



HAL
open science

Mean Curvature Flow and Heegaard Surfaces in Lens Spaces

Reto Buzano, Sylvain Maillot

► **To cite this version:**

Reto Buzano, Sylvain Maillot. Mean Curvature Flow and Heegaard Surfaces in Lens Spaces. 2023. hal-04342321

HAL Id: hal-04342321

<https://hal.science/hal-04342321>

Preprint submitted on 13 Dec 2023

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution 4.0 International License

Mean Curvature Flow and Heegaard Surfaces in Lens Spaces

Reto Buzano and Sylvain Maillot

Abstract

We prove that the moduli space of mean convex two-spheres embedded in complete, orientable 3-dimensional Riemannian manifolds with nonnegative Ricci curvature is path-connected. This result is sharp in the sense that neither of the conditions of (strict) mean convexity, completeness, and nonnegativity of the Ricci curvature can be dropped or weakened. We also study the number of path components of mean convex Heegaard tori, again in ambient manifolds with nonnegative Ricci curvature. We prove that there are always either one or two path components and this number does not only depend on the homotopy type of the ambient manifold. We give a precise characterisation of the two cases and also discuss what happens if the mean convexity condition is weakened to nonnegative mean curvature.

1 Introduction

Given a closed (smooth) manifold M^n , a classical problem in Riemannian geometry is to investigate the homotopy type of the moduli space of positive scalar curvature metrics

$$\mathcal{M}_{R>0}(M) := \{g \text{ Riemannian metric on } M \text{ with } R_g > 0\} / \text{Diff}(M).$$

This problem has its origins in an over 100 years old result of Weyl [25] who showed that for $n = 2$, this space is always path-connected. In fact, Rosenberg-Stolz [23] proved that it is contractible. When $n = 3$, it is only much more recently that Marques [22] proved path-connectedness and Bamler-Kleiner [4] proved contractibility. These results are in sharp contrast to the situation in higher dimensions where the moduli spaces can have very complicated topology – though their homotopy types are in general unknown. An illustrative example is given by spheres \mathbb{S}^{4k+3} , $k \geq 1$, for which the moduli spaces are known to have infinitely many path components [20].

In this article, we are concerned with the following extrinsic version of this problem: given a Riemannian three-manifold (M^3, g) and a closed surface Σ^2 , investigate the homotopy type of the moduli space of mean convex two-sided embeddings of Σ into M ,

$$\mathcal{M}_{H>0}(\Sigma, M) := \{\Sigma \hookrightarrow M \text{ smooth two-sided embedding with } H > 0\} / \text{Diff}(\Sigma).$$

This extrinsic problem was first studied by the first author in joint work with Haslhofer and Hershkovits in the case where $(M^3, g) = (\mathbb{R}^3, g_{Eucl.})$ is Euclidean space.¹ In [10], they obtained

¹In fact, they studied the moduli spaces of *two-convex* embeddings of hypersurfaces Σ^{n-1} into $(\mathbb{R}^n, g_{Eucl.})$, but we only focus on dimension $n = 3$ here.

that for $\Sigma = \mathbb{S}^2$ the moduli space is path-connected and in [11] they proved that for $\Sigma = T^2 = \mathbb{S}^1 \times \mathbb{S}^1$ the connected components of the moduli space are in bijective correspondence with the knot classes of closed embedded curves.

The aim of this article is to extend the study of the extrinsic problem to more general ambient three-manifolds (M^3, g) . In this context we say that a smoothly embedded, two-sided surface $\Sigma \subset M$ has *positive mean curvature* (respectively *nonnegative mean curvature*) with respect to g if there exists a choice of unit normal vector field ν on Σ such that $H > 0$ (resp. $H \geq 0$) everywhere. In other words, Σ has positive mean curvature if the mean curvature vector always points to the same side of Σ . Note that this latter definition only makes sense if Σ is two-sided. We also say that Σ is *mean convex* if it has positive mean curvature.

For $\Sigma = \mathbb{S}^2$, we obtain the following generalisation of the main result from [10].

Theorem 1.1. *Let (M, g) be a complete, orientable Riemannian three-manifold with nonnegative Ricci curvature. Then the moduli space $\mathcal{M}_{H>0}(\mathbb{S}^2, M)$ of mean convex embedded two-spheres in (M, g) is path-connected.*

Remark 1.2. Theorem 1.1 does not hold if the condition $H > 0$ is weakened to $H \geq 0$ or if $\text{Ric}_g \geq 0$ is weakened to $\text{Ric}_g \geq -C$ for some $C > 0$ (or replaced by $1 \geq \text{Ric}_g \geq -\varepsilon$ for some $\varepsilon > 0$). Indeed, in both cases, it is easy to construct examples where the moduli space contains two-spheres that do not bound a ball and others that do; thus the moduli space has at least two path components. Moreover, the completeness condition on M cannot be dropped: for example if M is a round three-sphere with a point removed, then every mean convex two-sphere bounds a ball, but for some of them the mean curvature vector points toward the ball while for others it points in the opposite direction. Hence the moduli space again has at least two path components. See Proposition 2.1 and Remark 2.2 below for more details.

For $\Sigma = T^2$, we restrict the problem to *Heegaard surfaces* in order to guarantee that Σ bounds a solid (unknotted) torus. It is well known that the moduli space of *all* Heegaard tori is path-connected, but surprisingly, in contrast to the Euclidean case, the moduli space of *mean convex* Heegaard tori can have more than one path component. Our main theorem says exactly how many there are depending on the ambient manifold M . Before stating the results, we give some definitions and fix notation.

In his PhD thesis, Heegaard [16] introduced an intriguingly simple way to construct complicated orientable three-manifolds: taking two handlebodies of the same genus, one can glue their boundaries together by a (possibly very complicated) diffeomorphism. Here, a handlebody of genus k is an orientable three-manifold with boundary obtained from a closed 3-ball \bar{B}^3 by attaching k handles, the process of attaching a handle being understood as removing two disjoint closed disks D_0 and D_1 in ∂B^3 and attaching $D^2 \times [0, 1]$ by identifying $D^2 \times \{j\}$ with D_j , $j = 0, 1$. In fact, as a consequence of Whitehead's triangulation theorem [26], in fact every closed, orientable three-manifold can be obtained in this way. Conversely, if M is a closed, orientable three-manifold, and Σ a closed, orientable surface smoothly embedded in M , we say that Σ is a *Heegaard surface* if the closure of each component of $M \setminus \Sigma$ is a handlebody. Three-manifolds admitting a Heegaard surface of genus one are called *lens spaces*. Let us now describe a general construction of them.

Let \mathbb{S}^1 and D^2 be the unit circle and the unit disk in \mathbb{C} , respectively. A *solid torus* V is a three-manifold diffeomorphic to $\mathbb{S}^1 \times D^2$. Its *core* is the simply closed curve $\mathbb{S}^1 \times \{0\}$ and the *longitude*

and *meridian* are defined as $\mathbb{S}^1 \times \{1\}$ and $\{1\} \times \mathbb{S}^1$, respectively. Note that these definitions imply that a diffeomorphism to $\mathbb{S}^1 \times D^2$ has been fixed, so these concepts are defined only up to isotopy.

Let p, q be coprime integers with $p \geq 0$. The *lens space* $L(p, q)$ is obtained by gluing together two solid tori V_1 and V_2 via the diffeomorphism $\theta : \partial V_1 \rightarrow \partial V_2$ defined by

$$\theta(u, v) = (u^r v^p, u^s v^q),$$

where r, s are integers such that $rq - sp = 1$. Put differently, $L(p, q)$ is the three-manifold obtained by gluing the boundaries of two solid tori in such a way that the meridian of the first torus goes to a curve wrapping around the longitude p times and around the meridian q times on the second torus. We denote by γ_k the core of V_k for $k \in \{1, 2\}$. By construction, the torus $\Sigma = \partial V_1$ is Heegaard.

Special cases of lens spaces include the 3-sphere $\mathbb{S}^3 = L(1, 0)$, real projective space $\mathbb{RP}^3 = L(2, 1)$, and $\mathbb{S}^1 \times \mathbb{S}^2 = L(0, 1)$. On the one hand, the standard product metric on $L(0, 1) = \mathbb{S}^1 \times \mathbb{S}^2$ has $\text{Ric}_g \geq 0$. On the other hand, in the case $p \neq 0$, $L(p, q)$ has an alternative description as a quotient of $\mathbb{S}^3 \subset \mathbb{C}^2$ by the (finite cyclic) group of isometries generated by

$$(z_1, z_2) \mapsto (e^{2\pi i/p} z_1, e^{2\pi i q/p} z_2)$$

and therefore carries a Riemannian metric with (constant) positive sectional curvature, thus in particular again with $\text{Ric}_g \geq 0$. The assumption of nonnegative Ricci curvature is therefore quite natural for lens spaces and throughout the article (unless otherwise mentioned), we will always assume that we have fixed a background metric g on $L(p, q)$ with $\text{Ric}_g \geq 0$.

Definition 1.3. We denote by

$$\mathcal{M}_{H>0}(p, q) = \mathcal{M}_{H>0}(T^2, L(p, q))$$

the *moduli space of Heegaard tori* in $L(p, q)$ that have positive mean curvature with respect to the background metric g on $L(p, q)$. Similarly, $\mathcal{M}_{H \geq 0}(p, q)$ denotes the moduli space of Heegaard tori in $L(p, q)$ with nonnegative mean curvature.

Our main results are the following two theorems.

Theorem 1.4. *Let p, q be coprime integers with $p \geq 0$. Fix a metric g on the lens space $L(p, q)$ with nonnegative Ricci curvature. If $q \not\equiv \pm 1 \pmod{p}$, then $\mathcal{M}_{H>0}(p, q)$ is path-connected. Otherwise it has exactly two path-components.*

Remark 1.5. An interesting aspect of Theorem 1.4 is that it is not true that the homotopy type of the moduli spaces $\mathcal{M}_{H>0}(p, q)$ depends only on the homotopy type of the ambient space $L(p, q)$. An illustrative example is given by $L(7, 1)$ and $L(7, 2)$. These two lens spaces have the same homotopy type (and therefore also isomorphic fundamental groups and the same homology). Nevertheless, by Theorem 1.4, $\mathcal{M}_{H>0}(7, 1)$ is path-connected while $\mathcal{M}_{H>0}(7, 2)$ is not.

By a result of Bonahon [5], any two Heegaard tori in a fixed lens space $L(p, q)$ are isotopic. In particular, when $q \not\equiv \pm 1 \pmod{p}$ we can pick two Heegaard tori T_1 and T_2 lying in different components of $\mathcal{M}_{H>0}(p, q)$; then T_1 and T_2 are isotopic, but no isotopy between them preserves mean convexity. By contrast, our next result says that we can always find an isotopy that preserves nonnegativity of mean curvature.

Theorem 1.6. *Let p, q be coprime integers with $p \geq 0$. Fix a metric g on the lens space $L(p, q)$ with nonnegative Ricci curvature. Then $\mathcal{M}_{H \geq 0}(p, q)$ is always path-connected.*

We finish this introduction with an outline of the proofs of the above theorems and of the article.

First, in Section 2, we show that every mean convex 2-sphere in a complete, orientable Riemannian three-manifold with nonnegative Ricci curvature as in Theorem 1.1 bounds a mean convex domain (a compact codimension 0 submanifold with mean curvature vector pointing inwards). This important step allows us to work with domains instead of surfaces in the following; it also highlights that the theorem is sharp in the sense described in Remark 1.2. We then continue by extending a gluing result of the first author with Haslhofer and Hershkovits [10] from Euclidean space to the ambient manifold setting, allowing us to connect mean convex domains along thin “strings” in a way that preserves mean convexity. This gluing result could be of independent interest.

Next, in Section 3, we recall the theory of mean curvature flow with surgery from Haslhofer-Kleiner [15] as extended by Haslhofer-Ketover [14] to the ambient manifold setting. In particular, for any mean convex sphere or Heegaard torus Σ , there is a mean curvature flow (MCF) with surgery starting at Σ that becomes extinct or converges to a union of minimal surfaces. If the ambient manifold has nonnegative Ricci curvature, then the only possibility is finite time extinction.

In Section 4, we show that any element of $\mathcal{M}_{H > 0}(\mathbb{S}^2, M)$ can be deformed to a *marble tree* which then itself can be deformed to a small geodesic sphere. Similarly, any element of $\mathcal{M}_{H > 0}(p, q)$ can be deformed to a *marble circuit* which then itself can be deformed further to the boundary of the ε -neighbourhood of a simple closed curve $T_\varepsilon = \partial N_\varepsilon(\gamma)$. These results follow again the ideas of the first author with Haslhofer and Hershkovits [10, 11], using the gluing construction of Section 2 to undo the surgeries.

The proofs of the three theorems above are then completed in Section 5. Theorem 1.1 follows immediately from the results in Section 4. Theorem 1.4 requires further topological arguments. First, we show that $\mathcal{M}_{H > 0}(p, q)$ has at most two components. This uses critically Bonahon’s uniqueness theorem [5]: since the initial torus is Heegaard, the simple closed curve γ from the above paragraph must be isotopic (up to orientation) to one of the cores γ_1 or γ_2 . In the case where $q \cong \pm 1 \pmod p$, path-connectedness then follows immediately from a Lemma of Bonahon that says that in this case γ_1 and γ_2 are isotopic. The second case is by contradiction: let T_1 and T_2 be the boundary of the ε -neighbourhood of γ_1 and γ_2 , respectively. Then T_1 (respectively T_2) has positive mean curvature and the mean curvature vector points towards γ_1 (respectively γ_2). If they were in the same component of the moduli space, along the isotopy the mean curvature would point in the same direction. The key point is that “same” direction is well-defined because γ_1 is not homologous to γ_2 . Finally, to prove Theorem 1.6, we show that every Heegaard torus with $H \geq 0$ is isotopic (via an isotopy preserving $H \geq 0$) to a strictly mean convex Heegaard torus. Therefore $\mathcal{M}_{H \geq 0}(p, q)$ cannot have more path-components than $\mathcal{M}_{H > 0}(p, q)$. Moreover, in the case where $\mathcal{M}_{H > 0}(p, q)$ has two path components there exists a minimal Heegaard torus Σ by [19] and we prove that it is isotopic *on each side* to a mean convex torus. The result follows.

Acknowledgements. RB thanks Gianmichele Di Matteo and Luciano Mari for interesting discussions. He has been partially supported by the EPSRC grant EP/S012907/1. SM thanks Giuseppe Pipoli for interesting discussions. He has been partially supported by Agence Nationale de la Recherche grant ANR-17-CE40-0034.

2 Mean convex gluing in ambient three-manifolds

Here and in the following two sections, (M, g) denotes a complete, orientable Riemannian three-manifold with nonnegative Ricci curvature $\text{Ric}_g \geq 0$ (unless mentioned otherwise). Since M is orientable, all orientable embedded surfaces in M are two-sided. This applies in particular to spheres and tori.

A *domain* in M is a (possibly disconnected) smooth *compact* codimension 0 submanifold with boundary. A domain K is *mean convex* if its boundary $\Sigma = \partial K$ satisfies $H \neq 0$ everywhere, with mean curvature vector pointing toward K .

In the following, instead of working with the mean convex surfaces Σ in (M, g) , it will be convenient to work with the mean convex domains bounded by them. In the situations relevant for us, this can always be done. For mean convex Heegaard tori in lens spaces this is easy. Indeed, they bound two solid tori and the condition of mean convexity means that the mean curvature vector always points to one of these two solid tori which will therefore be our mean convex domain. For mean convex spheres as in Theorem 1.1 this is less obvious; we therefore start with a proof of this fact.

Proposition 2.1. *Let Σ be a mean convex embedded two-sphere in (M, g) . Then Σ bounds a mean convex domain \bar{K} .*

Remark 2.2. Note that the proposition is sharp in the sense that the condition $H > 0$ cannot be weakened to $H \geq 0$, as seen by the counterexample $\Sigma = \mathbb{S}^2 \times \{0\}$ in $M = \mathbb{S}^2 \times \mathbb{S}^1$ (or $\mathbb{S}^2 \times \mathbb{R}$) with standard product metric. Similarly, the condition $\text{Ric}_g \geq 0$ on (M, g) cannot be weakened to $\text{Ric}_g \geq -C$ for some $C > 0$ (or replaced by $1 \geq \text{Ric}_g \geq -\varepsilon$ for some $\varepsilon > 0$). Also here the counterexample is $\Sigma = \mathbb{S}^2 \times \{0\}$ in $M = \mathbb{S}^2 \times \mathbb{S}^1$ (or $\mathbb{S}^2 \times \mathbb{R}$), this time with a warped product metric $g = w(x)g_{\mathbb{S}^2} + dx^2$ arbitrarily close to the standard product metric (i.e. with warping factor $w(x)$ close to 1) and with $w'(0) \neq 0$ to guarantee that Σ has $H > 0$. Finally, completeness of M cannot be removed as seen by the counterexample $M = \mathbb{S}^3 \setminus \{p\}$ with Σ a small geodesic two-sphere around p . The sphere Σ bounds a compact domain only on one side, but the mean curvature vector points to the complement of this domain. All these counterexamples also work as counterexamples for path-connectedness of the moduli space $\mathcal{M}_{H>0}(\mathbb{S}^2, M)$ in Theorem 1.1 if one of the assumptions is removed or weakened.

Proof. Assuming towards a contradiction that there exists an embedded 2-sphere in M with $H > 0$ that does not bound a domain, we can lift it to the universal cover \widetilde{M} (due to simply-connectedness). The lifted two-sphere is still two-sided, still has $H > 0$, and still does not bound a domain. Hence we can work directly in the universal cover.

Since M is complete, orientable, and has $\text{Ric}_g \geq 0$, a theorem of Liu [21] implies that its universal cover \widetilde{M} is either diffeomorphic to \mathbb{S}^3 or \mathbb{R}^3 or it splits isometrically as a Riemannian product $\widetilde{M} = \mathbb{S}^2 \times \mathbb{R}$, where the \mathbb{S}^2 factor has (possibly nonconstant) nonnegative curvature. We consider these cases separately.

Case 1. \widetilde{M} is diffeomorphic to \mathbb{S}^3 .

This is the easiest case. By Alexander's theorem [2], Σ bounds a 3-ball on either side. One of these 3-balls is a mean convex domain.

Case 2. \widetilde{M} is diffeomorphic to \mathbb{R}^3 .

Again, by Alexander's theorem [2], Σ bounds a 3-ball. We need to prove that the mean curvature vector points towards this ball. By Kasue's theorem [18, Theorem C], if $H \geq 0$ with respect to the *outward pointing normal*, i.e. the mean curvature vector points to the complement of the ball, then this complement is isometric to a Riemannian product $\Sigma \times [0, \infty)$ and Σ is totally geodesic, i.e. $H \equiv 0$. This shows that the mean curvature vector must point inwards. (See also Theorem 1.6. of [1] for a strengthening of Kasue's theorem that yields a direct proof that the mean curvature vector points towards the ball in our case). As seen in the remark above, completeness of \widetilde{M} is essential for this argument.

Case 3. \widetilde{M} is a Riemannian product $\mathbb{S}^2 \times \mathbb{R}$ with nonnegatively curved \mathbb{S}^2 factor.

Let $\pi_2 : \widetilde{M} = \mathbb{S}^2 \times \mathbb{R} \rightarrow \mathbb{R}$ denote the projection to the second factor.

Assume Σ does not bound a ball. Call A^- and A^+ the two connected components of $(\mathbb{S}^2 \times \mathbb{R}) \setminus \Sigma$. Both components are homeomorphic to $\mathbb{S}^2 \times \mathbb{R}$ and adjacent to exactly one end. We assume that $p(A^+)$ (respectively $p(A^-)$) is adjacent to $+\infty$ (resp. $-\infty$). Since $H > 0$, the mean curvature vector of Σ always points to one of the two components. Without loss of generality, we may assume that it always points to A^+ (otherwise invert the orientation of \mathbb{R} , swapping A^-, A^+).

Now consider a maximal point p of the restricted projection $\pi_2|_{\Sigma}$. On the one hand, by construction, the mean curvature vector of Σ at p points to A^+ . On the other hand, by definition of p , Σ lies on the "lower" side of the (minimal) two-sphere $\mathbb{S}^2 \times \{\pi_2(p)\}$ and touches it at p . The maximum principle thus forces the mean curvature vector of Σ at p to either vanish or point to A^- , yielding the desired contradiction.

Hence Σ bounds a ball B . Again looking at a maximum of the restriction of π_2 to Σ and using the maximum principle we see that the mean curvature vector points toward B at such a point. We conclude that B is a mean convex domain. \square

In what follows, we fix a compact domain \bar{K} (which in our applications with noncompact (M, g) will be the mean convex domain bounded by Σ) and only consider subsets of \bar{K} . In particular, everything below might depend on this choice of \bar{K} . By compactness, the curvature tensor and its derivatives are uniformly bounded over \bar{K} and $\text{inj}(\bar{K}) > 0$, where $\text{inj}(\bar{K})$ denotes the infimum of the injectivity radius of M at p over all points $p \in \bar{K}$.

We denote by \mathcal{D} the set of all mean convex domains contained in \bar{K} . For such a domain $K \in \mathcal{D}$, $p \in \partial K$, and $\alpha > 0$, we denote by $\bar{B}_{\alpha}^{\pm}(p)$ the closed metric balls of radius $\alpha H(p)^{-1}$ with center $\exp_p(\pm \alpha H(p)^{-1} \nu(p))$ in M where $\nu(p)$ is the unit normal vector to ∂K at p pointing towards K .

Definition 2.3 (Quantitative noncollapsing, see [3, 14, 15, 24]). Let $\alpha > 0$. We say that a mean convex domain $K \in \mathcal{D}$ is α -*noncollapsed* if for every $p \in \partial K$ the following conditions hold:

1. $H(p) \geq 4\alpha \text{inj}(\bar{K})^{-1}$;
2. $\bar{B}_{\alpha}^+(p) \subset K$;
3. $\bar{B}_{\alpha}^-(p) \cap K \subset \partial K$.

A *curve* in M is a (possibly disconnected) smooth compact 1-submanifold with boundary. The set of curves contained in \bar{K} is denoted by \mathcal{C} . In this section, we will glue certain *controlled* domains along certain *controlled* curves in \bar{K} .

Definition 2.4 (Controlled domains and curves, slight modification of [10]). Let $\mathbb{A} = (\alpha, c)$ be a pair of positive numbers. We say that a domain in \bar{K} is \mathbb{A} -controlled if it is mean convex with $H \geq c$, α -noncollapsed, and satisfies $|A| + |\nabla A| \leq c^{-1}$. The set of \mathbb{A} -controlled domains in \bar{K} is denoted by $\mathcal{D}_{\mathbb{A}}$.

Let $\mathbb{b} > 0$. We say that a curve γ in \bar{K} is \mathbb{b} -controlled if its curvature vector satisfies $|\kappa| \leq \mathbb{b}^{-1}$ and $|\nabla \kappa| \leq \mathbb{b}^{-2}$, each connected component of γ has normal injectivity radius $\geq \mathbb{b}/10$, and different connected components of γ are at least $10\mathbb{b}$ apart. Moreover, we require that the $\mathbb{b}/10$ -tubular neighbourhood of γ is still contained in \bar{K} . We denote by $\mathcal{C}_{\mathbb{b}}$ the set of \mathbb{b} -controlled curves.

An (\mathbb{A}, \mathbb{b}) -controlled configuration is a pair $(K, \gamma) \in \mathcal{D}_{\mathbb{A}} \times \mathcal{C}_{\mathbb{b}}$ satisfying the following properties:

1. The interior of γ is contained in the complement of K .
2. If $p \in \partial\gamma$, then $p \in \partial K$ and γ meets ∂K orthogonally at p .
3. $d(\gamma \setminus \cup_{p \in \partial\gamma} B_{\mathbb{b}/10}(p), \partial K) \geq \mathbb{b}/20$.

We denote the set of (\mathbb{A}, \mathbb{b}) -controlled configurations by $\mathcal{X}_{\mathbb{A}, \mathbb{b}}$.

We note that if $\mathbb{A}' = (\alpha', c')$ is such that $\alpha' \leq \alpha$ and $c' \leq c$ then $\mathcal{D}_{\mathbb{A}} \subseteq \mathcal{D}_{\mathbb{A}'}$ and if $\mathbb{b}' \leq \mathbb{b}$ we have $\mathcal{C}_{\mathbb{b}} \subseteq \mathcal{C}_{\mathbb{b}'}$. In particular also $\mathcal{X}_{\mathbb{A}, \mathbb{b}} \subseteq \mathcal{X}_{\mathbb{A}', \mathbb{b}'}$.

The main result of this section is the following theorem which can be seen as an ambient manifold version of the gluing result from [10, Theorem 4.1].

Theorem 2.5 (Mean convex gluing). *There exists a smooth gluing map*

$$\mathcal{G} : \mathcal{X}_{\mathbb{A}, \mathbb{b}} \times (0, \bar{r}) \rightarrow \mathcal{D}, \quad ((K, \gamma), r) \mapsto \mathcal{G}_r(K, \gamma),$$

where $\bar{r} = \bar{r}(\mathbb{A}, \mathbb{b}, \bar{K}, M, g) > 0$ is a constant, and a smooth increasing function $\delta : (0, \bar{r}) \rightarrow \mathbb{R}_+$ with $\lim_{r \rightarrow 0} \delta(r) = 0$, with the following properties:

1. For every $(K, \gamma) \in \mathcal{X}_{\mathbb{A}, \mathbb{b}}$ and every $0 < r < \bar{r}$, the resulting domain $\mathcal{G}_r(K, \gamma)$ deformation retracts to $K \cup \gamma$.
2. The nontrivial part of the gluing only happens near the points $p \in \partial\gamma$ in the following sense:

$$\mathcal{G}_r(K, \gamma) \setminus \bigcup_{p \in \partial\gamma} B_{2\delta(r)}(p) = K \cup N_r(\gamma) \setminus \bigcup_{p \in \partial\gamma} B_{2\delta(r)}(p),$$

where $N_r(\gamma)$ denotes the solid r -tubular neighbourhood of γ .

We will often refer to the regions $N_r(\gamma)$ as *strings* and call r the *string radius*.

Remark 2.6. For simplicity, we only state and prove the theorem for (uniformly) mean convex domains in ambient three-manifolds, but with some very minor modifications of the proof, one obtains the same result in all dimensions for (uniformly) two-convex domains. Moreover, the result in [10] also allows γ to have “loose end points”, that is, points $p \in \partial\gamma \setminus \partial K$, but we do not need them for the sake of this article.

The idea of the proof is as follows: If $p \in \partial\gamma \cap \partial K$, we lift the domain K and curve γ in a neighbourhood of p to the tangent space via the exponential map, use the gluing construction from Theorem 4.1 in [10] in Euclidean \mathbb{R}^3 , and then use the exponential map again to obtain a domain in M . The resulting domain is not yet exactly what we want, as the “string part” will be the image (under the exponential map) of a tubular neighbourhood in Euclidean space rather than a tubular neighbourhood of γ in (M, g) , thus requiring a further modification in the annular regions $B_{2\delta(r)}(p) \setminus B_{\delta(r)}(p)$. All of this is relatively easy if the mean curvature is very large (in which case the errors caused by the ambient curvature of (M, g) will essentially be irrelevant), but it becomes more tricky when gluing at points of “small” (but positive) mean curvature. To deal with this problem carefully, we will need the following technical lemma.

Recall that for each $\varrho < \text{inj}(M, p)$, the exponential map $\exp_p : T_p M \simeq \mathbb{R}^3 \rightarrow M$ is a diffeomorphism from the Euclidean ϱ -ball $B_\varrho(0) \subset \mathbb{R}^3$ to the geodesic ϱ -ball $B_\varrho(p) \subset M$ (where 0 is mapped to p).²

Lemma 2.7. *Let $p \in \bar{K}$ and $\varrho < \text{inj}(\bar{K})$ be fixed. Let $S \subset B_\varrho(0) \subset \mathbb{R}^3$ be a two-sided surface and $\Sigma = \exp_p(S) \subset B_\varrho(p) \subset M$ its image under the above diffeomorphism and let these surfaces be endowed with normal vector fields pointing “in the same direction”. We denote by A^S and H^S the second fundamental form and mean curvature of S in \mathbb{R}^3 and by A^Σ and H^Σ the second fundamental form and mean curvature of Σ in (M, g) with respect to these normal vector fields. Then the following holds.*

1. *If $H^\Sigma \geq c_\Sigma > 0$, $|A^\Sigma| + |\nabla A^\Sigma| \leq c_\Sigma^{-1}$, and Σ is α -noncollapsed everywhere in $B_\varrho(p)$, then there is $0 < \delta_1 \leq \varrho$ depending on the curvature of (M, g) in p and on c_Σ and α , such that $S \cap B_{\delta_1}(0)$ satisfies $H^S \geq \frac{9}{10}c_\Sigma$, $|A^S| + |\nabla A^S| \leq \frac{10}{9}c_\Sigma^{-1}$, and is $\frac{9}{10}\alpha$ -noncollapsed.*
2. *If $\gamma \subset B_\varrho(p)$ is a \mathbb{b} -controlled curve, then there is $0 < \delta_2 \leq \varrho$ depending on the curvature of (M, g) in p and on \mathbb{b} , such that $\exp_p^{-1}(\gamma) \cap B_{\delta_2}(0)$ is $\frac{9}{10}\mathbb{b}$ -controlled.*
3. *If $H^S \geq c_S > 0$ everywhere in $B_\varrho(0)$, then there is $0 < \delta_3 \leq \varrho$ depending on the curvature of (M, g) in p and on c_S , such that $\Sigma \cap B_{\delta_3}(p)$ has $H^\Sigma \geq \frac{9}{10}c_S$.*

Proof of Theorem 2.5 (assuming Lemma 2.7). As the construction is local, we first explain the gluing around one point $p \in \partial\gamma \cap \partial K$, making sure no constants will depend on the choice of this p (but possibly on the set \bar{K}). First, set $\varrho = \frac{1}{2} \min\{\text{inj}(\bar{K}), \min_{q_1, q_2 \in \partial\gamma} d(q_1, q_2)\}$, guaranteeing that \exp_p restricted to $B_\varrho(0) \subset \mathbb{R}^3$ is a diffeomorphism and p is the only gluing point contained in $B_\varrho(p)$. We denote $\Sigma = \partial K \cap B_\varrho(p)$ and $S = \exp_p^{-1}(\Sigma) \subset B_\varrho(0) \subset \mathbb{R}^3$.

By assumption, K is $\mathbb{A} = (\alpha, c)$ -controlled and γ is \mathbb{b} -controlled. We can therefore let δ_1 be as in Point 1 of Lemma 2.7 with $c_\Sigma = c$, and δ_2 as in Point 2 of Lemma 2.7. This ensures that for $\delta_0 = \min\{\delta_1, \delta_2\}$ the surface $S \cap B_{\delta_0}(0)$ has $H^S \geq \frac{9}{10}c$, $|A^S| + |\nabla A^S| \leq \frac{10}{9}c^{-1}$ and is $\frac{9}{10}\alpha$ -noncollapsed and moreover $\widehat{\gamma} = \exp_p^{-1}(\gamma) \cap B_{\delta_0}(0)$ is $\frac{9}{10}\mathbb{b}$ -controlled. We can therefore apply the Euclidean gluing result from Theorem 4.1 of [10] with these constants after extending $\exp_p^{-1}(K) \cap B_{\delta_0}(0)$ to some $\mathbb{A}' = (\frac{9}{10}\alpha, \frac{9}{10}c)$ -controlled domain \widehat{K} in \mathbb{R}^3 that satisfies $\partial\widehat{K} \cap B_{\delta_0}(0) = S$. This yields for every sufficiently small $r < \bar{r}$ a mean convex domain $\mathcal{G}_r(\widehat{K}, \widehat{\gamma})$ in \mathbb{R}^3 such that

$$\mathcal{G}_r(\widehat{K}, \widehat{\gamma}) \setminus B_{\delta(r)}(0) = \widehat{K} \cup N_r(\widehat{\gamma}) \setminus B_{\delta(r)}(0),$$

²More precisely, after picking an oriented orthonormal basis $\{E_1, E_2, E_3\}$ of $T_p M$ at p , we identify \mathbb{R}^3 with $T_p M$ via the linear isometry $\iota : \mathbb{R}^3 \rightarrow T_p M$ given by $\iota(x^1, x^2, x^3) = \sum_{k=1}^3 x^k E_k$.

where $\delta(r)$ and \bar{r} are given by Theorem 4.1 in [10]. Moreover, by possibly further decreasing \bar{r} we may also assume that $\delta(r) < \delta_0$. Tracing through the proof, we see that in fact $\mathcal{G}_r(\bar{K}, \bar{\gamma})$ is not only mean convex, but has an explicit lower bound $\widehat{c}(r) > 0$ for the mean curvature which only depends on $(\mathbb{A}', \mathbb{b}')$ (or equivalently on (\mathbb{A}, \mathbb{b})) and on r and does *not* degenerate as $r \rightarrow 0$, hence $\inf_{r \in (0, \bar{r})} \widehat{c}(r) =: \widehat{c} > 0$. We will also make sure that \bar{r} is small enough so that $\delta(r) < \frac{1}{4}\delta_3$, where $\delta_3 > 0$ is from Point 3 of Lemma 2.7 with $c_S = \widehat{c}$.

The region created this way is not yet exactly what we want, but the last adjustment is rather easy. For sufficiently small \bar{r} (depending only on \mathbb{b}), both $N_r(\bar{\gamma})$ and $\exp_p^{-1}(N_r(\gamma))$ have very large mean curvature, say at least $100\widehat{c}$. We can therefore easily interpolate from one region to the other in $B_{2\delta(r)}(0) \setminus B_{\delta(r)}(0)$ with the resulting region agreeing with $N_r(\bar{\gamma})$ on $B_{\frac{5}{4}\delta(r)}(0) \setminus B_{\delta(r)}(0)$ and with $\exp_p^{-1}(N_r(\gamma))$ on $B_{2\delta(r)}(0) \setminus B_{\frac{7}{4}\delta(r)}(0)$ and still having \widehat{c} as lower bound for the mean curvature everywhere.

Applying \exp_p in $B_{\delta_3}(0)$, by Point 3 of Lemma 2.7, we obtain the desired $\mathcal{G}_r(K, \gamma) \cap B_{\delta_3}(p)$, which will have $\frac{9}{10}\widehat{c}$ as a lower bound for its mean curvature and agree with $K \cup N_r(\gamma)$ outside $B_{2\delta(r)}(p)$. Finally, we note that if \bar{r} is small enough so that $\bar{r}, \delta(\bar{r}) < \mathbb{b}/10$, then the resulting domain is still contained in \bar{K} . This finishes the proof. \square

It remains to prove the technical lemma. The idea is that *at* p the second fundamental form and mean curvatures A^S and H^S of S in \mathbb{R}^3 and A^Σ and H^Σ of Σ in (M, g) agree and from there the difference between the two values will grow by a function of the distance to p – similarly for the curvature of curves. Hence restricting to sufficiently small balls, the errors will be small enough to conclude the claims of the lemma.

Proof of Lemma 2.7. Let us first focus on the Points 1 and 3 concerned with the second fundamental form and mean curvature of the two surfaces.

We use the diffeomorphism $\exp_p : \mathbb{R}^3 \supset B_\varrho(0) \rightarrow B_\varrho(p) \subset M$ (after identifying $T_p M$ with \mathbb{R}^3 via a linear isometry) to pull back the metric g on M to \mathbb{R}^3 , denoting the resulting metric with $\tilde{g} = \exp_p^* g$. This way, $\exp_p : (B_\varrho(0), \tilde{g}) \rightarrow (B_\varrho(p), g)$ is a Riemannian isometry.

We can then phrase the lemma in terms of the two different metrics on $B_\varrho(0) \subset \mathbb{R}^3$, proving that in sufficiently small balls all quantities we are interested in remain comparable in these different metrics. In particular, for a two-sided surface $S \subset B_r(0) \subset \mathbb{R}^3$, we denote by $A = h_{\alpha\beta}$ and H its second fundamental form and mean curvature in \mathbb{R}^3 with the Euclidean metric g_e , and by $\tilde{A} = \tilde{h}_{\alpha\beta}$ and \tilde{H} its second fundamental form and mean curvature in \mathbb{R}^3 with respect to the metric \tilde{g} . We then want to prove the following.

Claim 1. $\tilde{h}_{\alpha\beta}(q) = h_{\alpha\beta}(q) + O_p(\delta^2)$ and $\tilde{H}(q) = H(q) + O_p(\delta^2)$ for all $q \in S \cap B_\delta(0)$, where the errors $O_p(\delta^2)$ depend on the curvature of M and all its derivatives at the point $p \in \bar{K}$.

From this claim, the statement about the mean curvature in Point 1 of the lemma follows immediately by picking δ_1 sufficiently small so that $O_p(\delta_1^2) \leq \frac{1}{10}c_S$, while Point 3 follows by picking δ_3 small enough so that $O_p(\delta_3^2) \geq -\frac{1}{10}c_S$. Note that due to the compactness of \bar{K} , this can be done independently of the choice of p .

Proof of Claim 1. It is well known that our choice of (normal) coordinates implies that for all $x \in B_\rho(0) \subset \mathbb{R}^3$, we have

$$\tilde{g}_{ij}(x) = \delta_{ij} - \frac{1}{3}R_{ikj\ell} x^k x^\ell + O(|x|^3), \quad (2.1)$$

see for example [27]. We note that the expression $O(|x|^3)$ depends on the curvature in p and all its derivatives. In the following, we normally drop dependence on $x \in \mathbb{R}^3$ where obvious to keep notation shorter and we also do not always explicitly keep track of the quadratic terms.

Parametrise S by $\theta = (\theta^1, \theta^2, \theta^3) : U \rightarrow \mathbb{R}^3$, where $U \subset \mathbb{R}^2$. Using coordinates z^1, z^2 in $U \subset \mathbb{R}^2$, we have $\theta^i = \theta^i(z^1, z^2)$, for $i = 1, 2, 3$. We denote by $\theta_\alpha = (\theta_\alpha^1, \theta_\alpha^2, \theta_\alpha^3)$ the tangent vector $\frac{\partial}{\partial z^\alpha} \theta$, $\alpha = 1, 2$. In particular, we have the first fundamental form of S in (\mathbb{R}^3, g_e) given by

$$g_{\alpha\beta}^S = g_e(\theta_\alpha, \theta_\beta) = \delta_{ij} \theta_\alpha^i \theta_\beta^j,$$

where here and in the following calculations indices such as i, j, k, ℓ are in $\{1, 2, 3\}$ and Greek indices like α, β are in $\{1, 2\}$ and we use the Einstein summation convention. Working instead in the metric \tilde{g} on \mathbb{R}^3 , we find the first fundamental form of S at a point x in $(\mathbb{R}^3, \tilde{g})$ using (2.1),

$$\tilde{g}_{\alpha\beta}^S = \tilde{g}(\theta_\alpha, \theta_\beta) = \tilde{g}_{ij} \theta_\alpha^i \theta_\beta^j = g_{\alpha\beta}^S - \frac{1}{3}R_{ikj\ell} x^k x^\ell \theta_\alpha^i \theta_\beta^j + O(|x|^3).$$

We point out that here and in the following, the terms $O(|x|^k)$ might depend on the vectors θ_1, θ_2 up to at most quadratic order. The inverses also satisfy a similar relation. For us, it is not necessary to have an explicit expression for the terms that are quadratic in x , but it is sufficient to note that

$$(\tilde{g}^S)^{\alpha\beta} = (g^S)^{\alpha\beta} + O(|x|^2). \quad (2.2)$$

The unit normal vector to S for the Euclidean background metric is given by $\nu = (\theta_1 \wedge \theta_2) / \|\theta_1 \wedge \theta_2\|$. We note that for sufficiently small ρ , the unit normal vector $\tilde{\nu}$ with respect to the background metric \tilde{g} lies very close to ν , and we can therefore make the ansatz $\tilde{\mu} = \nu + a^\alpha \theta_\alpha$ to obtain a (not yet normalised) normal vector. Indeed, the condition that $\tilde{\mu}$ is normal to θ_β with respect to the metric \tilde{g} yields

$$0 = \tilde{g}(\tilde{\mu}, \theta_\beta) = \tilde{g}(\nu, \theta_\beta) + a^\alpha \tilde{g}(\theta_\alpha, \theta_\beta) = -\frac{1}{3}R_{ikj\ell} x^k x^\ell \nu^i \theta_\beta^j + O(|x|^3) + a^\alpha \tilde{g}_{\alpha\beta}^S,$$

where in the last step we used the expansion (2.1) as well as $g_e(\nu, \theta_\beta) = 0$. Solving for a^α , we find

$$a^\alpha = (\tilde{g}^S)^{\alpha\beta} \left(\frac{1}{3}R_{ikj\ell} x^k x^\ell \nu^i \theta_\beta^j + O(|x|^3) \right) = (g^S)^{\alpha\beta} \left(\frac{1}{3}R_{ikj\ell} x^k x^\ell \nu^i \theta_\beta^j + O(|x|^3) \right).$$

In order to normalise $\tilde{\mu}$, we note that

$$\begin{aligned} \|\tilde{\mu}\|_{\tilde{g}}^2 &= \tilde{g}(\tilde{\mu}, \tilde{\mu}) = \tilde{g}(\tilde{\mu}, \nu + a^\beta \theta_\beta) = \tilde{g}(\tilde{\mu}, \nu) = \tilde{g}(\nu, \nu) + a^\beta \tilde{g}(\theta_\beta, \nu) \\ &= \|\nu\|_{g_e}^2 - \frac{1}{3}R_{ikj\ell} x^k x^\ell \nu^i \nu^j + O(|x|^3) + a^\beta (g_e(\theta_\beta, \nu) - \frac{1}{3}R_{ikj\ell} x^k x^\ell \theta_\beta^i \nu^j + O(|x|^3)) \\ &= 1 - \frac{1}{3}R_{ikj\ell} x^k x^\ell \nu^i \nu^j + O(|x|^3). \end{aligned}$$

On the first line of this calculation, we used that $\tilde{g}(\tilde{\mu}, \theta_\beta) = 0$, for the second line we used the expansion (2.1), and finally in the very last step we used $\|\nu\|_{g_e}^2 = 1$, $g_e(\theta_\beta, \nu) = 0$, and the fact that a^α grows quadratically in x . Therefore

$$\|\tilde{\mu}\|_{\tilde{g}}^{-1} = 1 + \frac{1}{6}R_{ikj\ell} x^k x^\ell \nu^i \nu^j + O(|x|^3) \quad (2.3)$$

and we find for the unit normal vector

$$\tilde{\nu} = \tilde{\mu}/\|\tilde{\mu}\|_{\tilde{g}} = \left(\nu + \frac{1}{3}(g^S)^{\alpha\beta} R_{ikjl} x^k x^\ell \nu^i \theta_\beta^j + O(|x|^3) \right) \left(1 + \frac{1}{6} R_{ikjl} x^k x^\ell \nu^i \nu^j + O(|x|^3) \right),$$

in particular $\tilde{\nu} = \nu + O(|x|^2)$. We then calculate the second fundamental form and the mean curvature of S with respect to the two different metrics. For the Euclidean metric, we have

$$h_{\alpha\beta} = -g_e(\nabla_{\theta_\alpha} \nu, \theta_\beta), \quad H = (g^S)^{\alpha\beta} h_{\alpha\beta},$$

and similarly for the background metric \tilde{g} , we have

$$\tilde{h}_{\alpha\beta} = -\tilde{g}(\nabla_{\theta_\alpha} \tilde{\nu}, \theta_\beta), \quad \tilde{H} = (\tilde{g}^S)^{\alpha\beta} \tilde{h}_{\alpha\beta}.$$

Multiplying with $\|\tilde{\mu}\|_{\tilde{g}}$, we can calculate the second fundamental form in the latter case as

$$\|\tilde{\mu}\|_{\tilde{g}} \tilde{h}_{\alpha\beta} = -\tilde{g}(\nabla_{\theta_\alpha} \tilde{\mu}, \theta_\beta) = -\tilde{g}(\nabla_{\theta_\alpha} \nu, \theta_\beta) - a^\gamma \tilde{g}(\nabla_{\theta_\alpha} \theta_\gamma, \theta_\beta)$$

The first of the terms on the right hand side will give us $h_{\alpha\beta}$ as well as error terms at least quadratic in x . The second term on the other hand will always be at least quadratic in x because of the factor a^γ . We can therefore arrive at

$$\|\tilde{\mu}\|_{\tilde{g}} \tilde{h}_{\alpha\beta} = h_{\alpha\beta} + O(|x|^2).$$

Proving this precisely requires a rather lengthy calculation (involving also estimates for the Christoffel symbols with respect to the different background metrics), and in an attempt to keep this article short and because the calculation does not offer any useful further insight, we skip the details here. We refer the reader to [13] where such a calculation has been carried out in full detail in the more specific situation where S is a perturbation of a sphere.

Now using first our formula (2.3) we obtain $\tilde{h}_{\alpha\beta} = h_{\alpha\beta} + O(|x|^2)$ and then using (2.2), we deduce that $\tilde{H} = H + O(|x|^2)$, concluding the claim. \square

The statement about the second fundamental form in Point 1 of the lemma follows by first showing that the first derivatives of the second fundamental forms satisfy a similar relation as the second fundamental forms themselves. This again involves estimating Christoffel symbols and we leave the details to the reader. The claimed estimates involving norms of the second fundamental form and its derivative then follow from (2.2).

We also note that S is α -noncollapsed (with respect to the pulled back metric \tilde{g}) if for each point $x \in S$ the balls with radius $\tilde{r}(x) = \alpha/\tilde{H}(x)$ and center $\tilde{C} = x \pm \tilde{r}(x)\tilde{\nu}(x)$ touch S only in the point x . By Claim 1, for δ sufficiently small, the mean curvature $H(x)$ of S at x with respect to the Euclidean metric will be extremely close to $\tilde{H}(x)$. Moreover, as seen in the proof of Claim 1, the unit normal vectors $\nu(x)$ and $\tilde{\nu}(x)$ will be very close as well. Therefore, for sufficiently small δ , the above balls contain the balls with radius $r(x) = \frac{9}{10}\alpha/H(x)$ and center $C = x \pm r(x)\nu(x)$. It follows that with respect to the Euclidean metric, S is $\frac{9}{10}\alpha$ -noncollapsed, finishing the proof of Point 1 of the lemma.

Finally, the remaining Point 2 of the lemma follows in a similar way as the other two points and we leave the direct but lengthy calculation to the reader. \square

3 Mean curvature flow with surgery

In this section we provide a short overview of the main definitions and results from the theory of mean curvature flow with surgery, following the terminology of Haslhofer-Kleiner [15] (in Euclidean space) and its extension by Haslhofer-Ketover [14] to the ambient three-manifold setting. We also refer the interested reader to slightly earlier work of Brendle-Huisken [9], where a similar theory of mean curvature flow with surgery in three-manifolds has been developed, extending their prior results from [8] in Euclidean \mathbb{R}^3 (which itself extended results of Huisken-Sinestrari [17]).

As we have seen at the beginning of Section 2, in the situations we are interested in, rather than working with mean convex surfaces we can work with the mean convex domains bounded by them. A smooth family of mean convex domains $\{K_t\}_{t \in I}$ in a complete three-manifold (M, g) is said to move by *mean curvature flow* if every point in the boundary ∂K_t evolves with time derivative equal to the mean curvature vector of the boundary. It is a well known fact that mean convexity is preserved along the mean curvature flow, see also (3.1) below.

The reason to introduce a surgery procedure, in a nutshell, is that typically mean curvature flow develops singularities where it can no longer be smoothly extended in time. In the general case, singularities can be very complicated (for instance, singularity models can have arbitrary genus, see [12]), but in the mean convex case all local singularities have a (possibly degenerate) cylindrical neck structure. Hence, to avoid dealing with the singularities, one stops the flow shortly before the singularity occurs, replaces neck regions by pairs of opposing caps, discards regions of high curvature, and then restarts the flow. This surgery procedure allows to extend the flow until it becomes extinct. In the following, we will recall the rigorous definitions and existence results from [14, 15].

Remark 3.1. Haslhofer-Ketover [14] work in closed Riemannian manifolds, while our (M, g) is only complete. This however does not create any additional complications because our evolving domains K_t are all contained in the compact domain \bar{K} , hence we could change (M, g) into a closed manifold by modifying it outside \bar{K} .

We want to define a mean curvature flow with surgery starting from an element K_0 in \mathcal{D} (allowing in particular also $K_0 = \bar{K}$), but for simplicity, following [14, 15], we first define a more general class of flows called (α, δ) -flows.

Definition 3.2 ((α, δ) -flows, see [14, 15]). An (α, δ) -flow is a collection of smooth families of mean convex domains $\{K_t^i\}_{t \in [t_{i-1}, t_i]}$ ($i = 1, \dots, k$ and $0 = t_0 < t_1 < \dots < t_k < \infty$), such that:

1. For each t , the domain K_t is α -noncollapsed as in Definition 2.3³
2. For each i , the final time slices of some (possibly empty) collection of disjoint strong δ -necks are replaced by pairs of opposing standard caps as described in Definition 3.3 below, yielding a post-surgery domain $K_{t_i}^\sharp \subseteq K_{t_i}^i =: K_{t_i}^-$.
3. For each i , the initial time slice of the next flow $K_{t_i}^{i+1} =: K_{t_i}^+$ is obtained from the post-surgery domain $K_{t_i}^\sharp$ by discarding some (possibly empty) collection of connected components.

³While α -noncollapsedness is preserved in Euclidean space, this is not the case in ambient manifolds; but by an estimate of Brendle [7] the noncollapsing parameter can decay at most exponentially and therefore one could also assume this condition only for the initial domain.

4. All surgeries are at comparable scales, meaning there exists a radius $s_0 > 0$, such that all δ -necks in Point 2 of this definition have radius $s \in [\frac{1}{2}s_0, 2s_0]$ for each i .

Let us discuss the second point of this definition in more detail.

An (α, δ) -flow has a δ -neck with center p and radius s at time t_0 if $4s/\delta \leq \text{inj}(\bar{K})$ as well as $s^{-1} \exp_p^{-1}(K_{t_0} \cap B_{2s/\delta}(p))$ is δ -close in $C^{[1/\delta]}$ in $B_{1/\delta} \cap s^{-1} \exp_p^{-1}(B_{2s/\delta}(p)) \subset \mathbb{R}^3$ to a solid round cylinder $D^2 \times \mathbb{R}$ with radius 1. Moreover, it has a *strong* δ -neck with center p and radius s at time t_0 if $4s/\delta \leq \text{inj}(\bar{K})$ and $\{s^{-1} \exp_p^{-1}(K_{t_0+s^2t} \cap B_{2s/\delta}(p))\}_{t \in (-1, 0]}$ is δ -close in $C^{[1/\delta]}$ in $B_{1/\delta} \cap s^{-1} \exp_p^{-1}(B_{2s/\delta}(p)) \subset \mathbb{R}^3$ to the evolution of a solid round cylinder $D^2 \times \mathbb{R}$ with radius 1 at $t = 0$. Finally, a *standard cap* is a smooth, closed, convex, connected, and unbounded codimension 0 submanifold with boundary $K_{st} \subset \mathbb{R}^3$ that coincides with a solid round half-cylinder of radius 1 outside a ball of radius 10. We can create a rotationally symmetric standard cap by smoothing out the set obtained from a half-ball attached to a half-cylinder. In the following, we assume that some standard cap has been fixed.

We can now define what we mean by replacing a strong δ -neck with two opposing standard caps in the above definition.

Definition 3.3 (Surgery procedure, see [14, 15]). After fixing a large cap-separation parameter Γ , we say that the final time slice of a strong δ -neck ($\delta \leq 1/10\Gamma$) with center p and radius s is *replaced by a pair of standard caps*, if (in a neighbourhood of p) the pre-surgery domain K^- is replaced by a post-surgery domain $K^\sharp \subseteq K^-$ such that the following properties hold.

1. The replacement takes place inside $B = B_{5\Gamma s}(p)$.
2. For each integer $\ell \geq 0$ there exists a constant C_ℓ , such that $\sup_{\partial K^\sharp \cap B} |\nabla^\ell A| \leq C_\ell s^{-(1+\ell)}$.
3. For each $p^\sharp \in K^\sharp$ with negative lower principal curvature $\lambda_1(p^\sharp) < 0$ there is a corresponding $p^- \in K^-$ with $\lambda_1/H(p^\sharp) \geq \lambda_1/H(p^-) - o(s)$.
4. $s^{-1} \exp_p^{-1}(K^\sharp)$ is $o(\delta)$ -close in $B_{10\Gamma}(0) \subset \mathbb{R}^3$ to a pair of opposing standard caps at distance Γ from the origin. Moreover there is an ‘‘almost straight line’’ between the tips of the caps, i.e. a curve that is $o(\delta)$ -close to a straight line in \mathbb{R}^3 and satisfies the estimates on κ and $\nabla \kappa$ of an $o(\delta)^{-1}$ -controlled curve.

We are finally ready to give the main definition of mean curvature flow with surgery, refining the definition of (α, δ) -flow.

Definition 3.4 (MCF with surgery, see [14, 15]). Fix a large cap-separation parameter Γ and let $\mathbb{H} = (H_{\text{trig}}, H_{\text{neck}}, H_{\text{th}})$ be a triple of positive numbers with $H_{\text{trig}} > H_{\text{neck}} > H_{\text{th}}$, referred to as trigger, neck, and thick curvature scales. An $(\alpha, \delta, \mathbb{H})$ -flow (also called a mean curvature flow with surgery) is an (α, δ) -flow with the following properties.

1. $H \leq H_{\text{trig}}$ everywhere, and surgery/discarding occurs precisely at times when $H = H_{\text{trig}}$ somewhere.
2. The collection of necks in Point 2 of Definition 3.2 is a minimal collection of solid strong δ -necks of curvature H_{neck} which separate $\{H = H_{\text{trig}}\}$ from $\{H \leq H_{\text{th}}\}$ in K_t^- .

3. The collection of connected components that are discarded in Point 3 of Definition 3.2 contains precisely the components with $H > H_{\text{th}}$ everywhere.⁴
4. Replacements are done in an optimal way in the sense that if a strong δ -neck in Point 2 is in fact a strong $\widehat{\delta}$ -neck for some $\widehat{\delta} < \delta$, then Point 4 of Definition 3.3 also holds with this $\widehat{\delta}$.

Theorem 3.5 (Existence Theorem, Theorem 7.7 in [14]). *Let $K_0 \in \mathcal{D}$ be a mean convex domain in \bar{K} (allowing also $K_0 = \bar{K}$). Then for every $T < \infty$, choosing α and δ small enough as well as scales \mathbb{H} large enough⁵, there exists an $(\alpha, \delta, \mathbb{H})$ -flow $\{K_t\}_{t \in [0, T]}$ with initial condition K_0 .*

An important point is that \bar{K} is compact, thus the theory developed for closed manifolds also works, despite the fact that our (M, g) is only complete, as explained in Remark 3.1. Generally, some components of a mean curvature flow with surgery might exist for all time and converge to a minimal surface in the ambient three-manifold, but in our specific case where the ambient manifold M has nonnegative Ricci curvature $\text{Ric}_g \geq 0$, we always obtain a finite extinction result.

Proposition 3.6 (Finite extinction). *Let $K_0 \in \mathcal{D}$ be a mean convex domain in \bar{K} satisfying $\min_{p \in \partial K_0} H(p) \geq c > 0$. Then every mean curvature flow with surgery starting from K_0 as in Theorem 3.5 becomes extinct in time $T_{\text{ext}} < c^{-2}$.*

Proof. The evolution equation for the mean curvature (of ∂K_t) in the ambient manifold (M, g) with $\text{Ric}_g \geq 0$ yields

$$\partial_t H = \Delta H + |A|^2 H + \text{Ric}_g(\nu, \nu)H \geq \Delta H + \frac{1}{2}H^3$$

and hence in particular

$$\partial_t H_{\min}(t) \geq \frac{1}{2}H_{\min}^3(t), \tag{3.1}$$

where $H_{\min}(t) := \min_{p \in \partial K_t} H(p)$. The surgery procedure does not influence this bound until the entire flow becomes extinct. Therefore, if the initial domain K_0 satisfies $H_{\min}(0) \geq c > 0$, then by the comparison principle $H_{\min}(t) \geq (c^{-2} - t)^{-1/2} \nearrow \infty$ as $t \rightarrow c^{-2}$ and in particular the flow must become extinct in time $T_{\text{ext}} < c^{-2}$ (which will also depend on the choice of scales \mathbb{H}). Again, the comparison principle can be applied directly since we only work in the compact subset \bar{K} of (M, g) . \square

While the surgery procedure cuts our domain into several pieces, the gluing construction from Section 2 can be used to connect it again with a very thin tubular string, using in particular also that all domains and curves are contained in \bar{K} . We glue in a string along the image γ under the exponential map of the almost straight line between two opposing standard caps in a post-surgery region (as in Point 4 of Definition 3.3). We call γ an *almost geodesic*. An important result is that there is an isotopy between the pre-surgery domain and the domain obtained by first performing surgery and then gluing together again. Note that this latter domain is contained in the pre-surgery domain and hence in particular also in \bar{K} . It is important to preserve mean convexity, so we introduce the concept of *mean convex isotopy* which is simply a smooth family $\{K_t\}$ of mean convex domains indexed by a real interval. We then have the following result that generalises [10, Proposition 6.6] from the Euclidean case to the ambient three-manifold setting.

⁴In particular, this means that of each pair of opposing caps precisely one is discarded.

⁵To be precise, also the ratios $H_{\text{trig}}/H_{\text{neck}}$ and $H_{\text{neck}}/H_{\text{th}}$ need to be sufficiently large.

Proposition 3.7 (Isotopy from necks to glued caps). *Assume that K^\sharp is obtained from K^- by replacing a strong δ -neck by a pair of standard caps and assume that γ denotes the almost geodesic connecting the tips of the caps in M . Then for r small enough there is a mean convex isotopy between K^- and $\mathcal{G}_r(K^\sharp, \gamma)$.*

Proof. Denote by p_\pm the two gluing points and call their distance d . Set $\tilde{K}_\pm = \exp_{p_\pm}^{-1}(K^-)$ and $\tilde{G}_\pm^1 = \exp_{p_\pm}^{-1}(\mathcal{G}_r(K^\sharp, \gamma))$. If we work with H_{neck} sufficiently large and therefore the surgery happens in a sufficiently small ball around p_\pm , these regions are mean convex in \mathbb{R}^3 by Lemma 2.7. As in the proof of Theorem 2.5 we can easily interpolate between \tilde{G}_\pm^1 and the ‘‘Euclidean gluing’’ $\tilde{G}_\pm^2 = \mathcal{G}_r(\exp_{p_\pm}^{-1}(K^\sharp), \exp_{p_\pm}^{-1}(\gamma))$ inside balls of radius $\frac{2}{3}d$ around the gluing points, because $\exp_{p_\pm}^{-1}(N_r(\gamma))$ and $N_r(\exp_{p_\pm}^{-1}(\gamma))$ have very large mean curvature (and the domains otherwise agree by construction). This gives a mean convex isotopy between \tilde{G}_\pm^1 and \tilde{G}_\pm^2 in balls of radius $\frac{2}{3}d$ around p_\pm and brings us back to the Euclidean case where according to [10, Proposition 6.6] there also exists a mean convex isotopy between \tilde{K}_\pm and \tilde{G}_\pm^2 . Combining the two mean convex isotopies, applying Lemma 2.7 one more time to go back to M , and interpolating in the regions where the balls overlap proves the proposition. \square

The final ingredient that is needed is a description of the possible types of discarded components. For the following definition, we need the notion of a standard cap K_{st} given above, as well as of an (C, ε) -cap, by which we mean a closed, strictly convex, connected, and unbounded codimension 0 submanifold with boundary $K_{C, \varepsilon} \subset \mathbb{R}^3$ such that every point outside some compact subset of $\text{Int } K$ with diameter at most C is the center of some ε -neck of radius 1.

Definition 3.8 (Capped tubes and tubular loops). A *capped ε -tube* is a mean convex domain K diffeomorphic to \bar{B}^3 such that there is a curve $\gamma \subset K$ with $\partial\gamma = \gamma \cap \partial K = \{\bar{p}_+, \bar{p}_-\}$ and such that:

1. $H(\bar{p}_\pm)^{-1} \exp_{\bar{p}_\pm}^{-1}(K \cap B_{2CH(\bar{p}_\pm)^{-1}}(\bar{p}_\pm))$ is ε -close, after rescaling, to either a (C, ε) -cap $K_{C, \varepsilon}$ or a standard cap K_{st} .
2. For every $p \in \gamma$, if $d(p, \bar{p}_+) \geq CH(\bar{p}_+)^{-1}$ and $d(p, \bar{p}_-) \geq CH(\bar{p}_-)^{-1}$ then p is the center of some ε -neck of some radius s with axis tangent to γ , and such that γ is $\varepsilon^{-2}s$ -controlled in $B_{\varepsilon^{-1}s}(p)$.

An *ε -tubular loop* is a mean convex solid torus admitting a core γ with the following property: every point $p \in \gamma$ is center of an ε -neck of some radius s with axis tangent to γ , and such that γ is $\varepsilon^{-2}s$ -controlled in $B_{\varepsilon^{-1}s}(p)$.

A direct consequence of the canonical neighbourhood theorems (Theorem 1.22 in [15] and Theorem 7.6 in [14]) and Corollary 1.25 in [15] is the following result (see also [10, 11] for the corresponding statement in Euclidean space).

Proposition 3.9 (Discarded components). *All discarded components of the $(\alpha, \delta, \mathbb{H})$ -flow $\{K_t\}$ from Theorem 3.5 are diffeomorphic to \bar{B}^3 or to $D^2 \times \mathbb{S}^1$. Moreover, the components that are diffeomorphic to \bar{B}^3 are either (a) convex or (b) capped ε -tubes; the components diffeomorphic to $D^2 \times \mathbb{S}^1$ are (c) ε -tubular loops.*

4 Reduction to sphere or tubular neighbourhood of a curve

We still assume that \bar{K} is a compact domain in a complete orientable Riemannian three-manifold (M, g) with nonnegative Ricci curvature. The radii $r_s \ll r_m$ in the following definition have to be thought of as being extremely small.

Definition 4.1. A *marble graph* in \bar{K} with string radius r_s and marble radius r_m is a domain of the form $G = \mathcal{G}_{r_s}(K, \gamma)$ where (K, γ) is a (\mathbb{A}, \mathbb{b}) -controlled configuration such that there exists a finite set $P \subset \bar{K}$ such that

1. K is the disjoint union of the balls $\bar{B}_{r_m}(p) \subset \bar{K}$ for $p \in P$.
2. for every $p \in P$, $\gamma \cap \bar{B}_{10r_m}(p)$ is a union of geodesic segments.

We say that G is a *marble tree* (respectively a *marble circuit*) if $K \cup \gamma$ is contractible (respectively homotopy equivalent to a circle.) To shorten the proofs below a little bit, we sometimes work without the condition 2 and call the resulting object an *almost marble graph*.

In this section, we show the following results.

Theorem 4.2. *Every mean convex domain K_0 in \bar{K} is mean convex isotopic to a marble graph.*

Theorem 4.3. *Every marble tree in \bar{K} is mean convex isotopic to a closed geodesic ε -ball and every marble circuit is mean convex isotopic to the solid ε -tubular neighbourhood of a curve and this curve is a core of the marble circuit (which is a solid torus).*

Theorem 1.1 will essentially be a direct consequence of these results while the proofs of our other main results will require further topological arguments in the next section. We start with the proof of Theorem 4.2.

Proof of Theorem 4.2. The proof is essentially the same as in the work of the first author with Haslhofer and Hershkovits [10, 11] with only very minor modifications. We therefore try to be brief.

Given a mean convex domain K_0 in the fixed compact domain \bar{K} , fix α and δ small enough as well as scales \mathbb{H} and cap-separation parameter Γ large enough and then flow K_0 by mean curvature flow with surgery according to the existence result from Theorem 3.5. By Proposition 3.6 this flow will become extinct in finite time and by Proposition 3.9 we know how the discarded components look like. We note that everything happens in the compact domain \bar{K} , thus allowing a direct application of Theorem 2.5.

We first claim that for sufficiently small $r_s \ll r_m$ (much smaller than the neck size) all discarded components from this flow are mean convex isotopic to almost marble graphs, compare with Proposition 7.4 of [10] and Lemma 3.6 of [11]. While this is obvious for convex discarded components (that are isotopic to a single marble), for capped ε -tubes and ε -tubular loops we first perform a large number of surgeries (i.e. replacements of necks by opposing caps as in Definition 3.3) obtaining a domain K^\sharp , and then glue back in strings of radius r_s along almost geodesics γ using our gluing result from Theorem 2.5. By Proposition 3.7 the resulting domain $\mathcal{G}_{r_s}(K^\sharp, \gamma)$ is mean convex isotopic to the original discarded component. If this is done in such a way that the connected components of K^\sharp have sufficiently small diameter, then these are all very small perturbations of short capped off cylinders and hence isotopic to small balls of radius, say, $10r_m$. Denote by $\{K_i^\sharp\}$

the mean convex isotopy from K^\sharp to this union of balls and by $\{\gamma_t\}$ the family of curves starting with γ from above that are then extended such that the end points follow the isotopy $\{K_t^\sharp\}$ by normal motion. Then $\mathcal{G}_{r_s}(K_t^\sharp, \gamma_t)$ is a mean convex isotopy from $\mathcal{G}_{r_s}(K^\sharp, \gamma)$ to an almost marble graph.

We then use a backwards induction scheme described in more detail in [10] that essentially works as follows. At the extinction time t_k there are only discarded components. Each of them is isotopic to an almost marble graph as just described. This gives a mean convex isotopy from $K_{t_k}^-$ to an almost marble graph (possibly disconnected). We then follow the mean curvature flow back to the previous surgery time t_{k-1} yielding an extension to a mean convex isotopy between $K_{t_{k-1}}^+$ and the same almost marble graph. Each component that was discarded at t_{k-1} is also isotopic to an almost marble graph and thus $K_{t_{k-1}}^\sharp$ is mean convex isotopic to an almost marble graph (namely the union of all previously collected almost marble graphs). Since $K_{t_{k-1}}^\sharp$ has been obtained from $K_{t_{k-1}}^-$ by replacing necks by opposing caps, we can follow the exact same strategy as in the paragraph above: first glue back in a string along almost geodesics γ and use Proposition 3.7 to obtain a mean convex isotopy from $K_{t_{k-1}}^-$ to $\mathcal{G}_{r_s}(K_{t_{k-1}}^\sharp, \gamma)$. Next use the isotopy of $K_{t_{k-1}}^\sharp$ to an almost marble graph and extend γ following this isotopy by normal motion until all the end points arrive on some marble, yielding a mean convex isotopy of the pre-surgery domain $K_{t_{k-1}}^-$ to a (new) almost marble graph (obtained from the old one by gluing in additional strings). We can follow this procedure back in time until we arrive at a mean convex isotopy from the initial domain K_0 to an almost marble graph. Finally, we shrink down the marbles from radius $10r_m$ to radius r_m and extend the strings radially in order to obtain a proper marble graph rather than just an almost marble graph.

There is only one additional technical detail that we have ignored in the previous paragraph: if at some point γ is extended through regions that are modified by later surgeries, we artificially move it out of these surgery regions. This argument works in the exact same way as in the Euclidean case described in [10, Lemma 9.4]. The proof is then complete. \square

Proof of Theorem 4.3. Again, the proof follows [10, 11]. An important ingredient is a rearrangement result, which allows to move curves/strings meeting a fixed marble to meet that marble in points of our choice. Intuitively, we shrink the marble radius to create “free space” in which we can move the curves around. Shrinking for example by factor 20 will create enough space to first move the curves around (in the outer region of the free space created) and then let them be geodesic segments before meeting the marble (in the inner region of the new free space). This procedure makes our control parameters worse, but if we enlarge the whole manifold (M, g) by the same factor, then the control parameters remain the same (while ambient curvature terms decrease) and therefore Theorem 2.5 still works the same, showing that without rescaling everything goes through if we also decrease the string radius by factor 20. This rearrangement result was proved precisely in Euclidean space as Lemma 5.3 of [10] and the exact same argument also works in an ambient manifold using geodesic normal coordinates with origin at the center of the marble where we rearrange curves.

Let us now first consider the case of a *marble tree*. If there is only one marble, we are done. Else, fix a leaf of the tree, i.e. a marble \bar{B}_1 that has only one string attached and call \bar{B}_2 the unique marble that is connected to \bar{B}_1 . Say the curve $\gamma_{1,2}$ along which the two marbles have been glued together meets \bar{B}_2 at its north pole. By what we have just discussed above, we can move all other curves/strings that meet \bar{B}_2 to the southern hemisphere (a step that requires us to shrink both marble and string radius). We then move \bar{B}_1 along the curve $\gamma_{1,2}$ closer to \bar{B}_2 (possibly violating Point 2 of Definition 4.1 at marble \bar{B}_1 for some time along the way) until it has distance $10r_m$ to

\bar{B}_2 . Now the connecting curve $\gamma_{1,2}$ is a geodesic segment. We then use the exponential map based at the center of \bar{B}_2 to lift the domain to the tangent space (where $\gamma_{1,2}$ becomes a straight line) and then use the Euclidean marble reduction result from Proposition 5.5 of [10] to completely remove the leaf marble \bar{B}_1 and connecting curve $\gamma_{1,2}$ in a mean convex way. By Lemma 2.7, applying the exponential map to this mean convex isotopy, we find the desired mean convex isotopy in (M, g) . We can repeat this step until we are only left with one marble which is our desired ε -ball.

In the case of a *marble circuit*, we first use the above argument to get rid of all leaf marbles. In the resulting marble circuit there are two curves meeting each marble and by the rearrangement argument from the first paragraph of the proof, we can move them in such a way that they meet the marble at antipodal points. Let us now focus on one marble: there are two curves γ_1 and γ_2 meeting the marble at antipodal points as radial geodesics. We can then join these geodesics with a geodesic segment in the interior of the marble. We then lift the problem again to the tangent space via the exponential map based at the center of the marble. This brings us back to the Euclidean case (or more precisely, in fact to the model case of a tube along a straight line glued to a ball). By an easy argument (see Theorem 4.1 of [11]), we can find a mean convex isotopy deforming this domain to a solid tube. This way we get rid of the marble and are only left with the tubular neighbourhood of a geodesic where we previously had a marble. Performing this step with all marbles, we find our solid ε -tubular neighbourhood of a curve. It is clear that the curve is a core of this resulting solid torus and thus also of the original marble circuit. The proof of the theorem is complete. \square

5 Proofs of the main results

We start with the proof of Theorem 1.1 which follows directly from what we have done in the previous sections.

Proof of Theorem 1.1. Let (M, g) be a complete, orientable Riemannian three-manifold with non-negative Ricci curvature. Take two elements of $\mathcal{M}_{H>0}(\mathbb{S}^2, M)$, represented by two embedded mean convex two-spheres Σ_1 and Σ_2 . By Proposition 2.1, Σ_i bounds a mean convex domain \bar{K}_i ($i = 1, 2$) which by Theorems 4.2 and 4.3 is mean convex isotopic to a small ball $\bar{B}_{\varepsilon_i}(p_i)$ inside \bar{K}_i . (Note that all steps are performed inside the compact regions \bar{K}_i). Connect p_1 and p_2 with each other along a curve γ . Let $0 < \varepsilon \leq \varepsilon_i$ be small enough such that *all* balls $\bar{B}_\varepsilon(p)$ with $p \in \gamma$ are mean convex. Then we have another mean convex isotopy from $\bar{B}_{\varepsilon_1}(p_1)$ to $\bar{B}_{\varepsilon_2}(p_2)$ by first shrinking the ball to size ε , then moving it along γ and finally expanding it again to size ε_2 . Combining everything \bar{K}_1 is mean convex isotopic to \bar{K}_2 and therefore Σ_1 and Σ_2 are in the same path component of the moduli space, proving the theorem. \square

We are left with proving the two theorems about Heegaard tori. For this reason, we fix a lens space $L(p, q)$ and a metric g on it having nonnegative Ricci curvature. As $L(p, q)$ is compact, we can simply work with $\bar{K} = L(p, q)$. Moreover, we fix once and for all a Heegaard torus Σ^0 in $L(p, q)$ and denote by V_1^0 and V_2^0 the two solid tori whose common boundary is Σ^0 . We also fix for each $k \in \{1, 2\}$ a core γ_k^0 of V_k^0 that is a controlled curve.

For every $\varepsilon > 0$, we denote by V_k^ε the ε -neighbourhood of γ_k and by T_k^ε its boundary, $k \in \{1, 2\}$. We choose ε small enough so that V_k^ε is a mean convex domain, and for every $0 < \varepsilon' < \varepsilon$ the domain $V_k^{\varepsilon'}$

is mean convex isotopic to V_k^ε . In particular, the path components of T_1^ε and T_2^ε in $\mathcal{M}_{H>0}(p, q)$ are well-defined.

Proposition 5.1. *Let Σ be a mean convex Heegaard torus Σ and let V be the unique mean convex domain such that $\partial V = \Sigma$. Then there exists $k \in \{1, 2\}$ such that V is mean convex isotopic to V_k^ε . Moreover, if $q \not\equiv \pm 1 \pmod p$, then such a k is unique.*

Proof. Since Σ is Heegaard, V is a solid torus. By Bonahon's theorem (cf. [5, Théorème 1]), Σ is ambient isotopic to Σ^0 , which means that there is a continuous 1-parameter family ϕ_t of self diffeomorphisms of $L(p, q)$ such that ϕ_0 is the identity of $L(p, q)$ and $\phi_1(\Sigma) = \Sigma^0$. We choose $k \in \{1, 2\}$ so that $\phi_1(V) = V_k^0$.

By Theorems 4.2 and 4.3, V is mean convex isotopic to the solid ε -tubular neighbourhood V^ε of some controlled curve γ . It follows that V^ε is isotopic to V_k^0 . Hence their cores γ and γ_k^0 are also isotopic. If ε is chosen small enough, then the tubular neighbourhoods V^ε and V_k^ε are mean convex isotopic. Thus V is mean convex isotopic to V_k^ε .

If $q \not\equiv \pm 1 \pmod p$, then by a classical computation (cf. [6, Section 6]) the curves γ_1^0 and γ_2^0 are not homologous, hence not isotopic. This implies that V_1^ε is not isotopic to V_2^ε . \square

We now turn to the proof of Theorem 1.4, knowing already that $\mathcal{M}_{H>0}(p, q)$ has at most two components by Proposition 5.1.

Proof of Theorem 1.4. There are two cases.

Case 1. $q \equiv \pm 1 \pmod p$. In this case we have to prove that $\mathcal{M}_{H>0}(p, q)$ is path-connected.

In view of Proposition 5.1, there only remains to show that T_1^ε and T_2^ε are mean convex isotopic to each other. By [5, Lemme 4], γ_1^0 is isotopic to γ_2^0 . Again, choosing ε small enough it follows that T_1^ε is mean convex isotopic to T_2^ε .

Case 2. $q \not\equiv \pm 1 \pmod p$. Here we must show that T_1^ε and T_2^ε are *not* mean convex isotopic to each other.

We set $H = H_1(L(p, q), \mathbb{Z})/\pm 1$. Each simply closed curve γ in $L(p, q)$ defines a class $[\gamma]$ in H . As we recalled above, a homology computation shows that the hypothesis $q \not\equiv \pm 1 \pmod p$ implies that $[\gamma_1] \neq [\gamma_2]$.

Let T be a mean convex Heegaard torus in $L(p, q)$ and V be the unique mean convex solid torus such that $\partial V = T$. By Proposition 5.1, there exists a unique $k \in \{1, 2\}$ such that V is mean convex isotopic to V_k^ε . We set $\phi(T) = [\gamma_k]$.

This defines a continuous map ϕ from $\mathcal{M}_{H>0}(p, q)$ to the discrete space H with the following property: for each $k \in \{1, 2\}$, we have $\phi(T_k^\varepsilon) = [\gamma_k]$. This entails that T_1^ε and T_2^ε belong to two different path components of $\mathcal{M}_{H>0}(p, q)$. This completes the proof in Case 2, hence that of Theorem 1.4. \square

Finally, we prove Theorem 1.6.

Proof of Theorem 1.6. First note that if $H \geq 0$ on Σ , then unless $H \equiv 0$, we will *instantaneously* obtain positive mean curvature when starting the mean curvature flow from Σ . This follows by applying the maximum principle to the evolution equation

$$\partial_t H = \Delta H + |A|^2 H + \text{Ric}_g(\nu, \nu)H \geq \Delta H + \frac{1}{2}H^3.$$

This means that tori with $H \geq 0$ and $H \not\equiv 0$ lie on the boundary of the moduli space $\mathcal{M}_{H>0}(p, q)$ and therefore do not add to the number of path components.

It remains to consider the case of minimal Heegaard tori, i.e. Heegaard tori with $H \equiv 0$. For such a torus, denote by V_1 and V_2 the two solid tori bounded by Σ and let ν be the unit normal vector field on Σ pointing towards V_1 . We then consider the stability operator $L_\Sigma \varphi = \Delta_\Sigma \varphi + (|A^\Sigma|^2 + \text{Ric}_g(\nu, \nu))\varphi$, where $\varphi : \Sigma \rightarrow \mathbb{R}$. If the largest eigenvalue λ of L_Σ is positive (meaning that Σ is unstable) then we define for $t \in (-\varepsilon, \varepsilon)$ the family of tori

$$\Sigma_t = \exp_{x \in \Sigma}(t\varphi_\lambda(x)\nu(x)),$$

where φ_λ is a normalised, positive eigenfunction of the eigenvalue λ . Also consider unit vector fields ν_t on Σ_t , continuous in t and such that $\nu_0 = \nu$. We obtain

$$\left. \frac{\partial}{\partial t} \right|_{t=0} H_{\Sigma_t} = L_\Sigma \varphi_\lambda = \lambda \varphi_\lambda,$$

and therefore

$$H_{\Sigma_t} = t\lambda\varphi_\lambda + O(t^2).$$

In particular, for $t > 0$, Σ_t is mean convex with respect to the normal vector field ν_t and by our previous arguments Σ_t is mean convex isotopic to a tubular neighbourhood of a core γ_1 of V_1 . On the other hand, for $t < 0$, Σ_t is mean convex with respect to $-\nu_t$ and therefore mean convex isotopic to a tubular neighbourhood of a core γ_2 of V_2 . This means in particular that if $\mathcal{M}_{H>0}(p, q)$ has two path components, Σ is ($H \geq 0$)-isotopic to mean convex Heegaard tori in *both* components and thus $\mathcal{M}_{H \geq 0}(p, q)$ is path-connected.

If instead the largest eigenvalue is zero (i.e. Σ is weakly stable), then $A^\Sigma \equiv 0$ and $\text{Ric}_g(\nu, \nu) \equiv 0$. This means that if we consider

$$\Sigma_t = \exp_{x \in \Sigma}(t\nu(x)),$$

and extend ν correspondingly (with ν_t the unit normal vector field to Σ_t), then either for some $t_+ > 0$ we find $\text{Ric}_g(\nu_{t_+}, \nu_{t_+}) \not\equiv 0$ or $\text{Ric}_g(\nu_t, \nu_t) \equiv 0$ for all $t > 0$. In the latter case, V_1 splits isometrically as $\Sigma \times [0, \infty)$ which is a contradiction since V_1 is a solid torus. Therefore we must be in the former case. Now at t_+ , the stability operator $L_{\Sigma_{t_+}}$ has a positive eigenvalue and we can apply the above construction starting at Σ_{t_+} , eventually finding that Σ is ($H \geq 0$)-isotopic to a tubular neighbourhood of a core γ_1 of V_1 . By an analogous argument we also find $t_- < 0$ with $\text{Ric}_g(\nu_{t_-}, \nu_{t_-}) \not\equiv 0$ (otherwise V_2 splits as $\Sigma \times (-\infty, 0]$, a contradiction), and by the above argument starting from t_- , we find that Σ is also ($H \geq 0$)-isotopic to a tubular neighbourhood of a core γ_2 of V_2 . Also in this case $\mathcal{M}_{H \geq 0}(p, q)$ has only one path component.

The theorem now follows by noting that in the case where $\mathcal{M}_{H>0}(p, q)$ has two path components, there always exists a minimal Heegaard torus. This follows for example from [19] where the existence of a minimal Heegaard torus is proved for all lens spaces $L(p, q)$ except \mathbb{S}^3 and \mathbb{RP}^3 , therefore in particular for all lens spaces with $q \not\equiv \pm 1 \pmod p$. \square

References

- [1] V. Agostiniani, M. Fogagnolo, and L. Mazziari. *Sharp geometric inequalities for closed hypersurfaces in manifolds with nonnegative Ricci curvature*. *Invent. Math.* **222** (2020), 1033–1101.
- [2] J. W. Alexander. *On the subdivision of 3-space by a polyhedron*. *Proc. Nat. Acad. Sci. U.S.A.* **10** (1924), 6–8.
- [3] B. Andrews. *Noncollapsing in mean convex mean curvature flow*. *Geom. Topol.* **16**:3 (2012), 1413–1418.
- [4] R. Bamler and B. Kleiner. *Ricci flow and contractibility of spaces of metrics*. Preprint, ArXiv:1909.08710 (2019).
- [5] F. Bonahon. *Difféotopies des espaces lenticulaires*. *Topology* **22**:3 (1983), 305–314.
- [6] F. Bonahon and J.-P. Otal. *Scindements de Heegaard des espaces lenticulaires*. *Ann. Sci. Éc. Norm. Supér.* (4) **16**:3 (1983), 451–466.
- [7] S. Brendle. *An inscribed radius estimate for mean curvature flow in Riemannian manifolds*. *Ann. Sc. Norm. Super. Pisa Cl. Sci.* **16**:4 (2016), 1447–1472.
- [8] S. Brendle and G. Huisken. *Mean curvature flow with surgery of mean convex surfaces in \mathbb{R}^3* . *Invent. Math.* **203** (2016), 615–654.
- [9] S. Brendle and G. Huisken. *Mean curvature flow with surgery of mean convex surfaces in 3-manifolds*. *J. Eur. Math. Soc.* **20**:9 (2018), 2239–2257.
- [10] R. Buzano, R. Haslhofer, and O. Hershkovits. *The moduli space of two-convex embedded spheres*. *J. Differential Geom.* **118**:2 (2021), 189–221.
- [11] R. Buzano, R. Haslhofer, and O. Hershkovits. *The moduli space of two-convex embedded tori*. *Int. Math. Res. Not. IMRN* **2019**:2 (2019), 392–406.
- [12] R. Buzano, H. Nguyen, and M. Schulz. *Noncompact self-shrinkers for mean curvature flow with arbitrary genus*. Preprint, ArXiv:2110.06027 (2021).
- [13] G. Di Matteo and A. Malchiodi. *Double bubbles with high constant mean curvature in Riemannian manifolds*. *Nonlinear Anal.* **224** (2022), No. 113088, 41 pp.
- [14] R. Haslhofer and D. Ketover. *Minimal 2-spheres in 3-spheres*. *Duke Math. J.* **168**:10 (2019), 1929–1975.
- [15] R. Haslhofer and B. Kleiner. *Mean curvature flow with surgery*. *Duke Math. J.* **166**:9 (2017), 1591–1626.
- [16] P. Heegaard. *Forstudier til en topologisk Teori for de algebraiske Fladers Sammenhæng*. Det Nordiske Forlag, Copenhagen (1898).
- [17] G. Huisken and C. Sinestrari. *Mean curvature flow with surgeries of two-convex hypersurfaces*. *Invent. Math.* **175** (2009), 137–221.
- [18] A. Kasue. *Ricci curvature, geodesics and some geometric properties of Riemannian manifolds with boundary*. *J. Math. Soc. Jpn.* **35**:1 (1983), 117–131.
- [19] D. Ketover, Y. Liokumovich, and A. Song. *On the existence of minimal Heegaard surfaces*. Preprint, ArXiv:1911.07161 (2019).
- [20] M. Kreck and S. Stolz. *Nonconnected moduli spaces of positive sectional curvature metrics*. *J. Am. Math. Soc.* **6**:4 (1993), 825–850.
- [21] G. Liu. *3-manifolds with nonnegative Ricci curvature*. *Invent. Math.* **193** (2013), 367–375.
- [22] F. C. Marques. *Deforming three-manifolds with positive scalar curvature*. *Ann. Math. (2)* **176**:2 (2012), 815–863.
- [23] J. Rosenberg and S. Stolz. *Metrics of positive scalar curvature and connections with surgery*. *Ann. Math. Stud.* **149** (2001), 353–386.
- [24] W. Sheng and X. Wang. *Singularity profile in the mean curvature flow*. *Methods Appl. Anal.* **16**:2 (2009), 139–155.
- [25] H. Weyl. *Über die Bestimmung einer geschlossenen konvexen Fläche durch ihr Linienelement*. *Vierteljahrsschr. Naturforsch. Ges. Zur.*, **61** (1916), 70–72.
- [26] J. H. C. Whitehead. *On C^1 -complexes*. *Ann. Math. (2)* **41** (1940), 809–824.
- [27] T. J. Willmore. *Riemannian geometry*. Oxford science publications (1996).

RETO BUZANO

Università degli Studi di Torino, Dipartimento di Matematica, Torino, Italy

E-mail address: `reto.buzano@unito.it`

SYLVAIN MAILLOT

IMAG, Université de Montpellier, CNRS, Montpellier, France

E-mail address: `sylvain.maillot@umontpellier.fr`