MANIFOLDS WITH CONULLITY AT MOST TWO AS GRAPH MANIFOLDS

LUIS A. FLORIT AND WOLFGANG ZILLER

Abstract. We find necessary and sufficient conditions for a complete Riemannian manifold M^n of finite volume, whose curvature tensor has nullity at least $n-2$, to be a geometric graph manifold and classify the ones with nonnegative scalar curvature. In the process, we show that Nomizu's conjecture, well known to be false in general, is true for manifolds with finite volume.

The nullity space Γ of the curvature tensor R of a Riemannian manifold M^n is defined for each $p \in M$ as $\Gamma(p) = \{X \in T_pM : R(X,Y) = 0 \ \ \forall Y \in T_pM\}$, and its dimension $\mu(p)$ is called the *nullity* of M^n at p. It is well known that the existence of points with positive nullity has strong geometric implications. For example, on an open subset of M^n where μ is constant, Γ is an integrable distribution with totally geodesic leaves. In addition, if M^n is complete, its leaves are also complete on the open subset where μ is minimal; see e.g. [Ma]. Riemannian n-manifolds with conullity at most 2, i.e., $\mu \geq n-2$, which we call *CN2 manifolds* for short, appear naturally and frequently in several different contexts in Riemannian geometry, e.g.:

- Gromov's 3-dimensional graph manifolds admit a complete CN2 metric with nonpositive sectional curvature and finite volume whose set of flat points consists of a disjoint union of flat totally geodesic tori ([Gr]). These were the first examples of Riemannian manifolds with geometric rank one. Interestingly, any complete metric of nonpositive curvature on such a graph manifold is necessarily CN2 and quite rigid, as was shown in [Sch];
- A Riemannian manifold is called semi-symmetric if at each point the curvature tensor is orthogonally equivalent to the curvature tensor of some symmetric space, which is allowed to depend on the point. CN2 manifolds are semi-symmetric since they have pointwise the curvature tensor of an isometric product of a Euclidean space and a surface with constant curvature. Conversely, Szabó showed in [Sz] that a complete simply connected semi-symmetric space is isometric to a Riemannian product $S \times N$, where S is a symmetric space and N is, on an open and dense subset, locally a product of CN2 manifolds;
- Isometrically deformable submanifolds tend to have large nullity. In particular, by the classic Beez-Killing theorem, any locally deformable hypersurface in a space form has to be CN2. Yet, generically, CN2 hypersurfaces are locally rigid, and the classification of

The first author was supported by CNPq-Brazil, and the second author by a grant from the National Science Foundation, by IMPA, and CAPES-Brazil.

the deformable ones has been carried out a century ago in [Sb, Ca]; see [DFT] for a modern version and further results. The corresponding classification of locally deformable CN2 Euclidean submanifolds in codimension two is considerably more involved, and was obtained only recently in [DF] and [FF];
- A compact immersed submanifold $M^3 \subset \mathbb{R}^5$ with nonnegative sectional curvature not diffeomorphic to the 3-sphere $\mathbb{S}^3$ is necessarily CN2, and either isometric to $(\mathbb{S}^2 \times \mathbb{R})/\mathbb{Z}$ for some metric of nonnegative curvature on $\mathbb{S}^2$, or diffeomorphic to a lens space $\mathbb{S}^3/\mathbb{Z}_p$; see [FZ]. In the case of lens spaces, the set of points with vanishing curvature has to be nonempty with Hausdorff dimension at least two. However, it is not known yet if they can be isometrically immersed into $\mathbb{R}^5$;
- I. M. Singer asked in [Si] whether a Riemannian manifold is homogeneous if the curvature tensor at any two points is orthogonally equivalent. The first counterexamples to this question were CN2 manifolds with constant scalar curvature, which clearly have this property, and are typically not homogeneous; see [Se, BKV].

The most trivial class of CN2 manifolds is given by cylinders $L^2 \times \mathbb{R}^{n-2}$ with their natural product metrics, where L^2 is any (not necessarily complete) connected surface. More generally, we call a *twisted cylinder* any quotient

$$C^n = (L^2 \times \mathbb{R}^{n-2})/G,$$

where $G \subset \mathrm{Iso}(L^2 \times \mathbb{R}^{n-2})$ acts properly discontinuously and freely. The natural quotient metric is clearly CN2, and we call L^2 the *generating surface* of C^n, and the images of the Euclidean factor its *nullity leaves*. Observe that C^n fails to be complete only because L^2 does not need to be. Yet, what is important for us is that C^n is foliated by complete, flat, totally geodesic, and locally parallel leaves of codimension 2.

Our first goal is to show that these are the basic building blocks of complete CN2 manifolds with finite volume:

THEOREM A. *Let M^n be a complete CN2 manifold. Then each finite volume connected component of the set of nonflat points of M^n is globally isometric to a twisted cylinder.*

The hypothesis on the volume of M^n is essential, since complete locally irreducible Riemannian manifolds with constant conullity two abound in any dimension; see [Sz], [BKV] and references therein. These examples serve also as counterexamples to the Nomizu conjecture in [No], which states that a complete locally irreducible semi-symmetric space of dimension at least three must be locally symmetric. However, Theorem A together with Theorem 4.4 in [Sz] yield:

COROLLARY 1. *Nomizu's conjecture is true for manifolds with finite volume.*

For the 3-dimensional case, the fact that the set of nonflat points of a finite volume CN2 manifold is locally reducible was proved in [SW] with a longer and more delicate proof. Notice also that in dimension 3 the CN2 condition is equivalent to the assumption, called cvc(0) in [SW], that every tangent vector is contained in a flat plane, or to the condition that the Ricci endomorphism has eigenvalues $(\lambda, \lambda, 0)$. Furthermore, in [BS] it was shown that a complete 3-manifolds with finite volume and (geometric) rank one is a twisted cylinder.

Observe that we are free to change the metric in the interior of the generating surfaces of the twisted cylinders in Theorem A, still obtaining a complete CN2 manifold. Moreover, they are nowhere flat with Gaussian curvature vanishing at their boundaries. Of course, these boundaries can be quite complicated and irregular.

In general it is very difficult to fully understand how the twisted cylinders in Theorem A can be glued together through the set of flat points in order to build a complete Riemannian manifold. An obvious way of gluing them is through compact totally geodesic flat hypersurfaces. Indeed, when the boundary of each generating surface L^2 in the twisted cylinder $C = (L^2 \times \mathbb{R}^{n-2})/G$ is a disjoint union of complete geodesics γ_j along which the Gaussian curvature of L^2 vanishes to infinity order, the boundary of C is a disjoint union of complete totally geodesic flat hypersurfaces $H_j = (\gamma_j \times \mathbb{R}^{n-2})/G_j \subset M^n$, where G_j is the subgroup of G preserving γ_j. We can now use each H_j to attach another finite volume twisted cylinder C' to C along H_j, as long as C' has a boundary component isometric to H_j. Repeating and iterating this procedure with each boundary component we construct a complete CN2 manifold M^n. As we will see, the hypersurfaces H_j have to be compact if M^n has finite volume. This motivates the following concept of geometric graph manifold of dimension $n \geq 3$, which by definition is endowed with a CN2 Riemannian metric:

Definition. A connected Riemannian manifold M^n is called a *geometric graph manifold* if M^n is a locally finite disjoint union of twisted cylinders C_i glued together through disjoint compact totally geodesic flat hypersurfaces H_λ of M^n. That is,

$$M^n \setminus W = \bigsqcup_{\lambda} H_\lambda, \quad \text{where} \quad W := \bigsqcup_{i} C_i.$$

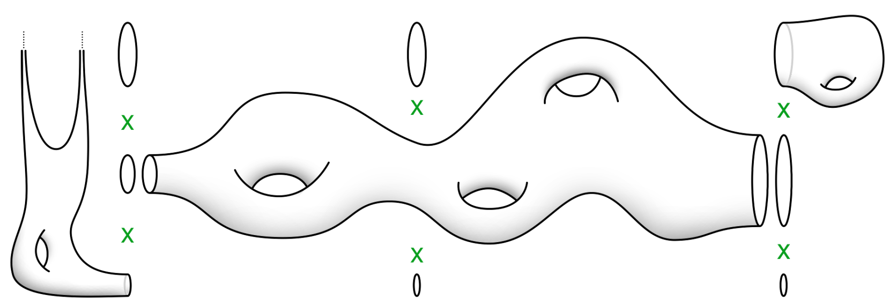

FIGURE 1. An irreducible 4-dimensional CN2 geometric graph manifold with three cylinders and two (finite volume) ends

Here we allow the possibility that a hypersurface H_λ is one-sided, even when M^n is orientable. We also assume, without loss of generality, that the nullity leaves of two cylinders C and C', glued along H_λ, have distinct limits in H_λ. This implies in particular that for each cylinder C, the Gauss curvature vanishes along ∂C to infinite order.

Observe that the complement of W is contained in the set of flat points of M^n, but we do not require that the generating surfaces of C_i have nonvanishing Gaussian curvature. In particular the sectional curvature of C_i, or equivalently its scalar curvature, can change sign. More importantly, W carries a well defined complete flat totally geodesic parallel distribution of constant rank $n-2$ contained in the nullity of M^n. Furthermore, W is dense and locally finite in the sense that it has a locally finite number of connected components (see Section 5 for a precise definition). These two topological properties will be crucial in what follows, so for convenience we say that a dense locally finite set is *full*.

A natural way to try to see if a Riemannian manifold M^n as in Theorem A is indeed a geometric graph manifold is the following. Theorem A implies that on the open set of nonflat points V we have the well defined parallel nullity distribution Γ of rank $n-2$, as in W above. Now, consider any open set $\hat{V} \supset V$ carrying a complete flat totally geodesic distribution $\hat{\Gamma}$ with $\hat{\Gamma}|_V = \Gamma$, which we call an *extension* of V. We will show that each connected component of $\hat{V}$ is still a twisted cylinder, and call $\hat{V}$ *maximal* if it has no larger extension. Clearly, by definition V always has a maximal extension, but it may not be unique. More importantly, all extensions of V may fail to be either dense, or locally finite, or both; see Examples 2 and 3 in Section 1.

Our second main goal is to prove that all we need to ask in order for M^n as in Theorem A to be a geometric graph manifold is that some extension of V is full:

THEOREM B. *Let M^n be a complete CN2 manifold with finite volume. Then M^n is a geometric graph manifold if and only if its set of nonflat points V admits a full extension. In particular, if V itself is full, then M^n is a geometric graph manifold.*

We point out that here we do not require a full extension $\hat{V}$ of V to be maximal, but clearly any maximal extension of $\hat{V}$ is also full. We can for example introduce complicated sets of flat points in the twisted cylinders, even as Cantor sets in the generating surfaces, but these flat sets will be absorbed by a maximal full extension. As we will show, any maximal full extension will satisfy the properties of W in the definition of geometric graph manifold, see Theorem 5.12. We expect that the methods developed to prove this can be extended for distributions of arbitrary rank.

Although complete 3-dimensional geometric graph manifolds are well studied in the context of nonpositive curvature (see e.g. [Gr, Sch]), they are also interesting for manifolds with nonnegative curvature, as e.g. in [FZ]. Notice that in this case the assumption of finite volume is equivalent to compactness. We first show that their study can be reduced to the 3-dimensional case:

Theorem C. *Let M^n be a compact geometric graph manifold with nonnegative scalar curvature. Then, the universal cover $\tilde{M}^n$ of M^n splits off an $(n-3)$-dimensional Euclidean factor isometrically, i.e., $\tilde{M}^n = N^3 \times \mathbb{R}^{n-3}$. Moreover, either $N^3 = \mathbb{S}^2 \times \mathbb{R}$ splits isometrically with a metric on $\mathbb{S}^2$ with nonnegative Gauss curvature, or $N^3 = \mathbb{S}^3$ with a geometric graph manifold metric.*

We will see that the first part follows, without any curvature conditions, already under the assumption that $M^n \setminus W$ is connected.

In dimension 3, the simplest nontrivial example with nonnegative scalar curvature is the usual description of $\mathbb{S}^3$ as the union of two solid tori endowed with a product metric, see Figure 2. If this product metric is invariant under $\mathrm{SO}(2) \times \mathrm{SO}(2)$, we can also take a quotient by the cyclic group generated by $R_p \times R_p^q$ to obtain a CN2 metrics on any lens space $L(p,q) = \mathbb{S}^3/\mathbb{Z}_p$. Here $R_p \in \mathrm{SO}(2)$ denotes the rotation of angle $2\pi/p$.

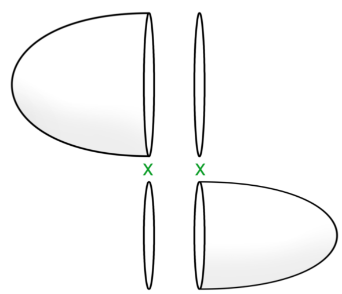

Figure 2. $\mathbb{S}^3 \subset \mathbb{R}^5$ as a nonnegatively curved CN2 manifold

There is a further family whose members also admit CN2 metrics with nonnegative scalar curvature: the so called *prism manifolds* $P(m,n) := \mathbb{S}^3/G_{m,n}$, which depend on two relatively prime positive integers m,n. A geometric graph manifold metric on $P(m,n)$ can be constructed as a quotient of a geometric graph manifold metric on $\mathbb{S}^3$ as above, with the two solid tori being in addition isometric, by the group generated by $R_{2n} \times R_{2n}^{-1}$ and $(R_m \times R_m) \circ J$, where J is a fixed point free isometry switching the two solid tori. Topologically $P(m,n)$ is thus a single solid torus whose boundary is identified to be a Klein bottle. Its fundamental group $G_{m,n}$ is abelian if and only if $m = 1$, and in fact $P(1,n)$ is diffeomorphic to $L(4n, 2n-1)$; see Section 7. Unlike in the case of lens spaces, the diffeomorphism type of a prism manifold is determined by its fundamental group.

Our last main result shows that these are the only nontrivial geometric graph manifolds with nonnegative scalar curvature. We will see that all twisted cylinders in this case are of the form $C = (D \times \mathbb{R})/\mathbb{Z}$, where D is a 2-disk of nonnegative Gaussian curvature whose boundary ∂D is a closed geodesic along which the curvature vanishes to infinite order. We fix once and for all such a metric $\langle\, , \rangle_0$ on a 2-disc D_0 whose boundary has length 1 and which is rotationally symmetric. We call a geometric graph manifold metric on a 3-manifold *standard* if the generating disk D of each twisted cylinder C as above is isometric to D_0

with metric $r^2\langle,\rangle_0$ for some $r > 0$. Observe that the projection of $\partial D \times \{s\}$ for $s \in \mathbb{R}$ is a parallel foliation by closed geodesics of the flat totally geodesic 2-torus $(\partial D \times \mathbb{R})/\mathbb{Z}$ of C.

We provide the following classification:

THEOREM D. *Let M^3 be a compact geometric graph manifold with nonnegative scalar curvature and irreducible universal cover. Then M^3 is diffeomorphic to a lens space or a prism manifold. Moreover, we have either:*

a) M^3 is a lens space and $M^3 = C_1 \sqcup T^2 \sqcup C_2$, i.e., M^3 is isometrically the union of two nonflat twisted cylinders over disks $C_i = (D_i \times \mathbb{R})/\mathbb{Z}$ glued together along their common totally geodesic flat torus boundary T^2. Conversely, any flat torus endowed with two parallel foliations by closed geodesics defines a standard geometric graph manifold metric on a lens space, which is unique up to isometry;

b) M^3 is a prism manifold and $M^3 = C \sqcup K^2$, i.e., M^3 is isometrically the closure of a single twisted cylinder over a disk $C = (D \times \mathbb{R})/\mathbb{Z}$ whose totally geodesic flat interior boundary is isometric to a rectangular torus T^2, and $K^2 = T^2/\mathbb{Z}_2$ is a Klein bottle. Conversely, any rectangular flat torus endowed with a parallel foliation by closed geodesics defines a standard geometric graph manifold metric on a prism manifold, which is unique up to isometry.

In addition, any geometric graph manifold metric with nonnegative scalar curvature on M^3 is isotopic, through geometric graph manifold metrics with nonnegative scalar curvature, to a standard one.

The diffeomorphism type of M^3 in Theorem D is determined by the relative (algebraic) slope between the parallel foliations by closed geodesics; see Section 7 for the precise definition. In case (a), M^3 is diffeomorphic to a lens space $L(p,q)$, where $q/p \in \mathbb{Q}$ is the relative slope between the foliations $[\partial D_i \times \{s\}]$, $i = 1, 2$. Analogously, the manifolds in case (b) are prism manifolds $P(m,n)$, where m/n is the relative slope between the foliation $[\partial D \times \{s\}]$ on the rectangular interior boundary torus $T^2 = S^1_{r_1} \times S^1_{r_2}$ and the foliation $S^1_{r_1} \times \{w\}$.

We can deform any geometric graph manifold metric in Theorem D to first be standard preserving the torus T^2, and then deform T^2 to be the unit square $S^1 \times S^1$, while preserving also the sign of the scalar curvature along the process. In case (a), we can also make one of the foliations equal to $S^1 \times \{w\}$. The metric is then determined by the remaining parallel foliation of the unit square by closed geodesics whose usual slope is equal to the relative slope. Since the diffeomorphism type of a lens space $L(p,q)$ is determined by $\pm q^{\pm 1} \bmod p$, we conclude:

COROLLARY 2. *The moduli space of geometric graph manifold metrics with nonnegative scalar curvature on a lens space $L(p,q)$ has infinitely many connected components, whereas on a prism manifold $P(m,n)$ with $m > 1$ it is connected.*

The assumption of local finiteness in Theorem B can be regarded as a mild regularity condition. But we believe that even without regularity conditions it should be possible to understand the gluing between the twisted cylinders. We state:

CONJECTURE 1. *If the set of nonflat points of a complete CN2 manifold with finite volume admits a dense (not necessarily locally finite) extension, then the complement of any maximal one is a disjoint union of compact totally geodesic flat hypersurfaces, possibly accumulating (see Example 3 in Section 1).*

Certainly more difficult, we can ask what happens if we remove all hypothesis on V. In particular, we do not know if the following is true:

QUESTION. *Does the set V of nonflat points of a complete CN2 manifold with finite volume admit a maximal (not necessarily dense or locally finite) extension $\hat{V}$ such that $\partial\hat{V}$ is a union of flat totally geodesic hypersurfaces, each of which has complete totally geodesic boundary (if nonempty)? (See Example 2 in Section 1).*

On the other hand, in the case of nonnegative or nonpositive curvature we believe one has a stronger statement:

CONJECTURE 2. *Every compact 3-dimensional CN2 manifold with nonnegative or nonpositive scalar curvature and finite volume is a geometric graph manifold.*

Another interesting question is to what extent complete CN2 manifolds with finite volume differ from geometric graph manifolds from a differentiable point of view:

QUESTION. *If M^n admits a complete CN2 metric with finite volume, does it also admit a geometric graph manifold metric?*

We caution that our definition of a graph manifold in dimension 3 is more special than the usual topological one, where the pieces are allowed to be nontrivial Seifert fibered circle bundles ([Wa]). Ours is similar, although more general, to the kind of graph manifolds one studies in nonpositive curvature.

The paper is organized as follows. In Section 1 we provide some examples in order to show that the two hypothesis in Theorem B are necessary. Section 2 is devoted to give some general facts about twisted cylinders and the action of their deck groups. A general semi-global version of the de Rham theorem is provided in Section 3 and will be used in Section 4 to prove Theorem A. The proof of Theorem B is carried out in Section 5. The last two sections are devoted to the classification of geometric graph manifolds of nonnegative scalar curvature, and therefore are independent from the rest of the paper. In Section 6 we prove Theorem C by showing that the manifold is a union of one or two twisted cylinders, while in Section 7 we classify their metrics.

1. Examples.

We now build some examples to help understand how geometric graph manifolds are linked with the CN2 property, and to what extent they differ. In particular, we exhibit CN2 metrics on the 3-torus T^3 which are C^∞ perturbations of the flat metric but that are not geometric graph manifold metrics.

1. *The 3-torus as a nontrivial geometric graph manifold.* Let $L^2 = [-1,1]^2$ with metric a C^∞ perturbation of the flat metric in a small open set $U \subset L^2$ whose closure is contained in the interior of L^2. The cube $C = L^2 \times [-1,1]$ with its product metric serves as a building block in all further examples, where the second factor will give rise to the nullity foliations. Depending on the example, we also adjust their sizes appropriately. Notice that the metric necessarily has scalar curvature of both signs. We now glue two such cubes along a common face in such a way that the nullity distributions are orthogonal, see Figure 3. Identifying opposite faces of the resulting larger cube defines a metric on T^3 with complete nullity foliations, making it into a nontrivial graph manifold. Figure 3 shows the nonflat points on the left, together with a full maximal extension and its two (un)twisted cylinders on the right.

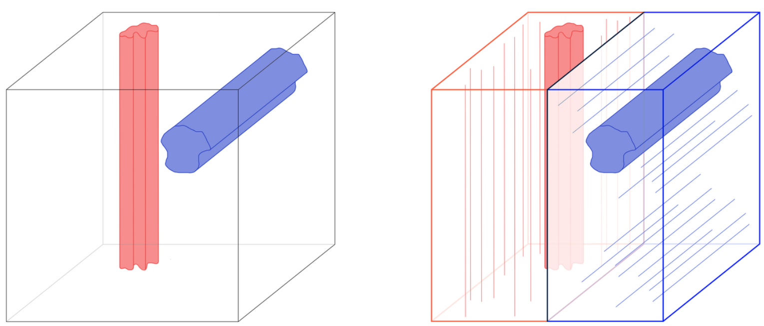

Figure 3. A CN2 3-torus with its set of nonflat points, and a full extension

2. *A CN2 3-torus failing to be a geometric graph manifold: no maximal dense extensions.* Here we take three basic building blocks and glue them together as in Figure 4. Adding two small flat cubes, we obtain a larger cube and identifying opposite faces defines a CN2 metric on T^3. But this is not a geometric graph manifold since the nullity distribution cannot be extended to a dense set of T^3. Figure 4 shows the set of nonflat points on the left, and a maximal extension of it on the right missing two octants of flat points. We point out that this example also shows that a CN2 manifold does not necessarily admit a T structure (see [CG]), as graph manifolds do.

3. *A CN2 3-torus failing to be a geometric graph manifold: no locally finite extension.* Take a sequence of building blocks $C_n = L_n \times [-1,1]$ with $L_n = [-1/2^n, 1/2^n] \times [-1,1]$. Glue one to the next along the squares $\{\pm 1/2^n\} \times [-1,1] \times [-1,1]$ as in Example 1, with

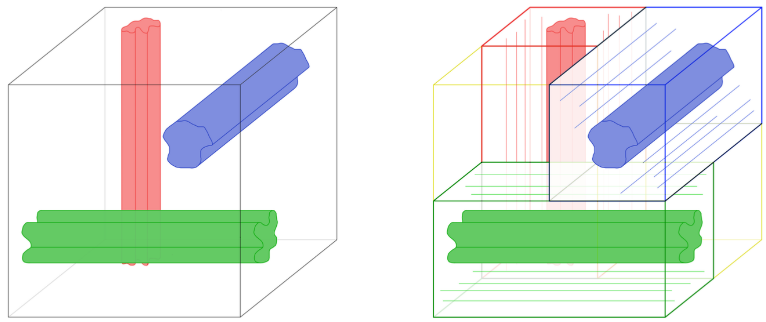

FIGURE 4. A CN2 3-torus with its nonflat points and a nondense maximal extension

nullity lines meeting orthogonally from one to the next, and accumulating at a two torus $T^2 = [-1,1] \times [-1,1]$. Now glue to this a copy of itself along T^2. Identifying opposite sides in the resulting cube defines a CN2 metric on T^3. The metric has two sequences of parallel totally geodesic flat 2-tori, approaching T^2 from both sides. It is not a geometric graph manifold since the number of connected components near T^2 of any extension of V is infinite, even though there exist dense maximal extensions of V. Notice that T^2 is disjoint from the union of the closures of the connected components of any extension of V.

4. *Drunken cylinders.* We can modify the previous examples to obtain a more complicated behavior of the twisted cylinders, illustrating another crucial difficulty in trying to prove Conjecture 1 in the introduction. For this, we start with a flat building block $C_n = L_n \times [-1,1]$ and identify two opposite faces to to obtain a flat metric on $[-1/2^n, 1/2^n] \times T^2$. Let γ_n be a closed geodesic in the two torus $\{0\} \times T^2$ and modify the metric in a small tubular neighborhood of γ_n in $(-1/4^n, 1/4^n) \times T^2$. The boundary of the resulting manifold consists of two flat square tori and we can hence glue one to the next as in the previous example. Gluing a mirror copy of the result and identifying the remaining two faces, we obtain a CN2 metric on T^3. We can now choose the infinite sequence of cylinders comprising the example such that the slopes of γ_n converge. Notice that in the previous examples the totally geodesic 2-tori separating components of a full extension of the set of nonflat points have the property that they are foliated by two different limits of nullity lines, whereas here the torus T^2 in the middle only has one family of limit nullity lines. The existence of two families is crucial in proving convexity properties of the boundary of a maximal full extension of the set of nonflat points; see Section 5.

2. Geometry of twisted cylinders

In this section we study some geometric properties of twisted cylinders with finite volume and nonempty complete flat totally geodesic boundary and set up some notation. We start with the flat case and afterwards we deal with the nonflat one.

We first need some preliminary results. Let us begin by characterizing the surfaces with boundary that have nondiscrete isometry group. Notice for the following that a rotationally symmetric 2-disc may be written as a singular warped product $[0,a] \times_\varphi S^1$ with $\varphi(0) = 0$.

LEMMA 2.1. *Let L^2 be a connected surface with a nonempty boundary consisting of complete geodesics, and assume that* $\mathrm{Iso}(L^2)$ *is not discrete. Then,* $\mathrm{Iso}(L^2)_0 \simeq S^1$ *or* $\mathbb{R}$*. Moreover, if γ is one of the boundary components of L^2, then L^2 is isometric to one of the following:*

a) A warped product $I \times_\varphi \gamma$, where $I = [0,a]$ or $I = [0,+\infty)$;
b) A rotationally symmetric 2-disc with $\gamma \simeq S^1$ as its boundary;
c) A nonorientable quotient of a warped product $([-a,a] \times_\varphi S^1)/\mathbb{Z}_2$, where φ is even and Z_2 acts as $(x,w) \to (-x,-w)$.

In particular, L^2 is diffeomorphic to a 2-disk, $I \times S^1$, $I \times \mathbb{R}$, or a Möbius strip, and φ' vanishes at the boundary of I.

Proof. If $\mathrm{Iso}(L^2)$ is not discrete, then $\dim \mathrm{Iso}(L^2) = 1$ since L^2 is a surface and $\mathrm{Iso}(L^2)_0$ acts transitively and freely on each connected component of its boundary. Hence $\mathrm{Iso}(L^2)_0$ is either S^1 or $\mathbb{R}$. In the first case, all the connected components of the boundary of L^2 are closed, while they are all open in the second case.

We claim that L^2 cannot have more than two boundary components. Indeed, if γ_i, $i = 1, 2, 3$, are three boundary components, let δ_1 be a minimal geodesic from γ_1 to γ_2 and δ_2 one from γ_1 to γ_3, which exist since $\mathrm{Iso}(L^2)_0$ acts transitively on each γ_i. But if $g \in \mathrm{Iso}(L^2)_0$ is such that $g\delta_1(0) = \delta_2(0)$, then $g\delta_1 = \delta_2$ and so $\gamma_2 = \gamma_3$.

If the boundary of L^2 consists of two geodesics $\gamma = \gamma_1$ and γ_2 (both either open or closed), then all geodesics orthogonal to γ_1 are minimizing until they meet γ_2 since $\mathrm{Iso}(L^2)_0$ acts transitively on such minimizing geodesics. Thus L^2 is a warped product $[0,a] \times_\varphi \gamma$. On the other hand, if the boundary of L^2 consists only of γ, we have that either all unit geodesics σ orthogonal to γ are minimizing for all time and L^2 is a warped product $[0,+\infty) \times_\varphi \gamma$, or they are minimizing up to a common time $a > 0$.

In the latter case, the set of their end points $\ell = \{\sigma(a)\}$ is the cut locus of γ. Furthermore, it is an orbit of $\mathrm{Iso}(L^2)_0$ and hence an embedded submanifold of L^2. We thus have two possibilities: either ℓ is a point, in which case L^2 is a rotationally symmetric disc, or ℓ is one-sided since the complement of the cut locus of γ is connected. Therefore its normal bundle is nontrivial, which implies that $\ell \simeq S^1$ and hence L^2 is a Möbius strip. Since the orientable double cover still satisfies the hypothesis of the lemma, it is isometric to $[-a,a] \times_\varphi S^1$. Finally, the only orientation reversing fixed point free involution is as claimed. $\square$

This also allows us to classify flat surfaces with totally geodesic boundary.

LEMMA 2.2. *If L^2 is a connected flat surface with complete totally geodesic boundary, then* $\mathrm{Iso}(L^2)$ *is not discrete. In particular, L^2 is isometric to either a flat Möbius strip or $I \times \gamma$, where $I = [0,a]$ or $I = [0,+\infty)$, and $\gamma \simeq S^1$ or $\gamma \simeq \mathbb{R}$.*

Proof. We first show that $\mathrm{Iso}(L^2)$ is not discrete. The double of L^2, i.e., $\hat{L^2} = L^2 \cup_{\partial L^2} L^2$, is a complete flat manifold whose universal cover is $\pi : \mathbb{R}^2 \to \hat{L^2}$ with a flat metric. The set K of all lifts of ∂L^2 into $\mathbb{R}^2$ is a union of lines, which are parallel since they cannot intersect. The deck group of π preserves K, and one easily sees that translations parallel to its lines commute with the full group of isometries preserving K. Hence they commute with the deck group as well and thus these translations descend to $\hat{L^2}$. So the isometry group of $\hat{L^2}$ has a subgroup isomorphic to either S^1 or $\mathbb{R}$ which preserves ∂L^2. This subgroup is then contained in $\mathrm{Iso}(L^2)$ and we can apply Lemma 2.1. The curvature of the warped product is given by $-\varphi''/\varphi$ and since φ' vanishes on the boundary, φ is constant. In particular, L^2 can not be a disc since in that case $\lim_{t\to 0+} \varphi(t) = 0$. The claim now follows from Lemma 2.1. $\square$

To finish our preliminary results, let us show the following elementary fact.

LEMMA 2.3. *A complete flat manifold with finite volume is compact.*

Proof. A Dirichlet fundamental domain U of the action of the deck group on the universal cover $\mathbb{R}^k$ is a convex set such that the translates of $\bar{U}$ fill all of $\mathbb{R}^k$, and the translates of U° are disjoint. Hence $U \subset \mathbb{R}^k$ is a convex set with finite volume. If $\bar{U}$ is not compact, it would contain a ray c and, if $B_\epsilon(c(0)) \subset U$, the subset $\{\overline{pc(t)} : p \in B_\epsilon(c(0)), t \geq 0\} \subset U$ would have infinite volume. $\square$

Recall that a twisted cylinder C^n is isometric to a quotient $(L^2 \times \mathbb{R}^{n-2})/G$, where $G \subset \mathrm{Iso}(L^2 \times \mathbb{R}^{n-2})$ acts properly discontinuously and freely. We will also assume in this section that C^n has finite volume and nonempty boundary which consists of a (possibly infinite) disjoint union of complete (not necessarily compact) flat totally geodesic hypersurfaces.

The flat case is particularly simple:

COROLLARY 2.4. *If C^n as above is flat, then it is isometric to either $[0,a] \times F$ or $([-a,a] \times F)/\mathbb{Z}_2$, where F is a connected compact flat hypersurface of C^n.*

Proof. By Lemma 2.2 we have that $C^n = (I \times \gamma \times \mathbb{R}^{n-2})/G$ for some G, with $I = [0,a]$ or $I = [0,+\infty)$, where in the case that L^2 is the Möbius strip, G also contains the orientation reversing isometry of $I \times \gamma$. Let $\hat{G} \subseteq G$ be the normal subgroup that preserves a boundary component of $I \times \gamma \times \mathbb{R}^{n-2}$. Then, either $\hat{G} = G$ or it has index two in G. Thus $\hat{G}$ does not act on I, and $\hat{C}^n = (I \times \gamma \times \mathbb{R}^{n-2})/\hat{G} = I \times F$ where $F = (\gamma \times \mathbb{R}^{n-2})/\hat{G}$ is flat and complete. Since $\hat{C}^n$ also has finite volume, $I = [0,a]$ and F has finite volume and is thus compact by Lemma 2.3. If $G \neq \hat{G}$, then $G/\hat{G} \simeq \mathbb{Z}_2$ acts as a reflection on I and freely on F. $\square$

Remark 2.5. We claim that we can ignore these flat twisted cylinders when building geometric graph manifolds. Let M^n be a geometric graph manifold with a flat compact totally geodesic boundary component F. If we attach $C = [0,a] \times F$ to M^n at F, the nullity leaves of M^n around F extend parallelly into C. Thus we can ignore flat cylinders that are products since a maximal full set will simply absorb them. On the other hand, if F has in addition a free $\mathbb{Z}_2$ action, we can also attach to M^n the flat twisted cylinder $C' = ([-a,a] \times F)/\mathbb{Z}_2$. The nullity leaves of M^n around F also extend parallelly to the dense subset $(([-a,0) \cup (0,a]) \times F)/\mathbb{Z}_2 \simeq (0,a] \times F \subset C'$, but not necessarily to $\{0\} \times F/\mathbb{Z}_2$. Therefore, instead of attaching C' to M^n, we can simply make a $\mathbb{Z}_2$ quotient of F eliminating it as a boundary component of M^n. In this way, $F/\mathbb{Z}_2$ becomes one of the (one-sided) totally geodesic hypersurfaces H_λ in the definition of geometric graph manifold.

We now assume until the end of this section that C^n is not flat. Hence, $\mathrm{Iso}(L^2 \times \mathbb{R}^{n-2}) = \mathrm{Iso}(L^2) \times \mathrm{Iso}(\mathbb{R}^{n-2})$ by the uniqueness in the local deRham decomposition theorem. We denote by $\pi\colon L^2 \times \mathbb{R}^{n-2} \to (L^2 \times \mathbb{R}^{n-2})/G$ the projection and by $p = [x,v] = \pi(x,v)$ points in C^n. Set

$$G_1 := \{g_1 \in \mathrm{Iso}(L^2) : (g_1,g_2) \in G\} \quad \text{and} \quad G_2 := \{g_2 \in \mathrm{Iso}(\mathbb{R}^{n-2}) : (g_1,g_2) \in G\}. \tag{2.6}$$

The stabilizer groups we denote by G_p, and for convenience we set

$$G^x = \{(g_1,g_2) \in G : g_1 x = x\} \quad \text{and} \quad G^v = \{(g_1,g_2) \in G : g_2 v = v\},$$

for $x \in L^2$ and $v \in \mathbb{R}^{n-2}$. Whenever necessary we can assume that G_2 acts effectively on $\mathbb{R}^{n-2}$, since we can otherwise first take the quotient of L^2 under the normal subgroup $\{g_1 \in G_1 : (g_1, \mathrm{Id}_{\mathbb{R}^{n-2}}) \in G\}$ of G_1. Notice that in this case, for every $g_2 \in G_2$ there exists a unique g_1 with $(g_1,g_2) \in G$, and $G \simeq G_2$. The nullity leaves are given by $\Gamma([x,v]) = \pi(\{x\} \times \mathbb{R}^{n-2}) = \mathbb{R}^{n-2}/G^x$, while the integral submanifolds of the distribution $\Gamma^\perp$ are $L^2_{[x,v]} = \pi(L^2 \times \{v\}) = L^2/G^v$. Observe that G^v acts properly discontinuously and freely on L^2, and G^x on $\mathbb{R}^{n-2}$ as well, since G does so on $L^2 \times \mathbb{R}^{n-2}$. The fixed point set of $g_2 \in G_2$ is a linear subspace of $\mathbb{R}^{n-2}$, and since G is countable, there exists a $v_0 \in \mathbb{R}^{n-2}$ with $G^{v_0} = \mathrm{Id}$. Therefore, $L^2_{[x,v_0]}$ is isometric to L^2 and covers all other leaves of $\Gamma^\perp$. In particular, we can assume that L^2 is a maximal leaf of $\Gamma^\perp$.

We first analyze the special case where G_1 is not discrete, but to do this we need the following notation. Let F be a compact flat manifold, and ξ a parallel line field on F. Given an interval $I \subset \mathbb{R}$ and $\varphi : I \to \mathbb{R}_+$ we denote by $I \times^\xi_\varphi F$ the manifold $I \times F$ together with the metric $dt^2 + \varphi(t)ds^2 + dx^2$, where (s,x) are local orthogonal coordinates on F with $\partial/\partial s \in \xi$. Observe that this is not a warped product metric on $I \times F$. In the case of Lemma 2.1 (b), the cylinder has the form $C^n = (D \times \mathbb{R}^{n-2})/G$ where D is a rotationally symmetric 2-disk which has a warped product metric $(0,a] \times_\varphi S^1$ almost everywhere. Therefore, for convenience, we write $C^n = [0,a] \times^\xi_\varphi F$ also in this case although with $\varphi(0) = 0$.

Proposition 2.7. *The subgroup $G_1 \subset \mathrm{Iso}(L^2)$ is not discrete if and only if there exists a nullity leaf which is not closed. In this situation, L^2 is isometric to one of the surfaces in Lemma* 2.1 *with $\int_I \varphi < +\infty$, and C^n is isometric to either $I \times^\xi_\varphi F$ or $([-a,a] \times^\xi_\varphi F)/\mathbb{Z}^2$,* *where F is a connected flat compact manifold. In particular, the boundary of C^n either has one or two connected components which are flat, compact, and totally geodesic.*

Proof. Let $\bar{G}_1$ be the closure of $G_1 \subset \mathrm{Iso}(L^2)$. We have $\overline{G_1 \cdot x} = \bar{G}_1 \cdot x$ for all $x \in L^2$ since by the Cartan-Ambrose-Hicks Theorem a sequence of isometries $\{f_k\}$ for which $\{f_k(p)\}$ converges, has a subsequence converging to an isometry. Hence $\Gamma(p) = ((G_1 \cdot x) \times \mathbb{R}^{n-2})/G$ and $\bar{\Gamma}(p) = ((\bar{G}_1 \cdot x) \times \mathbb{R}^{n-2})/G$. Furthermore, as is well known, G_1 is discrete in $\mathrm{Iso}(L^2)$ if and only if all orbits $G_1 \cdot x$ are discrete in L^2. This proves our first assertion.

Since G_1 is not discrete, $\mathrm{Iso}(L^2)$ is also not discrete and we apply Lemma 2.1. As we did in the proof of Corollary 2.4, let $\hat{G} \subseteq G$ be the normal subgroup that preserves a boundary component of $L^2 \times \mathbb{R}^{n-2}$. Thus $\hat{C}^n := (I \times_\varphi \gamma \times \mathbb{R}^{n-2})/\hat{G}$ is either C^n or a double cover of it, and $\hat{G}$ does not act on I. Hence $F = (\gamma \times \mathbb{R}^{n-2})/\hat{G}$ is a complete flat hypersurface, $\hat{C}^n = I \times_\varphi^\xi F$, and $Vol(\hat{C}^n) = Vol(F) \int_I \varphi$. Indeed, if w_0 is a volume element of F, then $\varphi dt \wedge w_0$ is a volume element of $\hat{C}^n$, wherever $\varphi > 0$. By Lemma 2.3 we have that F is compact. $\square$

Observe that, if one nullity leaf is not closed, then they are all not closed, except in the case where L^2 is a disc for which the nullity leaf through its center is closed.

As an application, we have:

COROLLARY 2.8. *If C^n as in Proposition 2.7 is not flat and its scalar curvature does not change sign, then L^2 is a rotationally symmetric disk with nonnegative Gaussian curvature.*

Proof. Since the boundary of C^n is totally geodesic, the derivatives of φ vanish at the boundary of I. The condition on the scalar curvature is equivalent to the Gaussian curvature $-\varphi''/\varphi$ of L^2 in Lemma 2.1 not changing sign. Since φ is thus either convex or concave, the result follows from $\int_I \varphi < +\infty$. $\square$

For the generic case where G_1 is discrete we have the following.

PROPOSITION 2.9. *If $G_1 \subset \mathrm{Iso}(L^2)$ is discrete, then L^2/G_1 is an orbifold with finite volume, and the projection onto the first factor induces an orbifold Riemannian submersion $C^n \to L^2/G_1$ whose fibers are compact nullity leaves.*

Proof. Discreteness implies that the action of G_1 on L^2 is properly discontinuous. The isotropy groups of G_1 are finite since they act on the tangent space effectively and hence are discrete subgroups of $\mathrm{O}(2)$. Thus L^2/G_1 an orbifold and the projection induces an orbifold Riemannian submersion $\sigma\colon (L^2 \times \mathbb{R}^{n-2})/G \to L^2/G_1$. The orbifold strata are points, and possibly complete geodesics which may intersect $\partial(L^2/G_1)$. The fibers of σ coincide with the nullity leaves $\sigma^{-1}([x]) = (\{x\} \times \mathbb{R}^{n-2})/G^x = \mathbb{R}^{n-2}/G^x$. Since C^n has finite volume, so does the base and the fibers since the set of orbifold points of L^2/G_1 has measure zero. Being complete and flat, the fibers must be compact by Lemma 2.3. $\square$

The last two propositions immediately imply the following result which will be needed in Section 5.

COROLLARY 2.10. *If the twisted cylinder C^n has finite volume, its leaves of nullity are all precompact, including their limits in the boundary of C^n.*

We finish this section by relating geometric graph manifolds to manifolds with nonpositive curvature.

PROPOSITION 2.11. *Let M^n be a geometric graph manifold such that none of the generating surfaces are diffeomorphic to a disc. Then M^n admits another geometric graph manifold metric with nonpositive curvature.*

Proof. For a twisted cylinder $C = (L^2 \times \mathbb{R}^{n-2})/G$ such that $G_1 \subset \mathrm{Iso}(L^2)$ is not discrete, Proposition 2.7 implies that C admits a flat metric. Thus we only need to consider twisted cylinders as in Proposition 2.9 and we will change the metric on L^2, while preserving the isometric action by G_1. Let $L^* = L^2 \cup L^2$ be the double of L^2 with its natural metric g induced by the one in L^2. Hence, L^* is a complete surface without boundary which admits an isometric involution τ switching the two copies, as well as an isometric action by G_1. Since we excluded flat surfaces and discs, L^* is a surface of higher genus. Thus, in the conformal class of g there exists a unique metric $\bar{g}$ with constant curvature -1. Since conformal maps of a hyperbolic metric are isometries, G_1 and τ act by isometries in $\bar{g}$ as well. Since $\partial L^2 \subset L^*$ is the fixed point set of τ, $\bar{g}|_L$ is a G_1-invariant metric on L^2 with constant curvature -1 whose boundary curves are still geodesics. So, near the boundary curves $\bar{g}|_L$ in Fermi coordinates has the form $ds^2 + f^2(s)dt^2$ with $f(s) = \cosh(s)$ for $s \in [0, \epsilon)$. Replacing $\cosh(s)$ with a convex function $f(s)$ with $f^{(k)}(0) = 0$ for all $k \geq 1$ we obtain a metric with nonpositive curvature with geodesic boundary along which the curvature vanishes to infinite order. □

We finish this section with some illustrative examples.

a) Let $C = (L^2 \times \mathbb{R})/G$ be a 3-dimensional cylinder such that $G_1 \subset \mathrm{Iso}(L^2)$ is discrete, and assume for simplicity that G acts orientation preserving on $\mathbb{R}$ and L^2. The nullity leaf through (x, v) is $\mathbb{R}/G^x$ and since it is compact, G^x is nontrivial for all x. Since the generic stabilizer group of G_1 is trivial, the normal subgroup $G^* \subset G$ acting as Id on L^2 is nontrivial, and since it acts freely on $\mathbb{R}$, we have $G^* \simeq \mathbb{Z}$. We can thus assume that the cylinder is of the form $C = (L^2 \times S^1)/G$. The action defines a homomorphism $\alpha\colon G \to \mathbb{R}$ such that for $g = (g_1, g_2) \in G$ the component g_2 acts on S^1 as a rotation with angle $\alpha(g)$. Now if $L^* = L/\ker(\alpha)$, then $C = (L^* \times S^1)/\mathbb{Z}^\ell$ for some ℓ since $\mathrm{Im}(\alpha)$ is a subgroup of $\mathbb{R}$. If $r_1, r_2, \cdots$ are generators of $\mathbb{Z}^\ell$, then the angles $\alpha(r_i)$ are linearly independent over $\mathbb{Q}$. Thus the leaves of $\Gamma^\perp$ are embedded only when $\ell = 1$, and are in fact dense in C otherwise. We can now change the action of G continuously, while remaining properly discontinuous, such that g_2 acts as $\epsilon\alpha(g)$ for some fixed ϵ. Letting $\epsilon \to 0$, we see that, since the stabilizer groups of G_1 are finite cyclic, C is homeomorphic to the product of a surface with S^1.

b) Conversely, let L^2 be a surface with boundary and $G = \pi_1(L^2)$. Then there exists a cover L' of L^2 such that the image of the fundamental group is $[G, G] \subset G$. Since $G/[G, G]$ is abelian, we can now choose a homomorphism $\alpha\colon G \to \mathbb{R}$ with $\alpha(G) = \mathbb{Z}^\ell$ where $\ell = b_2(L^2)$. We can then define the cylinder $C = (L/\ker(\alpha)) \times \mathbb{R})/\mathbb{Z}^\ell$. This includes the possibility that $\ell = \infty$. In fact, in [Gr] one finds a finite volume example with nonpositive sectional curvature where L^2 is an infinite connected sum of compact surfaces, with boundary consisting of infinitely many closed geodesics.

c) The boundary geodesics of the leaves of $\Gamma^\perp$ may not be closed. As an example, start with a complete flat strip $[-1, 1] \times \mathbb{R}$, remove infinitely many ϵ-discs centered at $(0, n)$, $n \in \mathbb{Z}$, and change the metric around them to make their boundaries totally geodesic in such a way that the metric remains invariant under $\mathbb{Z} = \langle h \rangle$ for $h(x, t) = (x, t + 1)$. If L^2 is this surface and R is a rotation of S^1 of angle $r\pi$ for some irrational number r, then $G = \mathbb{Z} = \langle (h, R) \rangle$ acts freely and properly discontinuously on $L^2 \times S^1$. Although L^2/G_1 has three closed geodesics as boundary, the boundary of the leaves of $\Gamma^\perp$ in $C^3 := (L^2 \times S^1)/G$ has one closed geodesic and also two complete open geodesics, each of which is dense in the corresponding boundary component of C^3. Observe that here the leaves of $\Gamma^\perp$ are also dense in C^3.

d) In general, a cylinder of finite volume does not necessarily have compact boundary. Indeed, consider the metric $ds^2 = dr^2 + e^{-\frac{1}{1-r^2} - t^2} dt^2$ on $L^2 = [-1, 1] \times \mathbb{R}$. Then $L^2 \times T^{n-2}$ has finite volume with two complete noncompact flat totally geodesic boundary components. Nevertheless, we will see that when we glue two twisted cylinders in a nontrivial way, Corollary 2.10 implies that their common boundary is in fact compact.

3. A semi-global de Rham theorem

The existence of a parallel smooth distribution Γ on a complete manifold M^n implies that the universal cover is an isometric product of a complete leaf of Γ with a leaf of $\Gamma^\perp$ by the global deRham theorem. In this section we prove a semi-global version of this fact to be used later on, and will concentrate on flat foliations for simplicity. Although this is a special case of Theorem 1 in [PR], our proof is simpler and more direct, so we add it here for completeness.

Proposition 3.1. *Let W be an open connected set of a complete Riemannian manifold M^n, and assume that Γ is a rank k parallel distribution on W whose leaves are flat and complete. If L is a maximal leaf of $\Gamma^\perp$, then the normal exponential map* $\exp^\perp : T^\perp L \to W$ *is an isometric covering, where $T^\perp L$ is equipped with the induced connection metric. In particular, W is isometric to the twisted cylinder $(\tilde{L} \times \mathbb{R}^k)/G$, where $\tilde{L}$ is the universal cover of L, and G acts isometrically in the product metric.*

Proof. Since Γ is parallel, its orthogonal complement $\Gamma^\perp$ is also parallel and hence integrable with totally geodesic leaves. Due to the local isometric product structure of W given by the local deRham Theorem, the normal bundle $T^\perp L$ of L is flat with respect to the normal connection of L, which in our case is simply the restriction of the connection on M^n since L is totally geodesic. Notice though that $T^\perp L$ does not have to be trivial. The normal connection defines, in the usual fashion, a connection metric on the total space. By completeness of the leaves of Γ, the normal exponential map $\exp^\perp\colon T^\perp L \to W$ is well defined on the whole normal bundle.

We first show that $\exp^\perp$ is a local isometry. Indeed, if $\alpha(s) = (c(s), \xi(s))$ is a curve in $T^\perp L$ with ξ parallel along c, then $\alpha'(0)$ is a horizontal vector in the connection metric, identified with $c'(0)$ under the usual identification $H_{\xi(0)} \simeq T_{c(0)}L$. Then $d(\exp^\perp)_{\alpha(0)}(\alpha'(0)) = J(1)$ where $J(t)$ is the Jacobi field along the geodesic $\gamma(t) = \exp^\perp(t\xi(0))$ with initial conditions $J(0) = c'(0)$ and $J'(0) = 0$ since ξ is parallel along c. Since $\gamma' \in \ker R$, the Jacobi operator $R(\,\cdot\,, \gamma')\gamma'$ vanishes and thus $J(1)$ is the parallel translate of $J(0) \in T_{c(0)}L$, which in turn implies that $\exp^\perp$ is an isometry on the horizontal space. On the vertical space it is an isometry since it agrees with the exponential map of the fiber which is a complete flat totally geodesic submanifold of W. The image of horizontal and vertical space under $\exp^\perp$ are also orthogonal, since the first is the parallel translate of TL and the second the parallel translate of $T^\perp L$ along γ. Hence $\exp^\perp$ is a local isometry. Notice that this also implies that if ξ is a (possibly only locally defined) parallel section of $T^\perp L$, then $\{\exp_p(t\xi(p)) \mid p \in L\} \subset W$ are integral manifolds of $\Gamma^\perp$ for all $t \in \mathbb{R}$.

Let $q \in W$ and denote by L_q the maximal leaf of $\Gamma^\perp$ containing q. Since the normal exponential map of L_q is also a local isometry, there exists an $\epsilon = \epsilon(q) > 0$ such that $B_\epsilon(q) \subset L_q$ is a normal ball and the set $\hat V_q := \{\xi \in T^\perp B_\epsilon(q) : \|\xi\| < \epsilon\} \subset T^\perp L_q$ is isometric to the Riemannian product $B_\epsilon(q) \times B_\epsilon \subset L_q \times \mathbb{R}^k \cong L_q \times T^\perp_q L_q$, and $\exp^\perp : \hat V_q \to V_q := \exp^\perp(\hat V_q)$ is an isometry. In particular, $q \in V_q \subset W$. We identify $B_\epsilon(q) \times B_\epsilon$, $\hat V_q$ and V_q via the normal exponential map of L_q. Accordingly, we denote the local leaf of $\Gamma^\perp$ through $x = (p, v) \in V_q$ by $L_{x,q} := B_\epsilon(q) \times \{v\} \subset L_x \cap V_q$. Moreover, for each $y \in V_q$ and $v \in T^\perp_y L_{y,q}$ there is a parallel vector field ξ in V_q with $\xi(y) = v$, and an isometric flow $\phi^\xi_t(x) = \exp_x(t\xi(x))$ for $x \in V_q$. Notice that this flow is defined for all $t \in \mathbb{R}$, and that the images of the leaves $L_{x,q}$ are again leaves of $\Gamma^\perp$ for all $t \in \mathbb{R}$, $x \in V_q$.

We now claim that $\exp^\perp$ is surjective onto W. Take a point $q \in W$ in the closure of the open set $U = \exp^\perp(T^\perp L)$ in W, and choose $y \in V_q \cap U$. Since the leaf of Γ containing y is also contained in the image of $\exp^\perp$, we can assume that $y \in L_{q,q}$. Then $y = \gamma(1)$, where $\gamma(t) = \exp^\perp(t\xi)$, for $x \in L$ and $\xi \in T^\perp_x L$. If η is the parallel vector field in V_q with $\eta(y) = \gamma'(1) \in T^\perp_y L_{q,q}$, then $\phi^\eta_{-1}(L_{q,q}) \subset L$ by maximality of L and hence $q \in U$. Hence, U is closed in W, so $U = W$.

In order to finish the proof that $\exp^\perp$ is a covering map, we show that it has the curve lifting property. Let $\alpha : [a, b] \to W$ be a smooth curve, and assume there exists a lift of $\alpha|_{[a,r)}$ to a curve $\tilde\alpha : [a, r) \to T^\perp L$. As usual, to extend $\tilde\alpha$ past r we only need to show that $\lim_{t\to r} \tilde\alpha(t)$ exists. Write $\tilde\alpha(t) = (c(t), \xi(t)) \in T^\perp L$. Since $\exp^\perp$ is a local isometry, $|\tilde\alpha'| = |\alpha'|$ and hence by the local product structure c and ξ have bounded length. Thus, $\lim_{t\to r} c(t) = x_\infty \in M$ and $\lim_{t\to r} \xi(t) = \xi_\infty$ exist by completeness of M^n and $\mathbb{R}^k$. We only need to show that $x_\infty \in L$ since then $\exp^\perp(x_\infty, \xi_\infty) = \alpha(r)$.

Consider $\delta > 0$ such that $\alpha((r-\delta, r]) \subset V_{\alpha(r)}$. For $t \in (r-\delta, r)$ let η_t be the parallel vector field in $V_{\alpha(r)}$ with $\eta_t(\alpha(t)) = \gamma_t'(1) \in T^\perp L_{\alpha(t)}$ where $\gamma_t(s) = \exp^\perp(s\xi(t))$. We then have $\lim_{t\to r} \eta_t = \eta_\infty \in T^\perp L_{\alpha(r)}$ with $\exp_{\alpha(r)}(-\eta_\infty) = x_\infty \in W$ and hence $\phi_{-1}^{\eta_\infty}(L_{\alpha(r),\alpha(r)}) \subset L_{x_\infty}$. Since we also have that $\phi_{-1}^{\eta_t}(L_{\alpha(t),\alpha(r)}) \subset L$ for $t < r$, it follows that $\lim_{t\to r} T_{c(t)}L = T_{x_\infty}L_{x_\infty}$ and thus $L \cup L_{x_\infty}$ is an integral leaf of $\Gamma^\perp$. By the maximality of L we conclude that $L_{x_\infty} \subset L$ and therefore $x_\infty \in L$, as we wished.

Finally, if $\pi\colon \tilde{L} \to L$ is the universal cover, then $\pi^*(T^\perp L) \to T^\perp L$ is also a cover and since $\pi^*(T^\perp L)$ is again a flat vector bundle over a simply connected base, it is isometric to $\tilde{L} \times \mathbb{R}^k$. This proves the last claim. □

4. The structure of the set of nonflat points

This section is devoted to the proof of the following stronger version of Theorem A.

Theorem 4.1. *Let M^n be a complete CN2 manifold, and V a connected component of the set of nonflat points of M^n. If V has finite volume, then its universal cover is isometric to $L^2 \times \mathbb{R}^{n-2}$, where L^2 is a simply connected surface whose Gauss curvature is nowhere zero and vanishes at its boundary.*

Proof. First, recall that, since $\Gamma = \ker R$ is a totally geodesic distribution, we have its *splitting tensor* $C : \Gamma \to End(\Gamma^\perp)$ defined as

$$C_T X = -(\nabla_X T)_{\Gamma^\perp},$$

where ∇ is the Levi-Civita connection of M^n, and a distribution as a subscript means to take the corresponding orthogonal projection. Clearly, $\Gamma^\perp$ is totally geodesic if and only if $C \equiv 0$, that is equivalent to the parallelism of Γ. Since the set of nonflat points in a CN2 manifold agrees with the points of minimal nullity, V is saturated by the flat complete leaves of Γ. Therefore, by Proposition 3.1, all we need to show is that C vanishes.

Let $U, S \in \Gamma$ and $X \in \Gamma^\perp$. Since Γ is totally geodesic,

$$\begin{aligned} C_{\nabla_U S} X &= -(\nabla_X \nabla_U S)_{\Gamma^\perp} = -(\nabla_U \nabla_X S)_{\Gamma^\perp} - (\nabla_{[X,U]} S)_{\Gamma^\perp} \\ &= (\nabla_U (C_S X))_{\Gamma^\perp} + C_S([X,U]_{\Gamma^\perp}) = (\nabla_U C_S) X + C_S(\nabla_U X) - C_S([U,X]_{\Gamma^\perp}) \\ &= (\nabla_U C_S) X + C_S(\nabla_X U) = (\nabla_U C_S) X - C_S C_U X, \end{aligned}$$

or

$$\nabla_U C_S = C_{\nabla_U S} + C_S C_U, \quad \forall\, U, S \in \Gamma. \tag{4.2}$$

We now consider the so called *nullity geodesics*, i.e. complete geodesics γ with $\gamma'(0) \in \Gamma$, which are hence contained in a leaf of Γ. Along such a geodesic γ, by (4.2) the splitting tensor $C_{\gamma'}$ satisfies the Riccati type differential equation $\nabla_{\gamma'} C_{\gamma'} = C_{\gamma'}^2$ over the entire real line. That is, with respect to a parallel basis,

$$C_{\gamma'}' = C_{\gamma'}^2, \text{ whose solutions are } C(t) = C_0(I - tC_0)^{-1}, \text{ for } C_0 := C_{\gamma'(0)}. \tag{4.3}$$

Therefore, along each nullity geodesic γ in V, all real eigenvalues of $C_{\gamma'}$ vanish. Since $\Gamma^\perp$ is 2-dimensional, for every $S \in \Gamma$ either all eigenvalues of C_S are complex and nonzero, or all eigenvalues are 0, i.e. C_S is nilpotent. In particular, C_S vanishes if it is self adjoint.

Let $W = \{p \in V : C \neq 0 \text{ at } p\}$, i.e. on $V \setminus W$ all splitting tensors vanish. Since the space of self-adjoint endomorphisms of $\Gamma^\perp$ is pointwise 3-dimensional and intersects $\operatorname{Im} C \subset End(\Gamma^\perp)$ only at 0, it follows that $\dim \operatorname{Im} C = 1$ in W, and hence $\ker C$ is a smooth codimension 1 distribution of Γ along W. Accordingly, write

$$\Gamma = \ker C \oplus^\perp \operatorname{span}\{T\},$$

for a unit vector field $T \in \Gamma$, which is well defined, up to sign, on W. By going to a two-fold cover of W if necessary, we can assume that T can be chosen globally on W.

Observe that if U, S are two sections of $\ker C$, then (4.2) implies that $\nabla_U S \in \ker C$, i.e. $\ker C$ is totally geodesic, and $\nabla_T U = 0$ as well. Since Γ is totally geodesic it follows that $\nabla_T T = 0$, that is, the integral curves γ of T are nullity geodesics. Therefore, from now on let for convenience $C = C_T$, $C(t) = C_{\gamma'(t)}$, and denote by $'$ the derivative in direction of T. In particular,

$$\operatorname{div} T = \operatorname{tr} \nabla T = -\operatorname{tr} C. \tag{4.4}$$

By (4.3) we have

$$\operatorname{tr} C(t) = \frac{\operatorname{tr} C_0 - 2t \det C_0}{1 - t \operatorname{tr} C_0 + t^2 \det C_0}, \quad \text{and} \quad \det C(t) = \frac{\det C_0}{1 - t \operatorname{tr} C_0 + t^2 \det C_0}. \tag{4.5}$$

Take $B \subset W$ a small compact neighborhood. Since either $\det C > 0$ or $\det C = \operatorname{tr} C = 0$ on W, by (4.5) there is $t_0 \in \mathbb{R}$ such that $\operatorname{tr} C(t)(q) \leq 0$ for every $q \in B$ and every $t \geq t_0$. In addition, defining $B_t := \phi_{t+t_0}(B)$ and $v(t) := \operatorname{vol} B_t$ we have that

$$v'(t) = \int_B \frac{d}{dt} \phi_t^*(dvol) = \int_B \operatorname{div} T = -\int_B \operatorname{tr} C \geq 0, \quad \forall\, t \geq 0.$$

So, the sequence of compact neighborhoods $\{B_n, n \in \mathbb{N}\}$ has nondecreasing volume in the set V of finite volume, and thus there is a strictly increasing sequence $\{n_k : k \geq 0\}$ such that $B_{n_k} \cap B_{n_0} \neq \emptyset$ for all $k \geq 1$. We will refer to this property as *weak recurrence*. In particular, there exists a sequence $p_k := \phi_{t_0+n_k}(q_k) \in B_{n_k} \cap B_{n_0}$, with $q_k \in B$, which has an accumulation point $p \in B_{n_0} \subset W$.

Consider the open subset $W' \subset W$ on which C has nonzero complex eigenvalues and notice that, by (4.3), W' is invariant under the flow ϕ_t of T. Using the above recurrence and sequence of points $p_k \to p$, (4.5) implies that

$$\det C_{T(p)} = \lim_{k \to +\infty} \det C_{T(p_k)} = \lim_{k \to +\infty} \frac{\det C_{T(q_k)}}{1 - (t_0 - n_k) \operatorname{tr} C_{T(q_k)} + (t_0 - n_k)^2 \det C_{T(q_k)}} = 0,$$

since $n_k \to +\infty$ and q_k lies in the compact set B. But this contradicts the fact that $p \in B_{n_0} \subset W'$, where $\det C > 0$. Thus C vanishes on W', which is a contradiction and shows that C is nilpotent on W.

We thus have a well defined 1-dimensional distribution on W spanned by the kernel of $C = C_T$, which is parallel along nullity lines by (4.3). Replacing W, if necessary, by the two-fold cover where this distribution has a section, and by a further cover to make W orientable, we can assume that there exists an orthonormal basis e_1, e_2 of $\Gamma^\perp$, defined on all of W, and parallel along nullity lines with

$$C(e_1) = 0, \quad C(e_2) = ae_1.$$

Hence

$$\nabla_T e_1 = \nabla_T e_2 = \nabla_T T = 0, \quad \nabla_{e_1} T = 0, \quad \nabla_{e_2} T = -ae_1,$$
$$\nabla_{e_1} e_1 = \alpha e_2, \quad \nabla_{e_2} e_2 = \beta e_1, \quad \nabla_{e_1} e_2 = -\alpha e_1, \quad \nabla_{e_2} e_1 = aT - \beta e_2,$$

for some smooth functions α, β on W. A calculation shows that

$$R(e_2, e_1)e_1 = (e_1(\beta) + e_2(\alpha) - \alpha^2 - \beta^2)e_2 + (a\beta - e_1(a))T,$$
$$R(e_1, e_2)e_2 = (e_1(\beta) + e_2(\alpha) - \alpha^2 - \beta^2)e_1 + \alpha aT,$$

and hence

$$\alpha = 0, \quad \mathrm{Scal}_M = e_1(\beta) - \beta^2, \quad \text{and} \quad e_1(a) = a\beta, \tag{4.6}$$

where Scal_M stands for the scalar curvature of M^n. The differential equation (4.3) implies that $a' = 0$, i.e. a is constant along nullity lines. Thus

$$e_2(a)' = e_2(a') + [T, e_2](a) = (\nabla_T e_2 - \nabla_{e_2} T)(a) = ae_1(a),$$

and $e_1(a)' = e_1(a') + [T, e_1](a) = 0$. So, $e_2(a) = ae_1(a)t + d$ for some smooth function d independent of t. By the weak recurrence property, $e_1(a)(p) = 0$ for $p \in B_{n_0}$ as above, and hence $e_1(a)(p') = 0$ as well, for $p' = \phi_{-n_0-t_0}(p) \in B$. Since B is arbitrary small, we have that $e_1(a) \equiv 0$ on W. But then (4.6) implies that $\beta = 0$ and thus $\mathrm{Scal}_M = 0$, contradicting that along V we have no flat points. Altogether, C vanishes everywhere on V and hence Γ is parallel. □

The proof becomes particularly simple for $n = 3$ since then Γ is one dimensional and only the differential equation $C' = C^2$ along the unique nullity geodesics is needed. The local product structure in this case was proved earlier in [SW] with more delicate techniques.

Let us finish this section with some observations about the geometric structure when the manifold is complete but the finite volume hypothesis is removed.

Remark. *a*) If C_T is nilpotent, then Scal is constant along nullity leaves. If this constant is positive, then the universal cover is isometric to $L^2 \times \mathbb{R}^{n-2}$. This follows since by (4.6) a^{-1} satisfies the Jacobi equation $(a^{-1})'' + a^{-1}\,\mathrm{Scal} = 0$ along integral curves of e_1, which cannot be satisfied for all $t \in \mathbb{R}$. For $n = 3$ this splitting was also proved in [AM].

b) If C_T does not vanish and has complex eigenvalues, then $\mathrm{Scal}(t) = \mathrm{Scal}(0)/(1 - t \operatorname{tr} C_0 + t^2 \det C_0)$ along nullity lines since one easily sees that $\mathrm{Scal}' = (\operatorname{tr} C_T)\,\mathrm{Scal}$. As was shown in [Sz], the universal cover of M^n splits off a Euclidean space of dimension $(n-3)$, i.e. all locally irreducible examples are 3-dimensional, c.f. Theorem C in the case of nonnegative curvature.

5. CN2 manifolds as geometric graph manifolds

The purpose of this section is to prove Theorem B in the introduction. As we will see, the proof is quite delicate and technical due to the lack of any *a priori* regularity of the boundary of a maximal full extension. We begin with some definitions.

Definition 5.1. We say that an open subset W of a topological space M is *locally finite* if, for every $p \in \partial W$, there exists an integer m such that, for every neighborhood $U \subset M$ of p, there exist a neighborhood $U' \subset U$ of p such that $U' \cap W$ has at most m connected components. We denote by $m(p)$ the minimum of such integers m.

In this situation, for every $p \in \partial W$, there exists a neighborhood $U_p \subset M$ of p such that $U_p \cap W$ has precisely $m(p)$ connected components. Notice that each component contains p in its boundary since by finiteness all other components have distance to p bounded away from 0. We call these components W_i the *local connected components of* W *at* p. Notice also that U_p can be chosen arbitrarily small. In fact, given any neighborhood U of p with $U \subset U_p$ we can construct a new neighborhood $U_p(U)$ of p as follows. Let X be the union of the connected components of $U \cap W$ that contain p in their boundary. Then $U_p(U) = (\overline{X})^\circ$ is a neighborhood of p (by local finiteness) and $U_p(U) \cap W$ has $m(p)$ connected components. Observe also that $U_p(U) \cap W_i$ are the connected components of $U_p(U) \cap W$ and all contain p in their boundary. Throughout this section U_p will always denote such a neighborhood of p and W_i the connected components of $U_p \cap W$.

In particular, by taking any $\delta > 0$ such that $B_\delta(p) \subset U_p(B_\epsilon(p))$ we get:

Lemma 5.2. *If M^n is a Riemannian manifold, $W \subset M^n$ is locally finite and $p \in \partial W$, then for every ball $B_\epsilon(p) \subset U_p$ there is $0 < \delta = \delta(\epsilon, p) < \epsilon$ such that $W_i \cap B_\delta(p)$ is arc-connected in $W_i \cap B_\epsilon(p)$, for all $1 \leq i \leq m$.*

Let M^n be a complete CN2 manifold whose nullity distribution is parallel along the set of nonflat points V of M^n, as is the case when M^n has finite volume by Theorem A. Suppose V has a full extension $W \subset M^n$, that is, $W \supset V$ is open, dense, locally finite, and W possesses a smooth parallel distribution Γ of rank $n-2$ whose leaves are flat and complete. Since any extension of a full extension is also full, we will assume in addition that W is maximal. Clearly, along $V \subset W$, Γ coincides with the nullity of M^n. Observe that maximal extensions always exist by definition, but they are not necessarily unique, and may fail to be dense or locally finite as shown in Examples 3 and 4 in Section 1. We call the leaves of Γ nullity leaves and, for simplicity, use $\Gamma(p)$ both for the distribution at p and for the leaf of Γ through p.

Observe that, for each sequence $\{p_n\}$ in W approaching a point $p \in \partial W$, $\Gamma(p_n)$ accumulates at an $(n-2)$-dimensional subspace of T_pM whose image under the exponential map gives a complete totally geodesic submanifold of M^n, by completeness of the leaves of Γ. We still denote the set of all these limit submanifolds by $\Gamma(p)$, and call each of them a *boundary nullity leaf* at p, or BNL for short. In addition, given $U \subset W$ with $p \in \partial U$, we denote by $\Gamma_U(p) \subset \Gamma(p)$ the BNL's at p that arise as limits of nullity leaves in U. In particular, if $W_1, \ldots, W_m$ are the local connected components of W at p, we have

$$\Gamma(p) = \Gamma_{W_1}(p) \cup \cdots \cup \Gamma_{W_m}(p).$$

We start with the following observation.

LEMMA 5.3. *Let $p \in \partial W$ such that $\Gamma(p)$ has only one BNL μ. Then, $\Gamma(q) = \{\mu\}$ for all $q \in \mu$.*

Proof. By definition, there is a unique BNL $\mu \in \Gamma(p)$. The hypersurface $B'_\epsilon(p) := \exp(\mu^\perp \cap B_\epsilon(0_p))$ is then transversal to Γ in $W'_\epsilon(p) = W \cap B'_\epsilon(p)$, which is thus dense in $B'_\epsilon(p)$.

Take $q \in \mu$, and write it as $q = \gamma_v(T)$ for some $v \in T_p\mu$, $T \in \mathbb{R}$, $\|v\| = 1$. By continuity of the geodesic flow at v, given $\delta > 0$, there is $\epsilon > 0$ such that all the nullity leaves of W through $W'_\epsilon(p)$ are C^0 δ-close to μ inside a compact ball of radius, say, $2T$, centered at p. In particular, these nullity lines of $W'_\epsilon(p)$ stay close to μ at q and form an open dense subset. This implies that there cannot be two different BNL's at q. Indeed, a second $\mu' \in \Gamma(p)$ is a limit of leaves of Γ of a local connected component W' at q. Thus all leaves in W' are close to μ' which implies that leaves of Γ on W, where it is an actual foliation, would intersect near q. □

For the next two lemmas we need a relationship between curvature bounds and local parallel transport for Riemannian vector bundles over surfaces.

LEMMA 5.4. *Let E^k be a Riemannian vector bundle with a compatible connection ∇ over a surface S, and let $D \subset S$ be a region diffeomorphic to a closed 2-disk with piecewise smooth boundary α. If the curvature tensor of E^k is bounded by $\delta > 0$ along D, then the angle between any vector $\xi \in E_{\alpha(0)}$ and its parallel transport along α is bounded by $(k-1)\,\delta Area(D)$.*

Proof. Let us consider polar coordinates on D through a diffeomorphism with a 2-disk. We can assume that ξ is a unit vector and we complete $\xi = \xi_1, \ldots, \xi_k$ to an orthonormal basis of $E_{\alpha(0)}$. By radially parallel transporting them first to p, and then radially to all of D, we get an orthonormal basis, which we again denote by $\xi_1, \ldots, \xi_k$, defined on D. If we consider the connection 1-forms $w_{ij}(X) = \langle \nabla_X \xi_i, \xi_j \rangle$ on D, then one easily sees that $dw_{ij} = \langle R_\nabla(\cdot,\cdot)\xi_i, \xi_j \rangle$ since $\dim D = 2$ and $w_{ij}(Y) = 0$ for the radial direction Y.

Let $\xi(t) = \sum_i a_i(t)\xi_i(\alpha(t))$ be the parallel transport of ξ along α between 0 and t. Then, since $\nabla_{\alpha'}\xi = 0$, we have $a_1' = \langle \xi_1, \xi \rangle' = \langle \nabla_{\alpha'}\xi_1, \xi \rangle = \langle \nabla_{\alpha'}\xi_1, \sum_{i=1}^k a_i\xi_i \rangle = \sum_{i=2}^k a_i w_{1i}(\alpha')$. Therefore, since $|a_i| \le 1$ we obtain

$$0 \le 1 - \langle \xi(0), \xi(1) \rangle = -\int_0^1 a_1' = -\sum_{i=2}^k \int_0^1 a_i\, w_{1i}(\alpha') \le \sum_{i=2}^k \int_\alpha |w_{1i}|.$$

For each i, choose a partition of α into countable many segments $\alpha_1, \alpha_2, \ldots$ with $w_{1i}|_{\alpha_{2j}} \ge 0$ and $w_{1i}|_{\alpha_{2j+1}} \le 0$. Then, α_j together with the radial curves (along which $w_{1i} = 0$) encloses a triangular region T_j where we apply Stokes Theorem to get

$$\int_\alpha |w_{1i}| = \sum_j \int_{T_{2j}} dw_{1i} - \sum_j \int_{T_{2j+1}} dw_{1i} \leq \int_D |\langle R_\nabla(\cdot,\cdot)\xi_1, \xi_i\rangle| \leq \delta Area(D). \qquad \square$$

In the following lemmas we study the behavior of the nullity leaves and BNL's in $U_p \subset B_\epsilon(p)$. To do this, we will be able to restrict the discussion to a single surface S transversal to the nullity leaves and BNL's near p due to the following result; see Figure 5.

LEMMA 5.5. *For each $p \in \partial W$ there exist a 2-plane $\tau \subset T_pM$, a sufficiently small convex ball $B_\epsilon(p)$, and a neighborhood $U_p \subset B_\epsilon(p)$ of p such that the surface $S := \exp(\tau \cap B_\epsilon(0_p)) \subset B_\epsilon(p)$ satisfies:*

a) S intersects all nullity leaves and BNL's in U_p, and does so transversely;
b) $W_i \cap S$ is connected and its closure is not contained in $U_p \cap S$;
c) $U_p \cap S$ is diffeomorphic to an open disc.

Proof. For a fixed $1 \leq i \leq m(p)$, choose some BNL $\mu \in \Gamma_{W_i}(p)$ and for some convex ball $B_\epsilon(p)$ consider the surface $L = \exp(\mu^\perp \cap B_\epsilon(0_p))$. Given $\delta > 0$, we will fix $\epsilon > 0$ such that $(n-1)\max\{|\operatorname{Scal}_M(x)| : x \in B_\epsilon(p)\} Area(L) < \delta$ and choose $U_p \subset B_\epsilon(p)$. Notice also that a bound on Scal_M gives a bound on the full curvature tensor since M^n is CN2. Take a sequence $p_i \in W_i$ converging to p such that $\Gamma(p_i) \to \mu$. For i large enough $\Gamma(p_i)$ is transversal to L and we can thus assume that $p_i \in W_i \cap L$. Furthermore, fix i large enough such that for $q := p_i$ the parallel translate of $\Gamma(q)$ along the minimal geodesic $\overline{qp}$ from q to p has angle less than δ with μ. Let W_i' be the arc-connected component of $W_i \cap L$ that contains q. Thus for any $q' \in W_i'$, we can choose a curve $\alpha \subset W_i'$ connecting q to q'. If $\beta = \overline{pq}$ and $\beta' = \overline{q'p}$, we form the closed curve $\varphi = \beta * \alpha * \beta' \subset L$. We can also choose α such that φ is simple and hence bounds a disc $D \subset L$. According to Lemma 5.4 the parallel transport of μ along φ forms an angle less than δ with μ. In other words, the parallel transport of $\Gamma(q')$ along $\overline{q'p}$ has angle less than 2δ with μ for any $q' \in W_i'$. We can thus choose $\epsilon' < \epsilon$ sufficiently small, and $U_p \subset B_{\epsilon'}(p)$ such that all nullity leaves in W_i' intersect L transversely at an angle bounded away from 0. We now claim that this implies that $W_i' = W_i$, i.e. $W_i \cap L$ is connected. Otherwise, there exists an $x \in W_i$ with $x \in \partial W_i'$. Since $\Gamma(x)$ still intersects L transversely, we can choose a small product neighborhood $U \subset W_i$ as in the proof of Theorem 3.1 such that $x \in U$ and all nullity leaves in U intersect L in a unique point and transversely. But then any two points in $U \cap L$ can be connected in U and then projected along nullity leaves to lie in L. Thus $U \cap L$ is also contained in W_i'.

Since there are only finitely many connected components, and all components of $U_p \cap W$ contain p in their boundary, there exists a common ϵ' and BNL's $\mu_i \in \Gamma_{W_i}(p)$ satisfying the above properties. We can now choose a 2-plane $\tau \subset T_pM$ transversal to all μ_i and set $S := \exp(\tau \cap B_\epsilon(0_p))$. Repeating the above argument for this surface S, we see that ϵ can be chosen sufficiently small such that all nullity leaves in $U_p = U_p(B_\epsilon(p))$ intersect S

transversely and that for all its connected components $S \cap W_i$ is connected as well. Since in addition we can assume that the angle between the nullity lines and S is bounded away from 0, all BNL's are transversal to S as well. Notice also that now the components of $W \cap U_p \cap S$ are precisely $W_i \cap S$.

So far S and U_p satisfy the properties in part (a) and the first part of (b). For the second part of (b), since W is locally finite, we simply choose $\epsilon' < \epsilon$ small enough such that the closure of the connected components $W_i \cap S$ are not strictly contained in $B_{\epsilon'}(p)$. But then $U_p(B_{\epsilon'}(p))$ is the desired neighborhood (see Figure 5).

We now claim that such a neighborhood is also simply connected. Let α be a closed curve in U_p which bounds a disc $D \subset S$. If D contains a point in another component W' of $W \cap S$, then W' is fully contained inside D since it does not touch α. Hence the closure of W' is contained in U_p, which contradicts (b). Thus D is contained in U_p, and we can find a null homotopy of α in $D \subset U_p$. □

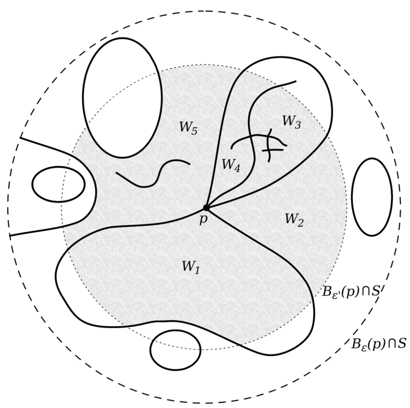

FIGURE 5. A point $p \in \partial W$ with $m(p) = 5$, the dark lines represent $\partial W \cap S$, while the shaded area corresponds to $U_p \cap S$

We will use $\epsilon > 0$, S, and U_p as in Lemma 5.5 for the remainder of this section. We point out though that the open sets $W_i \cap S$ can have quite complicated boundary. In fact, ∂W_i may not even be a Jordan curve and hence may not consist of a union of continuous arcs. Furthermore, for $p \in \partial W_i$ there may not even be a continuous curve in W_i with endpoint p. We thus carefully avoid using any such assumptions on properties of these boundaries.

It already follows from the proof of Lemma 5.5 that all BNL's in $\Gamma_{W_i}(p)$ form a small angle. We will now show that it is in fact unique.

LEMMA 5.6. *If W_i are the local connected components at p, then $\Gamma_{W_i}(p)$ is a single BNL, for each $1 \leq i \leq m$.*

Proof. Let $\mu_1, \mu_2 \in \Gamma_{W_i}(p)$ be two BNL's at p and two sequences $p_{r,k} \in W_i \cap S$, $r = 1, 2$, converging to p with $\Gamma(p_{r,k}) \to \mu_r$. For any $\epsilon' > 0$ choose $0 < \delta(\epsilon', p) < \epsilon'$ as in Lemma 5.2.

For k large enough $p_{r,k} \in B_{\delta'}(p)$ and we can choose a curve $\alpha_k \subset W_i \cap B_{\epsilon'}(p)$ connecting $p_{1,k}$ to $p_{2,k}$ and by Lemma 5.5 we can also assume that α_k lies in S. Now define the loop $\varphi_k = \beta_{2,k} * \alpha_k * \beta_{1,k}^{-1} \subset S \cap B_{\epsilon'}(p)$, where $\beta_{r,k} = \overline{p_{r,k}p}$. We can assume it is a simple closed curve and hence encloses a 2-disk $D \subset S \cap B_{\epsilon'}(p)$. Therefore, Lemma 5.4 implies that the angle between μ_1 and its parallel transport along φ_k is bounded by $(n-1)Area(S)s(\epsilon')$, where $s(\epsilon') := \max\{|\operatorname{Scal}_M(x)| : x \in B_{\epsilon'}(p)\}$. On the other hand, the parallel transport of $\Gamma(p_{r,k})$ along $\beta_{r,k}$ converges to μ_r as $k \to \infty$ and the parallel transport of $\Gamma(p_{1,k})$ along α_k is equal to $\Gamma(p_{2,k})$. Hence the angle between μ_2 and the parallel transport of μ_1 along φ_k goes to 0 as $k \to \infty$. Finally, $s(\epsilon') \to 0$ as $\epsilon' \to 0$ since $\operatorname{Scal}_M(p) = 0$ and we conclude that $\mu_1 = \mu_2$ as $\epsilon' \to 0$. $\square$

LEMMA 5.7. *For $q \in \partial W_i \cap \partial W_j \cap S$, both $\Gamma_{W_i}(q)$ and $\Gamma_{W_j}(q)$ also contain a unique BNL, and the angle between them coincides with the angle between $\Gamma_{W_i}(p)$ and $\Gamma_{W_j}(p)$.*

Proof. Fix $\delta > 0$ and let $M_\delta = \{x \in M^n : |\operatorname{Scal}_M(x)| < \delta\}$. Then $p \in \partial W \subset V \subset M_\delta$. Choosing U_p and S as in Lemma 5.5 we can study Γ in U_p in terms of its intersection with S, and in the following drop S for clarity. In addition, assume that $\pm\delta$ are regular values of Scal_M restricted to S, and observe that $M_\delta \cap U_p$ is an open neighborhood of $\partial W_i \cap \partial W_j$. We denote by r either i or j.

Since $\pm\delta$ are regular values, the set $\{|\operatorname{Scal}_M| \geq \delta\} \cap U_p$ is contained in the union of finitely many closed disjoint 2-discs (or half disks) D_ℓ. If we remove from W_r those discs D_ℓ which are contained in it, we obtain the open set $W_r' \subset W_r$, which is connected since W_r is. Let $\mu_r \in \Gamma_{W_r}(q)$ be a BNL and set $\nu_r = \Gamma_{W_r}(p)$, which contains only one element by Lemma 5.6. Choose two sequences $p_{r,k}, q_{r,k} \in W_r'$ such that $p_{r,k} \to p$, $q_{r,k} \to q$, with $\Gamma(q_{r,k}) \to \mu_r$ and $\Gamma(p_{r,k}) \to \nu_r$. Choose smooth simple curves $\alpha_{r,k} \subset W_r'$ joining $p_{r,k}$ to $q_{r,k}$, and let $\beta_{r,k} = \overline{p_{r,k}p}$ and $\gamma_{r,k} = \overline{q_{r,k}q}$, which, since ∂M_δ has positive distance to p and q, we can assume to lie in M_δ for k sufficiently large. Thus we get two curves $\varphi_{r,k} = \beta_{r,k}^{-1} * \alpha_{r,k} * \gamma_{r,k}$ from p to q, and hence a closed curve $\varphi_k := \varphi_{j,k}^{-1} * \varphi_{i,k} \subset M_\delta \cap U_p$, which we can also assume to be simple. By part (c) in Lemma 5.5, φ_k bounds a 2-disk $D \subset U_p$.

We claim that $\alpha_{r,k} \subset M_\delta \cap W_r$ can be modified in such a way that $D \subset M_\delta \cap U_p$. First observe that any closed disc $D_\ell \subset D$ as above must be contained in either W_i or W_j since $\partial D_\ell \cap \varphi_k = \emptyset$ and no component has its closure contained in U_p. For each $D_\ell \subset W_r$, by means of a smooth curve $\phi_\ell \subset W_r' \cap D$ connecting the boundary of D_ℓ with a point y_ℓ in $\alpha_{r,k}$ we can contour D_ℓ from the interior of D by following $\alpha_{r,k}$ up to y_ℓ, $\phi_\ell * \partial D_\ell * \phi_\ell^{-1}$, and the remaining part of $\alpha_{r,k}$. We can repeat this procedure for each D_ℓ and can also arrange this in such a way that all curves ϕ_ℓ are disjoint. Observe that this new curve, that we still call $\alpha_{r,k}$, is contained in $W_r \cap M_\delta$, and the claim is proved (see Figure 6).

As $k \to \infty$, the parallel transport of ν_r along $\varphi_{r,k}$ approaches μ_r since Γ is parallel in W_r. By Lemma 5.4, the angle between the parallel transport of ν_i along $\varphi_{i,k}$ and along $\varphi_{j,k}$ can be bounded by $(n-1)\delta Area(S)$. Since the angle between ν_i and ν_j and their parallel transport along ν_j is the same, the claim follows by taking $\delta \to 0$.

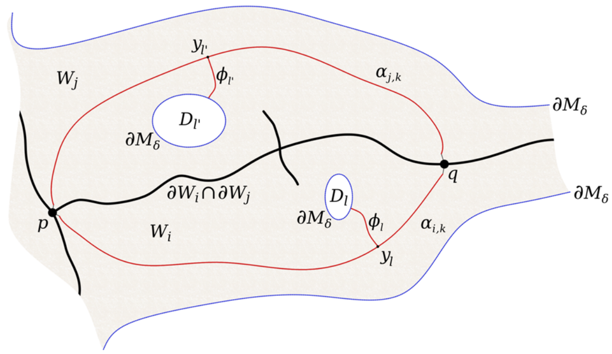

FIGURE 6. A neighborhood of $\partial W_i \cap \partial W_j$ in S with the shaded area representing M_δ

Finally, assume that there are two BNL's in $\Gamma_{W_r}(q)$. We can repeat the above argument with curves lying only in W_r since we did not assume that $i \neq j$, and it follows from Lemma 5.6 that the angle between them is 0. $\square$

LEMMA 5.8. *The distribution* Γ *does not extend continuously to any neighborhood of any* $p \in \partial W$.

Proof. Suppose that Γ extends continuously to $B_\delta(p)$. Let $\hat{\Gamma}$ be the smooth distribution in $B_\delta(p)$ obtained by parallel transporting $\Gamma(p)$ along geodesics emanating from p. Choosing a surface S centered at p as in Lemma 5.5, and contained in $B_\delta(p)$, we first claim that on S the distribution Γ agrees with $\hat{\Gamma}$. To see this, let $\alpha \subset S$ be a geodesic stating at p, and consider the angle function $\alpha(t)$ between $\Gamma(\gamma(t))$ and $\hat{\Gamma}(\gamma(t))$. An argument similar to the one in the proof of Lemma 5.7 shows that α is locally constant along the finitely many (not necessarily connected) sets $\gamma \cap \overline{W_i}$. Since by hypothesis α is continuous with $\alpha(0) = 0$, we conclude that $\alpha = 0$, as desired. Since the leaves of both Γ and $\hat{\Gamma}$ intersect S transversely, they must agree in a neighborhood of S on which Γ is thus smooth. By Lemma 5.3, this property also holds on the union of all complete leaves going through S, which contradicts the maximality of W. $\square$

LEMMA 5.9. *For every* $p \in \partial W$ *we have that* $2 \leq \#\Gamma(p) \leq m(p)$.

Proof. First, observe that Lemma 5.7 for $i = j$ shows that $\Gamma_{W_i}(q)$ contains a unique BNL for all $q \in \partial W_i \cap S$. We claim that this implies that $\partial W \cap S = \bigcup_{i \neq j}(\partial W_i \cap \partial W_j \cap S)$. If not, there exists a local component W_i, $q \in \partial W_i \cap S$ and a neighborhood U of q such that $U \subset \overline{W_i}$ and $U \cap \overline{W_j} = \emptyset$ for $j \neq i$ (see Figure 5). But then for every $r \in U \cap \partial W \cap S \subset U \cap \partial W_i \cap S$ we have that $\Gamma_{W_i}(r)$ contains a unique BNL. By Lemma 5.3 this is then also the case for any $r \in U \cap \partial W$, i.e. Γ is continuous in U, which contradicts Lemma 5.8.

We now show that $\#\Gamma(p) \geq 2$. So assume that $\Gamma(p)$ has a unique element, i.e. $\Gamma_{W_i}(p) = \Gamma_{W_j}(p)$ for all $i \neq j$. Then Lemma 5.7 implies that the same is true for any $q \in \partial W_i \cap \partial W_j \cap S$

and hence by the above for any $q \in \partial W \cap S$. Thus Γ is continuous in U_p, which again contradicts Lemma 5.8. The second inequality follows from Lemma 5.6. □

LEMMA 5.10. *There exists $\delta = \delta(p) > 0$ such that $W_i \cap B_\delta(p)$ is convex for all i.*

Proof. Let $\epsilon' > 0$ such that $B_{\epsilon'}(p) \subset U_p \subset B_\epsilon(p)$, and let $\delta = \delta(\epsilon', p)$ as in Lemma 5.2. Take points $q, q' \subset W_i \cap B_\delta(p)$ for which the minimizing geodesic segment $\overline{qq'}$ is not contained in $W_i \cap B_\delta(p)$. Take a curve $\alpha \subset W_i \cap B_{\epsilon'}(p)$ joining q with q', and set $s := \sup\{r : \sigma_t \subset W_i, \forall 0 \leq t < r\}$, where $\sigma_t = \overline{q\alpha(t)}$. Since $\overline{q\alpha(s)} \subset \overline{W_i} \cap B_{\epsilon'}(p)$, we have that $\overline{q\alpha(s)} \cap \partial W_i \neq \emptyset$. We claim that $m(x) = 1$ for all $x \in \overline{q\alpha(s)} \cap \partial W_i$, which contradicts the first inequality in Lemma 5.9.

To prove the claim, take $\mu \in \Gamma(x)$ and consider for each $0 < t < s$ the flat totally geodesic completely ruled hypersurface $H_t := \cup_{0<r<1}\Gamma(\sigma_t(r)) \subset W$ with limit $H := \lim_{t\to s} H_t$. If H intersects μ transversally, H_t would also for t close to s, which is a contradiction since $\mu \subset \partial W$. Therefore, $T_x\mu$ is a hyperplane contained in T_xH. Since H is foliated by complete flat hypersurfaces parallel to $\Gamma(q)$ along σ_s and $\mu \subset H$ is also a complete flat hypersurface, it follows that μ is parallel to $\Gamma(q)$ along σ_s as well. Thus μ is unique and hence $m(x) = 1$. □

We now come to the main result about the local structure of ∂W.

LEMMA 5.11. *The set $F_{ij} := \partial W_i \cap \partial W_j \subset \partial W$ is convex for all i, j, and along every geodesic in F_{ij} the two families of BNL's induced by W_i and W_j are parallel.*

Proof. Take two points $q, r \in F_{ij}$ and, for $k = i, j$, sequences $q_{k,n}, r_{k,n} \in W_k$ such that $q_{k,n} \to q$, $r_{k,n} \to r$. By convexity, $\overline{q_{k,n}r_{k,n}} \subset W_k$ and since both converge to $\overline{qr}$, it follows that $\overline{qr} \subset F_{ij}$. For the second assertion, simply observe that the parallel transport along $\overline{qr}$ of the BNL's agrees with the limits of the parallel transport along $\overline{q_{k,n}r_{k,n}}$. □

We are finally in a position to prove Theorem B, which follows from Theorem A and the following.

THEOREM 5.12. *Let M^n be a complete Riemannian manifold with a parallel rank $n-2$ distribution defined in a dense, locally finite and maximal open set W, whose leaves are complete and flat. Then $M^n \setminus W$ is a disjoint union of complete flat totally geodesic hypersurfaces. If, in addition, M^n has finite volume, then these hypersurfaces are compact and M^n is a geometric graph manifold.*

Proof. Consider F_{ij} as in Lemma 5.11 with $\Gamma_{W_i}(p) \neq \Gamma_{W_j}(p)$ which exists by Lemma 5.9. Then for $r = i, j$, each point in the interior F°_{ij} of F_{ij} is contained in a unique complete BNL of W_r, and we denote by $S_r \subset \partial W_r$ the union of such BNL's with $S_i \cap S_j \supset F^\circ_{ij}$. Observe in addition that S_r is a smooth flat totally geodesic hypersurface, and completely ruled since, as seen in the proof of Lemma 5.11, it arises as a limit of $H_n := \cup_{0<t<\epsilon}\Gamma(\sigma_n(t)) \subset W_r$ for a sequence of geodesic segments $\sigma_n \subset W_r$.

We now study the local connected components based at a point $q \in S_i \cap \partial S_j \subset \partial F_{ij}^{\circ}$. Clearly, W_i is one of those, with S_i smooth at q and W_i lying (locally) on one side of S_i. Let W_1', W_2' be any two other local components at q with BNL's $\mu_s \in \Gamma_{W_s'}(q)$ for $s = 1, 2$. Observe first that μ_s cannot be transversal to S_i since otherwise leaves of Γ in W_i and W_s' would intersect. Hence μ_1 and μ_2 are tangent to S_i, which implies that $\mu_1 = \mu_2$. Indeed, otherwise μ_1 and μ_2 would intersect transversally in S_i by dimension reasons, and then near q the leaves of W_1' and W_2' would again intersect since they are both locally on the same side of the hypersurface S_i. Therefore, all local connected components at q, apart from W_i, share the same BNL and thus $\#\Gamma(q) = 2$. Using a surface S at q and $U_q \subset S$ as in Lemma 5.5, we see that $W_i \cap S$ is a half disc with boundary a smooth geodesic containing q in its interior. For all remaining local connected components W_s' at q, Lemma 5.11 implies that the intersections $\partial W_1' \cap \partial W_2' \cap S$ are geodesics with endpoints at q. Since $\Gamma_{W_1'}(q) = \Gamma_{W_2'}(q)$, Lemma 5.7 implies that $\Gamma_{W_1'}(r) = \Gamma_{W_2'}(r)$ for all $r \in \partial W_1' \cap \partial W_2' \cap S$ and hence by Lemma 5.3 also in a neighborhood of S. Thus by Lemma 5.8 there can be only one such component, i.e. $m(q) = 2$, and hence W_j is the second component at q. But then S_j extends past q and hence F_{ij} is a complete flat hypersurface containing p in its interior. We conclude that the number of local connected components at ∂W is 2 everywhere and ∂W is a disjoint union of complete flat totally geodesic embedded hypersurfaces.

So far we have that F is a common boundary component of two global connected components W_r of W (that may even be the same). Each W_r induces a different complete totally geodesic codimension one foliation on F by means of its BNL's $\Gamma_{W_r} \subset F$, which are precompact by Corollary 2.10 if M^n has finite volume. If one such leaf is not compact, its closure is F, which is thus compact. If all leaves are compact, then F is easily seen to be compact using that the normal exponential map of $\mu \in \Gamma_{W_1}$ is a covering. $\square$

Remark. One of the difficulties in proving Conjecture 1 in the Introduction is that one needs to exclude the following situation when local finiteness fails. Let W' be a concave local connected component at p with $\partial W'$ consisting of two smooth hypersurfaces meeting along their common boundary BNL at p. Then one needs to show that the complement of W' near p cannot be densely filled with infinitely many disjoint twisted cylinders whose diameters go to 0 as they approach p.

6. A dichotomy and the proof of Theorem C

In this section we provide the general structure of geometric graph manifolds with nonnegative scalar curvature by showing a dichotomy: they are build from either one or two twisted cylinders over 2-disks. This will then be used to prove Theorem C.

Let M^n be a complete nonflat geometric graph manifold with finite volume and nonnegative scalar curvature. Since a complete noncompact manifold with nonnegative Ricci curvature has linear volume growth by [Ya], M^n is compact. We can furthermore assume that M^n is not itself a twisted cylinder since in this case the universal cover of M^n

is isometric to $\mathbb{S}^2 \times \mathbb{R}^{n-2}$, where $\mathbb{S}^2$ is endowed with a metric of nonnegative Gaussian curvature.

By assumption, there exists a collection of compact flat totally geodesic hypersurfaces in M^n whose complement is a disjoint union of (open) twisted cylinders C_i. Let $C = (L^2 \times \mathbb{R}^{n-2})/G$ be one of these cylinders whose boundary in M^n is a disjoint union of compact flat totally geodesic hypersurfaces. There is also an *interior boundary* $\partial_i C$ of C, which we also denote for convenience as ∂C by abuse of notation. This boundary can be defined as the set of equivalence classes of Cauchy sequences $\{p_n\} \subset C$ in the interior distance function d_C of C, where $\{p_n\} \sim \{p_n'\}$ if $\lim_{n\to\infty} d_C(p_n, p_n') = 0$. Since M^n is compact, such a Cauchy sequence $\{p_n\}$ converges in M^n, and we have a natural map $\sigma : \partial C \to M$ that sends $[\{p_n\}]$ into $\lim_{n\to\infty} p_n \in M^n$. This map is, on each component of ∂C, either an isometry or a local isometric two-fold cover since $H = \sigma(\partial C)$ consists of disjoint smooth hypersurfaces which are two-sided in the former case, and one-sided in the latter. Therefore, ∂C is smooth as well and $C \sqcup \partial C$ is a closed twisted cylinder with totally geodesic flat compact interior boundary as in Section 2, that by abuse of notation we still denote by C. Similarly, L^2 is a smooth surface with geodesic interior boundary components. Notice though that L^2 may be noncompact and the boundary geodesics may not be closed; see Section 2. By Remark 2.5, we can assume further that none of the cylinders C_i is flat.

We first show a crucial property implied by nonnegative curvature:

Proposition 6.1. *The surface L^2 is isometric to a 2-disk D with nonnegative Gaussian curvature, whose boundary is a closed geodesic along which the curvature of D vanishes to infinite order.*

Proof. Following the notation in Section 2, if $G_1 \subset \mathrm{Iso}(L^2)$ is not discrete this follows from Corollary 2.8. So assume G_1 is discrete. Proposition 2.9 implies that L^2/G_1 is a compact orbifold, and hence its boundary consists of closed geodesics. In particular, for each boundary geodesic γ of L^2 which is not closed, there exists $\mathbb{Z} \subset G_1$ which acts as translations along γ. Thus if γ_i, $i = 1, 2$, are two boundary geodesics of L^2, there exists a minimal geodesic δ connecting γ_1 with γ_2 with $\delta(0) = \gamma_1(t_1)$ and $\delta(r) = \gamma_2(t_2)$. If X is the parallel vector field along δ with $X(0) = \gamma_1'(t_1)$ then $X(r) = \gamma_2'(t_2)$. Let $\alpha : [0, r] \times \mathbb{R} \to M^n$ be defined as $\alpha(t, s) := \exp_{\delta(t)}(sX(t))$. Since the Gaussian curvature of L^2 is nonnegative, the Rauch comparison theorem implies that, for small s, the curves $t \to \alpha(t, s)$ from γ_1 to γ_2 have length at most that of δ, and are hence again minimizing geodesics from γ_1 to γ_2. By the rigidity in the Rauch comparison theorem $(t, s) \to \alpha(t, s)$ is a flat strip for small s. Repeating the argument, we see that it is a flat strip for all t and s, containing γ_i as its two boundary components. This implies that all of L^2 is flat since any point can be connected to γ_1 by a minimal geodesic. But then C is flat as well, a case that was excluded above.

So far, we have that L^2 is a smooth surface whose boundary is a single geodesic γ. Since G_1 acts effectively and isometrically on γ, $G_1 \simeq \mathbb{Z}_k$ or D_k, in which case γ is closed, or

$G_1 \simeq \mathbb{Z}$ or D_∞ if γ is not closed. In the first case, L^2 is compact since $L^2/G_1 = L^2/G_1$ is. But then the Gauss-Bonnet Theorem implies that L^2 is a disc. In the second case, $\mathbb{Z} \subset G_1$ and since its stabilizer groups are finite, $\mathbb{Z}$ acts freely on L^2 . Thus $L^2/\mathbb{Z}$ is smooth and compact with boundary a closed geodesic. But this must again be a disc, which is a contradiction since it cannot be nontrivially covered by L^2.

Finally, it is easy to see that the fact that the nullity distribution does not extends globally implies the curvature properties of L^2 at γ. $\square$

Remark 6.2. The properties at the boundary γ of a disk D as in Proposition 6.1 are easily seen to be equivalent to the fact that the natural gluing $D \sqcup (\gamma \times (-\epsilon, 0])$, $\gamma \cong \gamma \times \{0\}$, is smooth when we consider on $\gamma \times (-\epsilon, 0]$ the flat product metric. In fact, in Fermi coordinates $(s \geq 0, t)$ along γ, the metric is given by $ds^2 + f(t,s)dt^2$. The fact that γ is a (unparametrized) geodesic is equivalent to $\partial_s f(0,t) = 0$, while the curvature condition is equivalent to $\partial_s^k f(0,t) = 0$ for all t and $k \geq 2$. Therefore, $f(s,t)$ can be extended smoothly as $f(0,t)$ for $-\epsilon < s < 0$, which gives the smooth isometric attachment of the flat cylinder $\gamma \times (-\epsilon, 0]$ to D.

As a consequence of Proposition 6.1, $\partial C = (\gamma \times \mathbb{R}^{n-2})/G$ is connected, and so is $H = \sigma(\partial C)$. In particular, M^n contains at most two twisted cylinders with nonnegative curvature glued along H. We call such a connected compact flat totally geodesic hypersurface H a *core* of M^n. We conclude:

COROLLARY 6.3. *If M^n is not flat and not itself a twisted cylinder, then $M^n = W \sqcup H$ with core H, and either:*

a) H is two-sided, σ is an isometry, and $W = C \sqcup C'$ is the disjoint union of two open nonflat twisted cylinders as above attached via an isometry $\partial C \simeq H \simeq \partial C'$; or

b) H is one-sided, σ is a local isometric two-fold cover, $W = C$ is a single open nonflat twisted cylinder as above, and $M^n = C \sqcup H = C \sqcup (\partial C/\mathbb{Z}_2)$.

Furthermore, in case (a), if $H' \subset M^n$ is an embedded compact flat totally geodesic hypersurface then there exists an isometric product $H \times [0,a] \subset M^3$, with $H = H \times \{0\}$ and $H' = H \times \{a\}$. In particular, any such H' is a core of M^3, and hence the core is unique up to isometry. On the other hand, in case (b) the core H is unique.

Proof. We only need to prove the last assertion, and let us start with case (a). First assume that $H \cap H' \neq \emptyset$. Then the BNL Γ of C is contained in H'. Indeed if not, the product structure of the universal cover $\tilde{C} = L \times \mathbb{R}^{n-2}$, together with the fact that H' is totally geodesic and complete, would imply that L^2, and hence C, is flat. Analogously, the (distinct) BNL of C' is in H', and since H is the unique hypersurface containing both BNL's, we have that $H = H'$. If, on the other hand, $H \cap H' = \emptyset$, we can assume $H' \subset C = (L^2 \times \mathbb{R}^{n-2})/G$. Again by the product structure of $\tilde{C}$ and the fact that H' is embedded we see that $H' = (\gamma' \times \mathbb{R}^{n-2})/G'$ where $\gamma' \subset L^2$ is a simple closed geodesic and $G' \subset G$ the subgroup preserving γ'. Since the boundary γ of L^2 is also a closed geodesic and L^2 is a 2-disk with nonnegative Gaussian curvature, by Gauss–Bonnet there is a closed interval $I = [0, a] \subset \mathbb{R}$ such that the flat strip $\gamma \times I$ is contained in L^2, with $\gamma = \gamma \times \{0\}$ and $\gamma' = \gamma \times \{a\}$. Thus G' acts trivially on I, which implies our claim.

In case (b) we have that $H \cap H' = \emptyset$ as in case (a) since at any point $p \in H$ we have two different BNL's at $\sigma^{-1}(p)$. Hence as before $H' = (\gamma' \times \mathbb{R}^{n-2})/G' \subset C$ and $H \times [0, a] \subset M^3$, with $H = H \times \{0\}$ and $H' = H \times \{a\}$. But this contradicts the fact that H is one-sided. □

Remark 6.4. Any manifold in case (b) admits a two-fold cover whose covering metric is as in case (a). Indeed, we can attach to C another copy of C along its interior boundary $\partial_i C$ using the involution that generates $\mathbb{Z}_2$. Switching the two cylinders induces the two-fold cover of M^n.

We proceed by showing that the phenomenon is essentially 3-dimensional. Observe that we only use here that $M^n \setminus W$ is connected, with no curvature assumptions. In fact, the same proof shows that if $M^n \setminus W$ has k connected components, then M^n splits off an $(n - k - 2)$-dimensional Euclidean factor.

Claim. If $n > 3$, *the universal cover of* M^n *splits off an* $(n - 3)$*-dimensional Euclidean factor.*

Proof. Assume first that M^n is the union of two cylinders C and C' with common boundary H. Consider the nullity distributions Γ and Γ' on the interior of C and C', which extend uniquely to parallel codimension one distributions F and F' on H, respectively. If $F = F'$, then $\Gamma \cup \Gamma'$ is a globally defined parallel distribution, which implies that the universal cover is an isometric product $N^2 \times \mathbb{R}^{n-2}$. Otherwise $J := F \cap F'$ is a codimension two parallel distribution on H. We claim that J extends to a parallel distribution on the interior of both C and C'.

To see this, we only need to argue for C, so lift the distributions J and F to the cover $S^1 \times \mathbb{R}^{n-2}$ of H under the projection $\pi \colon L^2 \times \mathbb{R}^{n-2} \to C = (L^2 \times \mathbb{R}^{n-2})/G$, and denote these lifts by $\hat{J}$ and $\hat{F}$. They are again parallel distributions whose leaves project to those of J and F under π. At a point $(x_0, v_0) \in S^1 \times \mathbb{R}^{n-2}$ a leaf of $\hat{F}$ is given by $\{x_0\} \times \mathbb{R}^{n-2}$ and hence a leaf of $\hat{J}$ by $\{x_0\} \times W$ for some affine hyperplane $W \subset \mathbb{R}^{n-2}$. Since $\hat{J}$ is parallel, any other leaf is given by $\{x\} \times W$ for $x \in S^1$. Thus W is invariant under G_2 in (2.6) since G permutes the leaves of $\hat{F}$. Therefore $\pi(\{x\} \times W)$ for $x \in L^2$ are the leaves of a parallel distribution on the interior of C, restricting to J on its boundary.

Therefore, we have a global flat parallel distribution J of codimension three on M^n, which implies that the universal cover splits isometrically as $N^3 \times \mathbb{R}^{n-3}$.

Now, if M^n consists of only one open cylinder C and its one-sided boundary, by Remark 6.4 there is a two-fold cover $\hat{M}^n$ of M^n which is the union of two cylinders as above and whose universal cover splits an $(n - 3)$-dimensional Euclidean factor. □

We can now finish the proof of Theorem C. Since M^n is compact with nonnegative curvature, the splitting theorem implies that the universal cover splits isometrically $\tilde{M}^n = N^k \times \mathbb{R}^{n-k}$ with N^k compact and simply connected. According to the above claim, $k = 2$

and hence $N^2 \simeq \mathbb{S}^2$, or $k = 3$ and by Theorem 1.2 in [Ha] we have $N^3 \simeq \mathbb{S}^3$. In the latter case, the metric on $\mathbb{S}^3$ is also CN2 and V admits a full extension, which, by Theorem B, implies that the metric is a geometric graph manifold metric.

7. Geometric graph 3-manifolds with nonnegative curvature

In this section we classify 3-dimensional geometric graph manifolds with nonnegative scalar curvature, giving an explicit construction of all of them. As a consequence, we show that, for each lens space, the number of connected components of the moduli space of such metrics is infinite, while for each prism manifold, the moduli space is connected.

We first recall some properties of lens spaces and prism manifolds. One way of defining a lens space is as the quotient $L(p, q) = \mathbb{S}^3/\mathbb{Z}_p$, where $\mathbb{Z}_p \subset S^1 \subset \mathbb{C}$ acts as $g \cdot (z, w) = (gz, g^q w)$ for $(z, w) \in \mathbb{S}^3 \subset \mathbb{R}^4 = \mathbb{C}^2$. We can assume that $p, q > 0$ with $\gcd(p, q) = 1$. It is a well known fact that two lens spaces $L(p, q)$ and $L(p, q')$ are diffeomorphic if and only if $q' = q^{\pm 1} \mod p$. An alternative description we will use is as the union of two solid tori $D_i \times S^1$, with boundary identified by a diffeomorphism f such that $\partial D_1 \times \{p_0\} \in \pi_1(\partial D_1 \times S^1)$ is taken under f into $(q, p) \in \mathbb{Z} \oplus \mathbb{Z} \simeq \pi_1(\partial D_2 \times S^1)$ with respect to its natural basis.

A prism manifold can also be described in two different ways, see e.g. [ST, HK, Ru]. One way is to define it as the quotient $\mathbb{S}^3/(H_1 \times H_2) = H_1 \backslash \mathbb{S}^3/H_2$, where $H_1 \subset \mathrm{Sp}(1)$ is a cyclic group acting as left translations on $\mathbb{S}^3 \simeq \mathrm{Sp}(1)$ and $H_2 \subset \mathrm{Sp}(1)$ a binary dihedral group acting as right translations. A more useful description for our purposes is as the union of a solid torus $D \times S^1$ with the 3-manifold

$$N^3 = (S^1 \times S^1 \times I)/\langle (j, -Id) \rangle, \quad \text{where} \quad j(z, w) = (-z, \bar{w}). \tag{7.1}$$

Notice that N^3 is a bundle over the Klein bottle $K = T^2/\langle j \rangle$ with fiber an interval $I = [-\epsilon, \epsilon]$ and orientable total space. Thus ∂N^3 is the torus $S^1 \times S^1$, and we glue the two boundaries via a diffeomorphism. Here $\pi_1(N^3) \simeq \pi_1(K) = \{a, b \mid bab^{-1} = a^{-1}\}$ and $\pi_1(\partial N^3) \simeq \mathbb{Z} \oplus \mathbb{Z}$, with generators a, b^2, where a represents the circle on the right and b^2 the circle on the left. Then $P(m, n)$ is defined as gluing ∂C to ∂N^3 by sending $\partial D \times \{p_0\}$ to $a^m b^{2n} \in \pi_1(\partial N^3)$. We can again assume that $m, n > 0$ with $\gcd(m, n) = 1$. Furthermore, $\pi_1(P(m, n)) = G_{m,n} = \{a, b \mid bab^{-1} = a^{-1},\ a^m b^{2n} = 1\}$. This group has order $4mn$ and its abelianization has order $4n$. Thus the fundamental group determines and is determined by the ordered pair (m, n). In addition, $G_{m,n}$ is abelian if and only if $m = 1$ in which case $P(m, n)$ is diffeomorphic to the lens space $L(4n, 2n - 1)$. Surprisingly, unlike in the case of lens spaces, the diffeomorphism type of $P(m, n)$ is uniquely determined by (m, n). This follows, e.g., by using the Ricci flow together with the fact that the effective representations of $G_{m,n}$ in $\mathrm{SO}(4)$ are unique up to conjugation, and hence given by the one presented in the introduction. Prism manifolds can also be characterized as the 3-dimensional spherical space forms which contain a Klein bottle, which for $m > 1$ is also incompressible. Finally, observe that N^3 is diffeomorphic to an irreducible flat twisted cylinder as in Corollary 2.4, and hence $P(m, n)$ can be viewed metrically as the union of two twisted cylinders, one nonflat and the second one flat. If we let $\epsilon \to 0$, $P(m,n)$ is thus also a single solid torus whose rectangular flat torus boundary has been identified to a Klein bottle as in Corollary 6.3 (b).

Let M^3 be a compact nonflat geometric graph manifold with nonnegative scalar curvature. We first observe M^3 is orientable. Indeed, by Theorem C, $M^3 = \mathbb{S}^3/\Pi$ and if an element $g \in \Pi$ acts orientation reversing, the Lefschetz fixed point theorem implies that g has a fixed point. Thus every cylinder $C = (D \times \mathbb{R})/G$ is orientable as well, i.e. G acts orientation preserving.

Following the notation in Section 2, the group G_1 preserves the closed geodesic ∂D and hence $G_1 \subset \mathrm{O}(2)$ acts isometrically on the disk D. If G contains an element $g = (g_1, g_2)$ where g_1 reverses orientation, then so does g_2 and hence g would have a fixed point. Thus $G_1 \subset \mathrm{SO}(2)$. Furthermore, for any $g \in G$, g_1 has a fixed point by the Brouwer fixed point theorem and hence g_2 is a nontrivial translation, which implies that $G \simeq \mathbb{Z}$. Thus the twisted cylinders are of the form $C = (D \times \mathbb{R})/\mathbb{Z}$ with $\mathbb{Z}$ generated by $g = (g_1, g_2) \in \mathrm{Iso}(D \times \mathbb{R})$. If g_1 is nontrivial, and $x_0 \in D$ is the fixed point of g_1, then the isometry g_1 is determined by its derivative at x_0. After orienting D, $d(g_1)_{x_0}$ is a rotation R_θ of angle $2\pi\theta$, $0 \le \theta < 1$. We simply say that g_1 acts as a rotation R_θ on D. Thus g acts via

$$g(x, s) = (R_\theta(x), s + t) \in \mathrm{Iso}(D \times \mathbb{R}), \tag{7.2}$$

for a certain $t > 0$ after orienting $\Gamma \cong T^\perp D$.

In particular, we have that the interior boundary of C is a flat 2-torus. Notice also that the action of $\mathbb{Z}$ can be changed differentiably until $\theta = 0$, and hence C is diffeomorphic to a solid torus $D \times S^1$. According to Corollary 6.3, M^3 is thus either the union of two solid tori glued along their boundary, and hence diffeomorphic to a lens space, or it is a solid torus whose boundary is identified via an involution to form a Klein bottle, and therefore diffeomorphic to a prism manifold.

Remark 7.3. Let us clarify the role of orientations in our description of C in (7.2). Take a twisted cylinder C with nonnegative scalar curvature, and D a maximal leaf of $\Gamma^\perp$. Orienting Γ is then equivalent to orienting $T^\perp D$, which in turn is equivalent to choosing one of the two generators of $\mathbb{Z}$. On the other hand, orienting D is equivalent to choosing between the oriented angle θ above or $1-\theta$. In particular, these orientations are unrelated to the metric on C, i.e., changing orientations give isometric cylinders.

Next, we show that the geometric graph manifold metric on M^3 is isotopic to a standard one. In order to do this, fix once and for all a metric $\langle\ ,\ \rangle_0$ on the disc $D_0 = \{x \in \mathbb{R}^2 : |x| \le 1\}$ which is rotationally symmetric, has positive Gaussian curvature on the interior of D_0, and whose boundary is a closed geodesic of length 1 along which the Gaussian curvature vanishes to infinite order. We call the metric on M^3 *standard*, if for each twisted cylinder $C = (D \times \mathbb{R})/\mathbb{Z}$ in the complement of a core of M^3, the metric on D is isometric to $r^2\langle\ ,\ \rangle_0$ for some $r > 0$. Notice that such a metric on M^3 is unique up to isometry.

Proposition 7.4. *A geometric graph manifold metric with nonnegative scalar curvature is isotopic, through geometric graph manifold metrics with nonnegative scalar curvature, to a standard metric.*

Proof. We can define the isotopy separately on each cylinder $C = (D \times \mathbb{R})/\mathbb{Z}$, as long as the isometry type of the core $H = \partial C$, and the foliation of H induced by the nullity leaves of C, stays fixed. To do this, we first deform the metric $\langle\ ,\ \rangle$ on D induced from the metric on M^3.

Let $\langle\ ,\ \rangle'$ be the standard flat metric on D_0. By the uniformization theorem we can write $\langle\ ,\ \rangle = f_1^*(e^{2v}\langle\ ,\ \rangle')$ for some diffeomorphism $f_1 : D \to D_0$ and a smooth function v on D_0. The metric $e^{2v}\langle\ ,\ \rangle'$ is thus invariant under $C_{f_1}(G_1) = \{f_1 \circ g \circ f_1^{-1} : g \in G_1\}$ which fixes $f_1(x_0)$, where $x_0 \in D$ is the fixed point of the action of G_1. Equivalently, $h \in C_{f_1}(G_1)$ is a conformal transformation of $(D_0, \langle\ ,\ \rangle')$ with conformal factor $e^{2v-2v\circ h}$. Recall that the conformal transformations of $\langle\ ,\ \rangle'$ on the interior of D_0 can be viewed as the isometry group of the hyperbolic disc model. Hence there exists a conformal transformation j of D_0 with $j(f_1(x_0)) = 0$ and conformal factor $e^{2\tau}$. We can thus also write $\langle\ ,\ \rangle = f^*(e^{2u}\langle\ ,\ \rangle')$, where $f = j \circ f_1 : D \to D_0$ and $u := (v - \tau) \circ j$. Now the metric $e^{2u}\langle\ ,\ \rangle'$ is invariant under $C_f(G_1)$, which this time fixes the origin of D_0. So $k \in C_f(G_1)$ is a conformal transformation of $\langle\ ,\ \rangle'$ fixing the origin, with conformal factor $e^{2u-2u\circ k}$. But an isometry of the hyperbolic disc model, fixing the origin, is also an isometry of $\langle\ ,\ \rangle'$. Hence $e^{2u} = e^{2u\circ k}$, i.e. u is invariant under k. Altogether, $C_f(G_1) \subset \mathrm{SO}(2) \subset \mathrm{Iso}(D_0, \langle\ ,\ \rangle')$ and u is $C_f(G_1)$-invariant. Analogously, $r^2\langle\ ,\ \rangle_0 = f_0^*(e^{2u_0}\langle\ ,\ \rangle')$ with $f_0 \in \mathrm{Diff}(D_0)$ satisfying $f_0(0) = 0$ and u_0 being $\mathrm{SO}(2)$-invariant. In particular, u_0 is also $C_f(G_1)$-invariant.

We now consider the two metrics $e^{2u}\langle\ ,\ \rangle'$ and $e^{2u_0}\langle\ ,\ \rangle'$ on D_0. They both have the property that the boundary is a closed geodesic along which the curvature vanishes to infinite order. Notice that the assumption that the boundary is a closed geodesic, up to parametrization, is equivalent to the condition that the normal derivatives of u and u_0, with respect to a unit normal vector in $\langle\ ,\ \rangle'$, is equal to 1. Furthermore, since $Ke^{2u} = -\Delta u$, the curvature vanishes to infinite order if and only if Δu does. For each $0 \le s \le 1$, consider the $C_f(G_1)$-invariant metric on D_0 given by $\langle\ ,\ \rangle^s = e^{2(1-s)u_0+2su+a(s)}\langle\ ,\ \rangle'$, where $a(s)$ is the function that makes the boundary to have length r for all s. Clearly, for each s, the boundary is again a closed geodesic up to parametrization and K^s vanishes at the boundary to infinite order. Furthermore, since $K^s e^{2(1-s)u_0+2su+a(s)} = -(1-s)\Delta u_0 - s\Delta u$ and $\Delta u_0 < 0,\ \Delta u \le 0$, the curvature of $\langle\ ,\ \rangle^s$ is nonnegative and positive on the interior of D_0.

Now, the metrics $f^*\langle\ ,\ \rangle^s + dt^2$ on $D \times \mathbb{R}$ are invariant under the action of $\mathbb{Z}$ and hence induce the desired one parameter family of metrics on C, since $f^*\langle\ ,\ \rangle^0 + dt^2$ is isometric to $r^2\langle\ ,\ \rangle_0 + dt^2$ via the diffeomorphism $(f_0 \circ f^{-1}) \times \mathrm{Id}$. We then glue these metrics to the core H preserving the arc length parametrization of ∂D. □

We now discuss how C induces a natural marking on $\partial_i C$. For this, let us first recall some elementary facts about lattices $\Lambda \subset \mathbb{R}^2$, where we assume that the orientation on $\mathbb{R}^2$ is fixed. A *marking* of the lattice Λ is a choice of an oriented basis $\{v, \hat{v}\}$ of Λ, and we call such a marking *normalized* if $\theta := \langle v, \hat{v}\rangle/\|v\|^2 \in [0,1)$. Notice that for any $v \in \Lambda$, there exists a unique oriented normalized marking $\{v, \hat{v}\}$. Indeed, if $\{v, w\}$ is some basis of Λ, then $\langle v, w + nv\rangle/\|v\|^2 = \langle v, \hat{v}\rangle/\|v\|^2 + n$ and hence there exists a unique $n \in \mathbb{Z}$ such that $\{v, \hat{v}\}$ with $\hat{v} = w + nv$ is normalized. If T^2 is an oriented torus with base point z_0, then $T^2 = \mathbb{R}^2/\Lambda$ for some lattice Λ with $z_0 = [0]$. Then a (normalized) marking of T^2 is a basis of $T_{z_0}T^2 \simeq \mathbb{R}^2$ which is a (normalized) marking of the lattice Λ.

Consider an oriented twisted cylinder $C = (D \times \mathbb{R})/\mathbb{Z}$ with its standard metric, where the action of $\mathbb{Z}$ is given by (7.2). The totally geodesic flat torus $T^2 = \partial_i C$, which inherits an orientation from C, has a natural marking based at $z_0 = [(p_0, t_0)]$. For this, denote by $\gamma : [0,1] \to \partial D$ the simple closed geodesic of length $r = \|\gamma'(0)\| > 0$ with $\gamma(0) = p_0$ which follows the orientation of $D = [D \times \{t_0\}] \subset C$. Let $v = \gamma'(0)$, $\hat{v} = \theta v + t\partial/\partial s$, and notice that the geodesic $\sigma(s) = \exp(s\hat{v}), 0 \le s \le 1$, is simple and closed with length $\|\hat{v}\|$. Since $\theta \in [0,1)$, the basis $\mathcal{B} := \{v, \hat{v}\}$ is a normalized marking of T^2 based at z_0, which we denote by $\mathcal{B}(\gamma)$. We also have a parallel oriented foliation of T^2 by the closed geodesics $[\gamma \times \{t\}] \subset T^2$, which we denote by $\mathcal{F}(C)$.

It is important for us that the above process can be reversed for standard metrics:

Proposition 7.5. *Let T^2 be a flat oriented torus and $\mathcal{F}$ an oriented foliation of T^2 by parallel closed simple geodesics. Then, there exists a oriented twisted cylinder $C_{\mathcal{F}} = (D \times \mathbb{R})/\mathbb{Z}$ over a standard oriented disk D, unique up to isometry, such that $\partial C_{\mathcal{F}} = T^2$ and $\mathcal{F}(C_{\mathcal{F}}) = \mathcal{F}$. Moreover, different orientations induce isometric metrics.*

Proof. Choose $\gamma \in \mathcal{F}$, and set $z_0 = \gamma(0)$, $v = \gamma'(0)$, and let $\mathcal{B}(\gamma) = \{v, \hat{v}\}$ be the normalized marking of T^2 based at z_0 defined as above. Set $r = \|v\|$, $\theta = \langle v, \hat{v}\rangle/\|v\|^2$ and $t = \|\hat{v} - \theta v\|$. With respect to the oriented orthonormal basis $e_1 = v/r$, $e_2 = (\hat{v} - \theta v)/t$ of $T_{z_0}T^2$ we have

$$T^2 = \mathbb{R}^2/\Lambda = (\mathbb{R} \oplus \mathbb{R})/(\mathbb{Z}v \oplus \mathbb{Z}\hat{v}) = (S^1_r \times \mathbb{R})/\mathbb{Z}\hat{v},$$

where S^1_r is the oriented circle of length r. Since $v = re_1$ and $\hat{v} = \theta v + te_2$, we can also write $T^2 = (S^1_r \times \mathbb{R})/\langle g\rangle$ where $g(p,s) = (R_\theta(p), s + t)$. Now we simply attach $(D_0, r^2\langle\ ,\ \rangle_0)$ to S^1_r preserving orientations to build $C = (D_0 \times \mathbb{R})/\langle g\rangle$. Notice that any two base points of T^2 are taken to each other by an orientation preserving isometry of C, restricted to $\partial C = T^2$. Thus the construction is independent of the choice of z_0 and the choice of $\gamma \in \mathcal{F}$. By Remark 7.3, different choices of orientation induce the same metric on C, and hence $C_{\mathcal{F}}$ is unique up to isometry. □

Remark 7.6. If we do not assume that the metric on C is standard, then the construction of $C_{\mathcal{F}}$ depends on the choice of base point, and one has to assume that the metric on D is invariant under R_θ, where θ is the angle determined by the marking of T^2 induced by $\mathcal{F}$.

We can now easily classify standard geometric graph manifold metrics with two-sided core, proving case (a) of Theorem D.

THEOREM 7.7. *Let M^3 be a compact geometric graph manifold of nonnegative scalar curvature with irreducible universal cover, and assume that its core T^2 is two-sided. Then, $M^3 = C_1 \sqcup T^2 \sqcup C_2$, where $C_i = (D_i \times \mathbb{R})/\mathbb{Z}$ are twisted cylinders over 2-disks that induce two different foliations $\mathcal{F}_i = \mathcal{F}(C_i)$ of T^2 by parallel closed geodesics, $i = 1, 2$.*

Conversely, given a flat 2-torus T^2 with two different foliations $\mathcal{F}_i$ by parallel closed geodesics, there exists a standard geometric graph manifold $M^3 = C_1 \sqcup T^2 \sqcup C_2$ with irreducible universal cover whose core is T^2 and $C_i = C_{\mathcal{F}_i}$. Moreover, this data determines the standard metric up to isometries, i.e., if $h : T^2 \to \hat{T}^2$ is an isometry between flat tori, then $\hat{M}^3 = \hat{C}_1 \sqcup \hat{T}^2 \sqcup \hat{C}_2$ is isometric to M^3, where $\hat{C}_i = C_{h(\mathcal{F}_i)}$.

Proof. We only need to argue for the uniqueness. The core of a standard metric is unique since the set of nonflat points of a standard metric is dense, cf. Corollary 6.3. It is clear then that an isometry between standard geometric graph manifolds will send the core to the core, and the parallel foliations to the parallel foliations. Hence the core and the parallel foliations are determined by the isometry class of M^3.

Conversely, by uniqueness in Proposition 7.5 the standard twisted cylinders $C_{\mathcal{F}_i}$ and $C_{h(\mathcal{F}_i)}$ are isometric, which in turn induces an isometry between M^3 and $\hat{M}^3$. The only ambiguity is on which side of the torus to attach each of the twisted cylinders. But this simply gives an orientation reversing isometry fixing the core. $\square$

Now, let us consider the one-sided core case. Here we know that $M^3 = C \sqcup K$ and that K is a nonorientable quotient of the flat torus $\partial_i C$ and hence a flat Klein bottle. It is easy to see that, if a flat torus admits an orientation reversing fixed point free isometric involution j, then T^2 has to be isometric to a rectangular torus $S^1_r \times S^1_s$ along which j is as in (7.1). The irreducibility of the universal cover of M^3 is thus equivalent to $\mathcal{F}(C)$ not to coincide with one of the two invariant parallel foliations of j, $\mathcal{F}(j) = \{S^1_r \times \{w\} : w \in S^1_s\}$ and $\{\{z\} \times S^1_s : z \in S^1_r\}$.

As in the proof of Theorem 7.7, we conclude:

THEOREM 7.8. *Let M^3 be a compact geometric graph manifold of nonnegative scalar curvature with irreducible universal cover, and assume that its core K is one-sided. Then $M^3 = C \sqcup K$, where $C = (D \times \mathbb{R})/\mathbb{Z}$ is a twisted cylinder over a 2-disk with $\partial_i C = T^2$ isometric to a rectangular torus, and $\partial C = K = T^2/\mathbb{Z}_2$ a flat totally geodesic Klein bottle.*

Conversely, a rectangular flat torus $T^2 = S^1_r \times S^1_s$ and a foliation $\mathcal{F}$ of T^2 by parallel closed geodesics different from $S^1_r \times \{p\}$ or $\{p\} \times S^1_s$ define a standard geometric graph manifold with irreducible universal cover $M^3 = C_{\mathcal{F}} \sqcup K$ whose core K is one-sided. Moreover, T^2 and $\mathcal{F}$ determine M^3 up to isometry.

In order to determine the topological type of these geometric graph manifolds we introduce the concept of relative slope between two foliations of a flat torus. To define it, we first assume that the data is oriented. Let $\mathcal{F}_1$ and $\mathcal{F}_2$ be two different oriented parallel foliations by closed geodesics on the oriented flat torus T^2, and choose $\gamma_i \in \mathcal{F}_i$ such that $\gamma_1(0) = \gamma_2(0) =: z_0$. If $v_i = \gamma_i'(0)$ we have the normalized markings $\mathcal{B}(\gamma_i) = \{v_i, \hat{v}_i\}$ of T^2

based at z_0 defined as above. Since $\mathrm{SL}(2,\mathbb{Z})$ acts transitively on the set of oriented basis of a given lattice, there exist coprime integers $p, q,\ a, b \in \mathbb{Z}$ with $bq - ap = 1$ such that

$$v_2 = qv_1 + p\hat{v}_1, \quad \hat{v}_2 = av_1 + b\hat{v}_1.$$

We also have $p \neq 0$ since $v_1 \neq \pm v_2$. Notice that, since v_2 determines $\hat{v}_2$, the integers p and q determine a and b.

We call $q/p \in \mathbb{Q}$ the (algebraic) *slope* of γ_2 with respect to γ_1. Notice that the slope is independent of the choice of $\gamma_i \in \mathcal{F}_i$ since the foliations are parallel. The slope does not depend on the orientations of the geodesics either, since $\{-v, -\hat{v}\}$ is the oriented marking associated to $-\gamma$. Since $v_1 = bv_2 - p\hat{v}_2$, the slope of γ_1 with respect to γ_2 is $-b/p$. Observe though that the slopes depend on the orientation of the torus, since changing the orientation of the torus corresponds to changing $\hat{v}_i$ with $-\hat{v}_i$ which gives slopes $-q/p$ and b/p, respectively.

We denote by $[q/p] = \{\pm q/p, \mp b/p\}$ the *relative slope between the foliations*. Accordingly, if $M^3 = C_1 \sqcup T^2 \sqcup C_2$ has two-sided core, by Theorem 7.7 we have that the relative slope between $\mathcal{F}(C_1)$ and $\mathcal{F}(C_2)$ is an isometric invariant of M^3, that we call the *slope of* M^3. Analogously, if $M^3 = C \sqcup K$ has one-sided core $K = \partial_i C/\langle j\rangle$, we call the relative slope $[m/n] = \pm m/n$ of $\mathcal{F}(C)$ with respect to $\mathcal{F}(j)$ the *slope of* M^3. Notice that the slope of M^3 is well defined even when the geometric graph manifold metric is not standard.

We show next that slope of M^3 is precisely what determines its diffeomorphism type.

THEOREM 7.9. *Let M^3 be a compact geometric graph manifold of nonnegative scalar curvature with irreducible universal cover and slope $[m/n]$. Then, if the core of M^3 is two-sided M^3 is diffeomorphic to the lens space $L(n,m)$, while if the core of M^3 is one-sided M^3 is diffeomorphic to the prism manifold $P(m,n)$.*

Proof. Recall that the twisted cylinders C_i are diffeomorphic to $D_i \times S^1$ by deforming α_i continuously to 0. For two-sided core T^2, orient M^3 and T^2, choose $\gamma_i \in \mathcal{F}_i$, and let $\mathcal{B}(\gamma_i) = \{v_i, \hat{v}_i\}$ be the normalized markings of T^2 defined by C_i. Then the natural generators of $\pi_1(\partial(D_i \times S^1)) \simeq \mathbb{Z} \oplus \mathbb{Z}$ are represented by the closed geodesics γ_i and $\sigma_i(t) = \exp(t\hat{v}_i)$, $0 \le t \le 1$ since the marking $\{v_i, \hat{v}_i\}$ is normalized. According to the definition of slope, $v_2 = qv_1 + p\hat{v}_1$ which implies that under the diffeomorphism from $\partial D_2 \times S^1 \simeq \partial C_2$ to $\partial C_1 \simeq \partial D_1 \times S^1$, the element $(1,0) \in \pi_1(\partial(D_2 \times S^1))$ is taken to $(q,p) \in \pi_1(\partial(D_1 \times S^1))$. By definition this is the lens space $L(p,q)$.

To determine the topological type in the one-sided case, we view M^3 as the union of C with the flat twisted cylinder N^3 defined in (7.1). Then $\partial N^3 = T^2$ is a rectangular torus which we glue to $\partial_i C$. Taking $\epsilon \to 0$ (or considering $T^2 \times (0,\epsilon]$ as part of C instead), we obtain M^3. We can now use our second description of prism manifolds and the proof finishes as in the previous case. $\square$

Remark. Notice that if if the slope of a lens space is $[q/p] = \{\pm q/p, \mp b/p\}$ and hence $bq - ap = 1$, then $b = q^{-1} \bmod p$ which is consistent with the fact that $L(p,q)$ and $L(p,b)$ are diffeomorphic.

We finally classify the moduli space of metrics.

PROPOSITION 7.10. *On a lens space $L(p,q)$ the connected components of the moduli space of geometric graph manifold metrics with nonnegative scalar curvature are parametrized, up to sign, by its slope $[q/p]$, and therefore it has infinitely many components. On the other hand, on a prism manifold $P(m,n)$ with $m > 1$ the moduli space is connected.*

Proof. In Proposition 7.4 we saw that we can deform any geometric graph manifold metric into one which is standard. According to Theorem 7.7, the metric on a lens space can equivalently be defined by the triple $(T^2, \mathcal{F}_1, \mathcal{F}_2)$. We can now deform the flat metric on the torus, carrying along the foliations $\mathcal{F}_i$, which hence induces a deformation of the metric by standard metrics. In the proof of Proposition 7.5 we saw that, after choosing orientations, for $\gamma_i \in \mathcal{F}_i$ with $v_i = \gamma_i'(0)$ we have the normalized marking $\mathcal{B}(\gamma_i) = \{v_i, \hat{v}_i\}$ which represents a fundamental domain of the lattice defined by T^2. We can thus deform the flat torus to a unit square torus such that the first marking is given by $v_1 = (1,0),\ \hat{v}_1 = (0,1)$. Then $v_2 = (q,p) = qv_1 + p\hat{v}_1$, which in turn determines $\hat{v}_2$, and q/p is the slope of γ_2 with respect to γ_1. After changing orientations we can furthermore assume that $p, q > 0$. If we choose the second marking to make it standard, we instead obtain the second representative of the slope. Metrics with different slope can clearly not be deformed into each other since the invariant is a rational number.

For a prism manifold, we similarly deform the metric to be standard and the rectangular torus into a unit square. But then the slope $[m/n]$ already uniquely determines its diffeomorphism type. □

Remarks. *a*) Notice that the lens space $L(4n, 2n-1)$ has two types of CN2 metrics, one being the union of two nonflat twisted cylinders and the other being one twisted cylinder whose boundary is identified to a Klein bottle, or equivalently the union of a nonflat and a flat twisted cylinder. These clearly lie in different components of geometric graph manifold metrics.

b) One easily sees that the angle α between the nullity foliations of a lens space, i.e., the angle between v_1 and v_2, is given by

$$\cos(\alpha) = \frac{r_1}{r_2}(q + p\theta_1) = \frac{r_2}{r_1}(b - p\theta_2),$$

where $r_i = |v_i|$ and θ_i are the twists of the two cylinders. One can thus make the nullity leaves orthogonal if and only if $0 \le -q/p < 1$ and in that case $r_2 = pt_1,\ t_2 = t_1/p$ and $\theta_1 = -q/p,\ \theta_2 = b/p$. This determines the metric on the lens space described in the introduction as a quotient of Figure 2, and is thus the only component containing a metric with orthogonal nullity leaves.

c) We can explicitly describe the geometric graph manifold metrics on $\mathbb{S}^3 = L(1,0)$ up to deformation. We assume that the core is a unit square and that the first foliation is parallel to $(1,0)$, i.e. the first cylinder is a product cylinder. Then the second marking is given by $v_2 = (q,1),\ \hat{v}_2 = (q-1,1)$. By choosing the orientations appropriately, we can assume $q \geq 0$. According to the proof of Proposition 7.4, a marking $\{v, \hat{v}\}$ corresponds to a twisted cylinder as in (7.2) with $r = \|v\|$, $\theta = \langle v, \hat{v}\rangle / \|v\|^2$ and $t = \|\hat{v} - \theta v\|$. Thus in our case the second cylinder is given by

$$r = \sqrt{1+q^2}, \quad \theta = \frac{1+q^2-q}{1+q^2}, \quad t = \frac{1}{\sqrt{1+q^2}}.$$

The relative slope is $[q] = \{\pm q, \mp 1\}$, and the standard metric in Figure 2 corresponds to $q = 0$.

References

[AM] A. Aazami and C. Melby-Thompson, *On the principal Ricci curvatures of a Riemannian 3-manifold*. Preprint 2005, arxive:1508.02667v2.

[BS] R. Bettiol and B. Schmidt, *Three-manifolds with many flat planes*, Preprint 2014, arXiv:1407.4165.

[BKV] E. Boeckx, O. Kowalski and L. Vanhecke, *Riemannian manifolds of conullity two*. World Scientific, 1996.

[Ca] E. Cartan, *La déformation des hypersurfaces dans l'espace euclidien réel a n dimensions*. Bull. Soc. Math. France **44** (1916), 65–99.

[CG] J. Cheeger and M. Gomov, *Collapsing Riemannian manifolds while keeping their curvature bound I*, J. Dif. Geom. **23** (1986), 309-364.

[DF] M. Dajczer and L. Florit, *Genuine rigidity of Euclidean submanifolds in codimension two*. Geom. Dedicata **106** (2004), 195–210.

[DFT] M. Dajczer, L. Florit and R. Tojeiro, *On deformable hypersurfaces in space forms*. Ann. Mat. Pura ed Appl. **147** (1998), 361–390.

[FF] L. Florit and G. Freitas, *Classification of codimension two deformations of rank two Riemannian manifolds*. To appear in Comm. Anal. Geom.

[FZ] L. Florit and W. Ziller, *Nonnegatively curved Euclidean submanifolds in codimension two*. to appear in Comm. Math. Helv.

[Gr] M. Gromov, *Manifolds of negative curvature*. J. Differ. Geom. **13** (1978), 223–230.

[Ha] R. Hamilton, *Four-manifolds with positive curvature operator*. J. Diff. Geom. **24** (1986), 153–179.

[HK] S. Hong, J. Kalliongis, D. McCullough and J. Rubinstein, *Diffeomorphisms of Elliptic 3-Manifolds*. Lecture Notes in Mathematics 2055, Springer–Verlag, Berlin Heidelberg 2012.

[Ma] R. Maltz, *The nullity spaces of curvature-like tensors*. J. Diff. Geom. **7** (1972), 519–523.

[No] K. Nomizu, *On hypersurfaces satisfying a certain condition on the curvature tensor*. Tohoku Math. J. **20** (1968), 46–59.

[PR] R. Ponge and H. Reckziegel, *Twisted products in pseudo-Riemannian geometry*. Geom. Dedicata **48** (1993), no. 1, 15–25.

[Ru] J Rubinstein, *On 3-manifolds that have finite fundamental group and contain klein bottles*. Trans. AMS, **251** 129–137.

[SW] B. Schmidt and J. Wolfson, *Three manifolds with constant vector curvature*. Indiana Univ. Math. J. **63** (2014), 1757–1783.

[Sb] V. Sbrana, *Sulla varietá ad $n-1$ dimensioni deformabili nello spazio euclideo ad n dimensioni*. Rend. Circ. Mat. Palermo **27** (1909), 1–45.

[Sch] V. Schroeder, *Rigidity of Nonpositively Curved Graphmanifolds*. Math. Ann. **274** (1986), 19–26.

[ST] W. Threlfall and H. Seifert, *Topologische Untersuchung der Diskontinuitätsbereiche endlicher Bewegungsgruppen des dreidimensionalen sphärischen Raumes*. Math. Ann. **104** (1931), 1–70.

[Se] K. Sekigawa, *On the Riemannian manifolds of the form $B \times_f F$*. Kodai Math. Sem. Rep. **26** (1974), 343–347.

[Si] I. Singer, *Infinitesimally homogeneous spaces.* Comm. Pure Appl. Math. **13** (1960) 685–697.
[Sz] Z.I. Szabó, *Structure theorems on Riemannian spaces satisfying* $R(X,Y)\cdot R = 0$*, II, Global versions.* Geom. Dedicata **19** (1985), 65–108.
[Wa] F. Waldhausen, *Eine Klasse von 3-dimensionalen Mannigfaltigkeiten II.* Invent. Math. **4** (1967), 87–117.
[Ya] S. T. Yau, *Some Function Theoretic Properties of Complete Riemannian Manifold and Their Applications to Geometry.* Indiana University Mathematics Journal **25** (1976), 659–670.

IMPA: Est. Dona Castorina 110, 22460-320, Rio de Janeiro, Brazil
E-mail address: luis@impa.br

University of Pennsylvania: Philadelphia, PA 19104, USA
E-mail address: wziller@math.upenn.edu