

Hopf Arborescent Links, Minor Theory, and Decidability of the Genus Defect

Pierre Dehornoy*

Corentin Lunel[‡]

Arnaud de Mesmay[§]

Abstract

While the problem of computing the genus of a knot is now fairly well understood, no algorithm is known for its four-dimensional variants, both in the smooth and in the topological locally flat category. In this article, we investigate a class of knots and links called Hopf arborescent links, which are obtained as the boundaries of some iterated plumbings of Hopf bands. We show that for such links, computing the genus defects, which measure how much the four-dimensional genera differ from the classical genus, is decidable. Our proof is non-constructive, and is obtained by proving that Seifert surfaces of Hopf arborescent links under a relation of minors defined by containment of their Seifert surfaces form a well-quasi-order.

1 Introduction

A (tame) [knot](#) is a polygonal embedding of the circle S^1 into \mathbb{R}^3 , or equivalently, S^3 , and a [link](#) is a disjoint union of knots. Knot theory is both an old and very active mathematical field, yet from an algorithmic perspective, many problems arising naturally in knot theory are still shrouded in mystery. This can be illustrated with arguably the most fundamental algorithmic question in knot theory: in the KNOT EQUIVALENCE problem, we are given two knots K_1 and K_2 and are tasked with deciding whether they are equivalent, that is, whether one can deform one into the other continuously without creating self-intersections. The best algorithm for this problem, due to Kuperberg, is elementary recursive [25], yet the problem is not even known to be **NP**-hard (see for example [27, Conclusion]). We refer to Lackenby [28] for a survey on algorithmic problems in knot theory.

Given how seemingly hard testing the equivalence of knots is, a huge body of research has been devoted to designing and studying knot invariants in order to tell them apart. A classical invariant of a knot is its [genus](#): this is the smallest possible genus of an embedded oriented surface, called a [Seifert surface](#), having the knot as its boundary. Computing the genus of a knot turns out to be significantly more tractable: celebrated works of Hass, Lagarias and Pippenger [21] and Agol, Hass and Thurston [2], building on the normal surface theory of Haken [20], have shown that deciding if a knot has genus at most g is in **NP**, while Lackenby has proved that it is also in **co-NP** [29]. These algorithms run also well in practice within the software Regina [10].

There are, however, different notions of genus that are much less understood: considering S^3 as the boundary of the 4-dimensional ball B^4 , the 4-genus of a knot $g_4(K)$ is roughly the smallest possible genus of a surface in B^4 having the knot as its boundary. This comes in two flavours that are known to not be equivalent: the [topological locally flat 4-genus](#) and the [smooth 4-genus](#) depending on the regularity of the surface. We refer to the preliminaries for precise definitions. A knot is (topologically or smoothly) [slice](#) if it bounds a disc in B^4 , i.e., has 4-genus zero. One of the motivations for the study of such 4-dimensional invariants comes from algebraic geometry, as such surfaces arise naturally around singularities of algebraic curves in \mathbb{C}^2 [23, 42]. Another motivation is the slice-ribbon conjecture [14] which states that a knot is smoothly slice if and only if it is ribbon, i.e., it bounds an immersed disc with only ribbon-type singularities in S^3 . Unfortunately, no algorithmic framework at all is known to attack topological problems in 4-dimensional topology, and indeed many of these problems are known to be undecidable, e.g., the homeomorphism of 4-manifolds [31]. For some other problems, the decidability is a well-known open problem: this is the case for 4-sphere recognition [48] or embeddability of 2-dimensional complexes in \mathbb{R}^4 [32]. Similarly, no algorithm is known to decide the 4-genus of a knot or even to decide whether it is slice. To illustrate how hard these problems are, it is only in a recent breakthrough of Picirillo [40] that it was proved that the Conway knot is not smoothly slice, although it only has 11 crossings. From the perspective of lower bounds, recent work of de Mesmay, Rieck, Sedgwick and Tancer [12] has proved that an analogue of the 4-genus for links, the 4-ball Euler characteristic, is **NP**-hard to compute, but it is also not known to be decidable.

*Aix-Marseille Université, CNRS, I2M, Marseille, France, pierre.dehornoy@univ-amu.fr

[‡]LIGM, CNRS, Univ. Gustave Eiffel, ESIEE Paris, F-77454 Marne-la-Vallée, France, corentin.lunel2@univ-eiffel.fr

[§]LIGM, CNRS, Univ. Gustave Eiffel, ESIEE Paris, F-77454 Marne-la-Vallée, France, arnaud.de-mesmay@univ-eiffel.fr




Figure 1: Top: A positive Hopf band and a plumbing. Bottom: An Hopf arborescent link and an associated planar tree.

Our results. The goal of this paper is to investigate the structure of a particular class of links, which we call **Hopf arborescent links**, in order to prove the decidability of some of their 4-dimensional invariants. This family of links is informally defined as follows (we refer to Section 2 for more precise definitions). A **Hopf band** is the surface pictured in Figure 1 (top left), it can be either **positive** or **negative** depending on how it twists. If they are unlinked, i.e., there exists a sphere separating them, two Hopf bands can be **plumbed** together by identifying a square in one to a square in the other, as pictured in Figure 1 (top). The class of **Hopf arborescent links** is the class of links arising as the boundary of some iterated tree-like sequence of plumbings of Hopf bands. Hopf arborescent links can naturally be described using labelled trees, see Figure 1 (bottom), and are a subfamily of the more general arborescent links [8, 17]. The (topological or smooth) **genus defect** of a knot is defined as $\Delta g(K) = g(K) - g_4(K)$. It measures how much its 4-genus differs from its classical genus, and this definition can be extended to oriented links by considering surfaces having that oriented link as their boundary. Our main result is the following:

Theorem 1.1. *For any fixed k , deciding whether an Hopf arborescent link L has genus defect at most k is decidable. This holds both in the topological and smooth categories.*

The proof of Theorem 1.1 is not constructive. It is obtained as a corollary of another result, which establishes a well-behaved minor theory for Hopf arborescent links. A subsurface $\Sigma' \subseteq \Sigma$ of a Seifert surface Σ is **incompressible** if the complement $\Sigma \setminus \Sigma'$ has no open disc component. Given two Seifert surfaces Σ_1 and Σ_2 in \mathbb{R}^3 , we say that Σ_1 is a **surface-minor** of Σ_2 (or **minor** for short), denoted by $\Sigma_1 \preceq \Sigma_2$, if Σ_1 is isotopic to an incompressible subsurface of Σ_2 . This minor relation was introduced by Baader [3] (see also [4, 5]) with the goal of characterising those links that are closure of positive braids and whose signature is equal to twice their genus. The underlying question, still open today, is whether the canonical Seifert surfaces associated to positive braid closures form a well-quasi-order. A particularity of Hopf arborescent links is that they are **fibred** (see the definition in Section 2), which implies that to each link is associated a canonical Seifert surface (see Theorem 2.2), that we call a **Hopf arborescent surface**.

The notion of surface-minor naturally implies a minor relation for Hopf arborescent links. Our second result proves that Hopf arborescent surfaces are well-quasi-ordered under surface-minors.

Theorem 1.2. *The minor relation \preceq is a well-quasi-order for the set of Hopf arborescent surfaces, that is, for any infinite sequence $(\Sigma_n)_{n \in \mathbb{N}}$ of Hopf arborescent surfaces, there exists $i < j$ in \mathbb{N} such that $\Sigma_i \preceq \Sigma_j$.*

The idea behind the proof of Theorem 1.2 is to study a specific subset of the possible surface-minors that interacts nicely with an encoding of Hopf arborescent surfaces via labelled plane trees. We can then leverage the celebrated Kruskal tree theorem [24] to prove that the minor relation is a well-quasi-order. The connection from Theorem 1.2 to Theorem 1.1 follows from the fact the genus defect is minor-monotone, i.e., if $\Sigma_1 \preceq \Sigma_2$ are Seifert surfaces of minimal genus for links K_1 and K_2 , then $\Delta g(K_1) \leq \Delta g(K_2)$. This is not a new observation (see [6, Lemma 6]), we provide a proof in Proposition 4.1 for completeness. Therefore, for Hopf arborescent links, having genus defect at most k is characterized by a finite number of forbidden minors, and the algorithm of Theorem 1.1 proceeds by checking those. However, testing whether a surface is a minor of another one seems to be a very hard problem: even testing whether two tori are isotopic is already as hard a knot equivalence and an algorithm for genus two surfaces was only very recently found [7]. This problem is circumvented thanks to our restriction of the minor relation to one that is well-tailored to the arborescent structure of our links, which allows us to work entirely at the level of trees. In particular, Theorem 1.1 does not strictly follow from Theorem 1.2 but rather from its proof (see Proposition 3.7).

While the algorithms behind Theorem 1.1 are not explicit, we would like to offer three reasons to motivate our results. First, Theorem 1.1 proves that the corresponding problems are *not* undecidable, which is significant in the landscape of 4-dimensional topology. Second, this kind of existential algorithmic result has been a strong guiding light in algorithm design in the past decades: for a vast family of graph problems, the fact that an algorithm merely exists follows from Robertson-Seymour theory, and this has provided a strong impetus to actually look for explicit algorithms and optimise their complexity. This has been particularly influential in parameterized algorithms, we refer for example to the discussion in Chapter 6.3 in the book on parameterized algorithms [11], where it is conjectured that a result like our Theorem 1.1 precludes W[1]-hardness. Similarly, we are hopeful that our results can inspire future work aiming at developing explicit algorithms in 4-dimensional topology. Additionally, our framework directly proves that any property that is stable with respect to our link-minor relation (see Section 3.4) is decidable on the class of Hopf arborescent links. Finally, while it is certainly not the case that minor-based approaches can encompass the entirety of knot theory, it is fruitful to delineate exactly the classes which they can illuminate. In that respect, we find it interesting that our proof of Theorem 1.2 strongly relies on the structure of Hopf bands and does not seem to generalize to the wider family of arborescent knots, even when one bounds the number of twists in each band (see Remark 3.8).

Related work. It was observed by Baader and Dehornoy [4] that the natural Seifert surfaces for another class of knots, the positive braid knots with bounded braid index (we refer to the papers for the relevant definitions) also form a well-quasi-order. Furthermore, Liechti [30] proved that even without bounding the braid index, the set of positive braid knots of bounded genus defect is characterized by a finite number of forbidden minors. Since the minor relation in that setting simply amounts to removing letters in the braid presentation, this readily yields decidability as in our Theorem 1.1. While the two results are incomparable, we emphasize that our result also applies to links, and also features negative crossings (coming from negative Hopf bands): this extends the impact of our result to the smooth category, while for strongly quasipositive knots (and thus positive braid knots), the smooth defect is zero since the smooth 4-genus and the classical genus coincide [42].

All the knots and links we consider, as well as those considered by various authors in the context of surface-minor theory, are fibred (see again the definition in Section 2). This property is important as it brings control on the classical genus of the links. Also it is easy to construct infinite families of incomparable surfaces when dropping this assumption: the set $(A_n)_{n \in \mathbb{Z}}$ of those unknotted annuli in \mathbb{S}^3 with n twists forms an infinite antichain. In this direction, an optimistic conjecture would be that the collection of all fibred surfaces in \mathbb{S}^3 is a well-quasi-ordered set. If true, that would provide a strong generalisation of Theorem 1.2. However no strategy of proof is known to the authors for such a statement.

Also, it follows from a result of Giroux and Goodman [18] that *any* fibred link can be obtained from the unknot from a sequence of plumbings and deplumbings (a natural reverse operation to plumbing) of Hopf bands. While these (de)plumbings might not have the arborescent structure that characterizes ours, this shows that Hopf bands can be considered as basic building blocks for a wide class of three-dimensional objects.

Organization of this paper. After providing background and going through basic concepts of knot theory, fibred surfaces and well-quasi-orders in Section 2, we focus on Hopf arborescent links in Section 3. There, we define a precise construction of these links from plane trees, investigate this class of links, and prove Theorem 1.2. In Section 4, we prove Theorem 1.1. Finally, in Section 5, we first provide examples of Hopf arborescent links with non zero defect, and then explain how to combine them to obtain examples with arbitrarily large defect.

2 Preliminaries

Knot theory. We only recall the definitions that are critical to this paper, and refer to the textbooks of Burde and Zieschang [9] or Rolfsen [41] for a more general introduction to knot theory, and to Teichner [46] for a beginner-friendly introduction to its 4-dimensional aspects. A **knot**, respectively a **link**, is a polygonal embedding of the circle S^1 , respectively of a disjoint union of circles, into S^3 . Every knot and every component of a link inherits an orientation from the orientation of S^1 . For algorithmic purposes, we assume that an input link is given as a link diagram, i.e., a directed plane 4-valent graph with decorations at vertices indicating which strands are going over and under. Since our article focuses on decidability problems, switching to a different input (e.g., polygonal curves in R^3) makes no difference. A **Seifert surface** for a link L is a compact connected oriented surface Σ embedded in S^3 such that the oriented boundary of Σ is L . Throughout this article, we consider knots, links, and surfaces embedded in S^3 up to isotopies (continuous deformations without self-crossings). The **(classical) genus** of a link L , denoted by $g(L)$, is the smallest possible genus of a Seifert surface for L .

The 3-dimensional sphere can be seen as the boundary of the 4-dimensional ball B^4 . Being embedded in S^3 , a knot or a link can also be obtained as the boundary of surfaces embedded in B^4 . However, any knot K in S^3 can be used as a base that tapers to a point, the apex, inside B^4 to define a cone that bounds K . Hence, any knot in S^3 bounds a topological disc in B^4 . This motivates the following definition: a surface Σ embedded in B^4 is **locally flat** if for each point $x \in \Sigma$, there is a neighbourhood U in Σ and a neighbourhood V in B^4 such that the pair (U, V) is homeomorphic to the standard (B^2, B^4) . In the coning construction above, the latter condition is not satisfied at the apex, where the boundary of a disc is the knot instead of a standard S^1 . The **topological** (respectively **smooth**) **4-dimensional genus**, or simply **4-genus** of a link L , denoted by¹ $g_4(L)$ is the smallest possible genus of a compact connected oriented surface that is locally flat (respectively smoothly) embedded in B^4 , and that has L as its boundary.

A knot is **topologically** (resp. **smoothly**) **slice** if it bounds a locally flat (resp. a smooth) disc in B^4 . The **(topological or smooth) defect** of a link L is the quantity $\Delta g(L) = g(L) - g_4(L)$, where $g_4(L)$ denotes the topological or smooth 4-genus of L .

A **positive**, resp. **negative**, **Hopf band** is an unknotted annulus with a positive, resp. negative, full twist, as pictured in Figure 2. A **Hopf link** is the boundary of a Hopf band. Note that the two components of a positive Hopf link have linking number $+1$, while the two components of a negative Hopf link have linking number -1 . A Hopf band naturally retracts to a trivial knot, which we call its **core**.




Figure 2: A negative Hopf band on the left and a positive one on the right with its core in red.

A link is **fibred** if the complement $S^3 \setminus L$ **fibres** over S^1 , that is, if there exists a one-parameter continuous family of Seifert surfaces $(\Sigma_t)_{t \in S^1}$ for L which are disjoint except for their boundaries, and whose interiors together foliate $S^3 \setminus L$. These are called **fibres** or **fibre surfaces** (note that they are all isotopic by definition).

¹Our notation is intentionally ambiguous with respect to smooth or topological genus, because all our arguments will apply equally well in both categories.

A positive, resp. negative, Hopf link is fibred, with fibre the positive, resp. negative, Hopf band. Indeed, seeing \mathbb{S}^3 as the unit sphere $\{(z_1, z_2) \mid |z_1|^2 + |z_2|^2 = 1\}$ in \mathbb{C}^2 , a positive Hopf link is given by the equation $z_1 z_2 = 0$. For every argument $\theta \in \mathbb{S}^1$, the equation $\arg(z_1 z_2) = \theta$ describes a Seifert surface bounded by the Hopf link, and the collection of these surfaces describes the desired fibration.

It is a folklore result that goes back at least to Stallings [45] that a fibre surface of a fibred link is a Seifert surface of minimal (classical) genus, and moreover this surface is unique up to isotopy (see next paragraph). So, for fibred links (and all links in this paper will be fibred), it makes sense to speak of the *canonical* Seifert surface, by which we mean the unique (up to isotopy) fibre surface for that link.

Fibre surface, minimal genus, and uniqueness. Here, we comment on the following statements, which are considered as folklore in the knot theoretical community. The first one is fundamental for this paper since it brings control over the 3-genus of Hopf arborescent links. The second one justifies the notion of *canonical* Seifert surface for fibred links.

Theorem 2.1. *Fibred surfaces minimise the 3-genus over surfaces with the same oriented boundary.*

Theorem 2.2. *Two fibred surfaces with the same oriented boundary are isotopic, relatively to the boundary.*

Theorem 2.1 is often presented as a consequence of Stallings' Theorem [45]. Here we present a semi-elementary proof (for which we claim no novelty), in the sense that it only uses basic notions and results from homology theory.

Proof of Theorem 2.1. Consider $L = (L_1 \sqcup \dots \sqcup L_k)$ an oriented link in \mathbb{S}^3 . We consider the 3-manifold N_L which is the complement of an open tubular neighbourhood of L in \mathbb{S}^3 . It has k boundary components, which are all tori.

Note the following isomorphisms, from elementary algebraic topology: by excision [22, p119], one has $H_2(\mathbb{S}^3, L; \mathbb{Z}) \simeq H_2(N_L, \partial N_L; \mathbb{Z})$; in the long exact sequence $\dots \rightarrow H_2(\mathbb{S}^3; \mathbb{Z}) \rightarrow H_2(\mathbb{S}^3, L; \mathbb{Z}) \rightarrow H_1(L; \mathbb{Z}) \rightarrow H_1(\mathbb{S}^3; \mathbb{Z}) \rightarrow \dots$ the first and last term are trivial so the boundary map gives an isomorphism $\partial : H_2(\mathbb{S}^3, L; \mathbb{Z}) \rightarrow H_1(L; \mathbb{Z})$; and by Alexander duality [22, p254], $H^1(\mathbb{S}^3 \setminus L; \mathbb{Z})$ is isomorphic to $H_1(L; \mathbb{Z}) \simeq \mathbb{Z}^k$.

The second isomorphism states that a class in $H_2(\mathbb{S}^3, L; \mathbb{Z})$ is described by its boundary. That means that, when we restrict our attention to oriented Seifert surfaces for L , i.e., surfaces S so that $\partial S = L$, all such surfaces lie in the class $\partial^{-1}(1, \dots, 1)$. In particular they are homologous. Given two such surfaces, one can consider their restrictions S_1, S_2 to N_L , where they are also homologous. Their boundaries are curves on the tori ∂N_L , which therefore have the same homology class. With an isotopy it is thus possible to make S_1, S_2 parallel and disjoint in a neighbourhood of ∂N_L .

The class of the Seifert surfaces bounded by L is dual to the class $\text{Lk}(\cdot, L_1) + \dots + \text{Lk}(\cdot, L_k) \in H^1(\mathbb{S}^3 \setminus L; \mathbb{Z})$. We denote the latter by ℓ_L . We then consider the infinite cyclic covering $\hat{N}_L \rightarrow N_L$ associated to ℓ_L . This is the covering associated to the morphism $\ell_L : \pi_1(\mathbb{S}^3 \setminus L) \rightarrow \mathbb{Z}$, so that $\pi_1(\hat{N}_L) = \ker(\ell_L)$.

One way to construct \hat{N}_L is to consider an arbitrary Seifert surface S for L , cut \hat{N}_L along S , thus obtaining a 3-manifold with boundary $N_{L,S}$. This boundary consists of three parts: two copies of S that we denote by S^+ and S^- and call the “horizontal part” of the boundary, and a “vertical part” coming from ∂N_L which consists of k annuli (or sutures) connecting S^+ to S^- . Now consider \mathbb{Z} copies $(N_{L,S}^n)_{n \in \mathbb{Z}}$ of $N_{L,S}$. Call S_n^+ and S_n^- the horizontal part of the boundaries, and for every $n \in \mathbb{N}$ glue S_n^+ to S_{n+1}^- . The resulting 3-manifold $\cup_{n \in \mathbb{Z}} N_{L,S}^n$ has a natural projection to N_L and the loops in N_L that lift to closed loops in \hat{N}_L are exactly those in the kernel of ℓ_L . Therefore we indeed constructed \hat{N}_L .

Now suppose that $\mathbb{S}^3 \setminus L$ fibres over the circle with fibres having L as oriented boundary. Then the previous construction can be made by choosing for the surface S a fibre. In this case $N_{L,S}$ is the product manifold $S \times [0, 1]$, and \hat{N}_L is homeomorphic to $S \times \mathbb{R}$.

Consider now an arbitrary Seifert surface S' for L with oriented boundary L . Then S' is homologous to S , and in particular it is also dual to ℓ_L . So S' lifts in \hat{N}_L as \mathbb{Z} parallel copies $(S'_n)_{n \in \mathbb{Z}}$. Take any such component S'_0 . Since it lies in $\hat{N}_L \simeq S \times \mathbb{R}$, the projection on the first coordinate induces a map $f : S'_0 \rightarrow S$. Looking at the neighbourhood of the common boundary $\partial S = L = \partial S'_0$ we see that f has degree 1. Degree 1 maps induce surjections in homology, which implies the desired inequality. More concretely one can argue as follows: if there was a class $a \in H^1(S; \mathbb{Z})$ in the kernel of f^* , then there would be a class b such that the cup-product $a \smile b$ is the fundamental class $[S] \in H^2(S; \mathbb{Z})$, and therefore one would have $f^*([S]) = f^*(a) \smile f^*(b) = 0 \neq [S'_0]$, a contradiction. Therefore there is linear injection from $H^1(S; \mathbb{Z}) \simeq \mathbb{Z}^{2g(S)+k-1}$ into $H_1(S'_0; \mathbb{Z}) \simeq \mathbb{Z}^{2g(S'_0)+k-1}$. This implies $g(S) \leq g(S'_0)$, and proves the statement. \square

Theorem 2.2 is sometimes also attributed to Stallings, but we did not find a corresponding statement in Stallings' articles. However a proof is proposed by Whitten [49], which relies on the fundamental work of Waldhausen [47].

Indeed if S_1 is a fibre surface for a link L , then the infinite cyclic covering \hat{N}_L constructed in the proof of Theorem 2.1 is homeomorphic to $S_1 \times \mathbb{R}$. If S_2 is a Seifert surface for L that is also a fibred surface, then Theorem 2.1 implies that it is of minimal genus. Therefore any of its lifts in \hat{N}_L is incompressible (for otherwise it would not minimize the genus). One can pick a lift S'_2 that is disjoint of $S_1 \times \{0\}$ and assume it lives in $S_1 \times (0, N)$ for some large enough N . Proposition 3.1 of Waldhausen [47] then implies that S'_2 is isotopic to $S_1 \times \{0\}$. Projecting back in N_L yields an isotopy from S_2 to S_1 . The proof of [47, Proposition 3.1] seems too long to be detailed here. We only mention that it works by an induction on the complexity of the surface S_1 and by cutting it along curves until it is a disc.

Well-quasi-orders. We refer to Diestel [13, Chapter 12] for an introduction to well-quasi-orders and graph minor theory. An order \preceq on a set X is said to be a **well-quasi-order** if for every infinite sequence $(x_n)_{n \in \mathbb{N}}$ there exist $i, j \in \mathbb{N}$ such that $i < j$ and $x_i \preceq x_j$. Equivalently, \preceq is a well-quasi-order if it is well-founded and has no infinite **antichain** that is, no infinite sequence $(x_n)_{n \in \mathbb{N}}$ such that no two elements of (x_n) are comparable for \preceq . A property P is said to be stable for an order \preceq if for any x satisfying P and $y \preceq x$, then y satisfies P . It is well-known that if \preceq is a well-quasi-order and P is a property that is stable for \preceq , then there exists a finite family \mathcal{F} of elements of X such that $x \in X$ satisfies P if and only if there is no f in \mathcal{F} such that $f \preceq x$. The family \mathcal{F} is called a family of **excluded minors** for the property P . Therefore, if a parameter $p : X \rightarrow \mathbb{N}$ is monotone with respect to \preceq , i.e. $x \preceq y$ implies $p(x) \leq p(y)$, then for each $k \in \mathbb{N}$ the property $p(\cdot) < k$ is stable for \preceq and hence is characterized by such finite family of excluded minors for X .

In this paper, a **plane tree** is a rooted tree where each vertex v has a label $\ell(v)$ from an alphabet A , and the tree is provided with the combinatorial data of an embedding in the plane: each vertex is additionally given a permutation recording the ordering of the edges to its children. The root induces an orientation on the tree: every edge $\{u, v\}$ is directed from u to v , written $u \rightarrow v$, when u is closer to the root of the tree than v (i.e., edges go toward the leaves), we refer to the trees of Figure 4 for examples. In this paper, the alphabet A is $\{+, -\}$ endowed with the empty ordering \leq .

A plane tree T_1 has a **homeomorphic embedding into T_2** , written $T_1 \hookrightarrow T_2$, if T_1 can be obtained from T_2 by iteratively (i) removing a leaf and its adjacent edge, (ii) removing a root with a single child, its adjacent edge and rerooting at that child, (iii) reducing labels (with respect to \leq) and (iv) contracting paths into edges while preserving the labels of the endpoints, where all these operations must be consistent with the plane embedding. Throughout this paper, the relation \leq will be trivial, so that (iii) will never apply. Notice that this definition extends the notion of *topological minor* on graphs to a setting where vertices are labelled. In particular, it is more restrictive than the notion of *minor*, which allows to contract any edge to a point: in our case we can only contract paths to at least one edge. This property turns out to be critical in order to make our proofs work. The famous theorem of Kruskal [24] (see Nash-Williams [38] for a simpler proof) shows that this order is a well-quasi-order on the set of labelled plane trees.

Theorem 2.3 ([24, 38]). *The homeomorphic embedding on the set of labelled plane trees labelled by a well-quasi-order forms a well-quasi-order.*

3 Hopf arborescent links

Arborescent links are a class of knots and links originally defined and studied by Conway². This class has received much attention from knot theorists [8, 17, 44]. In this paper we study a subclass that we call Hopf arborescent links.

3.1 Hopf plumbing

The links that we investigate in this paper are boundaries of surfaces which are defined iteratively from Hopf bands using an operation called plumbing.

Let H be a Hopf band and Σ be an oriented surface with boundary, and let us assume that they are unlinked, that is, that there exists a sphere S in \mathbb{S}^3 separating them. To plumb H on Σ , pick an arc α on Σ whose endpoints lie on $\partial\Sigma$ and which is not boundary parallel (i.e., α is not isotopic relatively to its endpoints, to an arc in $\partial\Sigma$). Let D be a small neighbourhood of α in Σ that we see as a rectangle with two sides on $\partial\Sigma$ and two sides in the interior of Σ . Isotope Σ within $\mathbb{S}^3 \setminus S$ so that it intersects S exactly on D , see Figure 3, left. Then, define similarly D' a neighbourhood of the unique (up to isotopy) non boundary parallel arc in H with endpoints in ∂H . The

²Conway called them algebraic links, but this denomination is now more used for the links that come from algebraic curves in \mathbb{C}^2 .

orientations of Σ and H induce an orientation of the normal direction to the surface (so that concatenating the orientation of the surface with the positive normal direction gives a positive basis in \mathbb{S}^3). Finally, isotope H within its component of $\mathbb{S}^3 \setminus S$, so that D and D' are identified on S in a way that the sides of D that are on ∂S are matched with the sides of D' that are not on ∂H and the orientations of both rectangles match. The resulting surface is said to be obtained from Σ by **Hopf plumbing H on top of Σ along α** , see Figure 3.




Figure 3: A Hopf plumbing of a Hopf band H on top of a Seifert surface Σ along α .

Hopf plumbing is a special case of a more general operation called a Murasugi sum, see [37, 39]. A key property of Murasugi sums, proved by Gabai [16], is that it preserves fibredness. In the above setting, since Hopf bands are fibred, if Σ is fibred, then the surface obtained from Σ by Hopf plumbing H on Σ along any arc is also fibred.

3.2 From plane trees to Hopf arborescent surfaces and links

Recall that in this article, a plane tree is a rooted tree that is embedded in the plane and such that every vertex has a label $+$ or $-$. If v is a vertex, we denote by $\ell(v)$ its label. Let T be a plane tree. The associated surface $\Sigma(T)$ we construct is an oriented surface with boundary that retracts on the union of a finite set of oriented simple curves \mathcal{C}_T parametrized by the vertices of T , such that every $\alpha \in \mathcal{C}_T$ is the core of a Hopf band embedded on Σ and whose sign is the label of the corresponding vertex in T . For a vertex v in T , the curve $\alpha(v)$ intersects another curve $\alpha(v')$ if and only if vv' is an edge of T , and the two curves intersect exactly once. Moreover, following $\alpha(v)$ with its given orientation, the cyclic ordering of the intersection points with the curves $\alpha(v')$ coincides with the cyclic orderings of the neighbours of v in the plane tree T .

We now describe the construction inductively, see Figure 4 for an illustration.

1. Start from a Hopf band $H(v_r)$ where v_r is the root of T , and whose sign is the label $\ell(v_r)$.
2. For the induction step, assume that the tree T' is obtained from T by adding a leaf v a finite number of leaves v_1, \dots, v_k appearing in the plane in this order around v , and that the surface $\Sigma(T)$ is already constructed with its set of core curves \mathcal{C}_T .
 - (a) Since v is a leaf in T , the curve $\alpha(v)$ intersects only one curve $\alpha(v')$: the curve associated to v' , the parent of v in T .
 - (b) Starting from this intersection point we place k points p_1, \dots, p_k on $\alpha(v)$ in this order. Then we draw on $\Sigma(T)$ a family of k arcs β_1, \dots, β_k from ∂S to itself that correspond to those arcs that retracts on p_1, \dots, p_k . Each such arc β_i intersects the collection \mathcal{C}_T exactly at the point p_i .
 - (c) For $i = 1, \dots, k$, perform the Hopf plumbing of a Hopf band $H(v_i)$ of sign $\ell(v_i)$ on top of $\Sigma(T)$ along the arc β_i . The resulting surface is $\Sigma(T')$.
 - (d) Finally for every i orient the core of $H(v_i)$ so that when going from $\alpha(v)$ to $\alpha(v_i)$, we follow this rule: if $\ell(v)$ is positive, one turns to the left (with respect to the orientation of $\Sigma(T)$), and if $\ell(v)$ is negative, one turns to the right (once again with respect to the orientation of $\Sigma(T)$), see Figure 5. The set $\mathcal{C}_{T'}$ is the union of \mathcal{C}_T with $\alpha(v_1), \dots, \alpha(v_k)$.




Figure 4: A 3D-view of a Hopf arborescent link and its construction from two different plane trees. The chosen orientation of the root of each tree is indicated on the coloured core of the matching Hopf band. The orientation of the plane is counter-clockwise.




Figure 5: Orientation of the green core when its associated Hopf band (unsigned as it does not matter for the rule) is plumbed on top of the Hopf band with the red core.

Definition 3.1. A *Hopf arborescent surface* is a surface $\Sigma(T)$ obtained from a plane tree T by this construction, see Figure 4 for an example. A *Hopf arborescent link* is the boundary of a Hopf arborescent surface.

Since Hopf bands are fibred and this property is preserved under plumbing, Hopf arborescent surfaces are fibres for their boundaries, and are thus of minimal (classical) genus. The arbitrary-looking rule that we use to orient the cores of the Hopf band in Step 2d is new and will turn out to be key for our proofs of Theorem 1.2 and Proposition 3.2.

3.3 Connections to other classes of knots

Here, we provide some additional background on Hopf arborescent links, their plumbing structure and their relations to other classes (arborescent links and fibred links). In particular one may think that considering rooted trees and always plumbing the new Hopf bands on top of the surface is a strong restriction. We will show that this is not the

case, i.e. that the family of surfaces and links obtained with less restriction on the sides on which one plumbs is the same as the family considered here.

Arborescent knots. Our definition differs from the classical definition(s) of arborescent links (e.g., in [8, 17]) and their encoding via plane trees in three ways. Firstly, we restrict our attention to Hopf bands and not general unknotted annuli, hence the name Hopf arborescent links. This restriction is essential to our approach and the minor theory we develop does not extend when one drops this assumption, as Remark 3.8 shows.

Secondly, an arborescent surface is, in general, defined as a surface that retracts on a collection of core curves whose intersection pattern is a tree. One can show that in this case, the surface can be reconstructed using an algorithm similar to the one we propose in Section 3.2. The main difference is that when gluing a new band to the surface, it may be glued on top or below the surface. This difference could let one think that our construction is more peculiar than the standard one. However, it was remarked by Misev [35] that gluing a Hopf band on top or below an arc yields isotopic surfaces, see Figure 6. The only difference is that the orientation of the core curves of the glued band is then reversed. This change of orientation then implies that all vertices of the subtree subsequently glued to this band should be reversed to have the same surface. This argument shows that the set of surfaces we construct actually coincides with the set of surfaces obtained by the classical construction of arborescent links using Hopf bands only. The reason why we chose our presentation (i.e., always gluing on top when going away from the root of the tree) is that it behaves more nicely with respect to our minor theory, and in particular it simplifies the proof of Lemma 3.6.

Thirdly, the usual encoding of arborescent knots with plane trees [8, Chapter 12] is more detailed in that it puts the number of twists between each pair of adjacent edges around each vertex. Here, since bands have either 2 positive crossings or 2 negative, Figure 7 shows that we do not need to specify where the twists are on the band, and thus indicating the sign of the band on each vertex is sufficient.




Figure 6: Surfaces obtained by plumbing a Hopf band inside or outside a Seifert surface are isotopic.




Figure 7: The surface does not depend on where the twists are.

As noticed in Section 4.2 with the set $\mathcal{T}(L)$, it may happen that several trees yield the same (isotopy class of) oriented link and the same surface. In other words, our construction induces a well-defined map from plane trees to isotopy classes of surfaces in \mathbb{S}^3 , but this map may not be injective. This lack of injectivity is not an issue for the problems we address in this paper, for the preimage of a given link or surface is finite. However, it is carefully studied by Bonahon and Siebenmann [8] who give a recipe to detect (in the more general case of arborescent links) when two \mathbb{Z} -labelled plane trees yield isotopic links.

More general plumbing structures. By construction, the tree T involved in defining a Hopf arborescent surface $\Sigma(T)$ is the intersection graph of the set of core curves \mathcal{C}_T . It may well happen that a surface retracts on different collections of Hopf curves [36, 8] whose intersection pattern is very different from T , and in particular is not a tree. In that case, the surface can still be obtained by a sequence of Hopf plumbings, but some plumbing are

made on arcs that intersect the core of more than one Hopf band. On Figure 8 for example, the green curve on the right is indeed the core of a Hopf band, see top of Figure 9 to picture how a tubular neighbourhood of the core is a band with two positive crossings.




Figure 8: A Hopf arborescent plumbing with a cyclic plumbing graph.

3.4 Minors on surfaces, links, and plane trees

Since we focus our investigation on Hopf arborescent links, we define a stronger notion of minor that is well-tailored to these links. We say that a Hopf arborescent link L_1 is a **link-minor** of L_2 if there exist T_1 and T_2 , two labelled plane trees such that $\Sigma(T_1)$ and $\Sigma(T_2)$ are canonical Seifert surfaces of L_1 and L_2 respectively and $T_1 \hookrightarrow T_2$. The main result of this section is the following one, establishing that if L_1 is a link-minor of L_2 then the Seifert surface of L_1 is a surface-minor of the Seifert surface of L_2 .

Proposition 3.2. *Let T_1 and T_2 be two plane trees such that T_1 admits a homeomorphic embedding into T_2 . Then the Hopf arborescent surface $\Sigma(T_1)$ is an incompressible subsurface of $\Sigma(T_2)$.*

The proof relies on Lemmas 3.4, 3.5, and 3.6, which correspond respectively to the operations (i), (ii), and (iv) defining homeomorphic embeddings of trees. We first prove:

Lemma 3.3. *Let Σ be surface and γ be an arc that is not boundary-parallel in Σ with both extremities in $\partial\Sigma$. Then cutting Σ along γ yields a surface $\Sigma' = \overline{\Sigma \setminus \gamma}$ such that $\Sigma' \preceq \Sigma$.*

Proof. By definition of Σ' , there is a natural map $h : \Sigma' \rightarrow \Sigma$ that is injective except on $h^{-1}(\gamma) = \gamma_1 \cup \gamma_2$. Let T be a tubular neighbourhood of γ in Σ . Its boundary can be decomposed into t_1, t_2 , two arcs isotopic to γ in Σ and two open arcs of $\partial\Sigma$. Isotope h within Σ so that $h(\gamma_1) = t_1$, $h(\gamma_2) = t_2$, and $h(\Sigma') \cap T = t_1 \cup t_2$. It follows that $h(\Sigma')$ is a subsurface of Σ such that $\Sigma \setminus \Sigma'$ is not an open disc since γ is not boundary-parallel. \square

Lemma 3.3 essentially states that our surfaces behave well with respect to the surface-minor relation when cut along any essential arc. An important point is that cutting along an arc that is the diagonal of a plumbing rectangle merges two bands into one new band with two extra crossings that are either negative or positive depending on the diagonal, see Figure 9. So, cutting the plumbing of two positive Hopf bands along the diagonal that produces two negative crossings yields a positive Hopf band. Symmetrically, one can merge two negative Hopf bands into one negative by cutting along the other diagonal. Furthermore, when having a plumbing of two bands with opposite signs, one can merge them into a band with either sign depending on which cut one chooses.

Lemma 3.4. *Assume that T_1 is obtained from T_2 by deleting a leaf. Then $\Sigma(T_1)$ is an incompressible subsurface of $\Sigma(T_2)$.*

Proof. Let D be the plumbing rectangle of the Hopf band H associated to the additional leaf v of T_2 compared to T_1 . By definition, D has two sides γ_1, γ_2 in $\partial\Sigma(T_1)$. Thus γ_1 is also an arc of $\Sigma(T_2)$ with its extremities in $\partial\Sigma(T_2)$. By Lemma 3.3, $\Sigma(T_2)$ cut along γ_1 is an incompressible subsurface Σ' of $\Sigma(T_2)$. Furthermore, the remaining of H is a disc that can be isotoped into a neighbourhood of γ_2 so that $\Sigma(T_1) = \Sigma' \preceq \Sigma(T_2)$, see Figure 10. \square




Figure 9: Cutting a Hopf plumbing of two positive Hopf bands along the diagonal γ_2 (resp. γ_1) induces 2 positive (resp. negative) crossings. The Reidemeister moves are showed to help understand which crossings are obtained. The top right surface is a Hopf band while the bottom right one is a band with 3 positive full twists.

A very similar proof yields the following lemma.

Lemma 3.5. *Assume that T_1 is a plane tree whose root has only one child, T_2 is the subtree rooted at that child. Then $\Sigma(T_1)$ is an incompressible subsurface of $\Sigma(T_2)$.*

Proof. The proof is identical to the previous one: in that situation, $\Sigma(T_1)$ is obtained from $\Sigma(T_2)$ by plumbing a Hopf band, and cutting along one of the boundaries of the plumbing disc provides the needed incompressible subsurface, as in Figure 10. \square

Lemma 3.6. *Assume that T_2 is a plane tree in which $u \rightarrow v \rightarrow w$ are three consecutive vertices where v has degree 2, and that T_1 is obtained from T_2 by contracting $u \rightarrow v \rightarrow w$ into a single edge $u \rightarrow w$, while preserving the labels of the endpoints. Then $\Sigma(T_1)$ is an incompressible subsurface of $\Sigma(T_2)$.*

Proof. By the construction of Hopf arborescent surfaces, the edge between u and v in T_2 corresponds to a plumbing rectangle D . It is important to recall here our orientation convention: if u is labelled positively, the cores $\alpha(u)$ and $\alpha(v)$ are oriented so that one turns to the left when going from u to v at the rectangle D , while if v is labelled negatively, one turns to the right, see Figure 5. Now, we consider two diagonal arcs γ_1 and γ_2 on the plumbing rectangle D as pictured in Figure 9. When cutting along such a diagonal arc, we obtain a new surface in which the cores $\alpha(u)$ and $\alpha(v)$ merge into a single core. However, their orientations might mismatch, depending on whether we cut along γ_1 or γ_2 . We take the convention that γ_1 is the arc the preserves the orientations, while γ_2 induces an orientation mismatch, see Figures 9 and 11.

Now, let us first consider the case where the labels of u and v are the same. In this case, we consider the subsurface Σ' of $\Sigma(T_2)$ obtained by cutting along γ_1 . This has the effect of merging the core curves $\alpha(u)$ and $\alpha(v)$




Figure 10: Deleting a leaf yields a surface-minor.

in a way that respects their orientations. However, it might seem that since each curve $\alpha(u)$ and $\alpha(v)$ corresponds to a Hopf band, merging them like that yields a band that twists too much. But a key observation is that cutting along γ_1 adds a twist between these two bands, as pictured in Figure 9, and this twist is negative when the bands are positive, while it is positive when the bands are negative (indeed, this is the reason for our orientation convention). Therefore, the resulting surface Σ' is exactly the same as the one corresponding to the tree T_1 , and therefore $\Sigma(T_1)$ is an incompressible subsurface of $\Sigma(T_2)$ by Lemma 3.3. See the top and bottom pictures of Figure 11 for an illustration.

Now, let us consider the case where the label of u is $+$ while the label of v is $-$. In that case, we consider the surface Σ' of $\Sigma(T_2)$ obtained by cutting along γ_2 . This has the effect of merging the core curves $\alpha(u)$ and $\alpha(v)$ but with an orientation mismatch. We take the convention that the resulting core curve α' is oriented by $\alpha(u)$, and therefore disagrees with the orientation of $\alpha(v)$ while it follows it. Since u and v are labelled with opposite signs the two twists on their Hopf bands cancel out, but cutting along γ_2 adds a new positive twist, therefore we can consider α' as being the core curve of a positive Hopf band. Now, let us consider the plumbing rectangle D' corresponding to the edge between v and w . Due to the orientation mismatch, arriving at this rectangle from α' , we are oriented in the direction opposed to the one we would arrive with if we were arriving from $\alpha(v)$. But due to the orientation convention, when going from $\alpha(v)$ to $\alpha(w)$ in $\Sigma(T_2)$ we turn to the right since v is negative, while when going from α' to $\alpha(w)$ in Σ' we turn to the left since α' is a positive band. Therefore, this effect cancels out the orientation mismatch, and Σ' coincides exactly with the surface $\Sigma(T_1)$ corresponding to the tree T_1 . Therefore $\Sigma(T_1)$ is an incompressible subsurface of $\Sigma(T_2)$ by Lemma 3.3. See the third picture of Figure 11 for an illustration.

The same cancellation effect happens when the label of u is $-$ and the label of v is $+$: when cutting along γ_2 we have an orientation mismatch which is cancelled out by the fact that the new band is negative, and thus the orientation convention makes it turn in the opposite direction in the plumbing rectangle between v and w . Therefore, in that case $\Sigma(T_1)$ is also an incompressible subsurface of $\Sigma(T_2)$ thanks to Lemma 3.3. This is illustrated in the second picture of Figure 11. \square

As a corollary, contracting any path of T_1 into an edge whose labels match with the labels of the extremities of the path produces a tree T_2 such that $\Sigma(T_1) \preceq \Sigma(T_2)$.

Proof of Proposition 3.2. By definition, if T_1 admits a homeomorphic embedding into T_2 , it can be obtained iteratively from T_2 by (i) removing a child leaf, (ii) removing a parent leaf, (iii) reducing a label or (iv) contracting a path while preserving the labels of the endpoints. Since no two elements on the alphabet $\{+, -\}$ are comparable, case (iii) cannot happen. Then the cases (i), (ii) and (iv) are handled respectively by Lemma 3.4, Lemma 3.5 and Lemma 3.6. \square

On the other hand, the Kruskal Tree Theorem directly yields the following proposition.

Proposition 3.7. *Hopf arborescent links are well-quasi-ordered under the link-minor relation.*

Proof. Take an infinite sequence $(L_n)_{n \in \mathbb{N}}$ of Hopf arborescent links, let $(T_n)_{n \in \mathbb{N}}$ a sequence of plane trees such that for all $n \in \mathbb{N}$, $\Sigma(T_n)$ is a Seifert surface of L_n . Then, by Theorem 2.3, there exists $i < j$ such that T_i admits a homeomorphic embedding into T_j . Hence L_i is a link-minor of L_j . \square

We can deduce Theorem 1.2 as a direct corollary of Proposition 3.2 and Proposition 3.7.




Figure 11: All cases of contraction of a 3-path to an edge preserving the labels of the endpoints.

Proof of Theorem 1.2. Take an infinite sequence $(\Sigma_n)_{n \in \mathbb{N}}$ of canonical Seifert surfaces of Hopf arborescent links. Then by Proposition 3.7, $\partial \Sigma_i \hookrightarrow \partial \Sigma_j$ for some $i < j$. By Proposition 3.2 we have $\Sigma_i \preceq \Sigma_j$, i.e. the surface-minor order is a well-quasi-order on Hopf arborescent surfaces. \square

Remark 3.8. *The proof of Proposition 3.2 highlights that the minor relation on the set of Hopf arborescent surfaces is more subtle and fragile than one might expect. Indeed, the cuts involved when taking an incompressible subsurface in the proof of Lemma 3.6 inevitably merge Hopf bands and thus one needs to be careful in order to control the number of resulting twists. In particular, the proof does not seem to generalise to the more general classes of surfaces obtained by plumbing bands with a bounded number of twists (even though everything works well at the level of trees).*

4 Decidability of the genus defect for Hopf arborescent links

4.1 Monotonicity of the genus defect

Now that we proved that link-minor is a well-quasi-order on the set of Hopf arborescent links, we want to highlight a property that is stable for this minor relation. Recall that the genus defect $\Delta_g(L)$ of an oriented link L is the difference $g(L) - g_4(L)$ between its classical genus and its 4-dimensional genus. The latter can be either in the topological locally flat or in the smooth category. All statements in this section (and in particular Theorem 1.1) hold in both categories. We reprove Lemma 6 of [6] in the form of Proposition 4.1 using the fact that link-minor implies that the associated Seifert are surface-minors.

Proposition 4.1. *The genus defect Δ_g is monotone on the family of Hopf arborescent links with respect to the link-minor relation, i.e., if L_1 is a link-minor of L_2 , then $\Delta_g(L_1) \leq \Delta_g(L_2)$.*

We rely on the following lemma that highlights how the 4-genus behaves with respect to surface-minors. It uses a cut-and-paste construction and an Euler characteristic argument.

Lemma 4.2. *Let Σ be an oriented surface of \mathbb{S}^3 and Σ' be a surface-minor of Σ . If we write $L = \partial \Sigma$ and $L' = \partial \Sigma'$, then we have $g(\Sigma) - g_4(L) \geq g(\Sigma') - g_4(L')$.*

Proof. Seeing \mathbb{S}^3 as the boundary of the 4-ball \mathbb{B}^4 , consider a surface S' in \mathbb{B}^4 such that $\partial S' = L'$ and $S' \cap \mathbb{S}^3 = L'$. Gluing the remaining pieces of $\Sigma \setminus \Sigma'$ to S' along L' yields a surface S in \mathbb{B}^4 such that $\partial S = L$, see Figure 12.




Figure 12: Illustration of the construction of Σ , Σ' , L , S , and S' . Purple surfaces are in \mathbb{S}^3 while blue one are considered to be in \mathbb{B}^4 .

By definition of $g_4(L)$, we have $g_4(L) \leq g(S)$, and thus $g(\Sigma) - g_4(L) \geq g(\Sigma) - g(S)$. Furthermore, the genus of S is given by $g(S) = g(\Sigma) - g(\Sigma') + g(S')$. Thus one has $g(\Sigma) - g_4(L) \geq g(\Sigma') - g(S')$. Now if we assume S' to minimise the 4-genus over surfaces bounded by L' , we conclude: $g(\Sigma) - g_4(L) \geq g(\Sigma') - g_4(L')$. \square

As Hopf arborescent surfaces are of minimal genus for Hopf arborescent links, Lemma 4.2 can be used to prove Proposition 4.1.

Proof of Proposition 4.1. Let L and L' be two Hopf arborescent links such that L' is a link-minor of L . Consider Σ and Σ' the corresponding canonical Seifert surfaces. Since Hopf arborescent links are fibred, Σ is a Seifert surface of L of minimal genus (see Theorem 2.1), i.e. $g(L) = g(\Sigma)$ and similarly, $g(\Sigma') = g(L')$. By Proposition 3.2 one has $\Sigma \preceq \Sigma'$. Hence, by Lemma 4.2, one gets $\Delta_g(L) \geq \Delta_g(L')$. \square

4.2 Proof of Theorem 1.1

We first show that the link-minor relation can be decided using the decidability of link equivalence. For knots, equivalence can be tested as a combination of an algorithm that allows to decide whether two 3-manifolds with boundary are homeomorphic [25, 33] and the Gordon-Luecke Theorem [19] that states that two knots are equivalent if their complements, which are 3-manifolds with boundaries, are equivalent. In the case of links, we additionally need to keep track of a longitude of each torus boundary component in the complement of the link. We refer to the survey of Lackenby [26, Section 2] for a summary of the techniques that allow to prove the following theorem:

Theorem 4.3 (Link equivalence). *Given two links L_1 and L_2 , the problem of testing whether L_1 is ambient isotopic to L_2 is decidable.*

Given a Hopf arborescent link L , denote by $\mathcal{T}(L)$ the set of plane trees T whose associated Hopf arborescent surface $\Sigma(T)$ has L as oriented boundary. As a corollary, we obtain:

Lemma 4.4. *Given a Hopf arborescent link L , the set $\mathcal{T}(L)$ is computable.*

Proof. For increasing $k \in \mathbb{N}$, we enumerate all plane trees T_i with k vertices labelled by $\{-, +\}$ and store the trees T_i such that $\partial\Sigma(T_i)$ is isotopic to L , where we test isotopy using Theorem 4.3. If we find a k for which such a tree exists, we finish the enumeration for this value and return the stored trees. Indeed, plumbing n Hopf bands produces a surface with Betti number n . By Theorem 2.1, all the trees T_i such that $\partial\Sigma(T_i) = L$ produce surfaces $\Sigma(T_i)$ with the same genus, hence have the same number of vertices. As the entry is a Hopf arborescent link, there exists a tree T such that $\partial\Sigma(T) = L$. So the algorithm terminates. \square

Alternatively, and if one wants some control on the complexity of that algorithm, one can avoid blindly testing for increasing k by first computing the genus of the link [21, 33], or just computing an upper bound to it using, e.g., Seifert's algorithm, and then enumerate only the trees that produce surfaces up to that genus. From Lemma 4.4, we obtain:

Lemma 4.5. *Given two Hopf arborescent links L_1 and L_2 , testing if L_1 is a link-minor of L_2 is decidable.*

Proof. Using Lemma 4.4, we compute $\mathcal{T}(L_1)$ and $\mathcal{T}(L_2)$. The trees in $\mathcal{T}(L_1)$ (resp. $\mathcal{T}(L_2)$) all have the same number k_1 (resp. k_2) of vertices. Then we brute force every possible path contraction to an edge and iterated leaf deletion on trees of $\mathcal{T}(L_2)$ such that the result is a tree with k_1 vertices, and test whether it is equal to a tree of $\mathcal{T}(L_1)$. If such a test succeeds, we output yes, otherwise we return no. There is a finite number of trees in both $\mathcal{T}(L_1)$ and $\mathcal{T}(L_2)$ and a finite number of a trees with k_1 vertices that homeomorphically embed into a tree of $\mathcal{T}(L_2)$. Hence that algorithm eventually terminates. Its correctness follows directly from the definition of link-minor. \square

Finally we prove Theorem 1.1 by using the stability of the genus-defect by link-minor, the previous algorithms, and the well-quasi-order properties.

Proof of Theorem 1.1. By Proposition 3.7, the order defined by link-minors is a well-quasi-order on the set of Hopf arborescent links. Hence, the set of Hopf arborescent links \mathcal{H}_k that have defect at most k is characterized by a finite family \mathcal{F}_k of forbidden minors. It follows, by Proposition 4.1 that $\Delta_g(L) \leq k$ if and only if for all f in \mathcal{F}_k , f is not a link-minor of L . Using Lemma 4.5 we test for each $f \in \mathcal{F}_k$ if f is a link-minor of L . If such a test succeeds, output no, otherwise the input link has $\Delta_g(L) \leq k$. \square

As said in the introduction, our proof is not constructive as it relies at its core on the existence of a set of forbidden minors for having defect at most k . This set of forbidden minors is not explicit and hard-coded in the algorithm. Furthermore, the sets of excluded minors will be different for the two different notions of defect (smooth and locally flat). It is likely that computing them is a topological challenge requiring arguments of different nature.

Theorem 1.2 provides the existence of a set of forbidden minors for having defect at most k but for a different and stronger definition of minors on links that relies only on the surface-minor relation on the Seifert surface and not the trees. However deciding this relation, even by a brute force argument, seems challenging: in addition to the fact that no algorithm seems to be known for testing isotopy of surfaces, one would also need to control the complexity of the cutting arcs. Even with positive Hopf arborescent links only, this seems delicate [34].

5 Examples: Hopf arborescent links with non-trivial defect

It is not clear a priori that the topological and/or smooth defects of Hopf arborescent links are nonzero. For instance, our building block, the Hopf band, has both defects equal to 0 since it bounds an annulus in \mathbb{S}^3 ; which has genus 0. Furthermore, as mentioned in the introduction, when Hopf arborescent links are only made with positive Hopf bands, they belong to a class of links called *positive* links, which implies that they are *strongly quasi-positive* [43], and this implies in turn that their smooth 4-genus equal their 3-genus [42]. So for this class of links, the smooth defect is always zero. In contrast, in this section, we provide an example of a Hopf arborescent knot for which both the topological and the smooth defect are nonzero, and an argument explaining how to use this knot to provide examples with arbitrarily large defects.

Example 1. A Hopf arborescent link with non trivial topological defect is given by Baader, Feller, Lewark and Liechti [6, Example 4]. It consists in plumbing 6 positive Hopf bands in a latin cross-like pattern. Since it is positive, its smooth defect is zero. However, one can find a torus sub-surface of S_X such that the restricted Alexander polynomial vanishes. By Freedman's Disc Theorem [15], this implies that one can remove this subsurface and glue along its boundary a locally flat disc in the 4-ball. This means that the topological defect is at least 1. We refer to the article [6] for more details.

Example 2. By the aforementioned result of Rudolph, in order to obtain a Hopf arborescent link with nonzero smooth defect, we need to use both positive and negative Hopf bands. Here is such an example: the Hopf arborescent link obtained from the tree pictured in Figure 13 is in fact a knot, named 8_{10} in the Rolfsen table [41]. One can readily check from knot censuses (for example from the Knot Atlas [1]) the smooth and topological genera of this knot, and we have $\Delta_g^{smooth}(8_{10}) = \Delta_g^{top}(8_{10}) = 2$. For completeness, we provide here another argument to explain why the defects are nonzero.




Figure 13: The Hopf arborescent knot 8_{10} (from the Rolfsen table).

The genus-minimizing surface is the Hopf arborescent surface pictured in Figure 14, whose genus is 3, and with 1 boundary component. The red curve on the left of Figure 14 is the core of an embedded annulus and its two boundaries are unlinked. Replacing this annulus by two discs glued on its two boundary components yields an immersed surface with genus 1 less, still bounded by $K_{8_{10}}$. The singularities of this immersion are of a particular type called ribbon singularities [41, p.225] (right side of Figure 14). Such singularities can be resolved in the 4-ball, and thus this implies that the Hopf arborescent surface can be pushed into a smoothly embedded surface in the

4-ball so that its smooth defect is at least 1. Furthermore, for any link L , we have that $g_4^{\text{smooth}}(L) \geq g_4^{\text{top}}(L)$ so that its topological defect is also at least 1.




Figure 14: Left: A simple red curve on the surface associated to $\Sigma(8_{10})$. The intersection number in green show that the boundaries of the corresponding annulus are unlinked. Right : A ribbon intersection.

One can also check in the censuses that these two examples are hyperbolic.

Arbitrarily large defect. We can use the previous examples to create Hopf arborescent links with arbitrarily large (smooth or topological) defect. The following proposition proves that the defect coming from tree patterns add up if they are isolated enough.

Proposition 5.1. *Let L be an arborescent link associated to a plane tree T . Let T_1, \dots, T_n be disjoint subtrees of T that are pairwise at distance at least two within T . Denote $L_i = \partial(\Sigma(T_i))$, then $\sum_{1 \leq i \leq n} \Delta_g(L_i) \leq \Delta_g(L)$.*

Proof. For each i , identifying the Hopf bands of the vertices of T_i within $\Sigma(T)$ provides a family of embeddings of the $\Sigma(T_i)$. These embeddings are disjoint by the distance assumption.

We can now repeat the proof of Lemma 4.2 and Proposition 4.1 with a family of links to obtain the desired inequality. Since the $\Sigma(T_i)$ are disjoint subsurfaces of $\Sigma(L)$, $\bigcup \Sigma(T_i)$ is a surface-minor of $\Sigma(L)$. As in the proof of Lemma 4.2, we define for each L_i , a surface S'_i in \mathbb{B}^4 that has L_i as its boundary. Gluing each S'_i to the pieces of $\Sigma(T) \setminus \bigcup \Sigma(T_i)$ along its associated L_i yields a surface S in \mathbb{B}^4 such that $\partial S = L$.

Now we have $g(S) = g(\Sigma(T)) - \sum_i g(\Sigma(T_i)) + \sum_i g(S'_i)$. It follows, by definition of g_4 that $g(\Sigma(T)) - g_4(L) \geq g(\Sigma(T)) - g(S) = \sum_i g(\Sigma(T_i)) - g(S'_i)$. By Theorem 2.1, and taking each S'_i to be of minimal genus, $\Delta_g(L) = g(L) - g_4(L) \geq \sum_i g(L_i) - g_4(L_i) = \sum_i \Delta_g(L_i)$. \square

Consider a Hopf arborescent link L associated to some tree T such that L has non zero defect. By identifying the roots of n copies of T with the leaves of a n -star graph (one vertex of degree n , n vertices of degree 1), we create a tree T' such that the link $L' = \partial\Sigma(T')$ has defect at least $n\Delta_g(L)$ by Proposition 5.1, therefore obtaining links with arbitrarily large topological and smooth defect.

Acknowledgements. We would like to thank Sebastian Baader for helpful discussions, and the anonymous reviewers for their questions and suggestions which allowed us to significantly improve the paper.

References

- [1] The knot atlas: 8_{10} . http://katlas.org/wiki/8_10. Accessed: 2024-13-03.
- [2] Ian Agol, Joel Hass, and William Thurston. The computational complexity of knot genus and spanning area. *Transactions of the American Mathematical Society*, 358(9):3821–3850, 2006.
- [3] Sebastian Baader. Positive braids of maximal signature. *L’Enseignement Mathématique*, 59(3):351–358, 2014.
- [4] Sebastian Baader and Pierre Dehornoy. Minor theory for surfaces and divides of maximal signature. *arXiv preprint arXiv:1211.7348*, 2012.
- [5] Sebastian Baader, Pierre Dehornoy, and Livio Liechti. Minor theory for quasipositive surfaces. In Athanase Papadopoulos, editor, *Essays in geometry, dedicated to Norbert A’Campo*, pages 351–358. EMS Pres, 2023.
- [6] Sebastian Baader, Peter Feller, Lukas Lewark, and Livio Liechti. On the topological 4-genus of torus knots. *Transactions of the American Mathematical Society*, 370(4):2639–2656, 2018.
- [7] Filippo Baroni. Classification of genus-two surfaces in \mathbb{S}^3 . *arXiv preprint arXiv:2309.05387*, 2023.
- [8] Francis Bonahon and Laurence C Siebenmann. New geometric splittings of classical knots and the classification and symmetries of arborescent knots. Preprint available on the authors’ webpage, 1979.
- [9] Gerhard Burde and Heiner Zieschang. *Knots*. Walter de gruyter, 2002.
- [10] Benjamin A. Burton, Ryan Budney, William Pettersson, et al. Regina: Software for low-dimensional topology. <http://regina-normal.github.io/>, 1999–2023.
- [11] Marek Cygan, Fedor V Fomin, Łukasz Kowalik, Daniel Lokshantov, Dániel Marx, Marcin Pilipczuk, Michał Pilipczuk, and Saket Saurabh. *Parameterized algorithms*, volume 4. Springer, 2015.
- [12] Arnaud de Mesmay, Yo’av Rieck, Eric Sedgwick, and Martin Tancer. The unbearable hardness of unknotting. *Advances in Mathematics*, 381:107648, 2021.
- [13] Reinhard Diestel. *Graph theory*. Number 173 in Graduate texts in mathematics. Springer, New York, 5th edition, 2016.
- [14] Ralph Fox. Some problems in knot theory. In *Topology of 3-manifolds and related topics (Proc. The Univ. of Georgia Institute)*, pages 168–176, Englewood Cliffs, N.J, 1961. Prentice-Hall.
- [15] Michael Hartley Freedman. The topology of four-dimensional manifolds. *Journal of Differential Geometry*, 17(3):357–453, 1982.
- [16] David Gabai. The Murasugi sum is a natural geometric operation. *Contemp. Math*, 20:131–143, 1983.
- [17] David Gabai. *Genera of the Arborescent Links*. Memoirs of the American Mathematical Society. American Mathematical Society, 1986.
- [18] Emmanuel Giroux and Noah Goodman. On the stable equivalence of open books in three-manifolds. *Geometry & Topology*, 10(1):97–114, 2006.
- [19] Cameron McA. Gordon and John S Luecke. Knots are determined by their complements. *Journal of the American Mathematical Society*, 2:371–415, 1989.
- [20] Wolfgang Haken. Theorie der normalflächen. *Acta Mathematica*, 105(3):245–375, 1961.
- [21] Joel Hass, Jeffrey C. Lagarias, and Nicholas Pippenger. The computational complexity of knot and link problems. *Journal of the ACM (JACM)*, 46(2):185–211, 1999.
- [22] Allen Hatcher. *Algebraic topology*. Cambridge University Press, Cambridge ; New York, 2002.
- [23] Peter B Kronheimer and Tomasz S Mrowka. The genus of embedded surfaces in the projective plane. *Mathematical Research Letters*, 1(6):797–808, 1994.

- [24] Joseph B. Kruskal. Well-quasi-ordering, the tree theorem, and Vazsonyi's conjecture. *Transactions of the American Mathematical Society*, 95:210–225, 1960.
- [25] Greg Kuperberg. Algorithmic homeomorphism of 3-manifolds as a corollary of geometrization. *Pacific Journal of Mathematics*, 301(1):189–241, September 2019. arXiv: 1508.06720. URL: <http://arxiv.org/abs/1508.06720>.
- [26] Marc Lackenby. Elementary Knot Theory. 01 2017.
- [27] Marc Lackenby. Some conditionally hard problems on links and 3-manifolds. *Discrete & Computational Geometry*, 58:580–595, 2017.
- [28] Marc Lackenby. Algorithms in 3-manifold theory. *Surveys in Differential Geometry*, 2020.
- [29] Marc Lackenby. The efficient certification of knottedness and Thurston norm. *Advances in Mathematics*, 387:107796, 2021.
- [30] Livio Liechti. On the genus defect of positive braid knots. *Algebraic & Geometric Topology*, 20(1):403–428, 2020.
- [31] Andrei Andreevich Markov. The insolubility of the problem of homeomorphy. In *Doklady Akademii Nauk*, volume 121, pages 218–220. Russian Academy of Sciences, 1958.
- [32] Jiří Matoušek, Martin Tancer, and Uli Wagner. Hardness of embedding simplicial complexes in \mathbb{R}^d . *Journal of the European Mathematical Society*, 13(2):259–295, 2010.
- [33] Sergei Vladimirovich Matveev. *Algorithmic topology and classification of 3-manifolds*. Springer, 2007.
- [34] Filip Misev. Cutting arcs for torus links and trees. *Bulletin de la Société Mathématique de France*, 145:575–602, 2014.
- [35] Filip Misev. *On the plumbing structure of fibre surfaces*. PhD thesis, Universität Bern, 2016.
- [36] Filip Misev. Hopf bands in arborescent Hopf plumbings. *Osaka Journal of Mathematics*, 56(2):375 – 389, 2019.
- [37] Kunio Murasugi. On a certain subgroup of the group of an alternating link. *American Journal of Mathematics*, 85(4):544–550, 1963.
- [38] C. St. J. A. Nash-Williams. On well-quasi-ordering finite trees. *Mathematical Proceedings of the Cambridge Philosophical Society*, 59(4):833–835, 1963.
- [39] Burak Ozbagci and Patrick Popescu-Pampu. Generalized plumbings and murasugi sums. *Arnold Mathematical Journal*, 2(1):69–119, dec 2015.
- [40] Lisa Piccirillo. The conway knot is not slice. *Annals of Mathematics*, 191(2):581–591, 2020.
- [41] Dale Rolfsen. *Knots and links*. AMS Chelsea Pub, Providence, R.I, 2003. OCLC: ocm52901393.
- [42] Lee Rudolph. Quasipositivity as an obstruction to sliceness. *Bulletin of the American Mathematical Society*, 29(1):51–59, 1993.
- [43] Lee Rudolph. Positive links are strongly quasipositive. *Geometry & Topology Monographs, Volume 2: Proceedings of the Kirbyfest*, pages 555–562, 1998.
- [44] Makoto Sakuma. Minimal genus seifert surfaces for special arborescent links. *Osaka Journal of Mathematics*, 31:861–905, 1994.
- [45] John R. Stallings. Constructions of fibred knots and links. *Proc. Symp. Pure Math.*, AMS 27:315–319, 1975.
- [46] Peter Teichner. Slice knots: Knot theory in the 4th dimension, 2011. Lecture notes by Julia Collins and Mark Powell. Electronic version available from <https://www.maths.ed.ac.uk/~v1ranick/papers/sliceknots2.pdf>.
- [47] Friedhelm Waldhausen. On irreducible 3-manifolds which are sufficiently large. *Annals of Mathematics*, 87:56–88, 1968.
- [48] Shmuel Weinberger. Homology manifolds. *Handbook of geometric topology*, pages 1085–1102, 2002.
- [49] Wilbur Whitten. Isotopy types of knot spanning surface. *Topology*, 12:373–380, 1973.