

Localized Geometric Moves to Compute Hyperbolic Structures on Triangulated 3-Manifolds

Clément Maria   

Inria Sophia Antipolis-Méditerranée, France

Owen Rouillé  

Inria Sophia Antipolis-Méditerranée, France

Abstract

A fundamental way to study 3-manifolds is through the geometric lens, one of the most prominent geometries being the hyperbolic one. We focus on the computation of a complete hyperbolic structure on a connected orientable hyperbolic 3-manifold with torus boundaries. This family of 3-manifolds includes the knot complements.

This computation of a hyperbolic structure requires the resolution of gluing equations on a triangulation of the space, but not all triangulations admit a solution to the equations.

In this paper, we propose a new method to find a triangulation that admits a solution to the gluing equations, using convex optimization and localized combinatorial modifications. It is based on Casson and Rivin's reformulation of the equations. We provide a novel approach to modify a triangulation and update its geometry, along with experimental results to support the new method.

2012 ACM Subject Classification Theory of computation → Computational geometry

Keywords and phrases knots and 3-manifolds, triangulation, hyperbolic structure, Thurston equations

Digital Object Identifier 10.4230/LIPIcs.ESA.2022.78

Related Version *Full Version:* <https://arxiv.org/abs/2112.06360>

Funding *Clément Maria:* Partially funded by ANR project ANR-20-CE48-0007 (AlgoKnot).

Owen Rouillé: Fully funded by ANR project ANR-20-CE48-0007 (AlgoKnot).

1 Introduction

A main problem of knot theory is to tell whether two knots are equivalent or distinct. Equivalence between knots is defined by the existence of an isotopy of the ambient space that would turn one knot into the other, i.e., a continuous deformation of the space that preserves the entanglement.

Isotopies are too difficult to compute in practice, and practitioners use *invariants* to tackle the knot equivalence problem. A *topological invariant* is a quantity assigned to a presentation of a knot, that is invariant by isotopy.

An important family of knots are the *hyperbolic knots*, which are the knots whose complements admit a *complete hyperbolic metric*. They are the subjects of active mathematical research, which motivates the introduction of efficient algorithmic tools to study their geometric properties, and most notably their *hyperbolic volume*. The hyperbolic volume of a hyperbolic knot is a topological invariant which is powerful at distinguishing between non-equivalent knots, is non-trivial to compute, and is at the heart of several deep conjectures in topology [11].

If it exists, the complete hyperbolic metric on a 3-manifold is unique[10], and may be combinatorially represented by a *complete hyperbolic structure* (CHS). In order to compute geometric properties (such as volume) of a hyperbolic knot, one triangulates the knot complement, and try to assign hyperbolic shapes to its tetrahedra. If these *shapes* verify a set of non-linear constraints called the *gluing equations*, they form a CHS and encode the



© Clément Maria and Owen Rouillé;
licensed under Creative Commons License CC-BY 4.0
30th Annual European Symposium on Algorithms (ESA 2022).

Editors: Shiri Chechik, Gonzalo Navarro, Eva Rotenberg, and Grzegorz Herman; Article No. 78; pp. 78:1–78:13

Leibniz International Proceedings in Informatics




LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

78:2 Localized Geometric Moves to Compute Hyperbolic Structures

■ **Table 1** On all ~ 9.7 millions prime knots with crossing numbers ranging from 12 to 17, alternating and non-alternating, we indicate (*% Failure on first try*) the percentage of knot complements (after triangulation and simplification) on which **SnapPy** fails to compute a CHS on first try. We also indicate (*Highest expected nb of retriang.*) the highest, over all knots, expected number of random re-triangulations necessary for **SnapPy** to succeed finding a CHS.

	Alternating						Non-Alternating					
#crossings	12	13	14	15	16	17	12	13	14	15	16	17
% Failure on first try	0.9	1.6	2.4	3.3	4.4	5.7	0.8	0.7	1.2	1.7	2.4	3.3
Expected nb of retriang.	2.9	3.9	6.2	7.1	9.5	15.9	2.0	3.1	6.1	11.8	9.7	13.5




■ **Figure 1** Distribution of the number of randomizations required to find a CHS for the complement of the knot “17nh_2654001” of the census, for 10000 tries. The mean is 10.3 randomizations, and the standard deviation is 8.7. The minimum number of re-triangulations is 1, and the maximum 92.

complete hyperbolic metric of the space. A major issue with this approach is that a solution to the constraints may not exist on all triangulations of a manifold, even if, as topological object, the manifold can carry a complete hyperbolic metric. Worst, it is not known whether every hyperbolic 3-manifold admits a triangulation on which a solution as CHS exists.

In practice, given an input knot, software can construct a triangulation of the knot complement, and simplify it to have a CHS. But if this fails, the only implemented practical solution, proposed by **SnapPy**[3], is to randomly modify and simplify the triangulation before trying again until a CHS is found.

Table 1 provides data on the search for a CHS with **SnapPy**, on the ~ 9.7 millions prime knots with crossing numbers up to 17. As observed, **SnapPy** has a high rate of success in finding a CHS after a standard triangulation and simplification of the knot complement. However, this standard construction of a triangulation fails to admit a CHS on more than 350 thousand knots in the census, and the percentage of problematic triangulations needing re-triangulations tends to increase with the number of crossings. Additionally, we observe that some knot complements may require in expectation a high number of random re-triangulations (up to 15.9), and the number of re-triangulations may itself suffer a high variance, as illustrated in Figure 1.



■ **Figure 2** Left: ideal triangulation of the complement of the 8-knot with two tetrahedra. The ideal vertex is truncated, and the red surface gives the torus link of the ideal vertex after gluing of the tetrahedra following the edge identifications. Middle: ideal tetrahedron in the upper half-space model, the dihedral angles are denoted α , β and γ , the complex shape parameter z is associated to the edge between 0 and ∞ . Right: example of shearing singularity, each triangle corresponds to a ideal tetrahedron seen from above (gray arrow).


Checking for the existence of a CHS requires the resolution of the non-linear gluing equations (with, e.g., Newton optimization method). Reducing the number of re-triangulations is consequently critical for performance of computation in knot theory, most notably for knots on which the state-of-the-art methods implemented in `SnapPy` require large numbers of re-triangulations, and even more so when proceeding to very large scale experiments such as the computation of the knot censuses [1] that are of great use to practitioners, where “difficult” knots are many.

Contribution. This paper introduces a new heuristic based algorithm to improve on the random approach for re-triangulating. The method is inspired by Casson and Rivin’s reformulation of the gluing equations [15, 4]: the gluing equations are split into a linear part and a non-linear part, and the resolution reduces to a convex optimization problem on a polytope domain. If the triangulation does not admit a complete structure, the optimization problem will converge on the boundary of the polytope and we exploit this information to introduce new *localized moves* to modify combinatorially the triangulation while reusing the partially computed geometry. We introduce necessary background on triangulations and geometry in Section 2, and the computation of hyperbolic structures with optimization in Section 3. We analyze precisely the behavior of the optimization phase on triangulations not admitting a CHS in Section 4, and introduce a re-triangulation algorithm in Section 5 guided by the partially computed geometry of the optimization phase. We illustrate experimentally the interest of the approach in Section 6 and propose a hybrid method with `SnapPy` in Section 6.2, that outperforms the state-of-the-art.

Note that, in this article we focus on complements of hyperbolic knots. However, the techniques introduced extend to more general hyperbolic 3-manifolds with torus boundaries [4].

2 Background

In this article, we focus on knot complements, i.e., non-compact 3-manifolds obtained by removing a closed regular neighborhood of a knot K in the sphere S^3 .



■ **Figure 3** Illustration of the Pachner moves 2-3 and 3-2.

2.1 Generalized and ideal triangulations

A *generalized triangulation* T is a collection of n abstract tetrahedra whose triangular facets are identified (e.g. Figure 2 (left)), or *glued*, in pairs. Note that the facets of the same tetrahedron may be glued together, and generalized triangulations are more general than simplicial complexes. The link of a vertex in T is the frontier of a closed regular neighborhood, and is itself a closed triangulated surface embedded in the triangulation. If the link of a vertex v is a 2-sphere, we call v an *internal* vertex, otherwise (e.g., if the link is a torus), we call v an *ideal* vertex.

A 1-vertex ideal triangulation is a triangulation with exactly 1 ideal vertex, and no internal vertices. They represent non-compact 3-manifolds, that can be recovered from the 1-vertex ideal triangulation by considering their realization where the vertex has been removed. Every knot complement can be represented by a 1-vertex ideal triangulation, where the link of the ideal vertex gives the frontier of a closed regular neighborhood of the knot. Intuitively, this is a triangulation of the sphere S^3 where the knot has been shrunk into a single point, distorting its neighborhood. Such 1-vertex ideal triangulations of a knot complement can be computed in polynomial time [6, 5] from a planar drawing of a knot.

Any two ideal triangulations of the same 3-manifold can be connected by a sequence of *Pachner moves* [13]. For 1-vertex ideal triangulations, they consist of the moves 2-3 and 3-2, inverse of each other, pictured in Figure 3.

2.2 Combinatorial description of hyperbolic geometry

Certain topological 3-manifolds can be equipped uniformly with a complete hyperbolic metric, which is unique up to isometry. They are called *hyperbolic manifolds*. They include the vast and important family of complements of *hyperbolic knots*.

We use the upper half-space model to represent the hyperbolic space \mathbb{H}^3 . This representation corresponds to $\{(z, r) | z \in \mathbb{C}, r \in \mathbb{R}_+^*\}$ with $\partial\mathbb{H}^3 = \mathbb{C} \cup \{\infty\}$ (from now on denoted $\partial\mathbb{H}^3$) consisting of the bottom plane $\mathbb{C} \times \{0\}$, together with the point at infinity, where geodesics are arcs of circles orthogonal to $\partial\mathbb{H}^3$. This model is *conformal*, i.e., Euclidean angles in the upper half space give the values of the angles in the hyperbolic space.

An *ideal hyperbolic tetrahedron* is the hyperbolic convex hull of four distinct points of $\partial\mathbb{H}^3$. These points are the (ideal) vertices of the tetrahedron.

A vertex on $\partial\mathbb{H}^3$ is a vertex at infinity, thus it is not part of the tetrahedron. An example of ideal tetrahedron is shown in Figure 2 (middle). Up to isometry, the geometric shape of an ideal tetrahedron can be represented by a single complex number:

► **Definition 1** (Shape parameter). *Given an ideal hyperbolic tetrahedron, there exists an isometry sending three of its vertices to 0, 1, and ∞ , and ensuring that the fourth vertex has positive imaginary part in the complex plane in $\partial\mathbb{H}^3 = \mathbb{C} \cup \{\infty\}$. The coordinate z of this fourth vertex is the shape parameter of the tetrahedron.*

The shape parameter defines and illustrates the shape of an ideal tetrahedron. It depends on which vertices are sent to 0, 1, and ∞ . Other permutations of the vertices give other equivalent shape parameters $z' = \frac{z-1}{z}$ and $z'' = \frac{1}{1-z}$ but the underlying tetrahedron is the same. The construction is well defined as isometries of \mathbb{H}^3 are determined by their action on three vertices of $\partial\mathbb{H}^3$.

Another way of characterizing the shape of an ideal hyperbolic tetrahedron is to consider its dihedral angles, i.e., the angle formed by two faces meeting on a common edge; see Figure 2 (middle). In an ideal tetrahedron, opposite angles are equal, and the sum of the six dihedral angles is 2π . We denote in the following the dihedral angles of a tetrahedron by a triplet (α, β, γ) with $(\alpha + \beta + \gamma) = \pi$.

3 Angle structures and hyperbolic volume

In this section, we introduce notions connected to the computation of complete hyperbolic structures on triangulations, via optimization methods. The approach given in this section is another formulation of Thurston's gluing equations, which are non-linear in the complex shape parameters mentioned above, we refer the reader to [4, 15] for more details.

3.1 Linear equations and angle structures

Let T be a 1-vertex ideal triangulation of a knot complement M with n tetrahedra. Since opposite edges have the same dihedral angles, all possible shapes of the tetrahedra can be represented by a vertex in \mathbb{R}^{3n} . We define an *angle structure*:

► **Definition 2** (Angle structures). *Given an ideal triangulation T , an angle structure is a value assignment to the dihedral angles of the tetrahedra of T such that:*

1. *all the angles are strictly positive;*
2. *the three dihedral angles (α, β, γ) of a tetrahedron sum to π ;*
3. *the dihedral angles around each edge of T sum to 2π .*

The set of angle structures on T is denoted $\mathcal{A}(T)$.

Conditions 1 and 2 ensure the angles are in $(0, \pi)$ and that the tetrahedra have the same orientation. Condition 3 is necessary for points on the interior of edges to have a neighborhood isometric to a hyperbolic ball.

Constraints 1, 2, 3 are linear, hence the set $\mathcal{A}(T)$ of angle structures of a triangulation is the relative interior of a polytope in \mathbb{R}^{3n} . It is of dimension $n + |\partial M|$ where $|\partial M|$ is the number of cusps of M ; in the case of knot complements, there is a single cusp. This polytope satisfies:

► **Theorem 3** (Casson; see [8]). *Let T be an ideal triangulation of M , an orientable 3-manifold with toric cusps. If $\mathcal{A}(T) \neq \emptyset$, then M admits a complete hyperbolic metric.*

Furthermore, [4] describes precisely a generating family for the tangent space of $\mathcal{A}(T)$ that can be computed in polynomial time.

3.2 Maximizing the hyperbolic volume

The hyperbolic volume of an ideal hyperbolic tetrahedron with dihedral angles $(\alpha, \beta, \gamma) \in (0, \pi)^3$ is given by the function vol :

$$\text{vol}(\alpha, \beta, \gamma) = L(\alpha) + L(\beta) + L(\gamma)$$

where L is the Lobachevsky function $L(x) = -\int_0^x \log |2 \sin t| dt$. The volume functional can be extended to the whole polytope $\mathcal{A}(T)$ by summing the volumes of the hyperbolic tetrahedra. The following result is due independently to Casson and Rivin.

► **Theorem 4** (Casson, Rivin[14]). *Let T be an ideal triangulation with n tetrahedra of M , an orientable 3-manifold with boundary consisting of tori. Then a point $p \in \mathcal{A}(T) \subset \mathbb{R}^{3n}$ corresponds to a complete hyperbolic metric on the interior of M if and only if p is a critical point of the function vol .*

Additionally, the volume functional is concave on $\mathcal{A}(T)$ and the maximum can be computed via convex optimization methods.

We can finally define CHS that represent combinatorially complete hyperbolic metrics:

► **Definition 5** (Complete Hyperbolic Structure). *A Complete Hyperbolic Structure (CHS) is a triangulation T equipped with an angle structure corresponding to the maximum of vol over $\mathcal{A}(T)$.*

► **Remark 6.** Angle structures are not CHS as they do not prevent shearing singularities, see Figure 2 (right), and consequently may represent *non-complete* hyperbolic metrics.

4 Behavior of the optimization

Exploiting the formalism of the previous section, one can design an algorithm [4] to find and compute a CHS by first finding a point in the polytope $\mathcal{A}(T)$, then computing a basis of the tangent space of $\mathcal{A}(T)$, and then maximizing the (concave) hyperbolic volume functional on this subspace. If $\mathcal{A}(T) \neq \emptyset$ and the procedure finds the point maximizing the volume in the inside of the polytope, then the triangulation admits a complete hyperbolic structure represented by this point. Finally, if $\mathcal{A}(T) \neq \emptyset$ but the maximum of the volume functional is on the boundary, then one needs to re-triangulate the manifold in order to search for a triangulation admitting a CHS.

In this section, we study the outcome of the hyperbolic volume maximization on the space of angle structures. The result is given by the following lemma, which also appeared independently in [12]:

► **Lemma 7.** *Let T be an ideal triangulation of M , a non-compact orientable 3-manifold with toric cusps. Let p be the point maximizing vol over $\overline{\mathcal{A}(T)}$ (the topological closure of $\mathcal{A}(T)$), then at p , if a tetrahedron has an angle equal to 0, then all other angles of the tetrahedron are in $\{0, \pi\}$.*

In other words, either the maximization succeeds on the interior of $\mathcal{A}(T)$, or there is at least one tetrahedron with angles $(0, 0, \pi)$ in p . Such tetrahedron is called *flat*. It is not possible to have a tetrahedron with angles $(0, a, b)$, $(a, b) \in (0, \pi)$ in p .

This lemma is proven by looking at the derivative of the volume around the boundary of $\mathcal{A}(T)$, see [9].

We introduce in the next section an algorithm that performs localized combinatorial modifications on a triangulation equipped with an angle structure, in order to get rid of flat tetrahedra.

5 Localized combinatorial modifications of triangulations

According to Lemma 7, the volume maximization either leads to a solution to the gluing equations, or to flat tetrahedra. In this section, we discuss a method to get rid of flat tetrahedra by combinatorial modifications of the triangulation, while attempting to maintain the value of the volume functional. This would lead to a triangulation admitting an angle structure with larger volume than the previous one, allowing to resume the maximization. In order to maintain the value of the volume functional, we introduce *geometric Pachner moves*.

5.1 Geometric Pachner moves

We perform Pachner moves on the triangulation that preserves the partial geometric data computed. More precisely, we define:

► **Definition 8** (Geometric Pachner move). *A geometric Pachner move in a triangulation T with angle structure is a Pachner move in T such that the resulting triangulation admits an angle structure with identical dihedral angles for the tetrahedra not involved in the move.*

To check if a geometric Pachner move can be done, it must be valid combinatorially, and it should be possible to assign dihedral angles to the new tetrahedra without altering the rest of the angle structure. The combinatorial conditions are simple [13]: for the 2-3 moves (resp. 3-2 move) in Figure 3, the tetrahedra sharing the common triangle BDC (resp. common edge AE) must be distinct. The conditions on the angle structures are given by the following lemma:

► **Lemma 9.** *Given a triangulation with angle structure T , if the combinatorial conditions given above are satisfied:*

- *a 3-2 move is always geometric;*
- *a 2-3 move is geometric if and only if the sum of the two dihedral angles around each edge of the common face is smaller than π .*


The idea to prove the lemma is to study the system of equations given by the assumption that the angle structure is not altered outside the move. Furthermore there are formulas to compute the new angles. The proofs of Lemma 9 and of this later claim are technical and can be found in [9].

► **Remark 10 (Non-conservation of the volume).** In Lemma 9, several angle structures are possible after the 2-3 move. Conversely, several angle structures lead to the same result for the 3-2 move. This is because angle structures are blind to shearing singularities, see Figure 2 (right). If such a singularity exists, a 3-2 move can still be performed while maintaining the whole angle structure, but it will decrease the hyperbolic volume of the structure. It is due to the fact that the singularity, which maximized the volume, is fixed in the process.

5.2 Getting rid of flat tetrahedra

Let T be a triangulation with an angle structure admitting a flat tetrahedron t , and (e, e') the associated edges with dihedral angle π . By Lemma 9, it is not possible to get rid of a tetrahedron with a geometric 2-3 move. Indeed, if an edge of t has a dihedral angle equal to π , the second tetrahedron concerned by the 2-3 move must have a dihedral angle equal to 0, and consequently a 2-3 move will produce a flat tetrahedron (see [9]).

In order to get rid of t , our strategy is to turn either e or e' into an edge of degree three, at the center of three tetrahedra on which a 3-2 move can be performed (central edge of Figure 3, right). Note that edges of degree 2 prevent the existence of an angle structure



■ **Figure 4** A sequence of moves getting rid of a flat tetrahedron. The bold vertical edge is contained in 4 tetrahedra, three of which are represented behind the edge. The fourth tetrahedron is implicit, situated at the front, and flat. First, create the red edge of the first drawing with a 2-3 move. Then delete the red edge of the second drawing with a 3-2 move.

as they force the two tetrahedra sharing the edge to be flat. As a consequence, we assume these configurations are removed, which can be done by `SnapPy`'s simplification for instance, see Section 6. Consequently, all edges have degree at least 3, and our strategy focuses on *reducing* the incidence degree of angle π edges.

To reduce the degree of an edge e , the strategy is to perform 2-3 moves on the tetrahedra containing e , such that each move reduces the degree of the edge by one. Either for topological (tetrahedron glued to itself) or geometrical reasons, these moves will not always be possible, and the order in which they are done matters. This is described in Figure 4.

► **Remark 11.** Doing a 2-3 move does not always reduce the degree of the edge, e.g., when a tetrahedron is represented several times around an edge. However, doing a move that does not decrease the degree of the edge may delete the multiple occurrences of a tetrahedron around the edge and allow to continue with the simplification.


► **Remark 12.** When a 2-3 move is performed around e , the value of the dihedral angle of the new tetrahedron around e is equal to the sum of the previous two dihedral angles around e . Since the sum of the dihedral angles of a tetrahedron is equal to π , this means successfully reducing the degree of e may unlock previously forbidden moves.

Recursive moves. If no geometric Pachner move is possible on an edge e , we attempt to reduce the degree of a nearby edge e_f , i.e., an edge for which there exists a tetrahedron containing e and e_f ; see Figure 5 where e_f is represented by the red dot, as it is seen from above. More precisely, let f be a triangle containing e such that the associated 2-3 move is forbidden for geometric reasons, then by Lemma 9 there is an edge e_f of f for which the associated dihedral angle is at least π .

In consequence, reducing the degree of e_f to three without modifying the tetrahedra containing f and then performing a 3-2 move at e_f will reduce the degree of e by one; see Figure 5.

Note that in case a geometric Pachner move is not possible on any pair of tetrahedra containing e_f either, we can call the procedure recursively in a neighborhood of e_f , using the argument above to pass from e to e_f .

► **Remark 13.** The choice of e_f in f is unique: it is not possible to have two dihedral angles larger than π ; selecting a move to reduce the degree of e boils down to selecting a triangle containing e .



■ **Figure 5** Sequences of moves to reduce the degree of an edge when no 2-3 move is available, seen from above (triangles represent sections of tetrahedra, edge section of triangles and dots section of edges, see Figure 2 (Right)). In all drawings, the gray rectangle represents the flat tetrahedron, e (light blue dot) is the edge of the flat tetrahedron of which we want to reduce the degree. Assuming no 2-3 move is available, we pick another edge e_f (red dot), if a 3-2 move can be performed on e_f , we are in case (B) and the 3-2 move reduces the degree of e (case (C)). Otherwise, we are in case (A) we need to create the red dashed tetrahedron by decreasing the degree of e_f , this can be done by calling recursively our procedure on e_f .

Procedure summary. The procedure to get rid of a flat tetrahedron is a tree-like *backtracking search*, each branch corresponding to a choice of 2-3 or 3-2 move to perform. The moves used are geometric to preserve the volume maximization advancement. A flat tetrahedron can only be deleted with a 3-2 move: the first step is to *reduce the degree* of one of its edges to three. To reduce the degree of an edge, the solution is to perform 2-3 moves on two tetrahedra containing this edge. If a 2-3 move is not possible for geometric reasons, we are in the situation of Figure 5, and the procedure to *reduce the degree* can be called *recursively* on a neighboring edge. Performing a 3-2 move on this latter edge resulting in a decrease of the degree of the initial one. The procedure ends when the initial edge has degree three and the flat tetrahedron can be removed or when no move can be applied.

5.3 Implementation details

The implementation faces several practical challenges.

Breaking infinite loops. As such, the algorithm may loop infinitely on some instance, as 2-3 and 3-2 moves may reverse themselves. To counteract this phenomenon and avoid redundant modifications, a solution is to store at each modification the *isomorphism signature* of the triangulation [2], characterizing uniquely the isomorphism type of the triangulation. This allows us to recognize already processed triangulations and break branches of the backtracking algorithms. Note however that this is costly compared to the Pachner moves, and it makes the procedure no longer local.

Selecting the edge e . When attempting to remove a flat tetrahedron, one needs to choose between the two π -angled edges e and e' to reduce. In our implementation, we consider both edges, and select the one that has the larger smallest angle in its link. This performs better than a random choice in our experiments.

Pruning the backtrack search. Some triangulations may have edges of large degree (more than 10), which produces wide search trees. Additionally, the recursive calls to the degree reduction procedure may induce trees of large depth. Experimentally, the better strategy consists of exploring exhaustively the first levels of the tree. We set the width of exploration to 8 and explore the first 2 levels of the search tree.

Sequencing of Pachner moves. Different sequencings of Pachner moves lead to different triangulations. In practice, we favor 2-3 moves over recursive ones, which performs best. Additionally, among the 2-3 moves performed to get rid of a flat tetrahedron containing edge e of angle π , we prioritize moves that eliminate tetrahedra for which e has a small dihedral angle. Among the recursive ones, we used the same method as the one choosing between e and e' , indeed choosing a recursive move boils down to choosing an edge to minimize its degree.

6 Experiments

In this section, we study the experimental performance of our approach¹ to find a triangulation with a CHS, and compare its behavior with the software `SnapPy` [3]. `SnapPy` is the state-of-the-art software to study the geometric properties of knots and 3-manifolds, and is widely used in the low-dimensional topology community. Following this analysis, we propose and study a hybrid method with practical interest.

SnapPy. `SnapPy`'s method is based on a random re-triangulation followed by a simplification. The first step is constituted of $4n$ random 2-3 moves, where n is the number of tetrahedra in the input triangulation. The simplification performs non-deterministic modifications to decrease the number of tetrahedra in the triangulation. Notably it removes some configurations that prevent angle structures from existing. The verification of the existence of a CHS is based on a Newton's method to solve the gluing equations [17].

Data set. We apply our algorithms to the census of prime hyperbolic knots with up to 19 crossings. This census has been constituted by the efforts of many researchers in the field, and recently completed with the exhaustive enumeration of all knots with crossing number smaller than 20^2 [1].

For each knot, given by a knot diagram, we compute a triangulation of the knot complement using `Regina` [2], and simplify it with `SnapPy`. We then keep the triangulations admitting an angle structure but not a CHS, and hence requiring re-triangulation. All the others, without angle structure or admitting a CHS, were discarded. They constitute the *Failure on first try* data of Table 1.

The knots are grouped by crossing numbers (from 14 to 19), and on whether they are alternating or not. In Section 6.1, our experiments are run on the first 2500 knots of each of the 12 groups having an angle structure but no CHS, and in Section 6.2 we use the first 10.000 knots of the same data sets.


6.1 Success rate and combinatorial performance

In terms of Pachner moves. Figure 6 (left) represents the rate of knots on which our algorithm succeeds to find a triangulation with a CHS, for all the 12 groups of knots. It represents the success rate as a function of the number of Pachner moves, the higher success rate groups are the alternating knots.

While there is a significant difference between the efficiency on alternating and non-alternating knots, all the curves have the same behavior: the first few Pachner moves are very effective. An important point is that, compared to the number of moves done by `SnapPy`,

¹ Our implementation uses the SLSQP [7] optimization method from SciPy [16].

² The census is available at <https://regina-normal.github.io/data.html>



■ **Figure 6** Left: on the sample of triangulations, success rate in finding a CHS against the number of Pachner moves for different crossing numbers and alternabilities. The limit rate of success for alternating knots is 0.98, the limit rate for non-alternating knots is 0.87. Right: Number of re-triangulations required by our method over the number of re-triangulations required by SnapPy on the 12 groups of knots, gathered in bins (the color indicates the log of the number of points inside). Since our method is deterministic, the points have discrete ordinates. The diagonal $x = y$ appears in red.

at least four times the number of tetrahedra and then roughly the same number for the simplification times the number of re-triangulation, our method uses a much lower number of Pachner moves on average (see Figure 6, left). This is of *combinatorial interest*, as doing a small number of moves can mitigate the impact of the procedure on the properties of the triangulation, such as keeping a small treewidth for instance.

After the first few steps, the growth of the success rate slows down drastically. An interpretation of this phenomenon is that the tetrahedra created by the 2-3 moves tend to be more flat than the original ones. This leads to an increase in the number of required re-triangulations and issues with floating point arithmetic.

In terms of re-triangulation. In order to compare the behaviors of our method and of the state of the art: we want to know if the two methods struggle on the same triangulations. Both methods using sequences of re-triangulations, we use the length of these sequences as metrics.

We compare in Figure 6 (right) the average number of randomized re-triangulations required by SnapPy to find a triangulation with a CHS, and the (deterministic) number of re-triangulations of our method. The majority of the CHS are found within few re-triangulations on average: both methods manage to find 78% of them within two steps.

However, it appears that for the manifolds requiring a large number of re-triangulations with either method, the other one will perform well on it: the performance of the methods are substantially orthogonal for the difficult cases. This phenomenon highlights the fact that the methods are *complementary*.

Limitations. It is to be noted that our method suffers from several phenomena: there are manifolds on which our method fails, the reasons to this being convergence problems because of the floating point arithmetic, the creation of flatter tetrahedra or the volume modification mentioned in Remark 10. Furthermore some triangulations have several very high degree edges and looking for the correct sequence of geometric Pachner moves can be costly. In terms of time performance, the constraint convex optimization is an important bottleneck, as well as the computation and book-keeping of a large number of isomorphism signatures.

However, in light of the previous results, we design a hybrid algorithm, mixing `SnapPy`'s pipeline together with the heuristics based on localized geometric Pachner moves introduced in this article, in order to outperform both methods in terms of experimental timings.


6.2 Hybrid algorithm and time performance

As is, `SnapPy` does not produce angle structures as defined in Section 3, as it constructs *negative tetrahedra*, i.e., tetrahedra whose shape has a negative imaginary part. Our hybrid method consists of calling `SnapPy` to randomize and simplify the triangulation, and, in case the triangulation admits only few negative tetrahedra (less than four in our experiments), call the resolution of flat tetrahedra with geometric Pachner moves on these negative tetrahedra, as introduced in Section 5. Loop until a CHS is found.

Our implementation is done in C directly in `SnapPy`'s kernel and consists of a simplified version of Section 5's method. For each negative tetrahedron, we try to reduce the degree of one of its edges down to 3 in turn using only geometric 2-3 moves, we try all the edges, and we do not use recursive moves or heuristics to perform the geometric moves.

The results of this method are summarized in Figure 7, where the time to compute the CHS are compared for the new hybrid method and the usual `SnapPy` pipeline, on the first 10.000 triangulations of the datasets of Section 6. The hybrid method using localized geometric Pachner moves shows, at worst, similar time performance compared to `SnapPy`, and performs significantly better overall. On average, the hybrid method is 20% faster, however, it is to be noted that, on most cases, few re-triangulations are required in which case the hybrid method and `SnapPy` show naturally similar time performance. More interestingly, on cases requiring many more re-triangulations in the `SnapPy` pipeline, the hybrid methods performs much better than `SnapPy`, with running times up to 18 times faster.

As indicator of the global behavior of both algorithms, we indicate the linear regression of the data points (purple dashed line, slope = 0.4) in Figure 7, to illustrate that the dense region is not concentrated on the $x = y$ diagonal.



■ **Figure 7** Left: time comparison of the hybrid method and `SnapPy` in seconds for the 10.000 first triangulations of each crossing number of the data set of the previous section, gathered in bins (the color indicates the log of the number of points inside). Line $x = y$ is plain and red, a linear regression through the data set is dashed and black. For readability, four outliers favoring the hybrid method are omitted. Right: distribution, in logarithmic scale, of the time required by `SnapPy` over the time required by the hybrid method.

References

- 1 Benjamin A. Burton. The Next 350 Million Knots. In Sergio Cabello and Danny Z. Chen, editors, *36th International Symposium on Computational Geometry (SoCG 2020)*, volume 164 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 25:1–25:17, Dagstuhl, Germany, 2020. Schloss Dagstuhl–Leibniz-Zentrum für Informatik. doi:10.4230/LIPIcs.SoCG.2020.25.
- 2 Benjamin A. Burton, Ryan Budney, William Pettersson, et al. Regina: Software for low-dimensional topology. <http://regina-normal.github.io/>, 1999–2021.
- 3 Marc Culler, Nathan M. Dunfield, and Jeffrey R. Weeks. SnapPy, a computer program for studying the geometry and topology of 3-manifolds. <https://snappy.math.uic.edu/>, 1991–2021.
- 4 David Futer and François Guéritaud. From angled triangulations to hyperbolic structures, 2010. doi:10.1090/conm/541/10683.
- 5 Joel Hass, Jeffrey C. Lagarias, and Nicholas Pippenger. The computational complexity of knot and link problems. *J. ACM*, 46(2):185–211, 1999. doi:10.1145/301970.301971.
- 6 William H. Jaco and J. Hyam Rubinstein. 0-efficient triangulations of 3-manifolds. *Journal of Differential Geometry*, 65:61–168, 2002.
- 7 D. Kraft. *A Software Package for Sequential Quadratic Programming*. Deutsche Forschungs- und Versuchsanstalt für Luft- und Raumfahrt Köln: Forschungsbericht. Wiss. Berichtswesen d. DFVLR, 1988.
- 8 Marc Lackenby. Word hyperbolic dehn surgery. *Inventiones mathematicae*, pages 243–282, 2000. doi:10.1007/s002220000047.
- 9 Clément Maria and Owen Rouillé. Computing complete hyperbolic structures on cusped 3-manifolds, 2021. doi:10.48550/ARXIV.2112.06360.
- 10 George D. Mostow. *Strong Rigidity of Locally Symmetric Spaces*. Princeton University Press, 1973. doi:doi:10.1515/9781400881833.
- 11 Hitoshi Murakami and Yoshiyuki Yokota. *Volume Conjecture for Knots*. Springer, 2018. doi:10.1007/978-981-13-1150-5.
- 12 Barbara Nimershiem. Geometric triangulations of a family of hyperbolic 3-braids, 2021. arXiv:2108.09349.
- 13 Udo Pachner. P.l. homeomorphic manifolds are equivalent by elementary shellings. *European Journal of Combinatorics*, 12(2):129–145, 1991. doi:10.1016/S0195-6698(13)80080-7.
- 14 Igor Rivin. Euclidean structures on simplicial surfaces and hyperbolic volume. *Annals of Mathematics*, 139:553–580, 1994.
- 15 W. P. Thurston. *The geometry and topology of 3-manifolds*, volume 1. Princeton University Press, Princeton, N.J., 1980. Electronic version 1.1 - March 2002. URL: <http://library.msri.org/books/gt3m/>.
- 16 Pauli Virtanen, Ralf Gommers, Travis E. Oliphant, Matt Haberland, Tyler Reddy, David Cournapeau, Evgeni Burovski, Pearu Peterson, Warren Weckesser, Jonathan Bright, Stéfan J. van der Walt, Matthew Brett, Joshua Wilson, K. Jarrod Millman, Nikolay Mayorov, Andrew R. J. Nelson, Eric Jones, Robert Kern, Eric Larson, C J Carey, İlhan Polat, Yu Feng, Eric W. Moore, Jake VanderPlas, Denis Laxalde, Josef Perktold, Robert Cimrman, Ian Henriksen, E. A. Quintero, Charles R. Harris, Anne M. Archibald, Antônio H. Ribeiro, Fabian Pedregosa, Paul van Mulbregt, and SciPy 1.0 Contributors. SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python. *Nature Methods*, 17:261–272, 2020. doi:10.1038/s41592-019-0686-2.
- 17 Jeffrey Weeks. Computation of hyperbolic structures in knot theory. *Handbook of Knot Theory*, October 2003. doi:10.1016/B978-044451452-3/50011-3.