamber C_A of $\mathbb{C}$ and D_B lying in a B -chamber C_B .

Lemma 10.4. *Let $\mathbb{C}_A$ and $\mathbb{C}_B$ be the flagged chamber complexes that result from the decompositions $\mathbb{C} \xrightarrow{D_A} \mathbb{C}_A$ and $\mathbb{C} \xrightarrow{D_B} \mathbb{C}_B$ respectively. Suppose S is balanced or almost balanced for $\mathbb{C}$ so both $\mathbb{C}_A$ and $\mathbb{C}_B$ certify. Then $\mathbb{C}_A$ and $\mathbb{C}_B$ cocertify.*

Proof. Let c be the circle described before the statement of the lemma.

Case 1: c is essential in $F = F(\mathbb{C})$.

Then c divides F into two non-planar components and F is non-planar balanced. We show that our rule for certifying a non-planar balanced chamber complex cocertifies $\mathbb{C}_A$ and $\mathbb{C}_B$:

As usual, align $\mathcal{D}_A$ in $\mathbb{C}_A$ and $\mathcal{D}_B$ in $\mathbb{C}_B$, via an appropriate choice of spine for the compression bodies $B_{\mathbb{C}_A}$ in $\mathbb{C}_A$ and $A_{\mathbb{C}_B}$ in $\mathbb{C}_B$. Since c is essential in F it is essential in the Heegaard surface T of S^3 that the chamber complex supports. Hence S is a reducing sphere for T . In particular there is a homeomorphism $h : (S^3, T) \rightarrow (S^3, T_g)$ as described before Definition 3.3 in which $h(S) = S_i$ for some $1 \leq i \leq g-1$. But this coincides with the definition of the homeomorphism $h_S : (S^3, T) \rightarrow (S^3, T_g)$ in Proposition 3.7, and that is the homeomorphism which certifies S in both $\mathbb{C}_A$ and $\mathbb{C}_B$. So $\mathbb{C}_A$ and $\mathbb{C}_B$ cocertify in this case.

Case 2: c bounds a disk $D_F \subset F \cap X$ but does not bound a disk in $F \cap Y$ (or vice versa).

Let S_A be the sphere $D_A \cup D_F$, a component of $F(\hat{\mathbb{C}}_A)$ and, symmetrically, S_B be the sphere $D_B \cup D_F$, a component of $F(\hat{\mathbb{C}}_B)$. One possibility is that S_A is essential in $\mathbb{C}_A \cap X$, perhaps because the ball $B_{X \cap A}$ in X bounded by S_A contains other components of $F(\mathbb{C}_A)$ or because it is occupied, and so not a goneball in $\hat{\mathbb{C}}_A$. Then S_A , pushed slightly out of $B_{X \cap A}$, is also essential in $\mathbb{C}_A$ because it separates $B_{X \cap A}$ from the (non-spherical) component of $F(\mathbb{C}_A)$ that contains the circle c . Thus it certifies for $\mathbb{C}_A$. See Figure 28.




Figure 28: Case 2: c bounds a disk in $F \cap X$ but not in $F \cap Y$

If S_B similarly certifies for $\mathbb{C}_B$ then S_A and S_B define disjoint reducing spheres for T and the result follows from Lemma 3.5. On the other hand, if S_B does not certify for $\mathbb{C}_B$ then $B_{X \cap B}$ is a goneball. In this case, decomposing by D_B has no effect, that is $\mathbb{C}_B$ and $\mathbb{C}$ are isotopic, so $\mathbb{C}$ also certifies. By Proposition 8.7, $\mathbb{C}_A \sim \mathbb{C} \sim \mathbb{C}_B$ as required.

Finally, if S_A is inessential in $\mathbb{C}_A \cap X$ and S_B is inessential in $\mathbb{C}_B \cap X$, so $B_{X \cap B}$ and $B_{X \cap A}$ are goneballs, then $\mathbb{C}_A$ and $\mathbb{C}_B$ are isotopic to $\mathbb{C}$, so $\mathbb{C}_A$ and $\mathbb{C}_B$ again cocertify.

Case 3: c lies on a sphere component S_F of F , dividing it into disks $D_X \subset X$ and $D_Y \subset Y$.

Since S_F does not bound a goneball in $\mathbb{C}$, it either divides S^3 into two occupied balls or is parallel to an incompressible sphere in an adjoining chamber of $\mathbb{C}$. Thus S_F is a certificate for $\mathbb{C}$. Then, per Proposition 8.7, $\mathbb{C}_A \sim \mathbb{C} \sim \mathbb{C}_B$ as required. $\square$

This dispenses with the “first ambiguity” preceding Definition 9.4.

Suppose $\mathbb{C} \subset S^3$ is a flagged chamber complex, with $F = F(\mathbb{C})$ in general position with respect to the height projection $p : S^3 \rightarrow [-1, 1]$. Recall that $S^3 - \{\text{poles}\}$ is swept out by the family of level spheres

S_s , where $S_s = p^{-1}(s)$, $s \in (-1, 1)$. We will consider these level spheres as potential guiding spheres for decomposition of $\mathbb{C}$.

Definition 10.5. Let $F = F(\mathbb{C})$, $g = \text{genus}(F)$.

For each $s \in (-1, 1)$ that is a regular value of $p|F$, let $F_{a(\text{bove})}(s)$ be the part of F lying above S_s , that is $F_a(s) = F \cap p^{-1}([s, 1])$. Similarly, let $F_{b(\text{elow})}(s)$ be the part of F lying below S_s , that is $F_b(s) = F \cap p^{-1}([-1, s])$. Finally, let $g_a(s) = \text{genus}(F_a(s))$ and $g_b(s) = \text{genus}(F_b(s))$.

Lemma 10.6. The functions g_a, g_b have these properties:

1. As s ascends from -1 to 1 , g_b ascends from 0 to g and g_a descends from g to 0 .
2. As s ascends through a critical point of $p|F : F \rightarrow [-1, 1]$, $g_b(s)$ may increase by 1 but it will not decrease. Symmetrically, $g_a(s)$ may decrease by 1 but it will not increase.
3. As S_s ascends through a critical point of $p|F$, at most one of $g_a(s), g_b(s)$ will change.

Proof. The first statement follows from the definition.

For the second and third statements, observe that passing through a critical point on F of index 0 (a minimum) or 2 (a maximum) simply adds or subtracts a disk from F_a and F_b , which has no effect on genus. So the interest is in index 1 critical points, that is saddle points of tangency.


Suppose, with little loss in generality, that as s ascends through a saddle tangency, two circles in $F \cap S_s$ fuse into one. See Figure 29. (If instead one circle is divided in two, just turn the following argument upside down.) Thus a band is added to F_b with its ends on two different boundary circles. If the two circles lie on different components of F_b , the genus of F_b does not change. If the two lie on the same component, g_b ascends by 1 . This proves the second statement. For the third statement, consider F_a : the effect on F_a is to cut out a band between two boundary components, effectively replacing a pair of pants in F_a with an annulus. This has no effect on g_a . This proves the third statement. $\square$

Corollary 10.7. Let

$$s_b = \sup\{s \in [0, 1] | g_b(s) = 0\} \quad s_a = \inf\{s \in [0, 1] | g_a(s) = 0\}.$$

Then

1. $s_a \neq s_b$.
2. If $s_a < s_b$ then for all $s_a < s < s_b$, S_s is planar balanced.
3. If $s_b < s_a$ then for all $s_b < s < s_a$, S_s is non-planar balanced.
4. If $s_a < s_b$ then for s just below s_a and s just above s_b , S_s is almost balanced.
5. If $s_b < s_a$ and $p|F$ has no index 1 critical points for $s \in (s_b, s_a)$ then for s just below s_b and s just above s_a , S_s is almost balanced.


 Figure 29: s ascends through a saddle and two circles fuse

Proof. All but the last two statements follow directly from the definitions and Lemma 10.6; the last two require little more:

It follows from Lemma 10.6(2) that for s just above s_b , $g_b(s) = 1$ and, if $s_a < s_b$, $s \in \{s | g_s(a) = 0\}$ so also $g_a(s) = 0$. So S_s is almost balanced. A symmetric argument shows that it is also almost balanced for s just below s_a . This proves (4).

Suppose $s_b < s_a$. Then for s just above s_b , $g_b(s) = 1$ and $g_a(s) \geq 1$; similarly for s just below s_a , $g_a(s) = 1$ and $g_b(s) \geq 1$. Since $p|F$ has no index one critical points in (s_b, s_a) , neither g_a or g_b changes in that interval. Hence for any $s \in (s_b, s_a)$, $g_a(s) = 1 = g_b(s)$. It follows from Lemma 10.6(3) that for s just above s_a , $g_a(s) = 0$ and $g_b(s) = 1$ so S_s is almost balanced. The symmetric argument shows that for s just below s_b , S_s is almost balanced. This proves (5). $\square$

Suppose that the chamber complex $\mathbb{C}$ is not tiny. Then for any s between s_a and s_b it follows from Corollary 10.7 and Propositions 10.2 and 9.5 that the homeomorphism

$$h_s = h_{(\mathbb{C}, S_s)} : (S^3, T) \rightarrow (S^3, T_g)$$

is well-defined up to eyeglass equivalence. What remains unclear is whether the eyeglass equivalence class of h_s depends on s . The central issue is whether the eyeglass equivalence class can change as s passes through a critical point of $p|F$. We will show (Theorem 16.2) that in fact the eyeglass equivalence class does not change. The proof is technically complex; it occupies the next five sections.

11 Deflation and bullseyes

Suppose $\mathbb{C} \subset S^3$ is a flagged chamber complex supporting the genus g Heegaard splitting $S^3 = A \cup_T B$. We continue under the inductive Assumption 3.4.

Let H be an occupied handlebody chamber of $\mathbb{C}$ that is not a ball. In the following description we assume that H is an A -chamber, but everything can also be taken symmetrically, see Definition 11.1. Let C be the B -chamber adjacent to H and C_A be an A -chamber adjacent to C so that $C_A \neq H$. Since the associated Heegaard splitting of $C = A_C \cup_{T_C} B_C$ is pure, there is a spine for the compression body A_C in which an edge e is incident to ∂H on one end and ∂C_A on the other.

Let $(D_e, \partial D_e) \subset (A_C, T_C)$ be the disk dual to e in A_C , that is a meridian of the tube in A_C that corresponds to a regular neighborhood of the edge e , so the tube runs from H to C_A . Alter D_e by tube-summing with a maximal bubble b in H and let $\mathbb{C}_d$ be the resulting flagged chamber complex. (More precisely, after amalgamation of $\mathbb{C}$ to the original splitting, D_e becomes a disk in A and b a bubble for T ; it is these that are tube-summed to create the disk D'_e that replaces D_e in $\mathbb{C}_d$.) In effect, the bubble b has been moved from H to the chamber C_A . See Figure 30. Since b is maximal in H , H changes from an *occupied* handlebody in $\mathbb{C}$ to an *empty* handlebody in $\mathbb{C}_d$. It is natural to call this process a *deflation* of H . That is, the flagged chamber complex $\mathbb{C}_d$ is obtained from $\mathbb{C}$ by deflation (of H to C_A). Note that the process is well-defined up to eyeglass equivalence by a choice of H and C_A , since a bubble move is an eyeglass move.




Figure 30: Simple deflation

More generally, the process can be applied to a collection of edges in a spine of A_C , each of which is incident to ∂H and an A -chamber adjacent to C , but not necessarily the same A -chamber for all the edges. That is, redefine the meridian A -disks dual in A_C to the collection of edges by tube-summing them with disjoint non-trivial bubbles in H , bubbles which together constitute a maximal bubble in H . Generalizing further, bubbles from H may be sent to further A -chambers $\{C_{A_i}\}$ throughout $\mathbb{C}$, not just A -chambers adjacent to C , by further tube-summing to A -disks in B -chambers. See Figure 31. The only

requirement is that in the resulting flagged chamber complex $\mathbb{C}_d$, the chamber H is empty. In this general form we have:




Figure 31: General deflation

Definition 11.1. The flagged chamber complex $\mathbb{C}_d$ is obtained from $\mathbb{C}$ by handlebody deflation (of H to the collection of chambers $\{C_{A_i}\}$). Deflation of an occupied handlebody B -chamber of $\mathbb{C}$ that is not a ball is defined symmetrically.

As is true for C_A a single chamber, the change of flagged chamber complex from $\mathbb{C}$ to $\mathbb{C}_d$ is well-defined up to eyeglass equivalence by a choice of H , the collection $\{C_{A_i}\}$, and the genus of the bubbles moved into each. The effect of deflation on the flagging of $\mathbb{C}$ is to change H to an empty handlebody chamber in $\mathbb{C}_d$ (hence the term deflation) and, for those chambers in $\{C_{A_i}\}$ that are empty handlebodies, to change the flagging to occupied.

Proposition 11.2. Suppose the flagged chamber complex $\mathbb{C}_d$ is obtained from $\mathbb{C}$ by deflating an occupied non-ball handlebody H . If $\mathbb{C}_d$ certifies then $\mathbb{C}_d$ and $\mathbb{C}$ cocertify.

Proof. The occupied handlebody H is a certificate for $\mathbb{C}$ so $\mathbb{C}$ certifies. With no loss assume, as above, that H is an A -chamber deflated into the set of A -chambers $\{C_{A_i}\}$. Suppose $\mathbb{C}_d$ certifies and let C_d be a certifying chamber of $\mathbb{C}_d$. Since H is empty in $\mathbb{C}_d$, $C_d \neq H$. The only other chambers whose flagging is affected by the deflation of H are $\{C_{A_i}\}$ so, unless $C_d \in \{C_{A_i}\}$, the chamber C_d also certifies for $\mathbb{C}$, as required. So suppose $C_d \in \{C_{A_i}\}$, i. e. after the bubble pass one or more of the chambers in $\{C_{A_i}\}$, call it C_a , certifies for $\mathbb{C}_d$.

The only change in C_a from $\mathbb{C}$ to $\mathbb{C}_d$ is that a bubble b from H has been passed into it. Since the chamber C_a certifies in $\mathbb{C}_d$, it is either an occupied handlebody or it contains an incompressible sphere. If $\mathbb{C}_d$ is an occupied handlebody then, per Proposition 8.1(1), $\partial b \subset C_a$ is a certificate for $\mathbb{C}_d$ as well as $\mathbb{C}$ completing the proof. Suppose, on the other hand, that the chamber C_a contains an incompressible sphere

S . By [Sc1] S can be aligned with the Heegaard surface T_a for C_a in $\mathbb{C}$, so S becomes a certificate for $\mathbb{C}$. In constructing $\mathbb{C}_d$ pass the bubble b into the chamber C_a so that b remains disjoint from the aligned S . Then the sphere S certifies for both $\mathbb{C}$ and $\mathbb{C}_d$, as required. $\square$

The same construction can be done when H is an occupied ball, but an important difference is that once the occupied ball is deflated, so H has a genus 0 splitting, H must be declared a goneball, else $\mathbb{C}_d$ is not a flagged chamber complex. That is, the sphere ∂H is removed from the defining surface F and H is absorbed into its neighboring chamber (the chamber C in the construction above). This difference requires a change in strategy to arrive at the equivalent of Proposition 11.2.

A set of disjoint spheres in a ball B^3 is *nested* if only one complementary component is a ball.

Definition 11.3. Suppose F_b is a nested set of spheres in a 3-ball B^3 , with ∂B^3 the outermost sphere. Call the collar between any pair of spheres a shell. A bullseye is the flagged chamber complex in B^3 whose defining surface is F_b , each shell has a genus 0 (pure) Heegaard splitting, and the ball B_b bounded by the innermost sphere is occupied with maximal bubble b .

A bullseye in a chamber complex $\mathbb{C}$ is maximal if it is not contained in any larger bullseye.

We now extend the notion of deflation to include bullseyes, not just occupied non-ball handlebodies. Suppose F_b is a bullseye in a chamber complex $\mathbb{C}$ and, with no loss of generality, the chamber adjacent to F_b and outside it is a B -chamber C . Let $\{C_{A_i}\}$ (resp. $\{C_{B_i}\}$) be a collection of A -chambers (resp. B -chambers) in $\mathbb{C}$, other than those contained in F_b . Use the same process as described above for deflating a handlebody to deflate b either into the chambers $\{C_{A_i}\}$, if the ball chamber of F_b is an A -chamber, or into $\{C_{B_i}\}$, if the ball chamber in F_b is a B -chamber. (In the former case, the number of nested spheres in F_b is odd; in the latter it is even.) After the deflation, each sphere in F_b bounds a goneball, so the spheres in F_b disappear in the resulting flagged chamber complex.


Definition 11.4. The flagged chamber complex $\mathbb{C}_d$ is obtained from $\mathbb{C}$ by bullseye deflation (of F_b to $\{C_{A_i}\}$ or $\{C_{B_i}\}$) or, informally, deleting a bullseye. Similarly, we informally say that $\mathbb{C}$ is obtained from $\mathbb{C}_d$ by inserting a bullseye into C .

Proposition 11.5. Suppose the flagged chamber complex $\mathbb{C}_d$ is obtained from $\mathbb{C}$ by deflating a bullseye F_b . If $\mathbb{C}_d$ certifies then $\mathbb{C}_d$ and $\mathbb{C}$ cocertify.

Proof. The proof proceeds as for Proposition 11.2, with one difference. In the case of bullseye deflation, the chamber C changes, because all the spheres in F_b disappear. That is, the bullseye becomes a nested set of goneballs that are absorbed into C to become a chamber we denote C_d . See Figure 32. The proof then proceeds much as in Proposition 11.2 except, unlike the situation in Proposition 11.2, it is possible that C turns into the lone certifying chamber C_d for $\mathbb{C}_d$. We examine this possibility:

If C_d certifies because it contains an incompressible sphere, let $S \subset C_d$ be such an incompressible sphere. By general position (with a point in the center of the ball chamber of F_b) S can be isotoped to be disjoint from F_b and so lie in C . Align S with the Heegaard splitting of C . Then it is aligned also with the Heegaard splitting of C_d regardless of whether $C \in \{C_{B_i}\}$. The sphere S then certifies for both $\mathbb{C}$ and $\mathbb{C}_d$ as required.

Finally, suppose C_d certifies because it is an occupied handlebody. There are two possibilities: $C \in \{C_{B_i}\}$ and $C \notin \{C_{B_i}\}$. If $C \in \{C_{B_i}\}$ then a sub-bubble b' of b becomes a bubble in C_d . $\partial b'$ is disjoint


 Figure 32: A nested set of goneballs absorbed into C

from ∂b , which certifies for $\mathbb{C}$, so by Lemma 3.5 $\partial b'$ also certifies for $\mathbb{C}$. After deflation, the sphere $\partial b' \subset C_d$, may be taken to be disjoint from the boundary of the maximal bubble of C_d , by Lemma 7.9, so $\partial b'$ also certifies for $\mathbb{C}_d$. Hence $\mathbb{C}$ and $\mathbb{C}_d$ cocertify, as required.

On the other hand, if $C \notin \{C_{B_i}\}$, so the entire bubble b is dispersed away from C_d and still C_d is occupied, the maximal bubble b_d of C_d must have come from the chamber C in $\mathbb{C}$. Since, in $\mathbb{C}$, b_d lies in C , it is disjoint in C from b . Hence, by Lemma 3.5, ∂b_d and ∂b cocertify, as required. $\square$

Suppose $\mathbb{C}$ is a flagged chamber complex in S^3 and $\mathbb{C}_d$ is a flagged chamber complex obtained by (non-ball) handlebody deflation. Suppose $\mathcal{D}$ is a disk set in $\mathbb{C}$ in preferred alignment and

$$\mathbb{C} \xrightarrow{\mathcal{D}} \mathbb{C}_{\mathcal{D}}$$

is the resulting flagged chamber complex decomposition. Similarly construct the flagged chamber complex decomposition

$$\mathbb{C}_d \xrightarrow{\mathcal{D}} (\mathbb{C}_d)_{\mathcal{D}}$$

Proposition 11.6. *If $\mathbb{C}_d$ does not certify, then either $(\mathbb{C}_d)_{\mathcal{D}} = \mathbb{C}_{\mathcal{D}}$ or $(\mathbb{C}_d)_{\mathcal{D}}$ is obtained from $\mathbb{C}_{\mathcal{D}}$ by handlebody deflation or by bullseye deflation.*

Proof. With no loss of generality assume the deflated non-ball handlebody H is an A -chamber and the chambers C and $\{C_{A_i}\}$ are as described preceding Definition 11.1. The chambers $\{C_{A_i}\}$ cannot contain a handlebody because when H is deflated into $\{C_{A_i}\}$ the result would be an occupied handlebody in $\mathbb{C}_d$, contradicting the assumption that $\mathbb{C}_d$ does not certify. So the only difference between $\mathbb{C}$ and $\mathbb{C}_d$ as flagged chamber complexes is that H is occupied in the former but is empty in the latter. Thus the difference between $\mathbb{C}_{\mathcal{D}}$ and $(\mathbb{C}_d)_{\mathcal{D}}$, if any, is in the flagging of the remnants of H , perhaps including the disappearance of remnants as goneballs. So we examine these remnants:

Unless each remnant of H is a disk handlebody, the flagging of the remnants in both $\mathbb{C}_{\mathcal{D}}$ and $(\mathbb{C}_d)_{\mathcal{D}}$ is determined by Definition 7.2: Each handlebody remnant of H is empty if and only if it is disk. Hence in this case $\mathbb{C}_{\mathcal{D}} = (\mathbb{C}_d)_{\mathcal{D}}$ as flagged chamber complexes. On the other hand, if each remnant of H is a disk handlebody then, per Definition 7.2, in $\mathbb{C}_{\mathcal{D}}$ there is exactly one handlebody remnant H' (possibly a ball) that is occupied, while in $(\mathbb{C}_d)_{\mathcal{D}}$ each remnant is empty. We examine this situation more closely:

Since $\{C_{A_i}\}$ contains no handlebodies, some remnant in $\mathbb{C}_{\mathcal{D}}$ of each C_{A_i} is not a disk handlebody, by Proposition 4.8. Let R be such a remnant of some C_{A_i} . Consider the flagged chamber complex obtained from $\mathbb{C}_{\mathcal{D}}$ by deflating H' into R . Since R is not a disk handlebody it is either not a handlebody or it is an occupied handlebody; in either case moving a bubble into R will not change the flagging of R . Thus this deflation changes $\mathbb{C}_{\mathcal{D}}$ to the same flagged chamber complex as $(\mathbb{C}_d)_{\mathcal{D}}$, as required.

If H' is non-ball handlebody then the deflation is a handlebody deflation. If H' is a ball then its deflation turns it into a goneball, so $(\mathbb{C}_d)_{\mathcal{D}}$ is a bullseye deflation of $\mathbb{C}_{\mathcal{D}}$, where the bullseye is centered on the ball H' . (In fact, it is not hard to see that the bullseye consists only of H' .) $\square$

There is a similar result for bullseye deflation, but it is a bit more difficult to state and prove. Suppose $\mathbb{C}$ is a flagged chamber complex in S^3 and $\mathbb{C}_d$ is the flagged chamber complex obtained by deflating a bullseye F_b in $\mathbb{C}$. Suppose $\mathcal{D}$ is a disk set in $\mathbb{C}$ in preferred alignment and

$$\mathbb{C} \xrightarrow{\mathcal{D}} \mathbb{C}_{\mathcal{D}}$$

is the resulting flagged chamber complex decomposition. Similarly construct the flagged chamber complex decomposition

$$\mathbb{C}_d \xrightarrow{\mathcal{D}'} (\mathbb{C}_d)_{\mathcal{D}'},$$

where $\mathcal{D}' \subset \mathcal{D}$ is the subcollection of disks not incident to any sphere in F_b . (Since the spheres in F_b disappear in $\mathbb{C}_d$ it would not make sense to include disks incident to them among the decomposing disk set in $\mathbb{C}_d$.)

Proposition 11.7. *Continue with the notation just given. If $\mathbb{C}_d$ does not certify, then either $(\mathbb{C}_d)_{\mathcal{D}'} = \mathbb{C}_{\mathcal{D}}$ or $(\mathbb{C}_d)_{\mathcal{D}'}$ is obtained from $\mathbb{C}_{\mathcal{D}}$ by bullseye deflation.*

Proof. As usual, without loss of generality, assume the chamber in $\mathbb{C}$ adjacent to F_b is a B -chamber C which, after deflation and absorption of the nested set of goneballs into C , becomes a chamber we denote $C_d \in \mathbb{C}_d$. Let $\mathcal{D}_b = \mathcal{D} - \mathcal{D}'$, the subcollection of disks that are incident to spheres in F_b . The boundary ∂D of any disk $D \in \mathcal{D}_b$ divides the sphere $S \subset F_b$ in which it lies into two disks $E_{\pm}$. If D lies in the occupied ball or a shell of F_b then it is parallel to one or both of $E_{\pm}$ and so is inessential in C . If $(D, \partial D) \subset (C, \partial F_b)$ and neither sphere $D \cup E_{\pm}$ bounds a ball in C , then the two spheres (which become parallel spheres in C_d , where F_b is gone) would be essential in C_d , and so would certify for $\mathbb{C}_d$, contradicting our assumption. We conclude that each disk in $\mathcal{D}_b$ is inessential in the chamber of $\mathbb{C}$ in which it lies.

Claim: The remnants of F_b in $\mathbb{C}_{\mathcal{D}}$ consist of a bullseye F'_b that intersects C only in a collar of ∂F_b . The occupied ball component of F'_b still has b as a maximal bubble.

The proof of the claim is straightforward but a bit tedious - here is a sketch of the major steps: Since each disk of $\mathcal{D}$ that is incident to ∂F_b and lies in C is inessential, there is a collar of ∂F_b that contains them all. Then each remnant of F_b in $\hat{\mathbb{C}}_{\mathcal{D}}$ lies in a ball F_b^+ , the union of F_b and the collar, and is bounded

by spheres, so it is a punctured ball. Among the remnants of each shell in F_b is one containing an incompressible sphere, so that remnant is not a handlebody; similarly the adjacent remnant of the chamber C is not a handlebody. Hence, per Definition 7.2, each ball remnant in $\mathbb{C}_D$ of each shell and of C is a goneball. This implies that each remnant of the ball B_b in $\mathbb{C}_D$ is a ball and, per Definition 7.2, exactly one is occupied, and so must contain all of b .

Each component of F_D that lies in F_b^+ is a sphere; the ball it bounds in F_b^+ is a goneball in $\mathbb{C}_D$ unless it contains b . Hence the components of $F(\mathbb{C}_D)$ that lie in F_b^+ consist of a collection of nested spheres centered on a ball containing b ; in other words, they define a bullseye F'_b , as claimed. See Figure 33.




Figure 33: From bullseye to bullseye

Since each ball remnant of C in F'_b is a goneball, the Heegaard surface T_C for C (as well as the splittings of other chambers in $\mathbb{C}$) is unaffected by decomposition along $\mathcal{D}_b$. So the original preferred alignment of $\mathcal{D}$ in $\mathbb{C}$ remains an alignment of $\mathcal{D}'$ in $\mathbb{C}$. That alignment is easily made into a preferred alignment of $\mathcal{D}'$ in $\mathbb{C}_d$ as follows: $\mathcal{D}'$ is already in preferred alignment in chambers not C_d nor among the chambers $\{C_{A_i}\}$ (say) into which b is deflated, since these are the only chambers changed by the deflation. Whichever chamber, in $\{C_{A_i}\}$ or C_d , that b is deflated into cannot be a handlebody, since $\mathbb{C}_d$ does not certify, so there is some remnant of that chamber that is not a disk handlebody; choose the alignment of $\mathcal{D}'$ in that chamber so that b appears in a remnant that is not a disk handlebody. This does not change the flagging and thereby ensures that $\mathcal{D}'$ is in preferred alignment with $\mathbb{C}_d$.

Now consider the effect of $\mathcal{D}$ surgery away from the bullseye F_b , that is surgery via $\mathcal{D}'$ on the surface $F - F_b$. The resulting surface $\hat{F}$ can be viewed in two ways: Since $F - F_b$ is also the defining surface for $\mathbb{C}_d$, $\hat{F}$ can be described as the first (surgery) step in the construction of the decomposition $\mathbb{C}_d \xrightarrow{\mathcal{D}'} (\mathbb{C}_d)_{\mathcal{D}'}$. And $\hat{F}$ is also that part of the surface F_D that lies outside the new bullseye F'_b .

It remains only to understand the effect of the second step in disk decomposition - eliminating goneballs after surgery. We have already determined that eliminating goneballs from surgery on F_b results in the bullseye $F'_b \subset \mathbb{C}_D$. So we need only consider sphere components of $\hat{F}$. Any such sphere bounds a ball (in fact two balls) in S^3 . If such a ball is a goneball in $\mathbb{C}_D$ then that ball is unaffected by a deflation of F'_b (since F'_b is necessarily outside the goneball) so it remains a goneball in $(\mathbb{C}_d)_{\mathcal{D}'}$.

Consider the symmetric question: if a sphere component of $\hat{F}$ bounds a goneball in $(\mathbb{C}_d)_{\mathcal{D}'}$, does it bound a goneball in $\mathbb{C}_D$? The goneball in $(\mathbb{C}_d)_{\mathcal{D}'}$ may or may not contain the deflated F'_b . If the goneball does not contain F'_b then the goneball lies also in $\mathbb{C}_D$ and so the sphere does not appear in either $\mathbb{C}_D$ or $(\mathbb{C}_d)_{\mathcal{D}'}$. Suppose, on the other hand, that there is a sphere component of $\hat{F}$ that bounds a goneball in

$(\mathbb{C}_d)_{\mathcal{D}'}$ and that goneball does contain the deflated F'_b . Let S be an innermost one. Since the ball it bounds contains F'_b it is not a goneball in $\mathbb{C}_{\mathcal{D}}$, so S persists as a sphere in the defining surface of $\mathbb{C}_{\mathcal{D}}$. On the other hand, observe this: since a goneball must have genus 0 splitting, the region between S and $\partial F'_b$ in $\mathbb{C}_{\mathcal{D}}$, a collar of $\partial F'_b$, must have genus 0 splitting. Hence the region is a genus 0 shell. We can then regard S as yet another nested sphere component in F'_b and continue the argument. In the end, the only components of $\hat{F}$ that appear in $\mathbb{C}_{\mathcal{D}}$ but not in $(\mathbb{C}_d)_{\mathcal{D}'}$ are spheres that just add shells to F'_b . In particular, they are all eliminated by a bullseye deflation, in this case of a bullseye that properly contains F'_b . So $(\mathbb{C}_d)_{\mathcal{D}'}$ is obtained from $\mathbb{C}_{\mathcal{D}}$ by bullseye deflation, as required. $\square$

Digressing momentarily, Proposition 11.7 gives an easy and early example of a *guiding set of disks*. These will be further discussed in Section 13.

Definition 11.8. Suppose $\mathbb{C}$ is a flagged chamber complex in S^3 and $\overline{\mathcal{D}}$ is a finite set of disjoint disks in S^3 transverse to $F = F(\mathbb{C})$ so that $F \cap \overline{\mathcal{D}} = \partial \mathcal{D}$, for some $\mathcal{D} \subset \overline{\mathcal{D}}$. Then $\overline{\mathcal{D}}$ is a *guiding set of disks* for the flagged chamber complex decomposition $\mathbb{C} \xrightarrow{\mathcal{D}} \mathbb{C}_{\mathcal{D}}$. Slightly abusing notation we can then write $\mathbb{C} \xrightarrow{\overline{\mathcal{D}}} \mathbb{C}_{\mathcal{D}}$.

For example, in Proposition 11.7, $\mathcal{D}$ is a disk set for the decomposition $\mathbb{C} \xrightarrow{\mathcal{D}} \mathbb{C}_{\mathcal{D}}$ and $\mathcal{D}$ is also a *guiding disk set* for the decomposition $\mathbb{C}_d \xrightarrow{\mathcal{D}'} (\mathbb{C}_d)_{\mathcal{D}}$. So, with no loss of meaning, we could also write the latter $\mathbb{C}_d \xrightarrow{\mathcal{D}} (\mathbb{C}_d)_{\mathcal{D}}$.

In general we will write $\mathbb{C}' \twoheadrightarrow \mathbb{C}$ to denote that $\mathbb{C}'$ is obtained from $\mathbb{C}$ by either (non-ball) handlebody deflation or bullseye deflation. (The direction of the arrow is meant to indicate that the defining surface $F(\mathbb{C}')$ embeds in $F(\mathbb{C})$.) Thus, in both Propositions 11.6 and 11.7 above, $\mathbb{C}_d \twoheadrightarrow \mathbb{C}$. Applying Definition 11.8 these two propositions, together with Propositions 11.2 and 11.5, can then be summarized as follows:

Corollary 11.9. Suppose $\mathbb{C}, \mathbb{C}'$ are flagged chamber complexes for which $\mathbb{C}' \twoheadrightarrow \mathbb{C}$ and suppose $\mathcal{D}$ is a disk set in $\mathbb{C}$.

If $\mathbb{C}'$ certifies, then $\mathbb{C}'$ and $\mathbb{C}$ cocertify.

If $\mathbb{C}'$ does not certify, then the following square of flagged chamber complexes commutes:

$$\begin{array}{ccc} \mathbb{C}' & \xrightarrow{\mathcal{D}} & \mathbb{C}'_{\mathcal{D}} \\ \downarrow & & \downarrow \\ \mathbb{C} & \xrightarrow{\mathcal{D}} & \mathbb{C}_{\mathcal{D}} \end{array}$$

Broaden now to sequences of decompositions:

Definition 11.10. . Suppose

$$\vec{\mathbb{C}}: \quad \mathbb{C}_0 \xrightarrow{\mathcal{D}_0} \mathbb{C}_1 \xrightarrow{\mathcal{D}_1} \dots \xrightarrow{\mathcal{D}_{n-1}} \mathbb{C}_n$$

and

$$\vec{\mathbb{C}}': \quad \mathbb{C}'_0 \xrightarrow{\mathcal{D}_0} \mathbb{C}'_1 \xrightarrow{\mathcal{D}_1} \dots \xrightarrow{\mathcal{D}_{n-1}} \mathbb{C}'_n$$

are flagged chamber complex decomposition sequences so that for each $i \geq 0$, $\mathbb{C}'_i \twoheadrightarrow \mathbb{C}_i$. Then the decomposition sequence $\vec{\mathbb{C}}'$ is a deflation of the decomposition sequence $\vec{\mathbb{C}}$.

Suppose

$$\vec{\mathbb{C}}: \quad \mathbb{C}_0 \xrightarrow{\mathcal{D}_0} \mathbb{C}_1 \xrightarrow{\mathcal{D}_1} \dots \xrightarrow{\mathcal{D}_{n-1}} \mathbb{C}_n$$

is a flagged chamber complex decomposition sequence and $\mathbb{C}'_0$ is a flagged chamber complex such that $\mathbb{C}'_0 \dashrightarrow \mathbb{C}_0$. Iteratively define a flagged chamber complex decomposition sequence

$$\vec{\mathbb{C}}^{m(ax)}: \quad \mathbb{C}'_0 \xrightarrow{\mathcal{D}_0} \mathbb{C}'_1 \xrightarrow{\mathcal{D}_1} \dots \xrightarrow{\mathcal{D}_{m-1}} \mathbb{C}'_m$$

by the following process: If $\mathbb{C}'_i \dashrightarrow \mathbb{C}_i$ take $\mathbb{C}'_{i+1}$ to be the result of flagged chamber complex decomposition of $\mathbb{C}'_i$ by (the guiding set of disks) $\mathcal{D}_i$. Continue until either $m = n$ or $\mathbb{C}'_{m+1}$ is not a deflation of $\mathbb{C}_{m+1}$. The sequence $\vec{\mathbb{C}}^m$ is a deflation of the sequence

$$\mathbb{C}_0 \xrightarrow{\mathcal{D}_0} \mathbb{C}_1 \xrightarrow{\mathcal{D}_1} \dots \xrightarrow{\mathcal{D}_{m-1}} \mathbb{C}_m$$

so the following definition is natural:

Definition 11.11. $\vec{\mathbb{C}}^m$ is the maximal deflationary subsequence of $\vec{\mathbb{C}}'$.

Corollary 11.12. Suppose

$$\vec{\mathbb{C}}: \quad \mathbb{C}_0 \xrightarrow{\mathcal{D}_0} \mathbb{C}_1 \xrightarrow{\mathcal{D}_1} \dots \xrightarrow{\mathcal{D}_{n-1}} \mathbb{C}_n$$

is a flagged chamber complex decomposition sequence supporting $S^3 = A \cup_T B$ and $\mathbb{C}'_0$ is a flagged chamber complex such that $\mathbb{C}'_0 \dashrightarrow \mathbb{C}_0$. Let $\vec{\mathbb{C}}^m$ be the maximal deflationary subsequence of $\vec{\mathbb{C}}'$.

- If $\vec{\mathbb{C}}^m$ does not certify then $m = n$.
- If $\vec{\mathbb{C}}^m$ does certify then $\vec{\mathbb{C}}'$ and $\vec{\mathbb{C}}$ cocertify.

Proof. Iteratively apply Corollary 11.9. □

12 Reordering disks

12.1 Parallel vs sequential flagged chamber complex decompositions

Suppose $\mathbb{C}$ is a flagged chamber complex supporting $S^3 = A \cup_T B$, with $F = F(\mathbb{C})$ its defining surface. Suppose $\mathcal{D}_x, \mathcal{D}_y$ are disjoint disk sets in $\mathbb{C}$. In this section we compare the results of flagged chamber complex decomposition using these disk sets in series vs in parallel. That is, we compare the two flagged chamber complexes that appear in the bottom row of this diagram:

$$\begin{array}{ccc} \mathbb{C} & \xrightarrow{\mathcal{D}_x} & \mathbb{C}_x \\ \mathcal{D}_x \cup \mathcal{D}_y \downarrow & & \downarrow \mathcal{D}_y \\ \mathbb{C}_{x+y} & & \mathbb{C}_{xy} \end{array}$$

In general the answer is complex: for starters, notice that some disks in $\mathcal{D}_y$ may be incident to the boundary of goneballs in $\mathbb{C}_x$, so the disks $\mathcal{D}_y$ become not a disk set for the decomposition $\mathbb{C}_x \xrightarrow{\mathcal{D}_y} \mathbb{C}_{xy}$ but a guiding disk set. In light of such difficulties, eventually we will restrict our interest to the case in which $\mathcal{D}_x$ or $\mathcal{D}_y$ consists of a single disk. See Subsection 12.2.

We first identify circumstances in which $\mathbb{C}_{x+y} = \mathbb{C}_{xy}$. To that end:

Definition 12.1. Suppose H is a handlebody in S^3 such that $\partial H \subset F_{\mathcal{D}_x \cup \mathcal{D}_y}$, the surface obtained from F by surgery on $\mathcal{D}_x \cup \mathcal{D}_y$. We say that H is coherent for the decompositions by $\mathcal{D}_x \cup \mathcal{D}_y$ in the diagram above when

- H is a chamber in $\mathbb{C}_{x+y}$ if and only if it is a chamber in $\mathbb{C}_{xy}$ and
- if H is a chamber in both $\mathbb{C}_{x+y}$ and $\mathbb{C}_{xy}$ then it has the same flagging.

Here is an easy example:

Lemma 12.2. Suppose there is a component of $F_{\mathcal{D}_x \cup \mathcal{D}_y}$ that lies in $\text{int}(H)$ and is not a sphere. Then H is not a chamber in either $\mathbb{C}_{x+y}$ or $\mathbb{C}_{xy}$. Hence H is coherent.

Proof. Let F_{x+y} be a component of $F_{\mathcal{D}_x \cup \mathcal{D}_y}$ that lies in $\text{int}(H)$ and is not a sphere. Since F_{x+y} is not a sphere it is not contained in a goneball of $\mathbb{C}_{x+y}$ (cf Lemma 4.2(2)) so it is a component of the defining surface $F(\mathbb{C}_{x+y})$ that lies inside H . Hence H is not a handlebody chamber in $\mathbb{C}_{x+y}$.

The complement of the scars in F_{x+y} is a non-planar connected surface contained in a closed component $F_0 \subset F$. Similarly the complement of just the scars of $\mathcal{D}_y$ in F_{x+y} is contained in a non-planar closed surface F_x of $F_{\mathcal{D}_x}$. Since both F_0 and F_x contain a non-planar subsurface they are not spheres. Since F_x is not a sphere it is not contained in a goneball of $\mathbb{C}_x$, again by Lemma 4.2(2). In particular it is a component of the defining surface $F(\mathbb{C}_x)$. F_{x+y} is a component of the surface that results from surgery by $\mathcal{D}_y$ on F_x . Since it is not a sphere it also is not contained in a goneball of $\mathbb{C}_{xy}$. Hence F_{x+y} is a closed component of the defining surface for $\mathbb{C}_{xy}$ lying in the interior of H . Thus H is not a handlebody chamber in $\mathbb{C}_{xy}$. $\square$

Returning to the definition of coherent, note that H could fail to be a chamber in $\mathbb{C}_{x+y}$ for two reasons: H contains a component of the defining surface $F(\mathbb{C}_{x+y})$ in its interior, as in Lemma 12.2, or H could be a goneball. Similarly, H could fail to be a chamber in $\mathbb{C}_{xy}$ because it has a component of the defining surface $F(\mathbb{C}_{xy})$ in its interior, or because H is a goneball under the sequence of decompositions

$$\mathbb{C} \xrightarrow{\mathcal{D}_x} \mathbb{C}_x \xrightarrow{\mathcal{D}_y} \mathbb{C}_{xy}.$$

Suppose the latter. H could become a goneball under the second decomposition, so the complement ∂H_- of the scars of $\mathcal{D}_y$ in ∂H is part of a surface in $F(\mathbb{C}_x)$. On the other hand, ∂H_- may not be part of a surface in $F(\mathbb{C}_x)$ if ∂H_- is part of a sphere $G \subset F = F(\mathbb{C})$ that bounds a goneball in $\mathbb{C}_x$. In the last case, if H also has no other components of $\mathbb{C}_{xy}$ in its interior, it is natural to call H a goneball of the decomposition sequence $\mathbb{C} \xrightarrow{\mathcal{D}_x} \mathbb{C}_x \xrightarrow{\mathcal{D}_y} \mathbb{C}_{xy}$. As we will see, such a ball H need not be a goneball in $\mathbb{C}_{x+y}$.

The following lemma explains why coherence is useful:

Lemma 12.3. *Suppose every handlebody in S^3 whose boundary is in $F_{\mathcal{D}_x \cup \mathcal{D}_y}$ is coherent. Then $\mathbb{C}_{x+y} = \mathbb{C}_{xy}$ as flagged chamber complexes.*

Proof. Let F_0 be a component of $F_{\mathcal{D}_x \cup \mathcal{D}_y}$. If F_0 is not a sphere, then, as in the proof of Lemma 12.2, F_0 is a component of both defining surfaces $F(\mathbb{C}_{x+y})$ and $F(\mathbb{C}_{xy})$. If F_0 is a sphere, then it divides S^3 into two 3-balls, $H_{\pm}$. By assumption both $H_{\pm}$ are coherent, so each is a goneball in $\mathbb{C}_{x+y}$ if and only if it is a goneball in $\mathbb{C}_{xy}$. In particular, if either $H_{\pm}$ is a goneball in one of $\mathbb{C}_{x+y}$ or $\mathbb{C}_{xy}$ then it is a goneball in both and F_0 is not in either $F(\mathbb{C}_{x+y})$ and $F(\mathbb{C}_{xy})$. On the other hand, neither of $H_{\pm}$ is a goneball in $\mathbb{C}_{x+y}$ if and only if neither is a goneball in $\mathbb{C}_{xy}$ and, in this case, F_0 does lie in both $F(\mathbb{C}_{x+y})$ and $F(\mathbb{C}_{xy})$. Thus, ignoring flagging, $\mathbb{C}_{x+y} = \mathbb{C}_{xy}$ as chamber complexes.

Now consider flagging of a handlebody chamber H in one of $\mathbb{C}_{x+y}$ or $\mathbb{C}_{xy}$ and so in both. Suppose H is a handlebody chamber in one of $\mathbb{C}_{x+y}$ or $\mathbb{C}_{xy}$. Since H is coherent, it is then a handlebody chamber in both $\mathbb{C}_{x+y}$ and $\mathbb{C}_{xy}$, and the flagging of H is the same. Thus $\mathbb{C}_{x+y} = \mathbb{C}_{xy}$ as flagged chamber complexes. $\square$

Here is a helpful example of incoherence in a handlebody: Let C be a chamber in $\mathbb{C}$ that a disk D , properly embedded in C , divides into two components: U , not a handlebody, and V , the complement in a handlebody H of a regular neighborhood η of a properly embedded arc $\alpha \subset H$. Let E be a meridian disk in η , so E is properly embedded in η and meets α in a single point. Note that $F \cap \text{int}(H)$ is the annulus $\partial\eta$. See Figure 34.

Since $F \cap \text{int}(H)$ is an annulus and not a collection of disks, H is occupied after the decomposition $\mathbb{C} \xrightarrow{D \cup E} \mathbb{C}_{D+E}$ and after the sequence of decompositions


$$\mathbb{C} \xrightarrow{D} \mathbb{C}_D \xrightarrow{E} \mathbb{C}_{DE}.$$

This means that H is a coherent handlebody for $\mathbb{C}_{D+E}$ and $\mathbb{C}_{DE}$. On the other hand, after the decomposition $\mathbb{C} \xrightarrow{E} \mathbb{C}_E$ the disk D divides $\mathbb{C}_E$ into U and the handlebody H . For the next decomposition $\mathbb{C}_E \xrightarrow{D} \mathbb{C}_{ED}$, Definition 7.2 requires that E be aligned so that H is empty in $\mathbb{C}_{ED}$. Thus H is not a coherent handlebody for $\mathbb{C}_{D+E}$ and $\mathbb{C}_{ED}$.

Replacing D in this example of incoherence by a disk family consisting of D and a disk D' parallel to E has no effect on the above argument: For $\mathcal{D} = D \cup D'$, H is coherent for $\mathbb{C}_{\mathcal{D}+E}$ and $\mathbb{C}_{DE}$ but incoherent for $\mathbb{C}_{\mathcal{D}+E}$ and $\mathbb{C}_{ED}$. This leads to the following observation, using analogous notation:

Lemma 12.4. *Let $\mathcal{D} \cup E$ be a disk family in a flagged chamber complex $\mathbb{C}$ in S^3 . Suppose the disk E is parallel to a disk $D \in \mathcal{D}$ within a chamber C of $\mathbb{C}$. Then*

1. *Unless the chamber C_D of $\mathbb{C}_{\mathcal{D}}$ containing E is an occupied handlebody, $\mathbb{C}_{\mathcal{D}} = \mathbb{C}_{DE}$. If C_D is an occupied handlebody then for one of the two resulting siblings in $\mathbb{C}_{DE}$, $\mathbb{C}_{\mathcal{D}} = \mathbb{C}_{DE}$.*
2. *Unless the chamber C is an occupied handlebody with only disk handlebody remnants in $\mathbb{C}_{\mathcal{D}}$, the chamber complex $\mathbb{C}_{\mathcal{D}} = \mathbb{C}_{\mathcal{D}+E}$. The same is almost true in the remaining case: If C is an occupied handlebody with only disk handlebody remnants in $\mathbb{C}_{\mathcal{D}}$ then C has sibling remnants in $\mathbb{C}_{\mathcal{D}+E}$ and deleting a single such sibling gives the remnants of C in $\mathbb{C}_{\mathcal{D}}$.*
3. *either $\mathbb{C}_{\mathcal{D}+E} = \mathbb{C}_{DE}$ or C_D is an occupied handlebody chamber in $\mathbb{C}_{\mathcal{D}}$ and, in one or both of two resulting siblings in $\mathbb{C}_{DE}$, $\mathbb{C}_{\mathcal{D}+E} = \mathbb{C}_{DE}$.*


 Figure 34: $\mathbb{C}_{DE} = \mathbb{C}_{D+E} \neq \mathbb{C}_{ED}$

Proof. Since D and E are parallel, after surgery on $\mathcal{D} \cup E$ there is a sphere G in $F_{\mathcal{D} \cup E}$ bounding a disk ball B_G with one scar of each of the disks D and E on G . Moreover the other remnant $C_- = C_D - B_G$ of C_D in $\mathbb{C}_{DE}$ is homeomorphic to C_D , so it is a handlebody if and only if C_D is. If C_D is not a handlebody, so also C_- is not a handlebody, or if C_D is an empty handlebody, then per Definition 7.2 the ball B_G is empty and so a goneball in $\mathbb{C}_{DE}$. Then $\mathbb{C}_D = \mathbb{C}_{DE}$ (with C_- replacing C_D). If C_D is occupied then the same statement is true for one of the two siblings of C_D . This proves 1).

The remnants of C in $\mathbb{C}_{D+E}$ are the same as in $\mathbb{C}_D$, with the addition of the disk ball B_G . If B_G is a goneball, we are done as in the proof of 1). The only way in which B_G could not be a goneball, per Definition 7.2, is if C is an occupied handlebody with only disk remnants. In the latter case, only the sibling in which B_G is occupied does not also occur among the sibling remnants of C in $\mathbb{C}_D$. This proves 2).

The proof of 3) follows almost immediately. Following 1), we may as well assume that $\mathbb{C}_D = \mathbb{C}_{DE}$. The result follows from 2), unless C is an occupied handlebody with only disk handlebody remnants

in $\mathbb{C}_{\mathcal{D}}$. In this case, C_- is also disk, so C also has only disk handlebody remnants in $\mathbb{C}_{\mathcal{D}+E}$. Then per Definition 7.2 the remnants of C in both $\mathbb{C}_{\mathcal{D}+E}$ and $\mathbb{C}_{\mathcal{D}E}$ have the same description of siblings, one corresponding to exactly one of the remnants of $\mathbb{C}_{\mathcal{D}+E}$ being occupied, with remnant B_G a goneball except when it is the occupied remnant. $\square$

Note: Switching the order of $\mathcal{D}$ and E makes a difference. The example preceding Lemma 12.4 shows that there may be a handlebody in S^3 whose boundary is in $F_{\mathcal{D} \cup E}$ and which is not even coherent for $\mathbb{C}_{\mathcal{D}+E}$ and $\mathbb{C}_{E\mathcal{D}}$.

Return now to the discussion of decomposition by $\mathcal{D}_x$ and $\mathcal{D}_y$.

Definition 12.5. Suppose H is a handlebody in S^3 so that $\partial H \subset F_{\mathcal{D}_x \cup \mathcal{D}_y}$. Then a disk $D \in \mathcal{D}_x \cup \mathcal{D}_y$ is incident to H if $D \subset \text{int}(H)$ or D leaves a scar on ∂H . The handlebody H is an x -handlebody (resp. y -handlebody) if H is incident only to disks in $\mathcal{D}_x$ (resp. $\mathcal{D}_y$). An x -handlebody (resp. y -handlebody) H is an x -handlebody (resp. y -handlebody) chamber in $\mathbb{C}_{xy}$ if it is a chamber there. That is, ∂H is in the defining surface $F(\mathbb{C}_{xy})$ and no other chamber lies in $\text{int}(H)$.

Lemma 12.6. If H is an x -handlebody chamber in $\mathbb{C}_{xy}$ then H is also a chamber of $\mathbb{C}_x$.

Proof. Let C be the chamber of $\mathbb{C}_x$ which has H as a remnant. If any disks in $\mathcal{D}_y$ are incident to C then each remnant of C in $\hat{\mathbb{C}}_{xy}$ has a scar from $\mathcal{D}_y$, so the same is true after getting rid of goneballs. Since H has no scars from $\mathcal{D}_y$, no disks from $\mathcal{D}_y$ are incident to C , and C is unaffected by the decomposition by $\mathcal{D}_y$. That is, $C = H$ as required. $\square$

Suppose $\mathbb{C}$ contains no occupied handlebody chambers. Then by Corollary 7.3 a handlebody chamber in $\mathbb{C}_x$ or a handlebody chamber in $\mathbb{C}_{x+y}$ is empty if and only if it is disk.

Lemma 12.7. Suppose $\mathbb{C}$ contains no occupied handlebody chambers. Then:

- Each x -handlebody chamber of $\mathbb{C}_{xy}$ is empty if and only if it is a disk chamber under the decomposition $\mathbb{C} \xrightarrow{\mathcal{D}_x} \mathbb{C}_x$.
- If $\mathbb{C}_x$ contains no occupied handlebody chambers then each y -handlebody chamber of $\mathbb{C}_{xy}$ is empty if and only if it is disk under the decomposition $\mathbb{C}_x \xrightarrow{\mathcal{D}_y} \mathbb{C}_{xy}$.
- If $\mathbb{C}_x$ does contain occupied handlebody chambers, there is a preferred alignment of $\mathcal{D}_y$ in $\mathbb{C}_x$ so that each y -handlebody chamber of $\mathbb{C}_{xy}$ is empty if and only if it is disk under the decomposition $\mathbb{C}_x \xrightarrow{\mathcal{D}_y} \mathbb{C}_{xy}$. Any other preferred alignment of $\mathcal{D}_y$ gives a sibling decomposition $\mathbb{C}_x \xrightarrow{\mathcal{D}_y} \mathbb{C}_{xy}$.

Proof. The first statement follows from Corollary 7.3 and Lemma 12.6. The second statement follows from Corollary 7.3.

The third statement is more difficult to prove. Observe first that by the argument in Corollary 7.3, any handlebody remnant of a chamber in $\mathbb{C}_x$ that is not an occupied handlebody is empty if and only if it is disk. Moreover the argument extends, via Definition 7.2, even to occupied handlebody chambers in $\mathbb{C}_x$ so long as at least one remnant in $\mathbb{C}_{xy}$ is not a disk handlebody. So we need only describe how to align

those disks in $\mathcal{D}_y$ that are incident to each chamber C_x in $\mathbb{C}_x$ of the following type: C_x is an occupied handlebody and every remnant of C_x in $\mathbb{C}_{xy}$ is a disk handlebody.

This is how it is done: Since C_x is occupied, by Corollary 7.3 C_x is not disk. So there is some component F_x of $F \cap \text{int}(C_x)$ that is not a disk. F_x must be incident to some disk in $\mathcal{D}_x$ else it would be a subsurface of $F(\mathbb{C}_x)$ and not lie in the interior of C_x .

The subcollection of disks in $\mathcal{D}_y$ that are incident to ∂C_x is disjoint from F_x , since $\text{int}(\mathcal{D}_y)$ is disjoint from F . Hence the surface F_x will lie in the interior of some remnant C_{xy} of C_x in $\mathbb{C}_{xy}$, so that remnant too is incident to $\mathcal{D}_x$. In particular, C_{xy} is not a y -handlebody. Following Definition 7.2, choose a preferred alignment where every (disk) handlebody remnant of C_x other than C_{xy} is empty. Then in particular each y -handlebody remnant is both disk and empty, as required. $\square$

Lemma 12.8. *Suppose $\mathbb{C}$ does not certify and H is a handlebody in S^3 so that $\partial H \subset F_{\mathcal{D}_x \cup \mathcal{D}_y}$. If H is an x -handlebody it is coherent. If H is a y -handlebody then it is coherent in one of the sibling decompositions of $\mathbb{C}_x \xrightarrow{\mathcal{D}_y} \mathbb{C}_{xy}$.*

Proof. If appropriate (i. e. if $\mathbb{C}_x$ contains occupied handlebody chambers) choose the alignment for $\mathcal{D}_y$ in $\mathbb{C}_x$ given in Lemma 12.7 to ensure that each y -handlebody chamber of $\mathbb{C}_{xy}$ is empty if and only if it is disk under the decomposition $\mathbb{C}_x \xrightarrow{\mathcal{D}_y} \mathbb{C}_{xy}$. Also from Lemma 12.7 we have that each x -handlebody chamber of $\mathbb{C}_{xy}$ is empty if and only if it is disk under the decomposition $\mathbb{C} \xrightarrow{\mathcal{D}_x} \mathbb{C}_x$.

Case 1: x -handlebodies

The proof will be by contradiction. Suppose there is an x -handlebody that is not coherent and let H be an innermost one that is not coherent. Suppose first that there are no components of $F_{\mathcal{D}_x \cup \mathcal{D}_y}$ contained in the interior of H . Then $\mathbb{C} \xrightarrow{\mathcal{D}_x} \mathbb{C}_x$ is disk (for $F \cap \text{int}(H)$ has no components at all) and hence empty (possibly a goneball) in $\mathbb{C}_{xy}$. Similarly $\mathbb{C} \xrightarrow{\mathcal{D}_x \cup \mathcal{D}_y} \mathbb{C}_{x+y}$ is disk and hence H is empty in $\mathbb{C}_{x+y}$. Since H is empty in both $\mathbb{C}_{xy}$ and $\mathbb{C}_{x+y}$ it is coherent, a contradiction.

We deduce then that there must be components of $F_{\mathcal{D}_x \cup \mathcal{D}_y}$ in the interior of H . Suppose a subsurface of $F_{\mathcal{D}_x \cup \mathcal{D}_y}$ in the interior of H becomes in $\mathbb{C}_x$ a component F_x of $F(\mathbb{C}_x)$. Then H is not a chamber in $\mathbb{C}_x$ (since it contains a component of $F(\mathbb{C}_x)$ in its interior) and, since H is incident to no disks in $\mathcal{D}_y$, H remains not a chamber in $\mathbb{C}_{xy}$. Since no disk in $\mathcal{D}_y$ is incident to F_x , F_x is also a component of $\mathbb{C}_{x+y}$ unless it is a goneball there. But if it were a goneball there, and not in $\mathbb{C}_{xy}$ it would be a further in handlebody that is not coherent, contradicting our choice of H . We conclude that F_x remains as a surface in $\mathbb{C}_{x+y}$ so H is not a chamber there either. But this implies H is coherent, a contradiction.

We deduce then that $F_{\mathcal{D}_x \cup \mathcal{D}_y} \cap \text{int}(H)$ must become a collection of spheres bounding goneballs in $\mathbb{C}_x$, hence in $\mathbb{C}_{xy}$. Then, by our consistency hypothesis, the spheres bound goneballs also in $\mathbb{C}_{x+y}$. Hence in each of $\mathbb{C}_{xy}$ and $\mathbb{C}_{x+y}$, H is either a chamber or is itself a goneball. Whether H is occupied or empty (so perhaps a goneball) is determined in $\mathbb{C}_{x+y}$ by whether the components of $F \cap \text{int}(H)$ consist entirely of disks, by Corollary 7.3. By the first statement of Lemma 12.7 the same is true in $\mathbb{C}_{xy}$. Hence H is coherent a final contradiction.

Case 2: y -handlebodies

The proof will be by contradiction, similar to Case 1. Suppose there is a y -handlebody that is not coherent and let H be an innermost one that is not coherent. Suppose there is a component F_0 of $F_{\mathcal{D}_x \cup \mathcal{D}_y}$

that lies in $\text{int}(H)$ and remains in $F(\mathbb{C}_{xy})$, so H is not a chamber in $\mathbb{C}_{xy}$. By Lemma 12.2 F_0 is a sphere, so it bounds a ball in H . The ball it bounds is not a goneball in $\mathbb{C}_{xy}$ so, by the consistency hypothesis, F_0 does not bound a goneball in $\mathbb{C}_{x+y}$. This implies that H is not a chamber in either $\mathbb{C}_{xy}$ or $\mathbb{C}_{x+y}$, contradicting our assumption that H is not coherent.

We deduce then that $\text{int}(H)$ is disjoint from $F(\mathbb{C}_{xy})$ so H is a chamber in $\mathbb{C}_{xy}$. A symmetric argument shows that H is also a chamber in $\mathbb{C}_{x+y}$. Suppose $F \cap \text{int}(H)$ has a component F_0 that is not a disk, so H is an occupied handlebody in $\mathbb{C}_{x+y}$.

Subcase 2a: $F_0 \subset \text{int}(H)$ is a closed surface.

Since H is a y -handlebody, F_0 is unaffected by surgery on $\mathcal{D}_x$ so it persists as a component of $F_{\mathcal{D}_x}$. Since F_0 does not bound a goneball in $\mathbb{C}$, it does not bound one in $\mathbb{C}_s$, so F_0 persists as a component of $F(\mathbb{C}_x)$. In particular, H is not diskly under the decomposition $\mathbb{C}_x \xrightarrow{\mathcal{D}_y} \mathbb{C}_{xy}$. This implies that H is an occupied handlebody in $\mathbb{C}_{xy}$, contradicting the hypothesis that H is not coherent.

Subcase 2b: $\partial F_0 \neq \emptyset$.

Let $\partial_- H = \partial H \cap F$, that is ∂H with all scars removed. Note that since F_0 is not a disk it is either non-planar or has more than one boundary component; in either case $F_0 \cup \partial_- H$ is non-planar. In particular, the subsurface of $F_{\mathcal{D}_x}$ that contains it cannot be a sphere, so $F_0 \cup \partial_- H$ persists as a subsurface of $F(\mathbb{C}_x)$. In particular, F_0 persists, and, just as in Subcase 2a this implies that H is an occupied handlebody in $\mathbb{C}_{xy}$, contradicting the hypothesis that H is not coherent.

We are thus reduced to the case that $F \cap \text{int}(H)$ consists entirely of disks, so H is diskly in $\mathbb{C}_{x+y}$. Since we are assuming $\mathbb{C}$ does not certify, and therefore has no occupied handlebody chambers, this implies that H is empty (and so a goneball if H is a ball) in $\mathbb{C}_{x+y}$. It's possible that in $\mathbb{C}_x$ the surface $\partial_- H$ is part of a sphere $G \subset F$ that bounds a goneball in $\mathbb{C}_x$. This would imply that ∂H is planar, so H is a ball, and would make H a goneball of the decomposition sequence $\mathbb{C} \xrightarrow{\mathcal{D}_x} \mathbb{C}_x \xrightarrow{\mathcal{D}_y} \mathbb{C}_{xy}$, as described before Lemma 12.3.

On the other hand, if $F \cap \partial H$ does remain in $F(\mathbb{C}_{\mathcal{D}_x})$, what we can conclude is that H is diskly in the decomposition $\mathbb{C}_x \xrightarrow{\mathcal{D}_y} \mathbb{C}_{xy}$ and so, per our choice of preferred alignment of $\mathcal{D}_y$, H is empty. (In particular, it is again a goneball if H is a ball.) In either case, H is then coherent for $\mathbb{C}_{xy}$ and $\mathbb{C}_{x+y}$, a final contradiction. $\square$

12.2 When a disk set is a singleton

We take the discussion above one step further:

Definition 12.9. Suppose H is a handlebody in S^3 so that $\partial H \subset F_{\mathcal{D}_x \cup \mathcal{D}_y}$. H is an xe -handlebody (resp. ey -handlebody) if H is incident only to disks in $\mathcal{D}_x$ (resp. $\mathcal{D}_y$) except for a single disk E (exceptional) which is also incident to H but lies in $\mathcal{D}_y$ (resp. $\mathcal{D}_x$).

We proceed to two lemmas that are parallel in spirit and whose proofs will use the same figures. To keep their use disentangled, in Figures 36, 37 and 39 the subscripts $xe, x + e$ will refer to the case in which $\mathcal{D}_y = E$ and the subscripts $ey, e + y$ will refer to the case in which $\mathcal{D}_x = E$.

Lemma 12.10. Suppose $\mathbb{C}$ does not certify, H is an ey -handlebody and, if $\mathbb{C}_x$ contains an occupied handlebody, $\mathcal{D}_y$ is given the preferred alignment in $\mathbb{C}_x$ of Lemma 12.7. Then either H is coherent or

- E lies in $\text{int}(H)$ and
- either ∂E is non-separating in $F - \partial \mathcal{D}_y$ or ∂E is non-separating in F and E is parallel to a disk in $\mathcal{D}_y$ and
- H lies in a chamber¹ of $\mathbb{C}_{xy}$ and H is either
 - empty (and so a goneball if H is a ball), the deflation of a handlebody chamber in $\mathbb{C}_{x+y}$ or
 - the deflation of a single bullseye of $\mathbb{C}_{x+y}$ in H . (Possibly H is part of the bullseye.)

Proof. As usual, let $F_{E \cup \mathcal{D}_y}$ be the surface obtained from $F = F(\mathbb{C})$ by surgery on $E \cup \mathcal{D}_y$. Following Lemma 12.2 we will assume that each component of $F_{E \cup \mathcal{D}_y}$ that lies in $\text{int}(H)$ is a sphere and so bounds a ball in H . Let $\mathfrak{S}$ be this set of spheres

Special case: Each ball in H bounded by a sphere in $\mathfrak{S}$ is coherent.

Claim: In this special case, H lies entirely in a chamber of $\mathbb{C}_{x+y}$ if and only if it lies entirely in a chamber of $\mathbb{C}_{xy}$.

Proof of claim: For both $\mathbb{C}_{x+y}$ and $\mathbb{C}_{xy}$, H is contained in a chamber if and only if each ball in H bounded by a sphere in $\mathfrak{S}$ is a goneball, for this determines whether there is any chamber contained in $\text{int}(H)$. But in this special case, each ball bounded by a sphere in $\mathfrak{S}$ is coherent, so the ball it bounds is a goneball in $\mathbb{C}_{x+y}$ if and only if it is a goneball in $\mathbb{C}_{xy}$. This proves the claim.

Following the claim, but still in the special case, we will assume that H lies entirely in a chamber (possibly it is itself an entire chamber) of both $\mathbb{C}_{x+y}$ and $\mathbb{C}_{xy}$ and need to determine under what circumstances H is not coherent. That is,

- If H is a ball, when is it a goneball in one of $\mathbb{C}_{x+y}$ or $\mathbb{C}_{xy}$ but not the other?
- If H is not a ball, so it is itself a chamber, when is it empty in one of $\mathbb{C}_{x+y}$ or $\mathbb{C}_{xy}$ but not the other?

We examine possible subcases. For all but the last, H will turn out to be coherent:

Subcase 1a: There is a component F_0 of $F \cap \text{int}(H)$ that is not a disk and not incident to E .

Since F_0 is not a disk, H is not a disk handlebody in the decomposition $\mathbb{C} \rightarrow \mathbb{C}_{x+y}$. Since F_0 is not incident to E , F_0 is not affected by the decomposition $\mathbb{C} \xrightarrow{E} \mathbb{C}_x$. This implies that H is also not a disk handlebody in the decomposition $\mathbb{C}_x \rightarrow \mathbb{C}_{xy}$. H is then an occupied handlebody in both $\mathbb{C}_{x+y}$ or $\mathbb{C}_{xy}$, so H is coherent.

Subcase 1b: A component F_e of $F \cap \text{int}(H)$ is incident to E and surgery on E (in the decomposition $\mathbb{C} \xrightarrow{E} \mathbb{C}_x$) does not turn F_e into the union of disks and spheres.

Essentially the same argument applies: F_e can't be a disk, so H is occupied in $\mathbb{C}_{x+y}$; some non-disk component of the surgered F_e remains as a component of $F(\mathbb{C}_x)$ so H is occupied in $\mathbb{C}_{xy}$.

Subcase 2: E does not lie in $\text{int}(H)$, so it leaves either one or two external scars on ∂H .

Following Subcase 1a, we can assume every component of $F \cap \text{int}(H)$ is a disk, hence H is disk handlebody and so, by Corollary 7.3 H is an empty chamber in $\mathbb{C}_{x+y}$. H is also a disk handlebody remnant of the

¹The phrase ' H lies in a chamber' means that either H is itself a chamber, or H is a goneball in a chamber.

decomposition $\mathbb{C}_x \rightarrow \mathbb{C}_{xy}$ but this does not immediately imply that H is empty in $\mathbb{C}_{xy}$: per Definition 7.2 H might be the remnant of an occupied handlebody C_x in $\mathbb{C}_x$ all of whose other remnants are also diskly handlebodies. We explore this possibility further:

Suppose this were the case, and C is the chamber in $\mathbb{C}$ of which C_x is the occupied handlebody remnant. If the disk E lies outside C then the scars of E in ∂C_x would be internal scars. By the hypothesis of this subcase we may therefore assume that E lies in the interior of C . The one or two remnants of C in $\mathbb{C}_x$ are then subsets of C , so F is disjoint from these remnants and the remnant C_x in particular is diskly. Per Definition 7.2 the only way that $\mathbb{C}_x$ can be a diskly remnant of C and still be occupied is if C is an occupied handlebody with all remnants diskly. But a hypothesis of the Lemma is that $\mathbb{C}$ does not certify, so C is not an occupied handlebody. From the contradiction we conclude that H is empty in $\mathbb{C}_{xy}$ and so H is coherent.

Following Subcase 2, we henceforth assume that $E \subset \text{int}(H)$.

Subcase 3: $F \cap \text{int}(H)$ consists of disks, so H is either a goneball or an empty chamber in $\mathbb{C}_{x+y}$.

Let F_e be the component of $F \cap \text{int}(H)$ on which ∂E lies. Since H is irreducible, E is parallel in H to a subdisk F_0 of F_e . Then the decomposition $\mathbb{C} \xrightarrow{E} \mathbb{C}_x$ has essentially no effect on $\mathbb{C}$: after removing the goneball bounded by $E \cup F_0$, all that has changed is a proper isotopy of F_e that replaces F_0 with E (and the removal of any disks in $\mathcal{D}_y$ that were incident to F_0 , but these only give rise to goneballs in $\mathbb{C}_{x+y}$ and $\mathbb{C}_{xy}$). Then the decomposition $\mathbb{C}_x \rightarrow \mathbb{C}_{xy}$ also leaves H either a goneball or an empty chamber in $\mathbb{C}_{xy}$, so H is coherent.

Subcase 4: All that remains: $F \cap \text{int}(H)$ has a single non-disk component F_e , F_e is incident to E , and surgery on E turns F_e into the union of disks and spheres.

F_e could be a torus bounding an empty solid torus, a once-punctured torus, or an annulus, and in each case E is a meridian disk. (See the left column in Figures 36 and 37.) It follows that ∂E is non-separating in F and, unless $\mathcal{D}_y$ also contains a disk whose boundary is parallel to ∂E , ∂E is also non-separating in $F - \partial \mathcal{D}_y$. If $\mathcal{D}_y$ does contain a disk D whose boundary is parallel to ∂E then D and E are parallel, since otherwise between them would lie a closed component of $F \cap \text{int}(H)$, contradicting the hypothesis of this subcase. So, in the end, we are in the situation described in the statement of the Lemma: **∂E is non-separating in $F - \partial \mathcal{D}_y$ or ∂E is non-separating in F and E is parallel to a disk in $\mathcal{D}_y$** . In this case

H is an occupied handlebody in $\mathbb{C}_{x+y}$, since F_e is not a disk, but may be empty in $\mathbb{C}_{xy}$, since $\mathbb{C}_x \xrightarrow{\mathcal{D}_y} \mathbb{C}_{xy}$ is diskly. Since H could be incoherent, we explore this situation further:

Consider the chamber C_x in $\mathbb{C}_x$ which has H as a remnant. If every other remnant of C_x is also a diskly handlebody, then C_x is a handlebody and it must be occupied because F_e is not a disk. Since $E \subset \text{int}(H)$, each other remnant of C_x is a diskly y -handlebody, so per Lemma 12.7, each is empty. This implies that H must be occupied in $\mathbb{C}_{xy}$, and H is coherent. Finally, if any remnant of C_x in $\mathbb{C}_{xy}$ is not a diskly handlebody then Definition 7.2 says that a preferred alignment of $\mathcal{D}_y$ leaves H deflated to such a remnant, an outcome allowed in Lemma 12.10. This concludes the proof in the special case.

The general case:

Following the special case, the remaining possibility is that there is a sphere in $\mathfrak{S}$ bounding a ball in H that is not coherent. Let G be an innermost such sphere. By the special case, E lies in the ball B_G that G bounds in H and B_G is occupied in $\mathbb{C}_{x+y}$ but deflates into another chamber in $\mathbb{C}_{xy}$. Since B_G is

occupied in $\mathbb{C}_{x+y}$, it is a chamber, so H does not lie inside a chamber in $\mathbb{C}_{x+y}$. If any sphere in $\mathfrak{S}$ does not bound a goneball in $\mathbb{C}_{xy}$ then H is also not a chamber in $\mathbb{C}_{xy}$ and so is coherent. See Figure 35.





Figure 35: Bullseye deflation $\mathbb{C}_{ey} \rightarrow \mathbb{C}_{e+y}$

We are left with the case that every sphere in $\mathfrak{S}$ bounds a goneball in $\mathbb{C}_{xy}$. Any such sphere not containing G in the interior of the ball it bounds will be a y -handlebody, since $E \subset B_G$, so it follows from Lemma 12.8 that the ball will also be a goneball in $\mathbb{C}_{x+y}$. So all spheres in $\mathfrak{S}$ that don't bound goneballs in $\mathbb{C}_{x+y}$ are nested around G . Let G' be an outermost sphere in this nested collection of spheres $\mathfrak{S}_e \subset \mathfrak{S}$ bounding goneballs in $\mathbb{C}_{xy}$ but not in $\mathbb{C}_{x+y}$, or ∂H itself if H is a ball. Since G' bounds a goneball in $\mathbb{C}_{xy}$, the ball has trivial Heegaard splitting. It follows that if $\mathfrak{S}_e$ has more than one sphere, the collar between any two has trivial splitting. Thus $\mathfrak{S}_e$ is a bullseye in $\mathbb{C}_{x+y}$, the final possibility allowed by Lemma 12.10. $\square$

The next lemma is more complicated to state, mostly because it has to accommodate the possibility that even if $\mathcal{D}_y = E$, there may be two xe -handlebody chambers in $\mathbb{C}_{x+y}$ and $\mathbb{C}_{xy}$. Namely if E divides a handlebody chamber $C_x \subset \mathbb{C}_x$ into two handlebodies, then they each are xe -handlebody chambers, since E leaves an external scar on both of them. According to Definition 7.2 if C_x is an occupied handlebody, there are two preferred alignments of E : H is empty and H' is occupied, or vice versa. The different alignments produce sibling flaggings of $\mathbb{C}_{xy}$. One sibling is obtained from the other by deflation of H' into H (which is a bullseye deflation if H' is then a goneball) or vice versa. In this case it is natural to call replacing the flagging of the chambers H and H' in $\mathbb{C}_{x+y}$ with either of the flaggings in $\mathbb{C}_{xy}$ a *sibling deflation*.



 Figure 36: H a ball, F_e a punctured torus

Lemma 12.11. *Suppose $\mathbb{C}$ does not certify and H is an xe -handlebody. Then either*

- H is coherent, or
- $\mathbb{C}_x$ has an occupied handlebody chamber containing E and, in one of the two sibling deflations, H is coherent, or
- H is a chamber in $\mathbb{C}_{xy}$, either the deflation of an occupied chamber in $\mathbb{C}_{x+y}$ or the result of the deflation of a single bullseye of $\mathbb{C}_{x+y}$ in H .

Furthermore, in the third instance, when neither of the first two outcomes applies,

- E leaves exactly one scar on ∂H ,


 Figure 37: H a ball, F_e an annulus

- ∂E is separating in a component of $F - \partial \mathcal{D}_x$, and
- E is not parallel to any disk in $\mathcal{D}_x$.

Proof. If E is parallel to a disk in $\mathcal{D}_x$, Lemma 12.4 says one of the first two outcomes occurs, and we are done. So henceforth we assume that E is not parallel to a disk in $\mathcal{D}_x$.

Here are some preliminary observations:

- If H (or H' in the case in which sibling deflation may arise) is coherent we are done. So we assume H (and H' , if relevant) is not coherent.
- Following Lemma 12.2, unless every component of $F_{\mathcal{D}_x \cup E}$ lying in $\text{int}(H)$ is a sphere, H is not a chamber in either $\mathbb{C}_{xy}$ nor $\mathbb{C}_{x+y}$ and so is coherent. So we henceforth assume that every component of $F_{\mathcal{D}_x \cup E}$ lying in $\text{int}(H)$ is a sphere. Let $\mathfrak{S}$ denote the collection of all such spheres.
- Sibling deflation may arise when E is a separating disk in an occupied handlebody chamber of $\mathbb{C}_x$. If that chamber has H as one of the two remnants, each ball bounded by a sphere in $\mathfrak{S}$ is an x -handlebody and so, per Lemma 12.8 the sphere is coherent.

On the last point, sibling deflation, there is more to say:

The Sibling Claim: Suppose that E is a separating disk in an occupied handlebody chamber C_x of $\mathbb{C}_x$. In at least one of the two sibling deflations of C_x in $\mathbb{C}_{xy}$, either

- H is coherent, or
- H is the deflation of an occupied chamber in $\mathbb{C}_{x+y}$ or
- the only remnant of C_x in the interior of H is a coherent ball.

Proof of the Sibling Claim: Since $\mathbb{C}$ has no occupied handlebodies, the handlebody $\mathbb{C}_x$ is occupied if and only if $F \cap \text{int}(C_x)$ contains a non-disk component, by Corollary 7.3. Hence one or both of the handlebody chamber remnants $C_x - E$ are occupied in $\mathbb{C}_{x+y}$. It follows that either chamber (but not necessarily both) can be made coherent by a choice of sibling deflation in $\mathbb{C}_{xy}$. This proves the Sibling Claim if either of these chambers is H . Another possibility is that both handlebody chambers lie in the interior of H , and so are necessarily balls. We have seen then that at least one ball is occupied in $\mathbb{C}_{x+y}$ and exactly one in $\mathbb{C}_{xy}$. Thus H is not a chamber in either $\mathbb{C}_{x+y}$ or $\mathbb{C}_{xy}$ so H is coherent.

The final possibility is that one of the handlebody chamber remnants is a ball B_H in the interior of H and the other lies outside H (in fact it is then the complement $S^3 - H$ since its boundary is connected). B_H is either empty or occupied in $\mathbb{C}_{x+y}$. By choosing the right sibling deflation we can match that status for B_H in $\mathbb{C}_{xy}$. This makes B_H coherent and so proves the Sibling Claim.

It is possible that a sphere in $\mathfrak{S}$ bounds a ball that is not coherent - that is, it is a goneball in one of $\mathbb{C}_{xy}$ or $\mathbb{C}_{x+y}$ but not the other (regardless of the choice of sibling deflation, when E lies in an occupied handlebody of $\mathbb{C}_x$). If this occurs, let B_G be an innermost such ball, B_G can't be an x -handlebody, by Lemma 12.8 so E is incident to B_G . By the Sibling Claim, we may assume that if E lies in an occupied handlebody in $\mathbb{C}_x$, B_G is not one of its remnants in $\mathbb{C}_{xy}$.

In view of these possibilities, in the discussion below we will use the term *target manifold* to refer either to B_G , if there is a B_G as above, or to H itself if there is no B_G . In either case, any handlebody contained in the interior of the target manifold is a coherent ball.

Let $F_x = F(\mathbb{C}_x)$ denote the defining surface of $\mathbb{C}_x$. We consider possible scars that E might leave on the boundary of the target manifold in $F_{\mathcal{D}_x \cup E}$. Only in the last case, when there is exactly one external scar on the boundary of the target manifold, does the possibility of sibling deflations arise:

Claim 1: There is a scar of E on the boundary of the target manifold (either $G = \partial B_G$ or ∂H .)

The argument is essentially the same for either target manifold, so we assume here the target manifold is B_G . If there is no scar on G then G is among the components of $F_{\mathcal{D}_x}$, the surface obtained by $\mathcal{D}_x$ surgery on F . Notice that B_G must be a chamber in both $\mathbb{C}_{x+y}$ and $\mathbb{C}_{xy}$, for if there are occupied balls in the interior of one, they would be occupied in the other (because B_G is innermost among the non-coherent), making B_G not a chamber and therefore coherent, contradicting hypothesis.

If $F \cap \text{int}(B_G)$ consists of disks, then B_G is disky and hence a goneball in both $\mathbb{C}_{x+y}$ and $\mathbb{C}_x$, by Corollary 7.3. Hence B_G is also a goneball in $\mathbb{C}_{xy}$, again contradicting that B_G is not coherent. On the other hand, if $F \cap \text{int}(B_G)$ is not all disks. Then B_G is an occupied ball in both $\mathbb{C}_{x+y}$ and $\mathbb{C}_x$. The latter

implies that B_G is also an occupied ball in $\mathbb{C}_{xy}$, since E is not incident to B_G ; B_G is its own and only remnant in $\mathbb{C}_{xy}$. This would make B_G coherent, a final contradiction that proves Claim 1.

Claim 2: E does not leave two external scars on the target manifold.

Again the argument for target manifold B_G applies also to H , so we can suppose that B_G is the target manifold and consider what happens if E leaves two external scars on G . Then the chamber C_x in $\mathbb{C}_x$ of which B_G is a remnant is obtained from G by attaching a 1-handle dual to E , and so is a handlebody in $\mathbb{C}_x$. B_G is the sole remnant of C_x under the decomposition $\mathbb{C}_x \rightarrow \mathbb{C}_{xy}$ and its interior is disjoint from F_x , so B_G is occupied in $\mathbb{C}_{xy}$ if and only if C_x is occupied in $\mathbb{C}_x$. Since $\mathbb{C}$ has no occupied handlebodies, the handlebody $\mathbb{C}_x$ is occupied if and only if $F \cap \text{int}(C_x)$ contains a non-disk component, by Corollary 7.3. But $F \cap \text{int}(C_x)$ contains a non-disk component if and only if $F \cap \text{int}(B_G)$ contains a non-disk component and the latter is equivalent to B_G being occupied in $\mathbb{C}_{x+y}$. Hence B_G is occupied in $\mathbb{C}_{xy}$ if and only if it is occupied in $\mathbb{C}_{x+y}$. Hence B_G is coherent, a contradiction that proves the claim.

Claim 3: E does not leave two internal scars on the target manifold.

The case when the target manifold is B_G is representative. If E leaves two internal scars on G then the belt annulus is a non-disk component of both $F \cap \text{int}(B_G)$ and $F_x \cap \text{int}(B_G)$, so B_G is occupied in both $\mathbb{C}_{x+y}$ and $\mathbb{C}_{xy}$. Hence B_G is coherent, a contradiction that proves the claim.

Claim 4: E does not leave one internal scar on the target manifold

This is the most difficult claim to verify, for the argument involves some unusual computation. To avoid the danger of over-simplifying we will assume that the more complicated H is the target manifold; the argument when B_G is the target manifold is then just a special case. Suppose that E leaves one internal scar on ∂H . Then the other scar must be on a component of $F_{\mathcal{D}_x \cup E}$ that lies in the interior of H and so is a sphere $S \in \mathfrak{S}$. Put another way, in $F_{\mathcal{D}_x}$ the belt annulus of E connects the scar of E on ∂H to the disk complement in S of the scar of E on S . So the scar of E on ∂H is parallel in H to a disk component F_0 of $F_x \cap \text{int}(H)$ (the union of the belt annulus of E and a subdisk of S). Since E is the only disk in the decomposition $\mathbb{C}_x \rightarrow \mathbb{C}_{xy}$, the disk F_0 is the only component of $F_x \cap \text{int}(H)$.

By assumption the ball B_S that S bounds in H is coherent: unless it is a goneball in both, it persists in both $\mathbb{C}_{x+y}$ and $\mathbb{C}_{xy}$, so H would not be a chamber in either and therefore is coherent, a contradiction. So we henceforth assume that B_S is a goneball in both $\mathbb{C}_{x+y}$ and $\mathbb{C}_{xy}$. The former means that $F \cap \text{int}(B_S)$ consists of disks, and the latter means that the chamber C_x of $\mathbb{C}_x$ of which H is a remnant is just the complement in H of the ball (parallel to B_S) between the scar of E in ∂H and F_0 . That is, H is parallel to C_x via this ball. In particular, H is the only remnant of C_x so, via Definition 7.2, H is occupied in $\mathbb{C}_{xy}$ if and only if C_x is occupied in $\mathbb{C}_x$.

We wish now to examine $F \cap \text{int}(H)$. F_0 is a disk containing scars of $\mathcal{D}_x$; the complement of the scars is a subset of F . In particular, if $F \cap \text{int}(H)$ contains any closed components they would be disjoint from F_0 . A closed component of F can't lie in B_S , since B_S is empty in $\mathbb{C}_{x+y}$, so it would have to lie in the chamber C_x of $\mathbb{C}_x$, making C_x , hence H , occupied in $\mathbb{C}_{xy}$ as well as $\mathbb{C}_{x+y}$. Then H would be coherent, a contradiction. So we need only consider the case in which $F \cap \text{int}(H)$ contains no closed components.

Here is a useful observation about compact orientable surfaces with no closed components: For any closed connected orientable surface Q , $\chi(Q) \leq 2$ with equality if and only if Q is a sphere. It follows that for any compact connected surface Q , $\chi(Q) + |\partial Q| \leq 2$, with equality if and only if Q is planar, and so, when $\partial Q \neq \emptyset$, $\chi(Q) - |\partial Q| \leq 0$ with equality if and only if Q is planar and has a single boundary

component, i. e. is a disk. More generally, this is then true for any compact orientable surface, so long as there are no closed (in fact no sphere) components. We have thus verified this claim:

Claim: Suppose Q is a compact orientable surface with no closed components. Then $\chi(Q) - |\partial Q| \leq 0$ with equality if and only if Q consists entirely of disks.




Figure 38: $F \cap \text{int}(H) = (F \cap \text{int}(C_x)) \cup (F \cap F_0) \cup (F \cap \text{int}(B_S))$

To apply the Claim, begin with the observation

$$F \cap \text{int}(H) = (F \cap \text{int}(C_x)) \cup (F \cap F_0) \cup (F \cap \text{int}(B_S)). \quad (1)$$

See Figure 38. For the purpose of this computation, include all boundary circles on each surface, so each surface is compact. We then have

$$\chi(F \cap \text{int}(H)) - \chi(F \cap \text{int}(C_x)) = \chi(F \cap F_0) + \chi(F \cap \text{int}(B_S)). \quad (2)$$

Note that $F \cap F_0$ is the disk F_0 with all the scars of $\mathcal{D}_x$ that lie in F_0 removed. Hence $\chi(F \cap F_0) = 1 - |\text{scars on } F_0|$. These scars on F_0 are of two types: those that come from components of F lying in $\text{int}(B_S)$ (all of them disks, since B_S is a goneball) and those coming from components of F lying in $\text{int}(C_x)$. Each disk in $F \cap \text{int}(B_S)$ contributes 1 to $\chi(F \cap \text{int}(B_S))$ and -1 to $\chi(F \cap F_0)$ so

$$\chi(F \cap F_0) + \chi(F \cap \text{int}(B_S)) = 1 - |\text{scars on } F_0 \text{ left by } F \cap \text{int}(C_x)|. \quad (3)$$

Another way to characterize the number of scars left on F_0 by $F \cap \text{int}(C_x)$ is that it is the difference of the number of scars left on ∂C_x and those left on ∂H . Furthermore, the total number of scars left on ∂H is one more than the number left by ∂C_x , since it includes ∂E . We then obtain:

$$1 - |\text{scars on } F_0 \text{ left by } F \cap \text{int}(C_x)| = |\partial(F \cap \text{int}(H))| - |\partial(F \cap \text{int}(C_x))|. \quad (4)$$

Combining (2), (3) and (4) we have

$$\chi(F \cap \text{int}(H)) - \chi(F \cap \text{int}(C_x)) = |\partial(F \cap \text{int}(H))| - |\partial(F \cap \text{int}(C_x))| \quad (5)$$

which implies

$$\chi(F \cap \text{int}(H)) - |\partial(F \cap \text{int}(H))| = \chi(F \cap \text{int}(C_x)) - |\partial(F \cap \text{int}(C_x))|. \quad (6)$$

Now apply the Claim: The left side of (6) is zero if and only if H is diskly and hence empty; the right side of (6) is zero if and only if C_x is diskly and hence empty. Thus H is empty in $\mathbb{C}_{x+y}$ if and only if it is empty in $\mathbb{C}_{xy}$. That is, H is coherent, a contradiction that proves Claim 4.

Claim 5: If the target manifold contains exactly one external scar of E then H satisfies one of the three outcomes listed in Lemma 12.11.

The argument for this claim depends on whether H or B_G is the target manifold.

Case 1: The target manifold is B_G .

Let C_x be the chamber in $\mathbb{C}_x$ of which B_G is a remnant. Since E leaves exactly one external scar on G , E separates C_x into B_G and another remnant R_x , so B_G is diskly under the decomposition $\mathbb{C}_x \rightarrow \mathbb{C}_{xy}$.

Subcase 1a: C_x is a handlebody.

We first show that C_x cannot be an empty handlebody. Since $\mathbb{C}$ contains no occupied handlebodies, C_x can only be empty if $F \cap \text{int}(C_x)$ consists entirely of disks, by Corollary 7.3. Since $\text{int}(B_G) \subset \text{int}(C_x)$ this implies that $F \cap \text{int}(B_G)$ consists entirely of disks, so B_G is empty in $\mathbb{C}_{x+y}$. Since C_x is empty in $\mathbb{C}_x$, Definition 7.2 says B_G is also empty in $\mathbb{C}_{xy}$. This contradicts the requirement that B_G is not coherent.

We deduce that C_x is an occupied handlebody. In this case, apply the Sibling Claim: the first two outcomes of that claim imply the first two outcomes of the Lemma; the third only arises when the target manifold is H , and then it contradicts Claim 4, completing the proof of the Lemma in this subcase.

Subcase 1b: C_x is not a handlebody.


In this case R_x cannot be a handlebody and in particular cannot be a ball, and so it can't be a chamber in $\text{int}(H)$. Then $R_x \subset S^3 - H$. Moreover, Definition 7.2 says that in this case the preferred alignment of E leaves B_G empty in $\mathbb{C}_{x+y}$ and so a goneball. Our assumption is that B_G is not coherent, so B_G must be occupied in $\mathbb{C}_{x+y}$. The difference then between H in $\mathbb{C}_{x+y}$ and $\mathbb{C}_{xy}$ is that the latter is obtained from the former by deflation of B_G . This can be viewed as a bullseye deflation, the third outcome allowed. See the right side of Figure 39.

Case 2: The target manifold is H .

The proof is highly analogous to Case 1, so we merely sketch: As in Case 1, the component C_x of $\mathbb{C}_x$ of which H is a remnant is divided by E into two chambers, H and R_x . If C_x is an empty handlebody then, just as in Case 1, H is empty in both $\mathbb{C}_{x+y}$ and $\mathbb{C}_{xy}$ and so is coherent, completing the proof. If $\mathbb{C}_x$ is an occupied handlebody then the Sibling Claim applies; the first two outcomes of that Claim imply the first two claims of the Lemma, and the third outcome contradicts Claim 4. Finally, if C_x (so also R_x) is not a handlebody, then the preferred alignment given by Definition 7.2 leaves H empty in $\mathbb{C}_{xy}$. Thus if H is empty in $\mathbb{C}_{x+y}$ it is coherent; if it is occupied then $\mathbb{C}_{xy}$ is obtained from $\mathbb{C}_{x+y}$ by deflation into R_x , implying the third outcome of the Lemma.

Having established the claims, we now prove the Lemma: Suppose a ball $B_G \subset \text{int}(H)$ presents itself, as described before Claim 1. Claim 1 says that E leaves scars on G , claims 2 and 3 show that E leaves exactly one scar on G , and Claim 4 says that it cannot be an internal scar. Thus E must leave one external scar on G , in which case Claim 5 shows that one of the three outcomes of Lemma 12.11 occurs.

On the other hand, if no B_G satisfies the criteria described before Claim 1 then Claims 1, 2, 3 and 4 show that, unless ∂H contains a single external scar of E , H is coherent. Claim 5 shows that when ∂H does contain a single external scar of E , one of the three outcomes of Lemma 12.11 occurs.


 Figure 39: E leaves one external scar

We now examine when neither of the first two outcomes occurs, to verify the final three claims of Lemma 12.11. We have already concluded that E is not parallel to any disk in $\mathcal{D}_x$. The proof so far shows also that the target manifold must have a single external scar, and then Claim 5 shows that in this case ∂H contains exactly one scar of E , an internal scar when the target manifold is B_G and an external scar when the target manifold is H . All that remains is to show that ∂E is separating in $F - \partial \mathcal{D}_x$:

The definition of F_x implies that the complement of the scars of $\mathcal{D}_x$ in F_x is a collection of components of $F - \partial \mathcal{D}_x$. Not removing the scars of the disks in $\mathcal{D}_x$, that is, adding disks back, does not affect the number of components in the surface, so to show that ∂E is separating in a component of $F - \partial \mathcal{D}_x$ it suffices to show that ∂E is separating in a component of F_x . But that follows immediately from the fact that E leaves only one scar on ∂H : ∂H is a component of $F_x - \partial E$ that is distinct from the component on which the other scar of E lies. $\square$

Proposition 12.12. *Suppose $\mathbb{C}$ is a flagged chamber complex supporting $S^3 = A \cup_T B$, $\mathbb{C}$ does not certify and $\mathcal{D}_x, \mathcal{D}_y$ are disjoint disk sets in $\mathbb{C}$. If either $|\mathcal{D}_x|$ or $|\mathcal{D}_y| = 1$, then, after an appropriate choice of sibling if $\mathbb{C}_x$ contains occupied handlebody chambers, either $\mathcal{D}_{xy} = \mathcal{D}_{x+y}$ or $\mathcal{D}_{xy} \dashrightarrow \mathcal{D}_{x+y}$.*

In more detail: Suppose $\mathcal{D} \cup E$ is a disk set in $\mathbb{C}$, where E is a single disk. Let

- $\mathbb{C}_{DE}$ be the flagged chamber complex given by the composition $\mathbb{C} \xrightarrow{\mathcal{D}} \mathbb{C}_{\mathcal{D}} \xrightarrow{E} \mathbb{C}_{DE}$
- $\mathbb{C}_{ED}$ be the flagged chamber complex given by the composition $\mathbb{C} \xrightarrow{E} \mathbb{C}_E \xrightarrow{\mathcal{D}} \mathbb{C}_{ED}$ and $\mathcal{D}$ the preferred alignment described in Lemma 12.8
- $\mathbb{C}_{D+E}$ be the flagged chamber complex given by the decomposition $\mathbb{C} \xrightarrow{\mathcal{D} \cup E} \mathbb{C}_{D+E}$.

(Note that here D in the subscripts refers not to a single disk but to the family $\mathcal{D}$ of disks.)

Then either

- $\mathbb{C}_{DE} = \mathbb{C}_{D+E} = \mathbb{C}_{ED}$ or
- $\mathbb{C}_{DE} = \mathbb{C}_{D+E} \leftarrow \mathbb{C}_{ED}$ or
- $\mathbb{C}_{DE} \dashrightarrow \mathbb{C}_{D+E} = \mathbb{C}_{ED}$ or
- E lies in an occupied handlebody in $\mathbb{C}_{\mathcal{D}}$ and its sibling $\mathbb{C}'_{DE}$ satisfies $\mathbb{C}'_{DE} = \mathbb{C}_{D+E} = \mathbb{C}_{ED}$

If $\mathbb{C}_{D+E} \neq \mathbb{C}_{ED}$ then either ∂E is non-separating in $F - \partial \mathcal{D}$ or ∂E is non-separating in F and E is parallel to a disk in $\mathcal{D}$.

Proof. The proof breaks naturally into two cases:

Case 1: $\mathbb{C}_{ED} \neq \mathbb{C}_{D+E}$.

By Lemma 12.3 there is a handlebody H in S^3 whose boundary is in $F_{\mathcal{D} \cup E}$ and which is not coherent in $\mathbb{C}_{ED}$ and $\mathbb{C}_{D+E}$. Then applying Lemma 12.10 to H , either ∂E is non-separating in $F - \partial \mathcal{D}$ or ∂E is non-separating in F and E is parallel to a disk in $\mathcal{D}$. Then by Lemma 12.11 every handlebody in S^3 whose boundary is in $F_{\mathcal{D} \cup E}$ is coherent for $\mathbb{C}_{DE}$ and $\mathbb{C}_{D+E}$. So, by Lemma 12.3, $\mathbb{C}_{DE} = \mathbb{C}_{D+E}$. Thus what is left to show in this case is that $\mathbb{C}_{ED} \dashrightarrow \mathbb{C}_{D+E}$.

There cannot be two disjoint handlebodies in S^3 whose boundaries are in $F_{\mathcal{D} \cup E}$ and which are not coherent for $\mathbb{C}_{ED}$ and $\mathbb{C}_{D+E}$ since Lemma 12.10 says that the disk E lies in the interior of each such handlebody. Then there is a maximal such handlebody - that is, one that contains any other - so without loss take H to be the maximal one. By Lemma 12.10 H is contained in a chamber of $\mathbb{C}_{ED}$ and is obtained from a chamber of $\mathbb{C}_{D+E}$ by deflation of a chamber or of a bullseye in $\mathbb{C}_{D+E}$. All chambers of $\mathbb{C}_{D+E}$ outside of H are the same in $\mathbb{C}_{D+E}$ and $\mathbb{C}_{ED}$ by the same proof as that of Lemma 12.3. Thus $\mathbb{C}_{ED} \dashrightarrow \mathbb{C}_{D+E}$ as required.

Case 2: $\mathbb{C}_{DE} \neq \mathbb{C}_{D+E}$

The proof is analogous to that of Case 1, but a bit more complicated. By Lemma 12.3 there is a handlebody H in S^3 whose boundary is in $F_{\mathcal{D} \cup E}$ and which is not coherent in $\mathbb{C}_{DE}$ and $\mathbb{C}_{D+E}$. Then according to Lemma 12.11 E is not parallel to a disk in $\mathcal{D}$ and ∂E is separating in $F - \partial \mathcal{D}$. Then by Lemma 12.10 every handlebody in S^3 whose boundary is in $F_{\mathcal{D} \cup E}$ is coherent for $\mathbb{C}_{ED}$ and $\mathbb{C}_{D+E}$ so by Lemma 12.3 $\mathbb{C}_{ED} = \mathbb{C}_{D+E}$. Thus what is left to show in this case is that $\mathbb{C}_{DE} \dashrightarrow \mathbb{C}_{D+E}$.

According to Lemma 12.11 E leaves exactly one scar on ∂H . If it is an internal scar, then per Claim 5 in the proof of Lemma 12.11, E lies in the interior of a chamber C_x of $\mathbb{C}_x$ dividing C_x into a ball in $\text{int}(H)$ and a chamber R_x adjacent to and outside of H , as in Figure 38. Unless R_x is also a handlebody, the choice of preferred alignment of E has no effect on the flagging of chambers outside of H . In particular, there can be no handlebody chamber incoherent for $\mathbb{C}_{DE}$ and $\mathbb{C}_{D+E}$ that is disjoint from H . In this case, the argument concludes as for Case 1, using Lemma 12.11 instead of Lemma 12.10. On the other hand, if R_x is a handlebody (and so, as noted in Lemma 12.11 Claim 5, Subcase 1a, the only chamber of $\mathbb{C}_{x+y}$ that lies outside of H) then the choice of alignment of E in C_x does affect the flagging of R_x . The argument in this case is much like the argument in the case that E leaves an external scar on H , so we turn to that case, noting that the argument applies here as well.

If E leaves an external scar on H then it is possible that there is a disjoint handlebody H' in S^3 just like H : $\partial H'$ is in $F_{\mathcal{D} \cup E}$, H' is not coherent for $\mathbb{C}_{DE}$ and $\mathbb{C}_{D+E}$, and E also leaves an external scar on $\partial H'$. Then, just as is true for H , H' lies in a handlebody chamber in $\mathbb{C}_{DE}$. Thus there is a handlebody chamber $C_x \in \mathbb{C}_{\mathcal{D}}$ containing E , with E dividing C_x into H and H' .

We now reprise part of the argument of Claim 4a in Lemma 12.10: C_x must be an occupied handlebody, for if C_x is empty then H and H' would both be empty in $\mathbb{C}_{D+E}$ and $\mathbb{C}_{DE}$, contradicting the assumption that H is not coherent. Then, per Definition 7.2, C_x is the parent of two siblings in $\mathbb{C}_{DE}$: one in which H is occupied and H' is empty, the other in which H' is occupied and H is empty.

Since C_x is occupied, $F \cap \text{int}(C_x)$ has non-disk components by Corollary 7.3. There are three cases to consider, depending on where these non-disk components lie:

- If some non-disk components lie in H and some in H' , then H and H' are both occupied in $\mathbb{C}_{D+E}$. Since H is not coherent, it must then be empty in $\mathbb{C}_{DE}$. In this case the sibling for C_{DE} is the one in which H is empty and H' is occupied, that is the one that results from an alignment for E which deflates H into H' , leaving H' occupied in $\mathbb{C}_{DE}$ as well. Thus $\mathbb{C}_{DE} \rightarrow \mathbb{C}_{D+E}$, an allowed outcome.
- If each component of $F \cap \text{int}(H)$ is a disk, then H is empty in $\mathbb{C}_{D+E}$. Since H is not coherent it is then occupied in $\mathbb{C}_{DE}$. But in this case the sibling C'_{DE} of C_{DE} is one in which H is empty and H' is occupied. Then in this sibling $\mathbb{C}'_{DE} = \mathbb{C}_{D+E}$, an allowed outcome.
- If each component of $F \cap \text{int}(H')$ is a disk then H' is empty and H is occupied in $\mathbb{C}_{D+E}$. Since H is not coherent, it is empty in $\mathbb{C}_{DE}$. In this case the sibling C'_{DE} of C_{DE} has H occupied and H' empty, just as is true for $\mathbb{C}_{D+E}$. Thus $\mathbb{C}'_{DE} = \mathbb{C}_{D+E}$, an allowed outcome.

□

Corollary 12.13. *Suppose $\mathbb{C}$ is a flagged chamber complex supporting $S^3 = A \cup_T B$, $\mathbb{C}$ does not certify and $\mathcal{D} \cup E$ is a disk set in $\mathbb{C}$, where E is a single disk. In the notation of Proposition 12.12, one of these is true (for some choice of sibling, if siblings exist):*

- $\mathbb{C}_{DE} = \mathbb{C}_{ED}$,
- $\mathbb{C}_{DE} \rightarrow \mathbb{C}_{ED}$ or
- $\mathbb{C}_{DE} \leftarrow \mathbb{C}_{ED}$

We denote the three possibilities above by $\mathbb{C}_{DE} \leftrightarrow \mathbb{C}_{ED}$. Then one way of expressing the conclusion of Corollary 12.13 is to say that this diagram commutes:

$$\begin{array}{ccc}
 \mathbb{C}_0 & \xrightarrow{\mathcal{D}} & \mathbb{C}_{\mathcal{D}} \\
 \downarrow E & & \downarrow E \\
 & & \mathbb{C}_{DE} \\
 & & \uparrow \\
 & & \mathbb{C}_{ED} \\
 \mathbb{C}_E & \xrightarrow{\mathcal{D}} &
 \end{array}$$

13 Guiding disk sets and disk addition

Recall from Section 11 the notion of guiding disks (Definition 11.8) for flagged chamber decompositions. The example given there involved decomposition of two chamber complexes that are related by bullseye deflation. Here is another example: Suppose $\mathcal{D} \cup E$ is a disk set in a flagged chamber complex $\mathbb{C}$. Let $\mathbb{C}'$ be the result of the flagged chamber complex decomposition $\mathbb{C} \xrightarrow{E} \mathbb{C}'$ and let $\mathcal{D}'$ be the set of disks in $\mathcal{D}$ whose boundaries lie on $\mathbb{C}'$. The decomposing set $\mathcal{D}'$ is contained in $\mathcal{D}$, but may not be all of $\mathcal{D}$, because it does include those disks in $\mathcal{D}$ that lie on goneballs of the decomposition $\mathbb{C} \xrightarrow{E} \mathbb{C}'$. Nevertheless $\mathcal{D}$ remains as a guiding set of disks for the decomposition $\mathbb{C}' \xrightarrow{\mathcal{D}'} \mathbb{C}'_{\mathcal{D}'}$.

We have already encountered another example of guiding disk sets, in the context of guiding spheres in Section 9. Suppose $\mathbb{C}$ is a flagged chamber complex in S^3 with defining surface $F = F(\mathbb{C})$, and $S \subset S^3$ is an embedded sphere transverse to F . Consider a sequence $\vec{\mathbb{C}}$ of flagged chamber complex decompositions guided by S , as given in Definition 9.4 and described after the proof of Lemma 9.3: The circles $F \cap S$ are partitioned into sets $c_i, 0 \leq i \leq n$ and to each c_i a collection of disjoint disks $\overline{\mathcal{D}}_i \subset S$ is assigned. The associated sequence $\vec{\mathbb{C}}$ of flagged chamber complex decompositions guided by S has, as each disk set $\mathcal{D}_i$ those disks in $\overline{\mathcal{D}}_i$ that remain incident to $F(\mathbb{C}_i)$ after all goneballs in earlier decompositions have been removed. Thus each set of disks $\overline{\mathcal{D}}_i$ is a guiding set of disks for the decomposition $\mathbb{C}_i \xrightarrow{\mathcal{D}_i} \mathbb{C}_{i+1}$ not the disk set itself. Consistent with the notation of Definition 11.8 we could then write the sequence in Definition 9.4 as

$$\mathbb{C} = \mathbb{C}_0 \xrightarrow{\overline{\mathcal{D}}_0} \mathbb{C}_1 \xrightarrow{\overline{\mathcal{D}}_1} \mathbb{C}_2 \xrightarrow{\overline{\mathcal{D}}_2} \dots \xrightarrow{\overline{\mathcal{D}}_{k-1}} \mathbb{C}_k$$

Joining this example with the previous one, suppose E is a properly embedded disk in $\mathbb{C}$ that is disjoint from the guiding sphere S and, as above, let $\mathbb{C}'$ be the result of the flagged chamber complex decomposition $\mathbb{C} \xrightarrow{E} \mathbb{C}'$. Then it is natural to construct a sequence $\vec{\mathbb{C}}'$ of flagged chamber complex decompositions

$$\vec{\mathbb{C}}' : \quad \mathbb{C}' = \mathbb{C}'_0 \xrightarrow{\mathcal{D}'_0} \mathbb{C}'_1 \xrightarrow{\mathcal{D}'_1} \mathbb{C}'_2 \xrightarrow{\mathcal{D}'_2} \dots \xrightarrow{\mathcal{D}'_{n-1}} \mathbb{C}'_n \quad (13.1)$$

by setting each $\mathcal{D}'_i$ to be the collection of disks in $\mathcal{D}_i$ (and so in $\overline{\mathcal{D}}_i$) whose boundaries lie on $\mathbb{C}'_i$. Then, just as for $\vec{\mathbb{C}}$, each decomposition $\mathbb{C}'_i \xrightarrow{\mathcal{D}'_i} \mathbb{C}'_{i+1}$ has guiding disk set $\overline{\mathcal{D}}_i$. So the sequence could also be

written

$$\vec{\mathbb{C}}' : \mathbb{C}' = \mathbb{C}'_0 \xrightarrow{\overline{\mathcal{D}}_0} \mathbb{C}'_1 \xrightarrow{\overline{\mathcal{D}}_1} \mathbb{C}'_2 \xrightarrow{\overline{\mathcal{D}}_2} \dots \xrightarrow{\overline{\mathcal{D}}_{n-1}} \mathbb{C}'_n.$$

One way to describe the relation between the sequences above is that the sequence $\vec{\mathbb{C}}'$ is what the sequence $\vec{\mathbb{C}}$ would be if we had inserted decomposition with the disk E before beginning the sequence. Generalizing, we could decompose by E at a later stage of the decomposition sequence. That is, for each $0 \leq k \leq n$, define a flagged chamber complex decomposition sequence

$$\vec{\mathbb{C}}_E^k : \mathbb{C}_0 \xrightarrow{\mathcal{D}_0} \mathbb{C}_1 \xrightarrow{\mathcal{D}_1} \dots \xrightarrow{\mathcal{D}_{k-1}} \mathbb{C}_k \xrightarrow{E} \mathbb{C}_k^k \xrightarrow{\mathcal{D}_k^k} \mathbb{C}_{k+1}^k \xrightarrow{\mathcal{D}_{k+1}^k} \dots \xrightarrow{\mathcal{D}_{n-1}^k} \mathbb{C}_n^k \quad (13.2)$$

where the disks $\mathcal{D}_i^k \subset \overline{\mathcal{D}}_i, k \leq i \leq n$ are those disks whose boundaries lie on $F(\mathbb{C}_i^k)$. In particular, $\vec{\mathbb{C}}'$ of Sequence 13.1 is $\vec{\mathbb{C}}_E^0$ with the first decomposition $\mathbb{C} \xrightarrow{E} \mathbb{C}'$ dropped. As usual, a simplified way of writing each sequence is

$$\vec{\mathbb{C}}_E^k : \mathbb{C}_0 \xrightarrow{\overline{\mathcal{D}}_0} \mathbb{C}_1 \xrightarrow{\overline{\mathcal{D}}_1} \dots \xrightarrow{\overline{\mathcal{D}}_{k-1}} \mathbb{C}_k \xrightarrow{E} \mathbb{C}_k^k \xrightarrow{\overline{\mathcal{D}}_k} \mathbb{C}_{k+1}^k \xrightarrow{\overline{\mathcal{D}}_{k+1}} \dots \xrightarrow{\overline{\mathcal{D}}_{n-1}} \mathbb{C}_n^k$$

Lemma 13.1. *Suppose $\mathbb{C}$ and $\mathbb{C}'$ are flagged chamber complexes and their defining surfaces have the property that $F(\mathbb{C}') \subset F(\mathbb{C})$. Suppose $\mathbb{C} \xrightarrow{\mathcal{D}} \mathbb{C}_{\mathcal{D}}$ and $\mathbb{C}' \xrightarrow{\mathcal{D}'} \mathbb{C}'_{\mathcal{D}'}$ have a guiding disk set $\overline{\mathcal{D}}$ in common. Then $\mathcal{D}$ is a guiding disk set for $\mathcal{D}'$.*

Proof. Each $D' \in \mathcal{D}'$ is a disk in $\overline{\mathcal{D}}$ whose boundary lies on $F(\mathbb{C}')$. Since $F(\mathbb{C}') \subset F(\mathbb{C})$, $\partial D'$ also lies on $F(\mathbb{C})$, so $D' \in \mathcal{D}$. Thus D' is a disk in $\mathcal{D}$ whose boundary lies in $F(\mathbb{C}')$. Conversely, suppose $D \in \mathcal{D}$ has its boundary on $F(\mathbb{C}')$. Since $D \in \mathcal{D}$, $D \in \overline{\mathcal{D}}$ and, since $\overline{\mathcal{D}}$ is a guiding disk set for $\mathcal{D}'$ and $\partial D \subset F(\mathbb{C}')$, $D \in \mathcal{D}'$. Thus $\mathcal{D}'$ consists exactly of those disks in $\mathcal{D}$ whose boundaries lie on $F(\mathbb{C}')$, as required. $\square$

Proposition 13.2. *Let $\mathbb{C}$ be a flagged chamber complex that supports $S^3 = A \cup_T B$ and*

$$\vec{\mathbb{C}} : \mathbb{C} = \mathbb{C}_0 \xrightarrow{\mathcal{D}_0} \mathbb{C}_1 \xrightarrow{\mathcal{D}_1} \dots \xrightarrow{\mathcal{D}_{n-1}} \mathbb{C}_n$$

be a flagged chamber complex decomposition sequence. Suppose

- *there is a guiding disk set $\overline{\mathcal{D}}_i$ for each decomposition $\mathbb{C}_i \xrightarrow{\mathcal{D}_i} \mathbb{C}_{i+1}$.*
- *E is a proper disk in a chamber of $\mathbb{C}$, and E is disjoint from all disks in all $\overline{\mathcal{D}}_i$*
- *$\vec{\mathbb{C}}'$ is the sequence defined in Formula 13.1 above, so each $\overline{\mathcal{D}}_i$ is a guiding disk set for $\mathbb{C}'_i \xrightarrow{\mathcal{D}'_i} \mathbb{C}'_{i+1}$*
- *for each $0 \leq k \leq n$ and decomposition sequence $\vec{\mathbb{C}}_E^k$ defined in 13.2 above each $\overline{\mathcal{D}}_i$ is also a guiding disk set for the corresponding decomposition in $\vec{\mathbb{C}}_E^k$.*

If $\mathbb{C}_n$, $\mathbb{C}'_n$ and all $\mathbb{C}_n^k, 0 \leq k \leq n$ certify, then the sequences $\vec{\mathbb{C}}$ and $\vec{\mathbb{C}}'$ cocertify.

Proof. Case 1: $\mathbb{C}$ certifies.

Then the sequence $\vec{\mathbb{C}}$ and the sequence

$$\vec{\mathbb{C}}_E^0 : \mathbb{C} \xrightarrow{E} \mathbb{C}' = \mathbb{C}'_0 \xrightarrow{\mathcal{D}'_0} \mathbb{C}'_1 \xrightarrow{\mathcal{D}'_1} \dots \xrightarrow{\mathcal{D}'_{n-1}} \mathbb{C}'_n$$

cocertify because the chamber complex $\mathbb{C}$ certifies for both. (The same argument shows that all sequences $\vec{\mathbb{C}}_E^k$ cocertify.) Similarly the sequences $\vec{\mathbb{C}}_E^0$ and $\vec{\mathbb{C}}'$ cocertify because, by hypothesis, $\mathbb{C}'_n$ certifies and, as noted above, the latter sequence is contained in the former sequence (by dropping $\mathbb{C} \xrightarrow{E} \mathbb{C}'$). Thus we have $\vec{\mathbb{C}} \sim \vec{\mathbb{C}}_E^0 \sim \vec{\mathbb{C}}'$ as required.

Case 2: $\mathbb{C}$ does not certify, but $\mathbb{C}_1$ does.

Since $\mathbb{C}$ does not certify then per Corollary 12.13, with some choice of siblings if there are siblings, the following diagram commutes:

$$\begin{array}{ccccccc} \mathbb{C} & \xrightarrow{\mathcal{D}_0} & \mathbb{C}_1 & \xrightarrow{\mathcal{D}_1} & \dots & \xrightarrow{\mathcal{D}_{n-1}} & \mathbb{C}_n = \vec{\mathbb{C}} \\ & & \downarrow E & & & & \\ & & \mathbb{C}_1^1 & \xrightarrow{\mathcal{D}_1^1} & \dots & \xrightarrow{\mathcal{D}_{n-1}^1} & \mathbb{C}_n^1 = \vec{\mathbb{C}}_E^1 \\ & & \uparrow & & & & \\ & & \mathbb{C}' & \xrightarrow{\mathcal{D}'_0} & \mathbb{C}'_1 & \xrightarrow{\mathcal{D}'_1} & \dots \xrightarrow{\mathcal{D}'_{n-1}} \mathbb{C}'_n = \vec{\mathbb{C}}' \\ & \downarrow E & & & & & \end{array}$$

We henceforth ignore the issue of choosing siblings, since different choices of sibling give cocertifying sequences, via the certifying parent chamber. (See the proof of Proposition 9.5.)

We would like to apply Corollary 11.12 to the sequences $\vec{\mathbb{C}}'$ and $\vec{\mathbb{C}}_E^1$. The Corollary requires that if $\mathbb{C}'_1$, say, deflates to $\mathbb{C}_1^1$ (i. e. $\mathbb{C}'_1 \twoheadrightarrow \mathbb{C}_1^1$) that the disks in $\mathcal{D}_1^1$ are a subset of the disks in $\mathcal{D}'_1$, namely those that have their boundaries on the defining surface for $\mathbb{C}_1^1$. In our language, $\mathcal{D}'_1$ needs to be a guiding disk set for $\mathbb{C}_1^1$. Moreover, it is required that the same true for each $\mathcal{D}'_i$ and $\mathcal{D}_i^1$ throughout the maximal deflationary sequence. But since, by hypothesis, $\overline{\mathcal{D}}_i$ is a guiding set of disks for both $\mathcal{D}'_i$ and $\mathcal{D}_i^1$, this is true, via Lemma 13.1.

So we can apply Corollary 11.12: Since $\vec{\mathbb{C}}_E^1$ certifies by the hypothesis that $\mathbb{C}_n^1$ certifies, $\vec{\mathbb{C}}_E^1$ and $\vec{\mathbb{C}}'$ cocertify. A symmetric argument leads to the same conclusion in the case that $\mathbb{C}_1^1$ deflates to $\mathbb{C}'_1$ or, most simply, when $\mathbb{C}_1^1 = \mathbb{C}'_1$. Since, in this case, we are assuming that $\mathbb{C}_1$ certifies, it certifies for the sequences that contain it, namely $\vec{\mathbb{C}}$, and $\vec{\mathbb{C}}_E^1$. So we have $\vec{\mathbb{C}} \sim \vec{\mathbb{C}}_E^1 \sim \vec{\mathbb{C}}'$, as required.

We are left with the possibility that $\mathbb{C}_1$ does not certify. The argument just concluded suggests an inductive approach, which we now pursue:

Inductive Step: Suppose the sequence $\vec{\mathbb{C}}'$ and, for $0 \leq k \leq \ell$, all $\vec{\mathbb{C}}_E^k$ cocertify. (The proof of Cases 1 and 2 show that this is true for $\ell = 1$.) Then either $\vec{\mathbb{C}} \sim \vec{\mathbb{C}}'$ and we are done, or the same statement is true with ℓ replaced by $\ell + 1$.

If $\mathbb{C}_\ell$ certifies, then it certifies for sequences that contain it, including $\vec{\mathbb{C}}$, and $\vec{\mathbb{C}}_E^\ell$. Thus $\vec{\mathbb{C}} \sim \vec{\mathbb{C}}_E^\ell$. By inductive hypothesis, $\vec{\mathbb{C}}_E^\ell \sim \vec{\mathbb{C}}'$. Hence $\vec{\mathbb{C}} \sim \vec{\mathbb{C}}'$, as required.

If $\mathbb{C}_\ell$ does not certify, Corollary 12.13 says the following diagram commutes:

$$\begin{array}{ccccccc}
 \mathbb{C} & \xrightarrow{\mathcal{D}_0} & \mathbb{C}_1 & \xrightarrow{\mathcal{D}_1} & \dots & \xrightarrow{\mathcal{D}_{\ell-1}} & \mathbb{C}_\ell & \xrightarrow{\mathcal{D}_\ell} & \mathbb{C}_{\ell+1} & \xrightarrow{\mathcal{D}_{\ell+1}} & \dots & = \vec{\mathbb{C}} \\
 & & & & & & \downarrow E & & \downarrow E & & & \\
 & & & & & & \mathbb{C}_\ell^{\ell+1} & \xrightarrow{\mathcal{D}_{\ell+1}^{\ell+1}} & \dots & & = \vec{\mathbb{C}}_E^{\ell+1} \\
 & & & & & & \uparrow & & \uparrow & & \\
 & & & & & & \mathbb{C}_\ell^\ell & \xrightarrow{\mathcal{D}_\ell^\ell} & \mathbb{C}_{\ell+1}^\ell & \xrightarrow{\mathcal{D}_{\ell+1}^\ell} & \dots & = \vec{\mathbb{C}}_E^\ell
 \end{array}$$

Applying Corollary 11.12 much as in Case 2, $\vec{\mathbb{C}}_E^\ell$ and $\vec{\mathbb{C}}_E^{\ell+1}$ cocertify. Hence $\vec{\mathbb{C}}'$ and, for $0 \leq k \leq \ell + 1$, all $\vec{\mathbb{C}}_E^k$ cocertify, completing the inductive step.

Continue to iterate the inductive step, showing that $\vec{\mathbb{C}}'$ and each $\vec{\mathbb{C}}_E^k$ cocertify, until a k for which $\mathbb{C}_k$ certifies is reached. We know that we will eventually reach such a k , since $\vec{\mathbb{C}}$ certifies. At that point, $\vec{\mathbb{C}}_E^k$ cocertifies with both $\vec{\mathbb{C}}$ and $\vec{\mathbb{C}}'$, so $\vec{\mathbb{C}}$ and $\vec{\mathbb{C}}'$ cocertify, as required. $\square$

Theorem 13.3. *Suppose $\mathbb{C}$ is a flagged chamber complex in S^3 that is not tiny, and S is a balanced or almost balanced sphere transverse to $F(\mathbb{C})$. Let E be a disk disjoint from S properly embedded and essential in a chamber of $\mathbb{C}$ that is not an empty torus. Let $\mathbb{C}'$ be the flagged chamber complex obtained from $\mathbb{C}$ by decomposition along E . Suppose also that S is balanced or almost balanced for $\mathbb{C}'$. Then any sequence in $(\mathbb{C}, S)$ and any sequence in $(\mathbb{C}', S)$ cocertify.*

Note: the requirements that E be essential in its chamber and that C not be an empty torus will be removed later, see Corollary 15.11.

Proof. Let F, F' denote the defining surfaces $F(\mathbb{C}), F(\mathbb{C}')$ respectively. The circles $F \cap S$ and $F' \cap S$ determine decomposition sequences $\vec{\mathbb{C}}$ and $\vec{\mathbb{C}}'$ as described in Section 9. Since E is essential and not the meridian of an empty torus, surgery on E creates no new goneballs, so in fact $F \cap S = F' \cap S$. Thus the guiding disk sets $\vec{\mathcal{D}}_i$ for the two decomposition sequences $\vec{\mathbb{C}}$ and $\vec{\mathbb{C}}'$ are the same. Import the notation for $\vec{\mathbb{C}}$ and $\vec{\mathbb{C}}'$ from the paragraph that precedes Proposition 13.2 to these decomposition sequences $\vec{\mathbb{C}}$ and $\vec{\mathbb{C}}'$.

We are given that the chamber complexes $\mathbb{C}$ and $\mathbb{C}'$, related by the decomposition $\mathbb{C} \xrightarrow{E} \mathbb{C}'$, are each balanced or almost balanced. As noted in the proof of Proposition 10.2 during the decomposition sequence the genus of the defining surface above S and below S does not change. Surgery on E may or may not lower by one the genus of the defining surface on the side of S in which E lies, but we are given that S is balanced or almost balanced for $\mathbb{C}'$ as well as for $\mathbb{C}$. It follows that for any $0 \leq k \leq m$ the chamber complex $\mathbb{C}_k^k$ defined by the decomposition $\mathbb{C}_k \xrightarrow{E} \mathbb{C}_k^k$ in $\vec{\mathbb{C}}_E^k$ is also balanced or almost balanced.

Then Propositions 10.2 and 10.3 imply that the last term in each of the sequences $\vec{\mathbb{C}}, \vec{\mathbb{C}}'$ and each $\vec{\mathbb{C}}_E^k$ certifies. Proposition 13.2 then says that $\vec{\mathbb{C}}$ and $\vec{\mathbb{C}}'$ cocertify, as claimed. $\square$

14 Delaying a disk

In this section we consider what happens in a sequence of two flagged chamber complex decompositions if, when this is possible, a decomposition disk in the first decomposing disk set is transferred to the second decomposing disk set.

Suppose $\mathbb{C}_0$ is a flagged chamber complex supporting $S^3 = A \cup_T B$, and

$$\mathbb{C}_0 \xrightarrow{\mathcal{D}_0} \mathbb{C}_1 \xrightarrow{\mathcal{D}_1} \mathbb{C}_2$$

is a sequence of flagged chamber complex decompositions. Suppose E is a properly embedded disk in a chamber of $\mathbb{C}_0$ and E is disjoint from both $\mathcal{D}_0$ and $\mathcal{D}_1$. There are a variety of decomposition sequences that can be defined in this situation. Here is the notation we will use for them:

- $\mathbb{C}_0 \xrightarrow{\mathcal{D}_0} \mathbb{C}_1 \xrightarrow{E} \mathbb{C}_{0E}$
- $\mathbb{C}_1 \xrightarrow{E} \mathbb{C}_E \xrightarrow{\mathcal{D}'_1} \mathbb{C}_{E1}$
- $\mathbb{C}_0 \xrightarrow{\mathcal{D}_0 \cup E} \mathbb{C}_{0+E}$
- $\mathbb{C}_1 \xrightarrow{E \cup \mathcal{D}_1} \mathbb{C}_{E+1}$

In the second of these sequences $\mathcal{D}'_1$ has the usual meaning: those disks in $\mathcal{D}_1$ whose boundaries lie on the defining surface of $\mathbb{C}_E$ (and so do not lie on the boundary of goneballs after decomposition by E). Note that if ∂E lies on the boundary of a goneball in $\hat{\mathbb{C}}_1$, then it does not lie on the defining surface for $\mathbb{C}_1$, so decomposing on a disk set containing E is the same as decomposing on the disk set without E . So in this situation $\mathbb{C}_{E1} = \mathbb{C}_{E+1} = \mathbb{C}_2$.

Lemma 14.1. *Suppose neither $\mathbb{C}_0$ nor $\mathbb{C}_1$ certify. Then, up to a choice of siblings, if siblings exist, either*

1. $\mathbb{C}_{0E} = \mathbb{C}_{0+E}$ and $\mathbb{C}_{E1} = \mathbb{C}_{E+1}$ or
2. $\mathbb{C}_{0E} \twoheadrightarrow \mathbb{C}_{0+E}$ and $\mathbb{C}_{E1} = \mathbb{C}_{E+1}$ or
3. $\mathbb{C}_{0E} = \mathbb{C}_{0+E}$ and $\mathbb{C}_{E1} \twoheadrightarrow \mathbb{C}_{E+1}$.

Proof. Let F_0, F_1 be the defining surfaces for respectively $\mathbb{C}_0, \mathbb{C}_1$. Suppose first that ∂E is separating in a component of $F_0 - \partial \mathcal{D}_0$. Then it is separating in a component of F_1 , because F_1 is obtained from $F_0 - \partial \mathcal{D}_1$ simply by adding (scar) disks to $\partial \mathcal{D}_1$ and perhaps deleting some spheres. Then ∂E (if it is not on the boundary of a goneball) is also separating in a component of $F_1 - \partial \mathcal{D}_1$. It follows from the last sentence of Proposition 12.12 that $\mathbb{C}_{E1} = \mathbb{C}_{E+1}$ and, further applying Proposition 12.12, that outcome (1) or (2) of Lemma 14.1 occurs.

On the other hand, if ∂E is non-separating in $F_0 - \partial \mathcal{D}_0$ then by Lemma 12.11 $\mathbb{C}_{0E} = \mathbb{C}_{0+E}$. Then by Proposition 12.12, outcome (1) or (3) occurs. $\square$

We now put this in the context of decomposition sequences. Suppose $\overline{\mathcal{D}}_1, \dots, \overline{\mathcal{D}}_n$ are a sequence of guiding disk sets in S^3 so that $\overline{\mathcal{D}}_1$ determines the decomposition $\mathbb{C}_1 \xrightarrow{E \cup \overline{\mathcal{D}}_1} \mathbb{C}_{E+1}$ above. Then the guiding disk sets $\overline{\mathcal{D}}_i$ determine extensions of the three sequences that we label as follows:

$$\begin{aligned} \vec{\mathbb{C}}^{0+E} : \mathbb{C}_0 &\xrightarrow{\mathcal{D}_0 \cup E} \mathbb{C}_{0+E} \xrightarrow{\overline{\mathcal{D}}_1} \mathbb{C}_2^{0+E} \xrightarrow{\overline{\mathcal{D}}_2} \mathbb{C}_3^{0+E} \xrightarrow{\overline{\mathcal{D}}_3} \dots \xrightarrow{\overline{\mathcal{D}}_{n-1}} \mathbb{C}_n^{0+E} \\ \vec{\mathbb{C}}^{E1} : \mathbb{C}_0 &\xrightarrow{\mathcal{D}_0} \mathbb{C}_1 \xrightarrow{E} \mathbb{C}_{0E} \xrightarrow{\overline{\mathcal{D}}_1} \mathbb{C}_2^{0E} \xrightarrow{\overline{\mathcal{D}}_2} \mathbb{C}_3^{0E} \xrightarrow{\overline{\mathcal{D}}_3} \dots \xrightarrow{\overline{\mathcal{D}}_{n-1}} \mathbb{C}_n^{0E} \\ \vec{\mathbb{C}}^{E+1} : \mathbb{C}_0 &\xrightarrow{\mathcal{D}_0} \mathbb{C}_1 \xrightarrow{E \cup \overline{\mathcal{D}}_1} \mathbb{C}_{E+1} \xrightarrow{\overline{\mathcal{D}}_2} \mathbb{C}_3^{E+1} \xrightarrow{\overline{\mathcal{D}}_3} \dots \xrightarrow{\overline{\mathcal{D}}_{n-1}} \mathbb{C}_n^{E+1} \end{aligned}$$

Lemma 14.2. *Suppose all three decomposition sequences certify. Then all three decomposition sequences cocertify.*

Proof. If $\mathbb{C}_0$ certifies, this is obvious: $\mathbb{C}_0$ is a term in all three sequences, so if it certifies then all three sequences cocertify.

If $\mathbb{C}_1$ certifies, then $\vec{\mathbb{C}}^{E+1}$ and $\vec{\mathbb{C}}^{E1}$ cocertify for the same reason: $\mathbb{C}_1$ is in each sequence. If further $\mathbb{C}_0$ does not certify, we can apply Proposition 12.12 to conclude that either $\mathbb{C}_{0E} = \mathbb{C}_{0+E}$ or $\mathbb{C}_{0E} \dashrightarrow \mathbb{C}_{0+E}$. In the former case the sequences $\vec{\mathbb{C}}^{E1}$ and $\vec{\mathbb{C}}^{0+E}$ also cocertify because, starting with $\mathbb{C}_{0E} = \mathbb{C}_{0+E}$, they are the same sequence and before that in $\vec{\mathbb{C}}^{0+E}$ is only $\mathbb{C}_0$, which does not certify. In the latter case, when $\mathbb{C}_{0E} \dashrightarrow \mathbb{C}_{0+E}$, they cocertify by Corollary 11.12.

Suppose neither $\mathbb{C}_0$ nor $\mathbb{C}_1$ certify. Then following Lemma 14.1 the other decompositions can be arranged towards the left in exactly one of the following three diagrams, which differ only in which of two possible deflations (shown as vertical dashed arrows) are in fact equalities:

$$\begin{array}{ccccc} \mathbb{C}_0 & \xrightarrow{\mathcal{D}_0} & \mathbb{C}_1 & \xrightarrow{\overline{\mathcal{D}}_1} & \mathbb{C}_{E+1} \xrightarrow{\overline{\mathcal{D}}_2} \mathbb{C}_3^{E+1} \xrightarrow{\overline{\mathcal{D}}_3} \dots & = \vec{\mathbb{C}}^{E+1} \\ & \searrow \mathcal{D}_0 \cup E & \downarrow E & & \downarrow = & \\ & & \mathbb{C}_{0E} & \xrightarrow{\overline{\mathcal{D}}_1} & \mathbb{C}_{E1} \xrightarrow{\overline{\mathcal{D}}_2} \dots & = \vec{\mathbb{C}}^{E1} \\ & & \downarrow = & & & \\ & & \mathbb{C}_{0+E} & \xrightarrow{\overline{\mathcal{D}}_1} & \mathbb{C}_1^{0+E} \xrightarrow{\overline{\mathcal{D}}_2} \dots & = \vec{\mathbb{C}}^{0+E} \end{array} \quad (14.1.1)$$

$$\begin{array}{ccccc} \mathbb{C}_0 & \xrightarrow{\mathcal{D}_0} & \mathbb{C}_1 & \xrightarrow{\overline{\mathcal{D}}_1} & \mathbb{C}_{E+1} \xrightarrow{\overline{\mathcal{D}}_2} \mathbb{C}_3^{E+1} \xrightarrow{\overline{\mathcal{D}}_3} \dots & = \vec{\mathbb{C}}^{E+1} \\ & \searrow \mathcal{D}_0 \cup E & \downarrow E & & \downarrow = & \\ & & \mathbb{C}_{0E} & \xrightarrow{\overline{\mathcal{D}}_1} & \mathbb{C}_{E1} \xrightarrow{\overline{\mathcal{D}}_2} \dots & = \vec{\mathbb{C}}^{E1} \\ & & \downarrow \text{---} & & & \\ & & \mathbb{C}_{0+E} & \xrightarrow{\overline{\mathcal{D}}_1} & \mathbb{C}_1^{0+E} \xrightarrow{\overline{\mathcal{D}}_2} \dots & = \vec{\mathbb{C}}^{0+E} \end{array} \quad (14.1.1)$$

$$\begin{array}{ccccccc}
 \mathbb{C}_0 & \xrightarrow{\mathcal{D}_0} & \mathbb{C}_1 & \xrightarrow{\overline{\mathcal{D}}_1} & \mathbb{C}_{E+1} & \xrightarrow{\overline{\mathcal{D}}_2} & \mathbb{C}_3^{E+1} \xrightarrow{\overline{\mathcal{D}}_3} \dots & = \check{\mathbb{C}}^{E+1} \\
 & \searrow \mathcal{D}_0 \cup E & \downarrow E & & \uparrow & & & \\
 & & \mathbb{C}_{0E} & \xrightarrow{\overline{\mathcal{D}}_1} & \mathbb{C}_{E1} & \xrightarrow{\overline{\mathcal{D}}_2} & \dots & = \check{\mathbb{C}}^{E1} \\
 & & \downarrow = & & \downarrow & & & \\
 & & \mathbb{C}_{0+E} & \xrightarrow{\overline{\mathcal{D}}_1} & \mathbb{C}_1^{0+E} & \xrightarrow{\overline{\mathcal{D}}_2} & \dots & = \check{\mathbb{C}}^{0+E}
 \end{array} \tag{14.1.1}$$

In the first diagram it is obvious that all three sequences cocertify: because of the equalities, all three sequences are in fact the same sequence after the terms $\mathbb{C}_{E+1}$, $\mathbb{C}_{E1}$ and $\mathbb{C}_1^{0+E}$. Similarly the sequences $\check{\mathbb{C}}^{E+1}$ and $\check{\mathbb{C}}^{E1}$ become the same sequence after those terms in the second diagram and so cocertify; the same is true (even one step earlier) for $\check{\mathbb{C}}^{E1}$ and $\check{\mathbb{C}}^{0+E}$ in the third diagram.

Finally, by Corollary 11.12, both $\check{\mathbb{C}}^{E1}$ and $\check{\mathbb{C}}^{0+E}$ cocertify in the second diagram and $\check{\mathbb{C}}^{E+1}$ and $\check{\mathbb{C}}^{E1}$ cocertify in the third. $\square$

15 Ghost circles and timing of disks

Let $\mathbb{C}$ be a flagged chamber complex in S^3 supporting $A \cup_T B$, and $S \subset S^3$ be a sphere transverse to $\mathbb{C}$. It was shown in Section 9 how S defines a flagged chamber complex decomposition sequence for $\mathbb{C}$. We briefly recall the process, filling in further description, assisted by augmented notation: For F the defining surface for $\mathbb{C}$, the innermost in S circles c_0 of $F \cap S$ bound disks in S which we regard as a disk set $\mathcal{D}_0$ for $\mathbb{C}$. The set of circles $\bar{c}_1$ in $F \cap S$ (denoted simply c_1 in Section 9) consists of those circles which bound disks in S whose interiors may contain circles in c_0 but otherwise are disjoint from F . Call this set of disks $\overline{\mathcal{D}}_1$. Proceeding in this manner partition the entire collection of circles $F \cap S$ into sets $\bar{c}_0 = c_0, \bar{c}_1, \dots, \bar{c}_n$, where $n = \lfloor \frac{\text{diam}(Y)+1}{2} \rfloor$, so that each circle in $\bar{c}_k$ bounds a disk D in S , and any circle of $F \cap S$ that lies in the interior of D lies in some $\bar{c}_i, i < k$. Denote the collection of such disks in S bounded by the circles in $\bar{c}_k$ by $\mathcal{D}_k$.

The decomposition sequence

$$\check{\mathbb{C}}: \quad \mathbb{C} = \mathbb{C}_0 \xrightarrow{\mathcal{D}_0} \mathbb{C}_1 \xrightarrow{\mathcal{D}_1} \mathbb{C}_2 \xrightarrow{\mathcal{D}_2} \dots \xrightarrow{\mathcal{D}_{n-1}} \mathbb{C}_n$$

is then iteratively defined, the first step being the decomposition $\mathbb{C} = \mathbb{C}_0 \xrightarrow{\mathcal{D}_0} \mathbb{C}_1$. After this decomposition two important things happen to the sets of circles $\bar{c}_i, 1 \leq i \leq n$:

- The circles c_0 have been surgered away, so the interior of each disk in $\overline{\mathcal{D}}_1$ is disjoint from the defining surface F_1 for $\mathbb{C}_1$.
- Some of the circles in $\bar{c}_1$ may lie on the goneballs of the decomposition, and so no longer lie in $S \cap F_1$.

Then let $\mathcal{D}_1 \subset \overline{\mathcal{D}}_1$ be the subset of disks that are left, i. e. the subset of disks whose boundaries lie on F_1 .

Since, after the decomposition by $\mathcal{D}_0$, the interiors of the disks $\mathcal{D}_1$ are disjoint from F_1 , they are a disk set in $\mathbb{C}_1$. Use these disks to define the decomposition $\mathbb{C}_1 \xrightarrow{\mathcal{D}_1} \mathbb{C}_2$. Continue in this manner to get the entire decomposition sequence, noting that each $\mathcal{D}_i$ in the construction consists of those disks in $\overline{\mathcal{D}}_i$ whose boundaries lie on the defining surface $F_i = F(\mathbb{C}_i)$, and not on the boundaries of goneballs from earlier decompositions in the sequence. In other words:

Lemma 15.1. *For each $1 \leq i \leq n-1$, the disk set $\overline{\mathcal{D}}_i$ is a guiding disk set for the decomposition $\mathbb{C}_i \xrightarrow{\mathcal{D}_i} \mathbb{C}_{i+1}$.* $\square$

Definition 15.2. *For each $1 \leq i \leq n-1$, the circles $\bar{c}_i - c_i$ (that is $\partial\overline{\mathcal{D}}_i - \partial\mathcal{D}_i$) are called the ghost circles of such a decomposition sequence in $(\overline{\mathbb{C}}, S)$. The circles $c_i = F_i \cap S$, that is the circles that bound the disks $\mathcal{D}_i \subset S$, will be called the F -circles of the decomposition sequence.*

The disk set $\mathcal{D}_i$ may well contain ghost circles in its interior; these indicate where $\mathcal{D}_i$ intersects spheres bounding goneballs in earlier decompositions, and so can be ignored. Nevertheless, the ghost circles do play a role in the construction of the decomposition sequence, namely they are part of what determines, at the beginning of the decomposition sequence, before they are ghost circles, where in the sequence of decompositions a particular sub-disk of S appears.


Following up on this idea, it is reasonable to consider how introducing ghost circles in S at the beginning of the decomposition process (when the sets of circles c_i for the decomposition sequence are being defined) might affect the decomposition sequence. For example, if a ghost circle c' were to be added, at the beginning of the decomposition sequence, parallel to a circle $c \in c_{i-1}$ and just inside the disk $D \in \mathcal{D}_{i-1}$ that c bounds, one effect would be to redefine D as a disk in $\mathcal{D}_i$, because c is no longer i -th innermost in $F \cap S$ but $(i+1)$ -st innermost in the collection of circles $c \cup c'$. There could well be a similar delay to the use of other disks in the following $\mathcal{D}_k, i < k < n$ in a manner that may be a bit difficult to describe - see Figure 40. Our next goal is to understand to some extent how adding or removing ghost circles (and thereby changing the location of the decomposing disks in the sequence of disk decompositions) affects the certification of a decomposition sequence. The news is good: eventually we will show that adding or removing ghost circles may alter the decomposition sequence, as we have seen, but the new decomposition sequence cocertifies with the old.

To this end we generalize the discussion of Section 9 to spheres in S transverse to the defining surface F of a flagged chamber complex $\mathbb{C}$ in S^3 and consider a finite collection of circles $\bar{c}$ in S containing all of the circles $c = F \cap S$ (which we have distinguished as F -circles) but augmented by other possible circles $\bar{c} - F$ in S called ghost circles. The collection of circles $\bar{c}$ determines a sequence of flagged chamber complex decompositions guided by disk sets $\overline{\mathcal{D}}_i$, where each circle in $\bar{c}$ is the boundary of a disk in some $\overline{\mathcal{D}}_i$, as described above and in Section 9. (This setting naturally arises in the middle of a flagged chamber complex decomposition sequence, as we have seen.)

As in Section 9, let Y be the tree associated to the circles $\bar{c} \subset S$, and, as in Definition 9.4, let

$$\vec{\mathbb{C}}: \quad \mathbb{C} = \mathbb{C}_0 \xrightarrow{\mathcal{D}_0} \mathbb{C}_1 \xrightarrow{\mathcal{D}_1} \mathbb{C}_2 \xrightarrow{\mathcal{D}_2} \dots \xrightarrow{\mathcal{D}_{n-1}} \mathbb{C}_n$$

be the complete sequence of flagged chamber complex decompositions guided by the sphere S , where $n = \lfloor \frac{\text{diam}(Y)+1}{2} \rfloor$. It's perfectly possible that for some values of i , $\bar{c}_i$ consists entirely of ghost circles, so


 Figure 40: Adding a ghost circle, $i = 3$

none of the disks $\overline{\mathcal{D}_i}$ have their boundaries on F and so $\mathcal{D}_i = \emptyset$. In this case the decomposition $\mathbb{C}_i \xrightarrow{\mathcal{D}_i} \mathbb{C}_{i+1}$ is just the identity: $\mathbb{C}_i = \mathbb{C}_{i+1}$.

Recalling Definition 9.1, each edge e in Y (corresponding to a circle $c \in \bar{\mathcal{C}}$) is assigned a value

$$0 \leq \rho_Y(e) \leq \lfloor \frac{\text{diam}(Y) - 1}{2} \rfloor.$$

For each vertex $v \in Y$ we can similarly assign $\rho_Y(v) = \max\{\rho_Y(e) \mid e \text{ an edge incident to } v\}$.

Here are some elementary observations. For e an edge in Y , let $m_{\pm}^e$ be the numbers defined for e in Definition 9.1.

Lemma 15.3. *There is an edge f in Y such that $m_+^f = m_-^f$ if and only if $\text{diam}(Y)$ is odd. In this case $\rho_Y(f) = \frac{\text{diam}(Y)-1}{2}$ and f is unique: it is the mid-edge of each path in Y whose length is $\text{diam}(Y)$. Call f the center edge of the tree Y . See Figure 40.*

Proof. Let e be any edge in Y , with endpoints v_+ and v_- belonging to the trees Y_+ and Y_- respectively in $Y - e$, as discussed in Definition 9.1. If there is a leaf v_ℓ of Y that belongs to Y_+ and is a distance $m > \frac{\text{diam}(Y)-1}{2}$ from v_+ , then v_ℓ is a distance $m + 1 > \frac{\text{diam}(Y)+1}{2}$ from v_- . It follows from the definition of diameter that no leaf of Y that belongs to Y_- can be a distance more than

$$\text{diam}(Y) - m - 1 = \frac{\text{diam}(Y) - 1}{2} + \frac{\text{diam}(Y) + 1}{2} - m - 1 < \frac{\text{diam}(Y) - 1}{2}$$

from v_- . Hence $m_- \leq \frac{\text{diam}(Y)-1}{2} < m \leq m_+$ and e is not as f is described in the Lemma.

Thus an edge f as described in the Lemma has the property that any leaf of Y in $Y_{\pm}$ has distance at most $\frac{\text{diam}(Y)-1}{2}$ from respectively the vertices $v_{\pm}$. This implies that any path in $Y_{\pm}$ can have length at most $\text{diam}(Y) - 1$, so in particular no path γ in Y of length $\text{diam}(Y)$ can lie entirely in either of $Y_{\pm}$. This implies that f must be in γ and in fact be the mid-edge of γ , defining f uniquely and establishing that $\rho_Y(f) = m_{\pm}^f = \frac{\text{diam}(Y)-1}{2}$.

Conversely, suppose $\text{diam}(Y)$ is odd and f is the mid-edge of a path γ in Y of length $\text{diam}(Y)$. If any leaf in Y lying in Y_+ (resp Y_-) is a distance greater than $\frac{\text{diam}(Y)-1}{2}$ from v_+ (resp v_-) then conjoining a path from that leaf to v_+ with the part of $\gamma - v_+$ containing v_- would give a path of length greater than $\text{diam}(Y)$, a contradiction. Hence both $m_{\pm}^f = \frac{\text{diam}(Y)-1}{2}$ as required. $\square$

Quite similarly we have:

Lemma 15.4. *There is a vertex w in Y that is incident to more than one edge e with $\rho_Y(e) = \rho_Y(w)$ if and only if $\text{diam}(Y)$ is even. In this case w is unique and $\rho_Y(w) = \frac{\text{diam}(Y)}{2} - 1$. Call w the center vertex of the tree Y . See Figure 40.*

Proof. Consider any vertex v in Y and suppose there is a path γ in Y from v to a leaf of Y and that the length of γ is $m > \frac{\text{diam}(Y)}{2}$. Let e be the edge of γ that is adjacent to v and observe that, by definition of diameter, any path from v to a leaf of Y that is longer than $\text{diam}(Y) - m < m$ must pass through e . Definition 9.1 then implies that $\rho_Y(e)$ is the length of the longest path in Y not passing through e from v to a leaf of Y , and any edge e' incident to v other than e has $\rho_Y(e') \leq \rho_Y(e) - 1$. In particular, e is the only edge incident to v with $\rho_Y(e) = \rho_Y(v)$, so e is not an edge as described in the lemma.

It follows that if w is as described in the lemma, then the length of any path from w to a leaf in Y has length $\leq \frac{\text{diam}(Y)}{2}$. So, by definition of diameter, any path in Y whose length is $\text{diam}(Y)$ must pass through w and have its length bisected by w . Thus $\text{diam}(Y)$ is even and the distance from w to a most distant leaf will be $\frac{\text{diam}(Y)}{2}$. In particular, any edge e incident to w will find the vertex at its other end a distance of at most $\frac{\text{diam}(Y)}{2} - 1$ from a leaf. Thus $\rho_Y(e) \leq \frac{\text{diam}(Y)}{2} - 1$. Since this is true for all edges incident to w , $\rho_Y(w) \leq \frac{\text{diam}(Y)}{2} - 1$. Furthermore, since any path in Y whose length is $\text{diam}(Y)$ must pass through w , there are at least two edges $e_{\pm}$ incident to w lying on such a path, and then $\rho_Y(e_{\pm}) = \frac{\text{diam}(Y)}{2} - 1$. It follows that for both edges $\rho_Y(e_{\pm}) = \frac{\text{diam}(Y)}{2} - 1 = \rho_Y(w)$ as required.

Conversely, if $\text{diam}(Y)$ is even, γ is a path in Y of length $\text{diam}(Y)$, and w is the vertex that bisects γ , the same argument shows that for e either of the two edges in γ incident to w , $\rho_Y(e) = \rho_Y(w)$. $\square$

Lemma 15.5. *Suppose $P \subset S - \bar{c}$ is the planar surface corresponding to a vertex $v \in Y$. If $P \subset D \in \overline{\mathcal{D}}_k$, then $\rho_Y(v) \leq k$.*

Proof. For $D \in \overline{\mathcal{D}}_k$, $\partial D \in \bar{c}_k$ and from Lemma 9.2 each circle in $\bar{c} \cap \text{int}(D)$ lies in some $\bar{c}_j$, $j < k$. Hence if $P \subset D$, each circle $c \in \partial P$ lies in some $\bar{c}_j$, $j \leq k$. Corresponding to this, we have that each edge e in Y that is incident to v has $\rho_Y(e) \leq k$. Hence by definition $\rho_Y(v) \leq k$ as required. $\square$

Lemma 15.6. *Suppose a ghost circle is added to S in a planar surface $P \subset S - \bar{c}$ corresponding to a vertex $v \in Y$ and*

$$\vec{c}' : \quad \mathbb{C} = \mathbb{C}_0 \xrightarrow{\mathcal{D}'_0} \mathbb{C}'_1 \xrightarrow{\mathcal{D}'_1} \mathbb{C}'_2 \xrightarrow{\mathcal{D}'_2} \dots$$

is the resulting sequence of flagged chamber complex decompositions. Suppose $\rho_Y(v) \geq 1$. Then for each $1 \leq i \leq \rho_Y(v)$, $\mathcal{D}'_{i-1} = \mathcal{D}_{i-1}$ and $\mathbb{C}'_i = \mathbb{C}_i$.

Proof. If $i \leq \rho_Y(v)$ then $\rho_Y(v) \not\leq i-1$. It follows from Lemma 15.5 that any disk $D \in \overline{\mathcal{D}}_{i-1}$ does not contain P and so is unaffected by the addition of the ghost circle. Thus, for each $1 \leq i \leq \rho_Y(v)$, $\overline{\mathcal{D}}'_{i-1} = \overline{\mathcal{D}}_{i-1}$. Now argue iteratively: Since $\mathcal{C}'_0 = \mathcal{C}_0$ and $\overline{\mathcal{D}}'_0 = \overline{\mathcal{D}}_0$ it follows that $\mathcal{D}'_0 = \mathcal{D}_0$. Hence $\mathbb{C}'_1 = \mathbb{C}_1$, so if $\overline{\mathcal{D}}'_1 = \overline{\mathcal{D}}_1$ then $\mathcal{D}'_1 = \mathcal{D}_1$. The argument continues iteratively until $i = \rho_Y(v)$: For each $1 \leq i \leq \rho_Y(v)$, $\mathbb{C}'_{i-1} = \mathbb{C}_{i-1}$ and $\overline{\mathcal{D}}'_{i-1} = \overline{\mathcal{D}}_{i-1}$ imply $\mathcal{D}'_{i-1} = \mathcal{D}_{i-1}$, so $\mathbb{C}'_i = \mathbb{C}_i$, as required. $\square$

Corollary 15.7. Suppose $\text{diam}(Y)$ is even and a ghost circle is added to S in the planar surface $P \subset S - \bar{c}$ corresponding to the center vertex $v \in Y$. Then $\vec{\mathbb{C}}' = \vec{\mathbb{C}}$.

Proof. Since $\text{diam}(Y)$ is even, the complete sequence $\vec{\mathbb{C}}$ has last term $\mathbb{C}_n$ with $n = \text{diam}(Y)/2 = \rho_Y(v)$. If $\rho_Y(v) = 0$ then Y is a star graph, each circle in $\bar{c}$ bounds a unique disk in $\overline{\mathcal{D}}_0$ and both $\vec{\mathbb{C}}'$ and $\vec{\mathbb{C}}$ are simply the sequence $\mathbb{C}_0 \xrightarrow{\mathcal{D}_0} \mathbb{C}_1$. Indeed, after this decomposition the only circle remaining in S is the added ghost circle, so $S \cap F(\mathbb{C}_1) = \emptyset$ and there is no further decomposition in the sequence. If, on the other hand, $\rho_Y(v) \geq 1$ then by Lemma 15.6, for all $0 \leq i \leq n$, $\mathbb{C}'_i = \mathbb{C}_i$ and after these decompositions, the only circle that might remain is again the added ghost circle. $\square$

Corollary 15.8. Suppose $\text{diam}(Y)$ is odd and a ghost circle is added to S in a planar surface $P \subset S - \bar{c}$ corresponding to a vertex at the end of the center edge $e \in Y$. Then if $\vec{\mathbb{C}}$ certifies, $\vec{\mathbb{C}}'$ and $\vec{\mathbb{C}}$ cocertify.

Proof. One possibility is that $\text{diam}(Y) = 1$, that is Y is the single edge e . In this case $\bar{c}$ is a single circle c (corresponding to e) dividing S into two disks D_A and D_B , as in Lemma 10.4. Say the ghost circle is added to $D_A = P$. If c is also a ghost circle, then F is disjoint from S and there is no further decomposition. In particular, $\vec{\mathbb{C}}' = \vec{\mathbb{C}}$. If, on the other hand, c is an F -circle, so $c = F \cap S$ then it is shown in Lemma 10.4 that if $\vec{\mathbb{C}}$ certifies, the choice of whether to decompose $\mathbb{C}$ along D_A or D_B makes no difference: the two results cocertify. Adding a ghost circle to D_A has the sole effect of declaring that in $\mathbb{C}'$ the decomposition will be along D_B (since D_A is no longer a disk component of $S - \bar{c}$). Hence $\vec{\mathbb{C}}'$ and $\vec{\mathbb{C}}$ cocertify.

If, on the other hand, $\text{diam}(Y) \geq 3$ then, per Lemma 15.3 $\rho_Y(v) = \rho_Y(e) = \frac{\text{diam}(Y)-1}{2} \geq 1$. Then according to Lemma 15.6 the decomposition sequences $\vec{\mathbb{C}}$ and $\vec{\mathbb{C}}'$ are identical through $\mathbb{C}'_{\frac{\text{diam}(Y)-1}{2}} = \mathbb{C}_{\frac{\text{diam}(Y)-1}{2}}$. But for $\mathbb{C}_{\frac{\text{diam}(Y)-1}{2}}$, $\bar{c}$ has shrunk to a single circle, and Y is a single edge. Thus we can apply the argument just given for this case and conclude that $\vec{\mathbb{C}}' \sim \vec{\mathbb{C}}$ as required. $\square$

Proposition 15.9. Suppose $\mathbb{C}$ is a flagged chamber complex in S^3 that is not tiny. Let $S \subset S^3$ be a sphere, transverse to the defining surface $F = F(\mathbb{C})$, that is either balanced or almost balanced. Let $\bar{c}$ be a collection of circles in S containing the set $c = F \cap S$ and let $\bar{c}'$ be the union of $\bar{c}$ and a disjoint (ghost) circle c' in S . Let $\vec{\mathbb{C}}$ and $\vec{\mathbb{C}}'$ be the sequences of flagged chamber complex decompositions determined by $\bar{c}$ and $\bar{c}'$ respectively. Then $\vec{\mathbb{C}}$ and $\vec{\mathbb{C}}'$ cocertify.

Proof. Let Y be the tree associated to $\bar{c}$, $P \subset S - \bar{c}$ be the planar surface component to which the ghost circle c' is added, and v be the vertex in Y corresponding to P . The proof is by induction on the distance between v and the center vertex or edge of Y . Corollaries 15.7 and 15.8 show that the proposition is true when that distance is 0.

Following Lemma 15.3, as used in the proof of Corollary 15.7, the decomposition sequences are identical until the tree has been shrunk to the point that v is a leaf. So we may as well assume that v is a leaf of Y , so P is an innermost disk of $S - \bar{c}$. Let $\bar{c}''$ be the collection of circles $\bar{c} \cup c''$, where c'' is a circle parallel to ∂P but just *outside* P , that is in the component P'' of $S - \bar{c}$ that is adjacent to P . See Figure 41. Then the vertex v'' corresponding to P'' is closer to the center edge or vertex of Y so, by inductive assumption, the sequence of flagged chamber complex decompositions $\vec{\mathbb{C}}''$ determined by $\bar{c}''$ and $\vec{\mathbb{C}}$ cocertify. So all that remains is to show that $\vec{\mathbb{C}}''$ and $\vec{\mathbb{C}}'$ cocertify.




Figure 41: Moving a ghost circle, so $\partial P \in c'_3 \rightarrow \partial P \in c''_2$

If ∂P is a ghost circle, then $\bar{c}'' = \bar{c}'$ so in fact $\vec{\mathbb{C}}'' = \vec{\mathbb{C}}'$ and we are done. So assume that ∂P is an F -circle in $\bar{c}$ and note that the only difference between $\bar{c}''$ and $\bar{c}'$ is the side on which a ghost circle is added. This switch in sides has no effect on the tree and so no effect on the disks in the decomposition sequence $\vec{\mathbb{C}}'$ except this: In $\vec{\mathbb{C}}''$ the disk P is in $\mathcal{D}_0$, since P is an innermost disk, whereas in $\vec{\mathbb{C}}'$ the disk P is in $\mathcal{D}_1$ because the disk bounded by ∂P contains the single added ghost circle c' .

Apply Lemma 14.2 to this situation, where here the disk P plays the role of E in that lemma, $\vec{\mathbb{C}}''$ here corresponds to $\vec{\mathbb{C}}^{0+E}$ there and $\vec{\mathbb{C}}'$ here corresponds to $\vec{\mathbb{C}}^{E+1}$ there. Since S is balanced or almost balanced, both sequences $\vec{\mathbb{C}}''$ and $\vec{\mathbb{C}}'$ certify, as does the sequence $\vec{\mathbb{C}}^{P1}$, defined for P as $\vec{\mathbb{C}}^{E1}$ is defined for E in the preamble to Lemma 14.2. The conclusion of Lemma 14.2 then says that $\vec{\mathbb{C}}''$ and $\vec{\mathbb{C}}'$ cocertify, so $\vec{\mathbb{C}} \sim \vec{\mathbb{C}}'' \sim \vec{\mathbb{C}}'$, as required. $\square$

Corollary 15.10. Suppose $\mathbb{C}$ is a flagged chamber complex in S^3 that is not tiny and $S \subset S^3$ is a sphere transverse to $F = F(\mathbb{C})$ which is either balanced or almost balanced for $\mathbb{C}$. Let $\bar{c}$ and $\bar{c}'$ be collections

of circles in S , each containing the set $\mathfrak{c} = F \cap S$. Let $\vec{\mathbb{C}}$ and $\vec{\mathbb{C}}'$ be the sequences of flagged chamber complex decompositions determined by $\vec{\mathfrak{c}}$ and $\vec{\mathfrak{c}}'$ respectively. Then $\vec{\mathbb{C}}$ and $\vec{\mathbb{C}}'$ cocertify.

Proof. It suffices to prove the Corollary in the case that $\vec{\mathfrak{c}} = \mathfrak{c}$. In this case $\vec{\mathfrak{c}}'$ is obtained from $\vec{\mathfrak{c}}$ by adding some number m of ghost circles. But this case follows by applying Proposition 15.9 m times. $\square$

Corollary 15.11. *Theorem 13.3 remains true even when the chamber C containing E is an empty torus, or when E is inessential.*

Proof. Referring to the proof of Theorem 13.3, the requirement that E be essential and C not be an empty torus is to ensure that the original decomposition $C = C_0 \xrightarrow{E} C'$ creates no goneballs. The boundaries of such goneballs would intersect $F = F(\mathbb{C})$ in a family of circles $\vec{\mathfrak{c}}_G$ in $F \cap S$ that continue to determine which circles in $F \cap S$ bound disks in each $\mathcal{D}_i$ during the subsequent decomposition

$$\vec{\mathbb{C}}' : \quad \mathbb{C}' = \mathbb{C}'_0 \xrightarrow{\mathcal{D}'_0} \mathbb{C}'_1 \xrightarrow{\mathcal{D}'_1} \mathbb{C}'_2 \xrightarrow{\mathcal{D}'_2} \dots \xrightarrow{\mathcal{D}'_{n-1}} \mathbb{C}'_n \quad (13.1)$$

used in the proof of Theorem 13.3.

In contrast, it is the collection of circles $F(\mathbb{C}') \cap S$ that determines a decomposition sequence in $(\vec{\mathbb{C}}', S)$ and the goneballs of the decomposition $C = C_0 \xrightarrow{E} C'$ do not appear in the definition of $F' = F(\mathbb{C}')$. In other words, $F \cap S = (F' \cap S) \cup \vec{\mathfrak{c}}_G$, so $\vec{\mathfrak{c}}_G$ acts as a collection of ghost circles in the decomposition sequence $\vec{\mathbb{C}}'$. Corollary 15.10 says that the addition of such ghost circles makes no difference in certification. That is, we may as well assume that $F \cap S = F' \cap S$ and proceed with the proof of Theorem 13.3 as written. $\square$

16 Isotopies of guiding spheres

Definition 16.1. *Suppose $\mathbb{C}$ is a flagged chamber complex in S^3 and S, S' are spheres transverse to the defining surface $F = F(\mathbb{C})$. If any $\vec{\mathbb{C}} \in (\vec{\mathbb{C}}, S)$ and $\vec{\mathbb{C}}' \in (\vec{\mathbb{C}}, S')$ cocertify, then we say $(\vec{\mathbb{C}}, S)$ and $(\vec{\mathbb{C}}, S')$ cocertify and write $(\vec{\mathbb{C}}, S) \sim (\vec{\mathbb{C}}, S')$.*


Our goal in this section is to prove:

Theorem 16.2. *Suppose $\mathbb{C}$ is a flagged chamber complex in S^3 that is not tiny and supports the Heegaard splitting $S^3 = A \cup_T B$. Suppose further that S_0 and S_1 are balanced or almost balanced spheres for $\mathbb{C}$ that are isotopic in S^3 through balanced or almost balanced spheres $S_t, 0 \leq t \leq 1$. Then $(\vec{\mathbb{C}}, S_0) \sim (\vec{\mathbb{C}}, S_1)$.*

As usual, let F denote the defining surface $F = F(\mathbb{C})$ and put the isotopy in general position so that generically S_t is transverse to F , but for certain discrete values of t , S_t is tangent to F at a single point, a non-degenerate critical point of index zero (F has a local max at v in a collar of S), two (a local min) or one (a saddle tangency).

Lemma 16.3. *Theorem 16.2 is true if the isotopy S_t passes through only a single critical point, of index zero or two (a max or a min).*

Proof. The cases of index zero or index two are symmetric. Let $S_{m(iddle)}$ be the sphere on which the critical point v occurs, and, in a collar of S_m , we may suppose F has a minimum. Call the bottom of the collar $S_{b(elow)}$ and the top $S_{a(bove)}$. We can assume that other than the local minimum, F intersects the collar in vertical cylinders. Then $F \cap S_a$ is the union of $F \cap S_b$ and a single circle c , where c bounds a disk D_F in $F - S_a$ (the disk in F containing v) and a parallel disk D_S in $S_a - F$. Let $\vec{\mathbb{C}}_b$ (resp. $\vec{\mathbb{C}}_a$) be the sequence of flagged chamber complex decompositions determined by $\mathbb{C}$ and S_b (resp. S_a). Thus $\vec{\mathbb{C}}_b \in (\vec{\mathbb{C}}, S_1)$ (say) and $\vec{\mathbb{C}}_a \in (\vec{\mathbb{C}}, S_0)$. See Figure 42. Augment the collection of circles $F \cap S_b$ with the single (ghost) circle c' parallel in S_b to $c \subset S_a$ and call the resulting sphere with ghost circle S_c . Let $\vec{\mathbb{C}}_c$ be a sequence of flagged chamber complex decompositions determined by $\mathbb{C}$ and S_c . Then by Corollary 15.10 $\vec{\mathbb{C}}_b \sim \vec{\mathbb{C}}_c$.


 Figure 42: Passing S through a minimum

Now compare $\vec{\mathbb{C}}_a$ with $\vec{\mathbb{C}}_c$. The pattern of circles in $F \cap S_a$ is the same as in $F \cap S_c$ except that the F -circle c in S_a appears as a ghost circle c' in S_c . Thus the decomposition $\mathbb{C}_0 \xrightarrow{\mathcal{D}_0} \mathbb{C}_1$ is the same in each sequence, except in $\vec{\mathbb{C}}_c$ the ghost circle c' is just removed, whereas in $\vec{\mathbb{C}}_a$, F is surgered by D_S , creating an extra ball chamber B_G bounded by the sphere $D_S \cup D_F$. But our convention for flagging declares the handlebody B_G to be empty and therefore a goneball, so the result of the flagged chamber complex decomposition $\mathbb{C}_1$ is the same in both cases, as is the remaining pattern of circles $S_a \cap F(\mathbb{C}_1) = S_c \cap F(\mathbb{C}_1)$. It follows that the entire ensuing flagged chamber complex decomposition sequence is the same, so $\vec{\mathbb{C}}_a \sim \vec{\mathbb{C}}_c$. Combining, we have $\vec{\mathbb{C}}_a \sim \vec{\mathbb{C}}_c \sim \vec{\mathbb{C}}_b$ so $(\vec{\mathbb{C}}, S_0) \sim (\vec{\mathbb{C}}, S_1)$ as required. $\square$

Next suppose the isotopy S_t has a single critical point that is a saddle tangency with F at a point $v \in F$. As in the proof of Lemma 16.3 we consider how F intersects a collar of S_m , the sphere of the tangency at

v , as shown in Figure 43. A neighborhood in F of $S_m \cap F$ (a figure eight) is a pair of pants with three disjoint circles: c_b , whose parallel in S_m is incident to v in two points, and c_1, c_2 each of whose parallels in S_m is incident to v in a single point. With S_a and S_b as in Lemma 16.3, we can assume that c_1 and c_2 both lie in S_a , and c lies in S_b .




Figure 43: A neighborhood of critical point v : two viewpoints

As in Section 9 let Y_a (resp Y_b) be the tree defined by the collection of circles $F \cap S_a$ (resp $F \cap S_b$), e_1, e_2 be the edges in Y_a corresponding to c_1, c_2 and e_b be the edge in Y_b corresponding to c_b . As defined in 9.1, let $\rho_i = \rho_{Y_a}(e_i), i = 1, 2$ and $\rho_b = \rho_{Y_b}(e_b)$. Put another way, in the decomposition sequences of $\mathbb{C}$ guided by S_a and S_b , each of the two circles $c_i, i = 1, 2$ bounds a disk $D_i \in \mathcal{D}_{\rho_i}$ in the decomposition $\mathbb{C}_{\rho_i} \xrightarrow{\mathcal{D}_{\rho_i}} \mathbb{C}_{\rho_{i+1}}$ and c_b bounds a disk $D_b \in \mathcal{D}_{\rho_b}$ in the decomposition $\mathbb{C}_{\rho_b} \xrightarrow{\mathcal{D}_{\rho_b}} \mathbb{C}_{\rho_{b+1}}$. With no loss of generality, assume $\rho_1 \leq \rho_2$ (i. e. $\mathcal{D}_{\rho_1}$ does not appear after $\mathcal{D}_{\rho_2}$ in the decomposition sequence). There are two possibilities for the disks $D_1, D_2 \subset S_a$, and the two viewpoints in Figure 43 are chosen so that in either case each D_i is on the bounded side of c_i in S_a : either D_1 and D_2 are disjoint, as shown in the left side of the figure, or $D_1 \subset D_2$ as shown in the right side. That is, the two viewpoints in the figure are chosen so that in each the D_i appear as bounded regions. Similarly, D_b can either lie on the side of c_b that contains v (the bounded side in the left of the figure and unbounded side in the right) or the other side. We have:

Lemma 16.4. *If D_b lies on the side of c_b that contains v , then $\rho_1 \leq \rho_b$.*

Proof. Recall that, aside from the component that contains v , F intersects the collar of S only in vertical annuli. Hence the interior of the bounded side of c_b on the left of Figure 43 (or the the interior of the unbounded side of c_b on the right) contains a parallel copy of each circle of $F \cap \text{int}(D_1)$. $\square$

Lemma 16.5. *Theorem 16.2 is true if the isotopy S_t passes through only a single critical point, of index one (a saddle tangency).*

Proof. We make two claims:

Claim 1: With no loss of generality, we can assume that $\rho_1 \leq \rho_b$.

Proof of Claim 1: By Lemma 16.4 this is true when D_b lies on the side of c_b that contains v , so suppose D_b lies on the other side. That is, suppose D_b lies on the outside (unbounded side) of c_b in the left of Figure 43 or on the inside (bounded side) of c_b on the right. Suppose $\rho_b < \rho_1$ and consider what happens if we add $\rho_1 - \rho_b$ ghost circles to D_b , parallel to ∂D_b . This raises ρ_b so that $\rho_b = \rho_1$. Since there is a copy of the pattern of circles $F \cap D_1$ on the other side of c_b in S_b , that is in $S_b - D_b$, (see left panel in Figure 44), it remains true that the maximal distance in Y_b from e_b to a leaf of $Y_{\pm}$ is still minimized in D_b , so we do not change the side of c_b on which D_b lies. Let S_b^+ denote S_b with these ghost circles added, and denote by $\vec{\mathbb{C}}_b^+$ be the sequence of flagged chamber complex decompositions on $\mathbb{C}$ determined by S_b^+ . By Corollary 15.10 it suffices to show that $\vec{\mathbb{C}}_b^+$ and $\vec{\mathbb{C}}_a$ cocertify. This proves Claim 1.




Figure 44: When D_b lies on the side of c_b not containing v

Since, per the claim, $\rho_1 \leq \rho_b$ and by definition $\rho_1 \leq \rho_2$ there is no difference between the decomposition sequences for $\vec{\mathbb{C}}_a$ and $\vec{\mathbb{C}}_b$ until we reach the decomposition $\mathbb{C}_{\rho_1} \xrightarrow{D_{\rho_1}} \mathbb{C}_{\rho_1+1}$. $F_{\rho_1} = F(\mathbb{C}_{\rho_1})$ still appears as in Figure 43, but with the added information that $\text{int}(D_1)$ contains neither ghost circles nor F -circles. In particular, it is possible to define a new chamber complex $\mathbb{C}'$ obtained from $\mathbb{C}_{\rho_1}$ by decomposing along D_1 alone. That is

$$\mathbb{C}_{\rho_1} \xrightarrow{D_1} \mathbb{C}'.$$

Claim 2: Both S_a and S_b are still balanced or almost balanced for $\mathbb{C}'$.

Proof of Claim 2: Let F' be the defining surface $F(\mathbb{C}')$. See Figure 44. We name 4 regions and 6 genera:

- X is the region below S_b and g_X is the genus of $F_{\rho_1} \cap X = F' \cap X$.
- $X^+ \supset X$ is the region below S_a , g_X^+ is the genus of $F_{\rho_1} \cap X^+$, and $g_X'^+$ is the genus of $F' \cap X^+$.
- Y is the region above S_a , g_Y is the genus of $F_{\rho_1} \cap Y$, and g_Y' is the genus of $F' \cap Y$.
- $Y^+ \supset Y$ is the region above S_b , and $g_Y'^+$ is the genus of $F' \cap Y^+$.

Surgery on $D_1 \subset S_a$ will not affect the genus of the part of F_{ρ_1} lying below or above S_a , so

(0) $g_X'^+ = g_X^+$, $g_Y' = g_Y$ and S_a is balanced or almost balanced for $\mathbb{C}'$.

To prove the claim, we then need only focus on S_b . Here are two further observations:

1. Since $F_{\rho_1} \cap X^+$ is obtained from $F_{\rho_1} \cap X$ by attaching a (genus 0) pair of pants on a single boundary component, $g_X^+ = g_X$.
2. Since $F' \cap Y^+$ is obtained from $F' \cap Y$ by attaching a (genus 0) annulus on a single boundary component, $g_Y'^+ = g_Y'$.

By definition S_b is balanced or almost balanced for $\mathbb{C}'$ if and only if the pair of numbers $\{g_X, g_Y'^+\} \neq \{0, k\}, k \geq 2$. Indeed, by definition, the pair is $\{0, 0\}$ if and only if S_b is planar balanced, it is $\{0, 1\}$ if and only if S_b is almost balanced, and the pair is $\{j, k\}, j, k \geq 1$ if and only if S_b is non-planar balanced. Similarly, since S_a is balanced or almost balanced for $\mathbb{C}_{\rho_1}$, we have $\{g_X^+, g_Y\} \neq \{0, k\}, k \geq 2$. From (1) we have $g_X = g_X^+$ and combining (0) and (2) gives $g_Y'^+ = g_Y$. In particular, $\{g_X, g_Y'^+\} = \{g_X^+, g_Y\} \neq \{0, k\}, k \geq 2$, completing the proof of Claim 2.




Figure 45: Only annuli between the planes

F' intersects the collar between S_a and S_b entirely in spanning annuli (see Figure 45 and right panel of Figure 44), so the decomposition sequence on $\mathbb{C}'$ determined by S_a and S_b is the same; we denote it $\vec{\mathbb{C}}'$. As in the proof of Lemma 16.3 let $\vec{\mathbb{C}}_b$ (resp $\vec{\mathbb{C}}_a$) be the sequence of flagged chamber complex decompositions determined by $\mathbb{C}$ and S_b (resp. S_a). Thus $\vec{\mathbb{C}}_b \in \overrightarrow{(\mathbb{C}, S_1)}$ (say) and $\vec{\mathbb{C}}_a \in \overrightarrow{(\mathbb{C}, S_0)}$. Now apply, for both S_a and S_b , Theorem 13.3 as augmented by Corollary 15.11 to conclude

$$\vec{\mathbb{C}}_a \sim \vec{\mathbb{C}}' \sim \vec{\mathbb{C}}_b \quad \text{so} \quad \overrightarrow{(\mathbb{C}, S_0)} \sim \overrightarrow{(\mathbb{C}, S_1)}$$

□

Proof of Theorem 16.2. As noted before Lemma 16.3 a generic isotopy S_t from S_0 to S_1 consists of a sequence of isotopies, each containing a single critical point and hence one to which either Lemma 16.3 or Lemma 16.5 applies. □

17 All balanced or almost balanced spheres cocertify

In this section we show that the requirement in Theorem 16.2 that S_0 and S_1 are isotopic through balanced or almost balanced spheres is superfluous. In particular we show:

Proposition 17.1. *Suppose $\mathbb{C}$ is a flagged chamber complex in S^3 . Then any pair of balanced or almost balanced spheres for $\mathbb{C}$ are isotopic in S^3 through balanced or almost balanced spheres for $\mathbb{C}$.*

We defer the proof in order to quickly observe some consequences: Following Theorem 16.2 this immediately implies:

Corollary 17.2. *Suppose $\mathbb{C}$ is a flagged chamber complex in S^3 that supports the genus g Heegaard splitting (S^3, T) and is not tiny. Let S^x and S^y be balanced or almost balanced spheres for $\mathbb{C}$. Then $\overrightarrow{(\mathbb{C}, S^x)} \sim \overrightarrow{(\mathbb{C}, S^y)}$.* $\square$

In particular, any such flagged chamber complex $\mathbb{C}$ in S^3 determines a unique certificate. That is, the homeomorphism $h_{(\mathbb{C}, S)} : (S^3, T) \rightarrow (S^3, T_g)$, for S a balanced or almost balanced sphere (provided, say, by Corollary 10.7), does not in fact depend on the choice of S , up to eyeglass equivalence. We will denote (a choice of) such a homeomorphism simply $h_{\mathbb{C}} : (S^3, T) \rightarrow (S^3, T_g)$. This expands the use of the notation, first introduced following Corollary 8.3, beyond just those flagged chamber complexes that themselves certify, to all flagged chamber complexes that are not tiny. Following Proposition 5.18, it applies to flagged chamber complexes obtained from a Heegaard surface by weak reduction. It also allows the following natural definition and notation:

Definition 17.3. *Suppose $\mathbb{C}, \mathbb{C}'$ are flagged chamber complexes in S^3 that support the genus g Heegaard splitting (S^3, T) and are not tiny. Then $\mathbb{C}$ and $\mathbb{C}'$ cocertify (written $\mathbb{C} \sim \mathbb{C}'$) if the homeomorphisms $h_{\mathbb{C}}, h_{\mathbb{C}'} : (S^3, T) \rightarrow (S^3, T_g)$ (as just defined above) are eyeglass equivalent.*

Corollary 17.4. *With this expanded definition of $h_{\mathbb{C}}$, Corollary 8.4 remains true. That is, if $\mathbb{C}$ certifies and $\tau \in G(S^3, T)$ then $h_{\tau(\mathbb{C})} \tau \sim h_{\mathbb{C}}$.*

Proof. This follows immediately from Corollary 9.6. $\square$

The proof of Proposition 17.1 is brief (it appears just before Proposition 17.7) but the background needed is extensive. Since it makes use of elementary Morse and Cerf Theory [Mi], [Ce], we will work in the smooth category. That is, we will take the defining surface $F = F(\mathbb{C})$ to be a smooth submanifold of S^3 , and the pair of spheres S^x and S^y , to be smooth submanifolds, smoothly transverse to F . Let $p : S^3 \rightarrow [-1, 1]$ be the standard height function, denote by S^2 the standard equator $p^{-1}(0)$, and call the points $p^{-1}(1), p^{-1}(-1)$ the north and south poles of S^3 . Denote by B_{ϵ} a pair of small ball neighborhoods of the poles, that is $B_{\epsilon} = p^{-1}([-1, -1 + \epsilon] \cup [1 - \epsilon, 1])$.

Choose a pair of points q_n, q_s (to correspond to north and south poles) disjoint from $F \cup S^x \cup S^y$ so that q_n and q_s lie on opposite sides of each of the spheres S^x and S^y . This is easily done: If S^x and S^y are disjoint, choose a generic point in the interior of each of the two ball components of $S^3 - (S^x \cup S^y)$. If, on the other hand, S^x and S^y intersect, observe that a regular neighborhood of a curve c of intersection intersects $S^x \cup S^y$ in a copy of $X \times c$. (Here X means two crossed lines, e. g. the x - and y - axes in $\mathbb{R}^2$.) Then set q_n, q_s to be generic points in opposite quadrants of some $X \times \{point\}$.

By the Schoenflies Theorem there are diffeomorphisms $\phi^x, \phi^y : (S^3; q_n, q_s) \rightarrow (S^3; p^{-1}(1), p^{-1}(-1))$ so that $\phi^x(S^x) = \phi^y(S^y) = S^2 \subset S^3$. Choose a small neighborhood B_ϵ of the poles in S^3 that is disjoint from $\phi^x(F) \cup \phi^y(F)$ and adjust ϕ^x, ϕ^y near q_n, q_s so that p_x and p_y coincide on a neighborhood $U = (p^x)^{-1}(B_\epsilon) = (p^y)^{-1}(B_\epsilon)$ of q_n, q_s . It is elementary in this case (or use [Lau] and the fact that the space of orientation preserving diffeomorphisms of the ball is connected, indeed contractible) that there is an isotopy $\theta : (S^3, U) \times I \rightarrow (S^3, B_\epsilon)$ between ϕ^x and ϕ^y , fixed on U . That is, for $\theta_t : (S^3, U) \rightarrow (S^3, B_\epsilon)$ defined by $\theta_t(z) = \theta(z, t)$, we have $\theta_0 = \phi^x$ and $\theta_1 = \phi^y$ and for all t , and each $u \in U$, $\theta_t(u) = \phi^x(u) = \phi^y(u)$.

Now consider the function $p^x : F \rightarrow [-1, 1]$ defined by $p^x = p\phi^x|_F$ and similarly $p^y = p\phi^y|_F$. The functions p^x, p^y are homotopic via the homotopy $p_t = p\theta_t|_F$. Morse and Cerf theory tell us that for a generic copy of $F \subset S^3$ arbitrarily near the original F (a copy which we henceforth take as F), the functions p^x, p^y are Morse functions and the homotopy p_t between them is Cerf. Here are the salient points of what is meant by this:

p^x (and p^y) are Morse: The function $p^x : F \rightarrow [-1, 1]$ has only a finite number of critical values $\{s_i\} \subset (-1, 1)$ and each critical value is the image of a single non-degenerate critical point in F . Non-degenerate means that the Hessian of the function is non-singular (see [Mi]). Put another way, each sphere $S_s^x = (p\phi^x)^{-1}(s)$ is transverse to F for each regular value of s , and, for each critical value s_i , $S_{s_i}^x$ is transverse to F except at a single point of tangency, where F intersects a neighborhood of the tangency point in $S_{s_i}^x$ as either a maximum, a minimum, or a saddle point.


The homotopy p_t is Cerf: For generic t the function $p_t : F \rightarrow [-1, 1]$ is Morse. At a finite set $\tau = \{t_j\} \subset I$ the function p_t is not Morse, solely because either

- near a single exceptional critical point on $z \in F$, p_t has a "birth-death" singularity given by the local model $p_t(z_1, z_2) = p_t(z) + z_1^3 \pm tz_1 \pm z_2^2$ or
- there are two non-degenerate critical points of p_t on F with the same critical value

Beyond this, and perhaps clarifying it, is a description of the "Cerf graphic" in the square $(t, s) \in I \times [-1, 1]$ (see, for example, introductory remarks in [GK], from which Figure 46 is taken): There are a finite collection Γ of curves in $I \times [-1, 1]$ so that each $\gamma \in \Gamma$ is the graph of a smooth function to $[-1, 1]$ on a closed interval in I whose end points are among the set $\tau \cup \partial I$. See green in Figure 46. Each generic vertical arc $t \times [-1, 1]$ is transverse to the graphs Γ and intersects each curve $\gamma \in \Gamma$ in at most one point. Each vertical arc $t_i \times [-1, 1]$, $t_i \in \tau$ has the same property, with a single exceptional point $s \in [-1, 1]$ where either

- (t_i, s) is a left end point of exactly two such graphs, a birth point. (So $t_i \times [-1, 1] \in$ red vertical lines in Figure 46)
- (t_i, s) is a right end point of exactly two such graphs, a death point. (So $t_i \times [-1, 1] \in$ red vertical lines in Figure 46)
- At (t_i, s) two graphs cross transversally. (So $t_i \times [-1, 1] \in$ blue vertical lines in Figure 46)

The collection Γ of graphs constitute all points where the functions p_t have critical points. That is, $(t, s) \in I \times [-1, 1]$ lies in Γ if and only if s is a critical value for p_t . In particular, each complementary


 Figure 46: Graphic (green) with $t_i \in \tau$

component R of Γ in $I \times [-1, 1]$ (called a *region* of the graphic) has the property that for any $(t, s) \in R$, s is a regular value for p_t .

Given this as background, observe that the natural sweep-out of $S^3 - \{\text{poles}\}$ by the 2-spheres $S_s = p^{-1}(s)$, $s \in (-1, 1)$, in which S_0 is the equator S^2 induces a t -parameterized smooth family of sweep-outs $S_s^t = \theta_t^{-1}(S_s)$ beginning with a sweep-out $(\phi^x)^{-1}(S_s)$ that contains $S^x = (\phi^x)^{-1}(S^2)$ and ending with a sweep-out $(\phi^y)^{-1}(S_s)$ that contains $S^y = (\phi^y)^{-1}(S^2)$. Taking this sweep-out point of view, we can now incorporate the ideas of Section 10, in particular Definition 10.5 and Lemma 10.6, as we now describe.

Suppose $(t, s) \in R$, where R is a region of the graphic. Since s is a regular value of p_t , the sphere S_s^t is transverse to F , and, since R is connected, any other sphere $S_{s'}^{t'}$, $(t', s') \in R$ is isotopic to S_s^t through spheres transverse to F . Following Definition 10.5 we can then assign a label $\frac{a_R}{b_R}$ to R , where the non-negative integer a_R is the genus of the part of F lying above S_s^t and similarly b_R is the genus of the part of F lying below S_s^t . Here are some elementary observations:

1. S_s^t is balanced or almost balanced for $\mathbb{C}$ if one of these is true:

- $0 \leq a_R, b_R \leq 1$
- Both $a_R, b_R \geq 1$

In this case we say that the region itself is balanced or almost balanced.

2. Hence S_s^t is *not* balanced or almost balanced for $\mathbb{C}$ (say S_s^t is *akilter*) if one of these is true:

- $a_R = 0$ and $b_R \geq 2$ (say then *+akilter*) or
- $b_R = 0$ and $a_R \geq 2$ (say then *-akilter*).

In this case we say that the region itself is respectively *+akilter* or *-akilter*.

3. Suppose R and R' are adjacent regions in the graphic, both incident to a birth singularity (t_0, s_0) . That is, the regions are separated near (t_0, s_0) by the pair of curves in Γ that end at (t_0, s_0) . Then a

generic sweep-out to the right of (t_0, s_0) (that is via $S_s^{(t_0+\epsilon)}$, $s \in (-1, 1)$), passing briefly through R' , say, differs from a sweep-out just to the left of (t_0, s_0) , passing through R , by the brief introduction of a minimum (or maximum), followed by a cancelling saddle tangency. The consequent brief transfer of a disk in F from one side of S_s^t to the other has no effect on the genus of F on each side, so $\frac{a_R}{b_R} = \frac{a_{R'}}{b_{R'}}$.

4. The same is true for R and R' adjacent regions in the graphic, both incident to a death singularity (t_0, s_0) .

We now interpret Lemma 10.6 in this context.

Lemma 17.5. *Suppose, in the setting described above, $\text{genus}(F) \geq 2$ and t_0 is a generic value of t , that is $t_0 \in I - \tau$. The vertical arc $v = t_0 \times [-1, 1]$ has the following properties:*


1. *There is a unique curve $\gamma_- \subset \Gamma$, intersecting v in a point (t_0, s_-) , so that the region just below γ_- is $-akilter$ and the region just above γ_- is balanced or almost balanced.*
2. *There is a unique curve $\gamma_+ \subset \Gamma$ intersecting v in a point (t_0, s_+) , so that the region just below γ_+ is balanced or almost balanced and the region just above γ_+ is $+akilter$.*
3. *$s_- < s_+$ and the subinterval $t_0 \times [s_-, s_+]$ of v intersects at least three different regions of the graphic.*

Proof. We will apply Lemma 10.6 to the sweepout of F by the spheres $S_s^{t_0}$. (The lemma is applicable because it is equivalent to considering the sweep-out of $\theta_{t_0}(F)$ by the spheres S_s .) Note first that the bottom regions of the graphic (those incident to $I \times \{-1\}$) are all labeled $\frac{\text{genus}(F)}{0}$ and the top regions are all labeled $\frac{0}{\text{genus}(F)}$. Since $\text{genus}(F) \geq 2$ the regions at the bottom are $-akilter$ and those at the top are $+akilter$.

As noted in Lemma 10.6 as s rises in v , passing from region to region, a_R cannot rise and b_R cannot fall. To progress under these rules from the label $\frac{a}{0}, a \geq 2$ to the label $\frac{0}{b}, b \geq 2$ there is a first transition from a region R that is $-akilter$ to a region R' that is balanced or almost balanced, either because $\frac{a_R}{b_R} = \frac{2}{0}$ and $\frac{a_{R'}}{b_{R'}} = \frac{1}{0}$ or because $\frac{a_R}{b_R} = \frac{a}{0}, a \geq 2$ and $\frac{a_{R'}}{b_{R'}} = \frac{a}{1}$. Moreover, once that transition is made, there is no way to transition back to $-akilter$ without either lowering the denominator in $\frac{a}{1}$ or raising the numerator in $\frac{1}{0}$, neither of which is allowed. So this transition defines the curve γ_- and so s_- . The symmetric argument identifies the curve γ_+ and shows that it lies above γ_- , that is $s_+ > s_-$. See Figure 47.

More subtly, notice that according to Lemma 10.6 as s passes upward through a curve γ from region R to region R' the difference $a_R - b_R \in \mathbb{Z}$ does not increase (for this would require either a_R to rise or b_R to fall) and can decrease by at most one (because at most one of these occurs: a_R drops by 1 or b_R rises by 1). That is, $(a_{R'} - b_{R'}) \leq (a_R - b_R) \leq (a_{R'} - b_{R'}) + 1$. So it takes at least 4 transitions to go from a $-akilter$ region, where $a_R - b_R \geq 2$, to a $+akilter$ region, where $a_R - b_R \leq -2$. So, not counting the end points, there are at least two graphs in Γ that cross the interior of $t_0 \times [s_-, s_+]$ and so 3 distinct regions through which $t_0 \times [s_-, s_+]$ passes. $\square$

Lemma 17.6. *Suppose, in the setting described above, $\text{genus}(F) \geq 2$. Then there are continuous (indeed, piecewise smooth) functions $f_-, f_+ : I \rightarrow [-1, 1]$ so that*


 Figure 47: Spanning arcs γ_+ , γ_- and the gap between

- For all t , $f_-(t) < f_+(t)$.
- The graphs of both $f_{\pm}$ lie in Γ .
- Regions below the graph of f_- are all $-akilter$.
- Regions above the graph of f_+ are all $+akilter$.
- Regions between the two graphs are all balanced or almost balanced.
- $f_-(0) < 0 < f_+(0)$ and $f_-(1) < 0 < f_+(1)$.

Proof. For $t_0 \notin \tau$ define $f_{\pm}(t_0) = s_{\pm}$ as given in Lemma 17.5. The functions $f_{\pm}$ are smooth near t_0 because each curve $\gamma \in \Gamma$ is smooth. The same argument applies for $t_0 \in \tau$ with some (t_0, s) a birth-death point since, as we have seen above, $\frac{a_R}{b_R}$ does not change as s passes through such a point.

The situation is only a bit more complicated for $t_0 \in \tau$ with two curves in Γ crossing at (exactly one) point (t_0, s_0) in $t_0 \times [-1, 1]$. In this case, the region below (t_0, s_0) could be $-akilter$ and the region above (t_0, s_0) either balanced or almost balanced, and then it is natural to define $f_-(t_0) = s_0$. By examining generic nearby values of t we see that the interval $t_0 \times (s_0, s_+)$ must still pass through at least two distinct regions, so $s_0 < s_+$ and $f_-(t_0) < f_+(t_0)$ as required. The graph of f_- may not be smooth at t_0 since values of $f_-(t)$ for t just less than t_0 may lie on one of the two crossing curves in Γ and values of $f_-(t)$ for t just more than t_0 lie on the other. Still the graph of f_- is piecewise smooth at t_0 . Symmetric comments apply when the region above (t_0, s_0) is $+akilter$ and the region below (t_0, s_0) either balanced or almost balanced.

Finally, our initial definition was that the balanced or almost balanced sphere $S^x = S_0^0$ so $f_-(0) < 0 < f_+(0)$ and the balanced or almost balanced sphere $S^y = S_0^1$ so $f_-(1) < 0 < f_+(1)$. $\square$

Proof of Proposition 17.1. If $\text{genus}(F) \leq 1$ let $f : [0, 1] \rightarrow [-1, 1]$ be any smooth function whose graph is transverse to the curves Γ in the graphic and for which $f(0) = f(1) = 0$. If $\text{genus}(F) \geq 2$ pick such a function so that, in addition, $f_- < f < f_+$. Then the 1-parameter family of spheres $S_{f(t)}^t, 0 \leq t \leq 1$ is an

isotopy of S^x to S^y . Moreover each sphere $S_{f(t)}^t$ is balanced or almost balanced for $\mathbb{C}$, since the graph of f passes only through balanced or almost balanced regions of the graphic. $\square$

There is a related proposition about single sweep-outs, whose proof is similar to that of Lemma 17.5:

Proposition 17.7. *Suppose $\mathbb{C}$ is a flagged chamber complex in S^3 that is not tiny and $S_s, s \in [-1, 1]$ is a generic sweep-out of S^3 by spheres. Suppose, E is a properly embedded disk in a chamber of $\mathbb{C}$ so that, for some generic $-1 < e < 1$, ∂E lies entirely in S_e . That is, ∂E is a circle component of $S_e \cap F(\mathbb{C})$. Put E in preferred alignment and let $\mathbb{C}'$ be the flagged chamber complex obtained from $\mathbb{C}$ by decomposition along E . Then $\mathbb{C}' \sim \mathbb{C}$.*

Proof. Since S_s is a generic sweep-out, there is a finite set $\sigma \subset [-1, 1]$ so that for $s \notin \sigma$, S_s is transverse to $F = F(\mathbb{C})$, and for each $s \in \sigma$, F and S_s have a single non-degenerate point of tangency. Apply the notational convention for regions introduced before Lemma 17.5 to the (interval) components of $[-1, 1] - \sigma$. That is, for each such interval $i \subset [-1, 1] - \sigma$ pick $s \in i$ and let $a_i \geq 0$ be the genus of the part of F lying above S_s and b_i be the genus of the part of F lying below S_s , and then assign to i the symbol $\frac{a_i}{b_i}$. Following Lemma 10.6, and similar to the argument used for a generic value of t in the proof of Lemma 17.5, we have:

- The lowest interval in $[-1, 1] - \sigma$ is assigned $\frac{\text{genus}(F)}{0}$
- The highest interval $[-1, 1] - \sigma$ is assigned $\frac{0}{\text{genus}(F)}$
- As s rises through each critical level $s \in \sigma$ the integer $a_i - b_i$ does not increase and can decrease by at most 1.

Since, as we ascend through the interval components of $[-1, 1] - \sigma$, the difference $a_i - b_i$ begins at $\text{genus}(F)$ and ends at $-\text{genus}(F)$ it follows that there is an interval i_0 for which the difference is zero, that is $a_{i_0} = b_{i_0}$. In particular, for $s \in i_0$, S_s is balanced for $\mathbb{C}$, either planar balanced (if $a_{i_0} = b_{i_0} = 0$) or non-planar balanced (if $a_{i_0} = b_{i_0} \geq 1$). Pick a generic $s_0 \in i_0$ ($s_0 \neq e$) so that S_{s_0} is transverse to E .

Claim: The sphere S_{s_0} is balanced or almost balanced for $\mathbb{C}'$ as well as for $\mathbb{C}$.

Proof of Claim:. Define the pair of whole numbers a', b' in direct analogy to the pair a_{i_0}, b_{i_0} : Namely, a' is the genus of the part of F' lying above S_{s_0} and b' is the genus of the part of F' lying below S_{s_0} . We will now show by construction that there is a close relation between the pairs a', b' and a_i, b_i .

We can assume, with no loss, that $s_0 > e$. (If $s_0 < e$ reverse the rolls of a and b in the following argument.). Let $\mathfrak{c} \subset S_{s_0}$ be the collection of circles $S_{s_0} \cap E$. Decomposing $\mathbb{C}$ along E to obtain $F' = F(\mathbb{C}')$ can be viewed as a 3-stage process:

- first remove a small annular neighborhood $\eta(\partial E)$ of ∂E from F , then
- attach a pushed-off copy of E to each boundary component of $F - \eta(\partial E)$; call this added pair of disks $2E$. The resulting surface F_η intersects the level sphere S_{s_0} in the curves $(F \cap S_{s_0}) \cup 2\mathfrak{c}$, where $2\mathfrak{c}$ denotes two parallel copies of $\mathfrak{c}$, one copy for each copy of E in $2E$.
- Finally, remove any resulting sphere components of F_η that bound goneballs.

Observe the behavior of the genus through each of these three stages of the construction. We consider the stages in reverse order: Removing spheres from F_η does not change the genus of the part of the surface lying above S_{s_0} or below, since each component removed is a subsurface of spheres and is therefore planar. Similarly, adding $2E$ does not affect the genus of the surface above or below, since each component added is a subsurface of $2E$ and therefore planar, and the ends of the annulus $\eta(\partial E)$ are attached to different planar subsurfaces of $2E$. Finally, removing $\eta(\partial E)$ from F does not affect the genus of the part of F above S_{s_0} , and, though it may lower the genus of the part of F below S_{s_0} (if it is non-separating in the component in which it lies), it lowers it by at most one.

We conclude that $a' = a_{i_0}$ and $b_{i_0} - 1 \leq b' \leq b_{i_0}$. In particular

- If $\frac{a_{i_0}}{b_{i_0}} = \frac{0}{0}$ then $\frac{a'}{b'} = \frac{0}{0}$.
- If $\frac{a_{i_0}}{b_{i_0}} = \frac{1}{1}$ then $\frac{a'}{b'} = \frac{1}{1}$ or $\frac{1}{0}$.
- If $\frac{a_{i_0}}{b_{i_0}} = \frac{m}{m}, m \geq 2$ then $\frac{a'}{b'} = \frac{m}{m}$ or $\frac{m}{m-1}$.

In any case, we see that the sphere S_{s_0} is balanced or almost balanced for $\mathbb{C}'$, proving the Claim.

The proof now proceeds by induction on $|c|$, the number of curves of intersection of E with S_{s_0} . If E is disjoint from S_{s_0} , so $|c| = 0$ the result follows from Theorem 13.3, as augmented by Corollary 15.11.

Assume then that $|c| \geq 1$ and let

$$\overrightarrow{(\mathbb{C}, S_{s_0})}: \quad \mathbb{C} \xrightarrow{\mathcal{D}_0} \mathbb{C}_1 \xrightarrow{\mathcal{D}_1} \dots \xrightarrow{\mathcal{D}_{n-1}} \mathbb{C}_n$$

and

$$\overrightarrow{(\mathbb{C}', S_{s_0})}: \quad \mathbb{C}' \xrightarrow{\mathcal{D}'_0} \mathbb{C}'_1 \xrightarrow{\mathcal{D}'_1} \dots \xrightarrow{\mathcal{D}'_{n-1}} \mathbb{C}'_n$$

be complete flagged chamber complex decomposition sequences guided by S_{s_0} for $\mathbb{C}$ and $\mathbb{C}'$ respectively. (First add the circles c as ghost circles for the decomposition sequence for $\mathbb{C}$, so that both sequences use the same set of guiding disks throughout, and so run in parallel. By Corollary 15.10 the addition of these ghost circles does not change how $\mathbb{C}$ certifies.) Focus attention on the first decomposition $\mathbb{C}'_i \xrightarrow{\mathcal{D}'_i} \mathbb{C}'_{i+1}$ in which a circle in c appears as the boundary of a disk in $\mathcal{D}'_i$. Let $c_0 \in c$ be such a circle, chosen among all candidates to be innermost in $2E$; let $D_0 \in \mathcal{D}'_i$ be the disk that c_0 bounds in S_{s_0} ; and let $E_0 \subset 2E$ be the disk that c_0 bounds in $2E$.

Although E_0 may intersect S_{s_0} in other components of c , by construction none bound disks in any $\mathcal{D}'_j, j \leq i$, so E_0 remains intact as part of the defining surface $F'_i = F(\mathbb{C}'_i)$. Consider whether D_0 is an essential disk in the chamber of $\mathbb{C}'_i$ in which it lies:

If D_0 is an essential disk, then the ball $B_0 \subset S^3$ that the sphere $S = D_0 \cup E_0$ bounds (on the side in S^3 not containing ∂E) contains components of F'_i . By construction, B_0 is disjoint from $2E$ so these components of F'_i must also be components of the defining surface $F_i = F(\mathbb{C}_i)$. Thus S is an incompressible sphere in a chamber of both $\mathbb{C}'_i$ and $\mathbb{C}_i$. This implies that $\mathbb{C}'_i \sim \mathbb{C}_i$, hence $\overrightarrow{(\mathbb{C}, S_{s_0})} \sim \overrightarrow{(\mathbb{C}', S_{s_0})}$ and so, following Definition 17.3, $\mathbb{C} \sim \mathbb{C}'$ as required.

On the other hand, if D_0 is an inessential disk, then $\text{int}(B_0)$ is disjoint from F'_i so B_0 is diskly in $\hat{\mathbb{C}}'_{i+1}$ and so (perhaps with an appropriate choice of sibling) becomes a goneball in $\mathbb{C}'_{i+1}$. In other words, after

decomposition by $\mathcal{D}'_i$ (and the similar elimination of the parallel copy of E_0 in $2E$ at the next stage of the decomposition sequence) we obtain the same chamber complexes as if E_0 had been replaced by D_0 in the disk E . But this replacement lowers $|c|$ by at least 1. By our inductive assumption on $|c|$, $\mathbb{C} \sim \mathbb{C}'$, as required.

To provide more detail in this last case, when D_0 is inessential, let E' be the disk that results from replacing the disk $E_0 \subset E$ by D_0 as just described. Recall the following diagram for the inductive step in the proof of Proposition 13.2, which in turn enables the proof of Theorem 13.3:

$$\begin{array}{ccccccc}
 \dots & \xrightarrow{\mathcal{D}_{i-1}} & \mathbb{C}_i & \xrightarrow{\mathcal{D}_i} & \mathbb{C}_{i+1} & \xrightarrow{\mathcal{D}_{i+1}} & \dots & = \vec{\mathbb{C}} \\
 & & \downarrow E & & \downarrow E & & & \\
 & & \mathbb{C}_i^i & \xrightarrow{\mathcal{D}_i^i} & \mathbb{C}_{i+1}^i & \xrightarrow{\mathcal{D}_{i+1}^i} & \dots & = \vec{\mathbb{C}}_E^i \\
 & & \downarrow E & & \downarrow E & & & \\
 & & \mathbb{C}_i^{i+1} & \xrightarrow{\mathcal{D}_i^{i+1}} & \mathbb{C}_{i+1}^{i+1} & \xrightarrow{\mathcal{D}_{i+1}^{i+1}} & \dots & = \vec{\mathbb{C}}_E^{i+1}
 \end{array}$$

As the argument there would ultimately be applied here, the decomposing disks are all guided by S_{s_0} . (One consequence is that the notation for the decomposing disks in the diagram can all be simplified to those for the guiding disks in the top row.) The problem with further use of that inductive step here is that for $j > i$ the decomposing disks $\mathcal{D}_j$ might pass through E . This means that their interiors may intersect $F(\mathbb{C}'_j)$, making later inductive steps potentially nonsensical. However, since B_0 is a goneball in the decomposition $\mathbb{C}_i \xrightarrow{\mathcal{D}_i} \mathbb{C}'_{i+1}$ nothing is lost by replacing E with E' to get the simplified diagram:

$$\begin{array}{ccccccc}
 \dots & \xrightarrow{\mathcal{D}_{i-1}} & \mathbb{C}_i & \xrightarrow{\mathcal{D}_i} & \mathbb{C}_{i+1} & \xrightarrow{\mathcal{D}_{i+1}} & \dots & = \vec{\mathbb{C}} \\
 & & \downarrow E' & & & & & \\
 & & \mathbb{C}_i^i & \xrightarrow{\mathcal{D}_i^i} & \mathbb{C}_{i+1}^i & \xrightarrow{\mathcal{D}_{i+1}^i} & \dots & = \vec{\mathbb{C}}_E^i
 \end{array}$$

Then observe that since $|E' \cap S_{s_0}| < |E \cap S_{s_0}| = |c|$ the inductive hypothesis on $|c|$ implies that $\overrightarrow{(\mathbb{C}_i, S_{s_0})} \sim \overrightarrow{(\mathbb{C}_i^i, S_{s_0})}$. This eliminates the need for the further inductive steps used in the proof of Proposition 13.2. $\square$

Because Proposition 17.7 requires that $\mathbb{C}$ not be tiny, it cannot be directly applied to the flagged chamber complex defined by a Heegaard surface T , since both chambers are empty handlebodies. But it does tell us something crucial about flagged chamber complexes that are obtained from T by weak reduction:

Proposition 17.8. *Suppose $S^3 = A \cup_T B$ is a Heegaard splitting and $S_s, s \in [-1, 1]$ is a generic sweep-out of S^3 by spheres. Suppose $\mathcal{D}$ is a weakly reducing collection of disks for T and E is a properly embedded disk disjoint from $\mathcal{D}$, lying in either A or B . Suppose further that, for some generic $-1 < e < 1$, ∂E lies entirely in S_e . Let $\mathbb{C}$ be the flagged chamber complex obtained from T by decomposition along $\mathcal{D}$ and $\mathbb{C}'$ be the flagged chamber complex obtained from T by decomposition along $\mathcal{D} \cup E$. Then $\mathbb{C}' \sim \mathbb{C}$.*

Proof. Let $\mathbb{C}_{DE}$ denote the flagged chamber complex obtained from $\mathbb{C}$ by decomposition along E . Proposition 17.7 does apply to this decomposition since, by Proposition 5.18, $\mathbb{C}$ is not tiny. Thus $\mathbb{C} \sim \mathbb{C}_{DE}$. It remains to prove a similar relation between $\mathbb{C}'$ and $\mathbb{C}_{DE}$.

We are in a position to apply Proposition 12.12, where the Heegaard splitting here plays the role of $\mathbb{C}$ in Proposition 12.12; $\mathbb{C}'$ here corresponds to $\mathbb{C}_{D+E}$ there; and the notation $\mathbb{C}_{DE}$ here was chosen to be consistent with the same notation there. Then, according to Proposition 12.12, either $\mathbb{C}_{DE} = \mathbb{C}'$ or $\mathbb{C}_{DE} \twoheadrightarrow \mathbb{C}'$. In the former case we are done, so we consider the case $\mathbb{C}_{DE} \twoheadrightarrow \mathbb{C}'$, and turn to Corollary 11.12.

Following Corollary 10.7 let S_{s_0} be a sphere in the sweep-out that is balanced for $\mathbb{C}'$ and let

$$\vec{\mathbb{C}}' : \quad \mathbb{C}' \xrightarrow{\mathcal{D}_0} \mathbb{C}'_1 \xrightarrow{\mathcal{D}'_1} \dots \xrightarrow{\mathcal{D}'_{n-1}} \mathbb{C}'_n$$

be a complete flagged chamber complex decomposition sequence guided by S_{s_0} . Similarly, let $\vec{\mathbb{C}}_{DE}$ be a complete flagged chamber complex decomposition sequence for $\mathbb{C}_{DE}$ guided by S_{s_0} . Since $\mathbb{C}_{DE} \twoheadrightarrow \mathbb{C}'$ their defining surfaces F_{DE} and F' differ by at most the insertion of a bullseye. Since the components of the bullseye are all spheres, the insertion does not affect the genus of the parts of the surface above and below S_{s_0} , so S_{s_0} is balanced for $\mathbb{C}_{DE}$ as well. In particular, by Proposition 10.2, both $\vec{\mathbb{C}}'$ and $\vec{\mathbb{C}}_{DE}$ certify.

Apply Corollary 11.12 to $\mathbb{C}_{DE} \twoheadrightarrow \mathbb{C}'$: Let $\vec{\mathbb{C}}_{DE}^m$ be the maximal deflationary sequence of $\mathbb{C}_{DE}$. If $\vec{\mathbb{C}}_{DE}^m = \vec{\mathbb{C}}_{DE}$ so the entire sequence is a deflationary sequence, then by the second outcome of Corollary 11.12, $\vec{\mathbb{C}}'$ and $\vec{\mathbb{C}}_{DE}$ cocertify. If, on the other hand, $\vec{\mathbb{C}}_{DE}^m$ is not the entire sequence $\vec{\mathbb{C}}_{DE}$ then by the first outcome of Corollary 11.12, $\vec{\mathbb{C}}_{DE}^m$ certifies, so by the second outcome, again $\vec{\mathbb{C}}'$ and $\vec{\mathbb{C}}_{DE}$ cocertify. In other words, $(\overrightarrow{\mathbb{C}_{DE}, S_{s_0}}) \sim (\overrightarrow{\mathbb{C}', S_{s_0}})$ or, following Definition 17.3, $\mathbb{C}_{DE} \sim \mathbb{C}'$ as required. $\square$

18 The Goeritz group is the eyeglass group

Suppose (S^3, T_g) is the standard genus $g \geq 2$ Heegaard splitting of Section 3 and τ is an element of the Goeritz group $G(S^3, T_g)$, as described in [JM]. Let $T_\theta \subset S^3, 0 \leq \theta \leq 2\pi$ be a representative of τ in $\pi_1(\text{Img}(S^3, T))$, with $S^3 = A_\theta \cup_{T_\theta} B_\theta$ and $T_0 = T_{2\pi} = T_g$.

Put another way, τ is represented by an isotopy $\Theta_\theta : S^3 \rightarrow S^3, 0 \leq \theta \leq 2\pi$ with Θ_0 the identity and $\Theta_{2\pi}(T_g) = T_g$. T_θ then denotes the surface $\Theta_\theta(T_g) \subset S^3$ and $\Theta_{2\pi}$ represents $\tau : (S^3, T_g) \rightarrow (S^3, T_g)$. Similarly A_θ, B_θ denote respectively $\Theta_\theta(A), \Theta_\theta(B)$. Finally, for any $\theta, \theta' \in [0, 2\pi]$, denote the composition $\Theta_\theta \Theta_{\theta'}^{-1}$ by $\Theta_\theta^{\theta'} : (S^3, T_{\theta'}) \rightarrow (S^3, T_\theta)$. This definition implies that, for any $\theta'' \in [0, 2\pi]$, $\Theta_\theta^{\theta''} \Theta_{\theta''}^{\theta'} = \Theta_\theta^{\theta'}$.

We briefly review some of the results of [FS1], where S^3 is swept out by level spheres S_s of the standard height function $p : S^3 \rightarrow [-1, 1]$.

There are values $0 < \theta_1 < \theta_2 < \dots < \theta_n < 2\pi$ so that for each $\theta \notin \{\theta_i, 1 \leq i \leq n\}$ there is a pair of weakly reducing disks (a_θ, b_θ) associated to T_θ so that:

- $a_\theta \subset A_\theta$ and $b_\theta \subset B_\theta$
- The isotopy class of the pair (a_θ, b_θ) does not change throughout each interval in $[0, 2\pi] - \{\theta_i, 1 \leq i \leq n\}$. By this we mean, if θ, θ' are in the same interval of $[0, 2\pi] - \{\theta_i, 1 \leq i \leq n\}$ then the pair of disks $(\Theta_\theta^{\theta'}(a'_\theta), \Theta_\theta^{\theta'}(b'_\theta))$ is properly isotopic to the pair (a_θ, b_θ) in (S^3, T_θ) .

A shorthand notation for this is $(a_\theta, b_\theta) = (a_{\theta'}, b_{\theta'})$.

- $(a_{2\pi}, b_{2\pi}) = (a_0, b_0)$. That is, the pair of disks $(a_{2\pi}, b_{2\pi})$ is properly isotopic to the pair (a_0, b_0) in (S^3, T_g) .
- For every $0 \leq \theta \leq 2\pi$ each of the disks a_θ and b_θ have level boundaries. That is, each the circles ∂a_θ and ∂b_θ lies in a single sphere of the sweep-out by S_s .

The last property is used for the application of Proposition 17.8 in the proof of Lemma 18.2 below. It follows from the construction described in [FS1, Appendix], as applied in [FS1, Subsection 4.5]. The disks a_θ and b_θ are chosen in [FS1] by how their boundaries lie in $T_\theta \cap S$, for S one of the level spheres S_s . (The argument in [FS1] requires that $\text{genus}(T) \geq 2$.) Note that, unlike the collection of disks used in disk decomposition sequences that we have been long discussing, it is not part of the construction for the sweep-outs in [FS1] that the entire disk a_θ or b_θ lies in a single sphere of the sweep-out, only that the boundary of each disk does.

The third property above, that $(a_{2\pi}, b_{2\pi}) = (a_0, b_0)$ (the notation here differs somewhat from that in [FS1]) follows from the fact that $T_{2\pi} = T_0 = T_g$. Indeed, we might as well take the sweep out by copies of the Heegaard surface used in the construction to be the same in both cases, which means that the weakly reducing pair of disks arising from the construction will be given by the same rule.

For each $\theta \notin \{\theta_i, 1 \leq i \leq n\}$, define the flagged chamber complex $\mathbb{C}_\theta$ as that obtained from T_θ by weak reduction along the pair of disks (a_θ, b_θ) .

Observe some properties:

- By Proposition 5.18 no chamber complex $\mathbb{C}_\theta$ is tiny.
- Suppose θ, θ' are in the same interval of $[0, 2\pi] - \{\theta_i, 1 \leq i \leq n\}$. Since the pair of disks $(\Theta_\theta^{\theta'}(a'_\theta), \Theta_\theta^{\theta'}(b'_\theta))$ is properly isotopic to the pair (a_θ, b_θ) , the chamber complex $\Theta_\theta^{\theta'}(\mathbb{C}_{\theta'})$ is also properly isotopic to the chamber complex $\mathbb{C}_\theta$ in (S^3, T_θ) .
- Since $(a_{2\pi}, b_{2\pi}) = (a_0, b_0)$ we have $\mathbb{C}_{2\pi} = \mathbb{C}_0$.

Finally, for each $\theta \notin \{\theta_i, 1 \leq i \leq n\}$ define $h_\theta : (S^3, T_\theta) \rightarrow (S^3, T_g)$ to be $h_{\mathbb{C}_\theta}$ as defined following Corollary 17.2. (Technically, h_θ is only defined up to eyeglass equivalence, since that is the case for $h_{\mathbb{C}_\theta}$.) We have

Lemma 18.1. *Suppose θ, θ' are in the same interval of $[0, 2\pi] - \{\theta_i, 1 \leq i \leq n\}$. Then $h_\theta \sim h_{\theta'} \Theta_{\theta'}^\theta$.*

Proof. As just observed the chamber complex $\Theta_\theta^{\theta'}(\mathbb{C}_{\theta'})$ is properly isotopic to the chamber complex $\mathbb{C}_\theta$ in (S^3, T_θ) so $h_\theta = h_{\mathbb{C}_\theta} = h_{\Theta_\theta^{\theta'}(\mathbb{C}_{\theta'})}$. Now apply Corollary 17.4: $h_{\Theta_\theta^{\theta'}(\mathbb{C}_{\theta'})} \Theta_\theta^{\theta'} \sim h_{\mathbb{C}_{\theta'}}$ so $h_\theta \sim h_{\mathbb{C}_{\theta'}} (\Theta_\theta^{\theta'})^{-1} = h_{\theta'} \Theta_{\theta'}^\theta$ $\square$

Further following [FS1], for each $1 \leq i \leq n$ there are three disjoint disks, either

- $a_{\theta_i} \subset A_{\theta_i}, b_{\theta_i}^\pm \subset B_{\theta_i}$ with the pairs $(a_{\theta_i}, b_{\theta_i}^+)$ and $(a_{\theta_i}, b_{\theta_i}^-)$ each weakly reducing or
- symmetrically $a_{\theta_i}^\pm \subset A_{\theta_i}, b_{\theta_i} \subset B_{\theta_i}$ with the pairs $(a_{\theta_i}^+, b_{\theta_i})$ and $(a_{\theta_i}^-, b_{\theta_i})$ each weakly reducing.

with the property that (respectively for the two cases) for small ϵ

- $\Theta_{\theta_i \pm \epsilon}^{\theta_i}(a_{\theta_i}) = a_{\theta_i \pm \epsilon}$ and $\Theta_{\theta_i \pm \epsilon}^{\theta_i}(b_{\theta_i}^{\pm}) = b_{\theta_i \pm \epsilon}$
- symmetrically $\Theta_{\theta_i \pm \epsilon}^{\theta_i}(a_{\theta_i}^{\pm}) = a_{\theta_i \pm \epsilon}$ and $\Theta_{\theta_i \pm \epsilon}^{\theta_i}(b_{\theta_i}) = b_{\theta_i \pm \epsilon}$

Lemma 18.2. *Suppose θ, θ' are any two points in $[0, 2\pi] - \{\theta_i, 1 \leq i \leq n\}$. Then $h_\theta \sim h_{\theta'} \Theta_{\theta'}^\theta$.*

Proof. It suffices to prove the case in which θ, θ' are in adjacent intervals, for then we can just proceed around the circle. (See Figure 48 for a highly schematic picture.) So suppose that they are on adjacent intervals sharing the end point $\theta_i \in (0, 2\pi)$. Following Lemma 18.1 it suffices to prove the case in which $\theta = \theta_i - \epsilon$ and $\theta' = \theta_i + \epsilon$, for small ϵ .

With no loss of generality assume that the three weakly reducing disks in (S^3, T_i) are $a_{\theta_i} \subset A_{\theta_i}, b_{\theta_i}^+ \subset B_{\theta_i}$ and let $\mathbb{C}_i$ be the (not tiny) chamber complex obtained by weakly reducing T_{θ_i} along the three disks $a_{\theta_i}, b_{\theta_i}^{\pm}$. Let $\mathbb{C}_i^+$ be that obtained by weakly reducing along just the pair $a_{\theta_i}, b_{\theta_i}^+$ and $\mathbb{C}_i^-$ be that obtained by weakly reducing along just $a_{\theta_i}, b_{\theta_i}^-$. Apply Proposition 17.8 to deduce that $\mathbb{C}_i^- \sim \mathbb{C}_i$. Similarly, $\mathbb{C}_i^+ \sim \mathbb{C}_i$, so as a result $\mathbb{C}_i^+ \sim \mathbb{C}_i^-$.

The argument of Lemma 18.1 further shows that $h_{\mathbb{C}_i^+} \sim h_{\theta_i + \epsilon} \Theta_{\theta_i + \epsilon}^{\theta_i}$ and $h_{\mathbb{C}_i^-} \sim h_{\theta_i - \epsilon} \Theta_{\theta_i - \epsilon}^{\theta_i}$ so, combining, $h_{\theta_i - \epsilon} \Theta_{\theta_i - \epsilon}^{\theta_i} \sim h_{\theta_i + \epsilon} \Theta_{\theta_i + \epsilon}^{\theta_i}$. Hence $h_{\theta_i - \epsilon} \sim h_{\theta_i + \epsilon} \Theta_{\theta_i + \epsilon}^{\theta_i} \Theta_{\theta_i - \epsilon}^{\theta_i} \sim h_{\theta_i + \epsilon} \Theta_{\theta_i + \epsilon}^{\theta_i - \epsilon}$ as required. $\square$




Figure 48: Schematic for Lemmas 18.1 and 18.2

Corollary 18.3. $\tau \in \mathcal{E}$

Proof. Choose, in Lemma 18.2 the values $\theta = 0, \theta' = 2\pi$. Then, per that lemma, $h_0 \sim h_{2\pi} \Theta_{2\pi} \Theta_0^{-1}$. Furthermore, since $\mathbb{C}_0 = \mathbb{C}_{2\pi}$, Corollary 8.3 has $h_0 \sim h_{\mathbb{C}_0} \sim h_{\mathbb{C}_{2\pi}} \sim h_{2\pi}$, and so $\tau = \Theta_{2\pi} \sim \Theta_0 = id_{(S^3, T_g)}$. $\square$

Remark: It is perhaps not surprising that it is possible to position T_g with respect to the height function on S^3 so that h_0 and $h_{2\pi} : (S^3, T_g) \rightarrow (S^3, T_g)$ are not just eyeglass equivalent, but in fact are both the identity (under the standard inductive Assumption 3.4. We do not need this here.

Corollary 18.4. $G(S^3, T) = \mathcal{E}$.

Index

- $F_{\mathcal{D}} \subset M$, 226
- $\leftrightarrow$, 305
- $\rightarrow$, 285
- b-wise alignment, 251
- $\hat{\mathbb{C}}_{\mathcal{D}}$, 226
- $\sim$, 218
- Akilter sphere for F , 324
- Aligned chamber complex decomposition, 250
- Aligned disk or sphere, 213
- Almost balanced sphere, 274
- Ancestral bubble pass, 253
- Balanced sphere, 274
- Bubble, 211
- Bubble exchange, 211
- Bubble move, 211
- Bubble pass, 213
- Bullseye, 281
- center edge of an odd diameter tree, 313
- center vertex of an even diameter tree, 314
- Certify; certificate, 266
- Chamber complex, 226
- Chamber complex decomposition, 226
- cocertify, 267
- Deflation, 279
- Deflation of decomposition sequence, 285
- Designated handlebody chamber, 232
- Disk decomposition, 226
- Disk set, 226
- Disky, 227
- Empty chamber, 256
- ey-handlebody, 292
- Eyeglass, 212
- Eyeglass group, 213
- Eyeglass twist, 212
- F-circles, 312
- Flagged chamber complex, 256
- Flip, 211
- Ghost circles, 312
- Goneball, 226
- Guiding set of disks, 285
- Heegaard split chamber complex decomposition, 243
- Maximal bubble, 256
- Maximal deflationary subsequence, 286
- Occupied chamber, 256
- Powell group, 213
- Preferred alignment, 256
- Pure Heegaard splitting, 235
- Remnant of chamber, 226
- Scars, internal and external, 226
- Sibling decompositions, 261
- Target manifold, 298
- Tiny chamber complex, 232
- Tiny Heegaard split chamber complex, 248
- Tube sum, 213
- Weak reduction, 240
- weakly reducing disk family, 233
- x-handlebody, 290
- xe-handlebody, 292
- y-handlebody, 290

Acknowledgments

Michael Freedman has been an avid source of support for this project; his notion of a ‘cycle of weak reductions’ underpins it all. Aoife McCormick inspired the notion of ‘chamber complex’ - the method here for translating Freedman’s idea into a recipe for Heegaard surface reduction.

References

- [Bo] F Bonahon, *Cobordism of automorphisms of surfaces*, Ann. Sci. École Norm. Sup.(4) **16** (1983) 237–270
- [Ce] J Cerf, *La stratification naturelle des espaces de fonctions différentiables réelles et le théorème de la pseudo-isotopie*, Publications mathématiques de l’I.H.É.S., tome 39 (1970) 5-173.
- [CG] A. Casson and C. Gordon, *Reducing Heegaard splittings*, Topology and its applications, **27** (1987) 275-283. [210](#), [220](#), [233](#)
- [EE] C. Earle and J. Eells, *The diffeomorphism group of a compact Riemann surface*, Bull. Amer. Math. Soc. **73** (1967) 557–559 [216](#)
- [FS1] M. Freedman and M. Scharlemann, *Powell moves and the Goeritz group*, Arxiv preprint 1804.05909 [209](#), [210](#), [211](#), [213](#), [214](#), [225](#), [330](#), [331](#)
- [FS2] M. Freedman and M. Scharlemann, *Uniqueness in Haken’s Theorem*, Mich. Math. J. **74** (2024) 119-142. [213](#), [219](#), [222](#), [223](#), [225](#), [234](#), [249](#), [251](#), [252](#), [253](#), [261](#), [262](#), [263](#), [269](#)
- [GK] D. Gay and R. Kirby, *Fiber-connected, indefinite Morse 2-functions on connected n -manifolds*, Proc. Natl. Acad. Sci. USA **108** (2011), no. 20, 8122–8125. [323](#)
- [GP] V. Guillemin and A. Pollack, *Differential Topology*, Prentice-Hall, 1974. [215](#)
- [Ha] A Hatcher, *Homeomorphisms of sufficiently large P^2 -irreducible 3-manifolds*, Topology **15** (1976), 343–347.
- [JM] J. Johnson and D. McCullough, *The space of Heegaard splittings*, J. Reine Angew. Math **679** (2013) 155-179 [209](#), [212](#), [330](#)
- [La] M. Lackenby, *An algorithm to determine the Heegaard genus of simple 3-manifolds with nonempty boundary*, Alg. Geom. Top. **8** (2008) 911-934. [219](#), [220](#), [233](#), [235](#), [237](#), [245](#)
- [Lau] F. Laudenbach, *Sur les 2-sphères d’une variété de dimension 3*, Ann. of Math. (2) **97** (1973) 57–81. [323](#)
- [Mi] J. Milnor, *Morse Theory (Based on lecture notes by M. Spivak and R. Wells)*, Annals of Mathematics Studies, **51** Princeton University Press, Princeton, N.J., 1963. [323](#)

- [Moi] E. E. Moise., *Geometric topology in dimensions 2 and 3*, Graduate Texts in Mathematics, **47** Springer-Verlag, New York - Heidelberg - Berlin, 1977. 211
- [Pa] R. Palais, *Local triviality of the restriction map for embeddings*, Comment. Math. Helv. **34** (1960) 305-312. 215
- [Po] J. Powell, *Homeomorphisms of S^3 leaving a Heegaard surface invariant*, Trans. Amer. Math. Soc. **257** (1980) 193–216. 209, 211
- [Sc1] M. Scharlemann, *A Strong Haken’s Theorem*, to appear in Alg. Geom. Top, Arxiv preprint 2003.08523. 213, 219, 221, 223, 224, 234, 242, 243, 245, 246, 251, 252, 258, 281
- [Sc2] M. Scharlemann, *One Powell generator is redundant*, Proc. Amer. Math. Soc., Series B **7** (2020) 138-141. 209
- [Sc3] M. Scharlemann, *Powell’s Conjecture on the Goeritz group of S^3 is stably true*, to appear in Alg. Geom. Top, Arxiv preprint arXiv:2210.13629. 210
- [Sp] M. Spivak, *Calculus on manifolds*, Avalon Publishing, 1965. 215
- [Wa] F. Waldhausen, *Heegaard-Zerlegungen der 3-Sphäre*, Topology **7** (1968) 195–203. 209, 211, 218, 222, 251, 261
- [Wh] J. H. C. Whitehead, *On equivalent sets of elements in a free group*, Ann. of Math. **37** (1936), 782–800.

AUTHOR

Martin Scharlemann
 Mathematics Department
 University of California
 Santa Barbara, CA 93106-3080 USA
mgscharl@math.ucsb.edu
<https://web.math.ucsb.edu/~mgscharl>