The enumeration and classification of prime 20–crossing knots

MORWEN B. THISTLETHWAITE

An account is given of the compilation of the $1,847,319,428$ prime knots with 20 crossings.

57K10; 57-04

1 Introduction

In the summer of 2018 the author tabulated the knots of 20 crossings. An independent tabulation was made simultaneously by B. Burton [Burton 2018], using the software *Regina* developed by him and others [Burton et al. 2023]; as the results of the two tabulations agree there is some confidence that the results are correct, despite the quantity and complexity of the data.

The knots are listed up to unoriented equivalence, that is to say we regard knot pairs (S^3, K), (S^3, L) as *equivalent* if there is a homeomorphism of pairs sending (S^3, K) to (S^3, L), and we list one representative of each equivalence class. The issue of determining which knots are amphicheiral or reversible will be addressed as a separate project.

A short historical note: knot tabulations began in earnest in the late nineteenth century with the work of P.G. Tait [Tait 1896], T.P. Kirkman [Kirkman 1885] and C.N. Little [Little 1885], Tait having being motivated by the (then current) Kelvin theory of vortex atoms. Initially, as Tait was aware, techniques were not available for distinguishing knot types rigorously; these techniques arrived shortly afterwards with the advent of the fundamental group [Poincaré 1895], whereupon M. Dehn and O. Schreier initiated the rigorous classification of knots, beginning with torus knots [Dehn 1914, Schreier 1924]. A fuller account of the history, up to the classification of 16–crossing knots, is given in [Hoste et al. 1998], and to complete the picture B. Burton [Burton 2020] pioneered the classification of knots of 17, 18 and 19 crossings.

A table listing the numbers of prime knots from 3 up to 20 crossings is given in the Appendix of this article.

Theorem 1.1 *The number of equivalence classes of prime knots that can be projected with* 20 *crossings, but not with fewer crossings, is* $1,847,319,428$*. Of these, all but 921 are hyperbolic, the remainder comprising* 915 *satellites of the trefoil knot,* 5 *satellites of the figure-eight knot, and the* $(3,10)$*–torus knot.*

The issue of primality is one that is easy to overlook, but it is important, as one has to guard against "imposter" knots that might be composite in some hidden way and are thus masquerading as prime knots. For this reason a section of this article is devoted to justifying the claim that all listed knots are prime.

2 The tabulation

2.1 Obtaining the raw list of knots

For the most part the method is the same as that employed in [Hoste et al. 1998], albeit with some minor differences. Traditionally all tabulations of knots with given crossing-number begin with a listing of the prime alternating knots with that number of crossings. This is one of the few steps in the process that is truly algorithmic: from the solutions of the various Tait conjectures, it is known that a knot with a reduced alternating n–crossing diagram cannot be projected with fewer than n crossings. Also, an alternating knot is guaranteed to be prime if its reduced alternating diagrams have the property that they do not admit a simple closed curve in the projection plane meeting the knot projection transversely in two points on distinct edges of the projection (Fig. 1). [Menasco 1984]. Furthermore, any two reduced alternating diagrams represent the same link type if and only if one can transform one to the other by means of a finite sequence of *flypes* (Fig. 2).

It is precisely the failure of non-alternating knots to adhere to such desirable properties that renders their classification a challenge.

It is relatively straightforward to write a program that generates all possible reduced alternating diagrams of a given crossing-number n, choosing a representative from each flype equivalence class, although skill is required in devising a program that will run in a reasonable time. This has indeed been accomplished quite dramatically for $n \leq 23$ [Rankin et al. 2004]. The present author has written a program that generates all prime alternating links with a given number of crossings, and for the case $n = 20$ the number of prime alternating knot types turns out to be $199,631,989$.

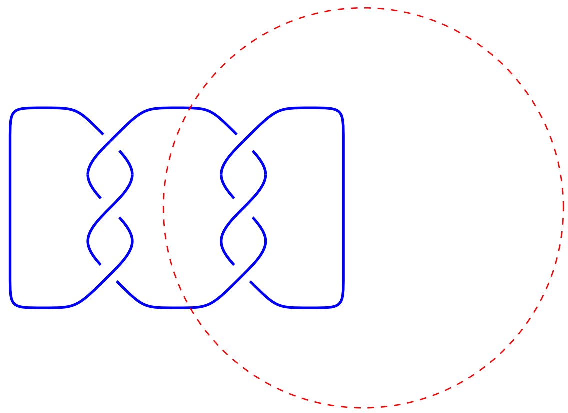

Figure 1: A composite alternating knot

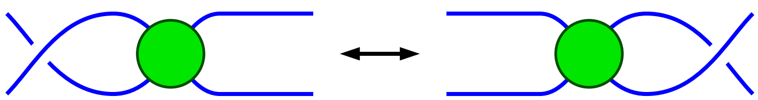

Figure 2: The "flype" transformation

Once one has the list of alternating diagrams to hand, non-alternating diagrams can be obtained from them by means of crossing switches. It is only necessary to take one alternating diagram from each flype equivalence class, as if alternating diagrams D_1, D_2 are flype-equivalent, then the diagrams obtained from crossing switches of D_1 are flype-equivalent to those obtained from D_2. Switching all crossings of a knot diagram produces a knot equivalent to the original on account of being its reflected image in the projection plane, so a crude estimate of the number of diagrams to be generated in this way from a single alternating diagram is $2^{19} = 524288$. However, it is only on rare occasions that this number is needed, as can be seen from the following observation: if a rational tangle diagram [Conway 1967] is not alternating, then there exists an isotopy of the tangle that reduces the number of crossings while keeping the four ends of the tangle fixed (Fig. 3). Therefore built into the program is a procedure that detects all nontrivial rational tangle substituents, and then we only allow crossing switches of the "base" alternating knot diagram that keep each of these tangles alternating.

The resulting non-alternating diagrams are subjected to a number of rapid viability

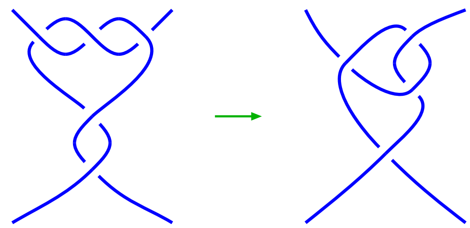

Figure 3: Reducing a non-alternating diagram of a rational tangle

tests to check whether the number of crossings can be reduced, and are immediately discarded upon failing any such test. A surviving diagram is then subjected to a different kind of test, specifically to see whether it can be transformed by flypes and passes (Figs. 2, 4) to a diagram whose DT code [Dowker and Thistlethwaite 1983] is lexicographically less. As the size of an equivalence class generated by these moves can be very large, even in the tens of thousands, we declare that the diagram passes the test if it is still lexicographically minimal once some fixed number k of diagrams has been generated by the moves. Smaller values of k will entail larger redundancy, but it makes sense to keep k quite small on account of the time that would be spent on processing a large set of diagrams.

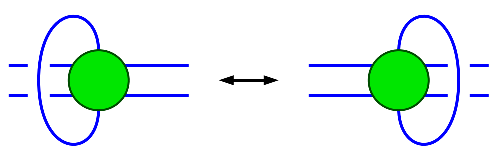

Figure 4: The "pass" transformation on diagrams

In [Hoste et al. 1998] some diagram moves more "exotic" than flypes and passes were used, but this approach was avoided here as it was deemed unnecessary, quite apart from the increased danger of introducing bugs into the program. In practice the value chosen for k was 200, and from experimentation with known tabulations with fewer crossings it was estimated that this resulted in roughly 25% redundancy overall.

2.2 Removing duplicate knots

The next task is to augment the current list of over 2 billion non-alternating diagrams by all tabulated knots with fewer crossings, and then to remove as many duplicates as possible. The extreme difficulty of achieving this "house cleaning" simply by inspecting or manipulating knot diagrams is wonderfully exhibited by the celebrated *Perko pair* [Perko 1974], a pair of knots with only 10 crossings declared to be inequivalent in C.N. Little's 1900 table, this status persisting until 1974 when K. Perko finally spotted the equivalence, thereby obtaining the first correct table of 10–crossing knots.

The 1970's also saw R. Riley's discovery of a hyperbolic structure on the complement of the figure-eight knot [Riley 1975], this being one of the inspirations for W. Thurston's breakthrough work on geometric structures on 3–manifolds. This in turn led to J. Weeks's extensive program SnapPea [Weeks 1989] and its more recent Python implementation SnapPy [Culler et al. 2007], one of whose many features is the ability to compute the *canonical cell decomposition* [Epstein and Penner 1988, Sakuma and Weeks 1995] of a hyperbolic 3–manifold with genus 1 cusps.

The preimage of this cell structure in the universal cover can either be seen in the upper half-space model as dual to the Ford domain, or it can be seen by means of a convex hull construction in the Minkowski model. SnapPea performs a very rapid computation of a purported canonical cell decomposition by starting with a known ideal triangulation of the manifold and then implementing a heuristic optimization process that applies combinatorial moves on the triangulation without affecting the underlying topology. Because of inevitable accumulation of roundoff error, the resulting cell decomposition might on occasion not be the canonical one, but nonetheless if two hyperbolic knots produce isomorphic cell decompositions, their respective complements are proved to be homeomorphic, and from the fact that knots are determined by their complements [Gordon and Luecke 1989], the knots are equivalent. In practice, even at the level of 20 crossings this is an effective way of removing duplicates, which otherwise could be very hard to spot.

Indeed, the current list of over 2 billion 20–crossing non-alternating diagrams was fed through SnapPea's canonical cell decomposition procedure, and the few hundred million diagrams producing duplicate cell decompositions were discarded. During this process approximately a mere 549491 were declared (with due caution) by SnapPea to be "apparently not hyperbolic", and these were copied to a separate list. The next stage is to try to distinguish by means of invariants the knots in the filtrate. There is no algorithm at work here, as there does not seem to be any way of predicting which invariants will distinguish which knots: we just throw invariants at the knots and hope for the best. However, the method is rigorous, as all computations of invariants are integer based.

2.3 Application of invariants to the list of knots

2.3.1 Description of the invariants

The first invariant applied to the remaining diagrams was the Jones polynomial, for which we are fortunate in having the very fast program of [Ewing and Millett 1991], which on a single processor of my workstation will process a million 20-crossing knots in $2\frac{1}{2}$ minutes. This partitions the set of diagrams into relatively small equivalence classes, such that within each equivalence class all knots have the same Jones polynomial. The knots in equivalence classes of size 1 are extracted and placed in the "resolved" folder, and the remaining knots are subjected to a sequence of further invariants, the aim being to subject the partition to successive refinements so that eventually all equivalence classes have size 1.

The remaining invariants are classical, and occur in papers of R.H. Fox [Fox 1962] and K. Perko [Perko 1976]. They rely on the fact that knot groups seem always to have an abundance of subgroups of small index. It follows from the work of W. Thurston that knot groups are residually finite, but this alone does not explain why the knots in our 20–crossing list are so rich in subgroups of index less than 10. Given a subgroup H of index n of a knot group G, the group G acts transitively by left multiplication on the set of n left cosets of H, giving rise to a transitive permutation representation of degree n of G. Conversely, given a transitive permutation representation of degree n of G, the stabilizers of the n symbols are conjugate subgroups of G of index n. Using the Reidemeister–Schreier rewriting process we can obtain a presentation of such a subgroup H, and by Abelianization obtain a finitely generated Abelian group, which is essentially the first homology group of the covering space of the knot complement corresponding to H. We can also glue in solid tori to this covering space so that the components of the preimage of a meridian curve are spanned by cross-sectional disks, thus obtaining the first homology group of the so-called *branched covering space*.

The technique is to choose a transitive permutation representation of some group, for example the natural representation of degree 5 of the alternating group A_5, and then find all homomorphisms of the knot group onto that group of permutations, up to composition with inner automorphisms of the image group. The multiset of Abelian groups thus obtained is then an invariant of that knot type, and amazingly, together with the Jones polynomial it was possible in this way to distinguish almost all listed 20–crossing knots from one another and from knots with fewer crossings, using only subgroups of the symmetric group S_7.

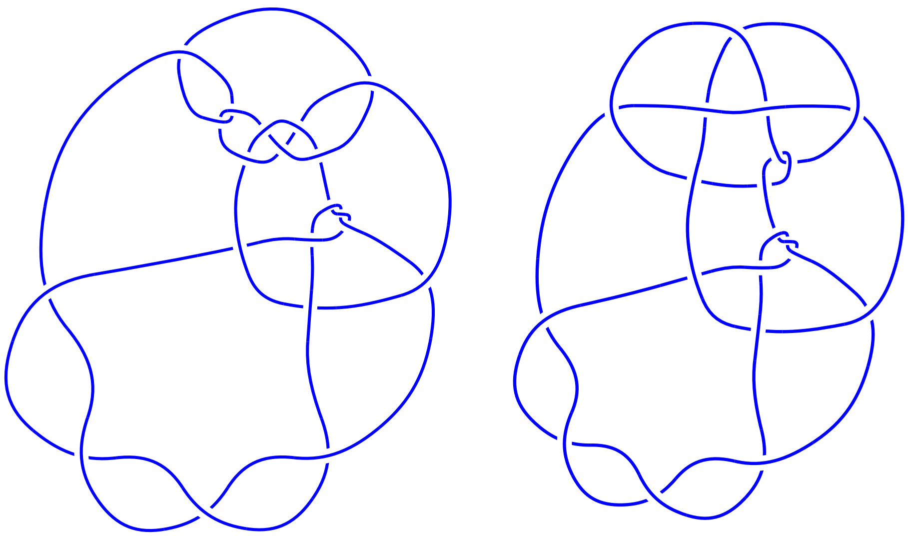

Figure 5: A pair K_1, K_2 of mutant 20–crossing knots

Here are two examples of this type of invariant applied to a fairly resistant mutant pair of 20–crossing knots.

Each diagram in Figure 5 consists of an upper tangle glued to a lower tangle along four strands; the second diagram can be obtained from the first by excising the upper tangle, rotating it through a half turn in the projection plane and then gluing it back to the lower tangle.

Being related by mutancy, these knots cannot be distinguished by the Jones polynomial nor indeed by the HOMFLYPT 2–variable polynomial; also they resisted homology groups associated with permutation representations of degrees 5 and 6. However they did succumb to permutation representations mapping meridians to one of the two conjugacy classes of 7–cycles in the alternating group A_7. For K_1 there were 14 representations, producing homology groups with torsion numbers as follows:

$$[1083964, 14], [10873394], [117987912], [1308356, 2, 2], [13423592, 8],$$
$$[155682849, 3], [2496669], [30245222, 2], [353577, 7], [477902327],$$
$$[4832310], [58290239, 7], [8694588], [909657, 7]$$

and for K_2 just 13 representations:

$$[10007522, 2], [1339604, 14], [20281751], [21298634], [24072097], [2742502, 2], [304197488, 2], [40220460], [46137, 21], [4719806], [53620280], [56118930, 3], [6282066, 2]$$

Since the group A_7 admits an automorphism sending each 7–cycle to its inverse, this conjugacy class cannot be used to detect nonreversibility of knots. It was observed by H. Trotter [Trotter 1964] that a more careful choice of target group can be effective for this purpose; indeed this was the first occasion that the existence of nonreversible knots was proved. Later R. Hartley [Hartley 1983] used (solvable) groups of functions $x \mapsto ax + b \ \ (a \neq 0)$ over finite fields to establish nonreversibility of many knots of up to 10 crossings. For the knots K_1, K_2 of Figure 5, the sporadic Mathieu simple group M_{11} is effective in showing that they are not reversible. Specifically, we can use the irreducible permutation representation of M_{11} of degree 11, and map meridians to one of the two conjugacy classes of size 990 containing elements of cycle type $(ab)(cdefghij)$. Here are the results, with torsion numbers for each representation enclosed in square brackets as above:

K_1	[1394030, 2]
Reverse of K_1	[287520], [65322]
K_2	[14118592], [5682, 2]
Reverse of K_2	[1598572], [4161904]

It is expected that one can determine reversibility in this way for the list of 20–crossing knots, although it could be very time consuming. Determining amphicheirality is in practice easier, as almost all instances of non-amphicheirality are detected by the Jones polynomial.

2.3.2 Details of the invariants' performance

There now follows details of the efficacy of the invariants used for distinguishing the 20–crossing knots from one another and from knots with fewer crossings. The enumeration stage of the classification process generated 2229828372 20–crossing nonalternating diagrams; barring programming error there were no omissions in this list, and the procedure was in three main stages as described below.

Stage 1. The raw list of 20–crossing nonalternating diagrams was augmented by the list of all 352151858 hyperbolic knots with fewer than 20 crossings, resulting in an enlarged list containing 2581980230 knots. SnapPea's canonical cell decomposition procedure was applied to each knot in the enlarged list, and the data was sorted so that duplicate cell decompositions became evident. Each cell decomposition was encoded by a string of approximately 300 bytes on average, so the amount of data involved in this step was around 775 GB. As explained in Section 2.2, it is not guaranteed that the cell decompositions output by this procedure are canonical in the sense of [Epstein and Penner 1988, Sakuma and Weeks 1995], but knots with duplicate cell decompositions have homeomorphic complements, so are equivalent owing to the fact that knots are determined by their complements [Gordon and Luecke 1989].

The canonical cell decomposition procedure declared that 549491 knots from the list were "apparently not hyperbolic" and these were put in a separate list for further treatment. Aggressive diagram moves revealed that out of these knots 200 were the unknot, 547611 were composite knots, and a further 482 could be drawn with fewer than 20 crossings. This left a residue of 1198 knots, which on being treated to still more stringent diagram moves were shown to belong to 921 knot types, distinct from one another and distinct from all nonhyperbolic knot types with fewer than 20 crossings. The proof that this list consisted of a single 20–crossing torus knot and 920 20–crossing satellite knots is given in Section 3.

After removing the 549491 knots declared to be "apparently not hyperbolic" and the knots whose complements had duplicate cell decompositions, the number of knots in the refined list was 1999847149. These were input into the next stage, it being expected that the only duplications were those arising from roundoff error in application of the canonical cell decomposition procedure.

Stage 2. From an accounting point of view this was the easiest stage. The Jones polynomials of the 1999847149 knots output by the previous stage were computed, and the 336548774 knots with unique polynomials were extracted and placed in the store of "resolved" knots . The remaining 1663298375 knots were input into Stage 3, which subjected them to the invariants described above, namely first homology groups of branched covering spaces corresponding to permutation representations of the knot groups.

Stage 3. Tables 1, 2 below summarize the results of this stage. The first table uses representations into alternating or symmetric groups of degrees $5, 6$, and the number of unresolved knots was reduced from 1663298375 to 728749. The column labelled "# unique" gives the number of knots distinguished from all others and placed into the "resolved" store, and the column labelled "# nonunique" gives the number of unresolved knots requiring further treatment. The column labelled "cycle type" gives the cycle type of the conjugacy class to which meridians of the knot group were mapped.

The machine used for these computations had 160 GB of memory and 20 processing cores.

The remaining knots were then subjected to permutation representations in various specific groups, as set out in Table 2. Each of these substages took less than a day of runtime.

degree	cycle type	# input	# unique	# nonunique	runtime
5	$(abcde)$	1663298375	906980266	756318109	20 days
5	$(abc)(de)$	756318109	317431388	438886721	4 days
5	$(abcd)$	438886721	309112549	129774172	5 days
6	$(abc)(def)$	129774172	66784736	62989436	6 days
6	$(abcde)$	62989436	57189475	5799961	6 days
6	$(abcdef)$	5799961	5071212	728749	1 day

Table 1

group	degree	cycle type	# input	# unique	# nonunique
PSL(2, 7)	7	7–cycles	728749	572093	156656
PSL(2, 11)	11	11–cycles	156656	117446	39210
PSL(2, 13)	14	13–cycles	39210	15364	23846
PSL(2, 17)	18	17–cycles	23846	5245	18601
A_7	7	7–cycles	18601	1071	17530
PSL(2, 19)	20	19–cycles	17530	2	17528

Table 2

At this point the list of 17528 unresolved knots were partitioned into 8755 equivalence classes, where knots within each equivalence class had resisted all invariants applied to date. It was suspected that each of these in fact represented a single knot type, and this was confirmed by a more persistent application of the canonical cell decomposition procedure: SnapPea has a convenient “random retriangulation” feature, and from this a small number of different contenders for canonicity were obtained, amongst which matching cell decompositions were found in each of the outstanding cases.

Surprisingly the last two knots to be distinguished, in the last row of Table 2, were a pair of 14–crossing two-bridged knots, with associated fractions $\frac{505}{192}$, $\frac{505}{212}$ and respective Conway codes 2111221112, 2211111122. These are easily distinguished by the fact that they are alternating, and also by their lens space two-fold branched covers, but for some reason they resisted polynomial invariants and the homology invariants of Tables 1, 2 until the very last step.

This concluded the task of obtaining a list of 20–crossing knots with no omissions or duplications, but it was still necessary to check that there were no "poseur" composite knots in the list. One almost hoped that some would materialize, as such examples would be noteworthy.

3 Establishing primality

A fundamental property of a hyperbolic 3–manifold is that it cannot contain an essential torus. The software *Regina* [Burton et al. 2023] confirmed that all presumed $1,847,318,507$ hyperbolic 20–crossing knots in our list are indeed hyperbolic (*i.e.* there are no false positives with respect to hyperbolicity) so all are immediately known to be prime.

A different approach is needed for showing that the 921 (apparently) non-hyperbolic knots are prime. The single torus knot was easily identified, and since torus knots are prime we may restrict our attention to the remaining 920 knots in this list.

We recall the terminology of [Lickorish 1981]. In that paper a *tangle* is defined to be a pair $(B\,,\,T)$ where B is a 3–ball, *i.e.* a manifold with boundary homeomorphic to the standard 3–ball B^3, and T is a proper 1–submanifold of B consisting of two disjoint arcs (naturally we assume that we are in the piecewise linear or smooth category). Thus the boundary of T consists of four points on ∂B. This definition has some obvious generalizations, for example we might allow the number of arcs in T to be greater than 2, but the definition as given is sufficient for our purposes. Tangles $(B_1\,,\,T_1)$, $(B_2\,,\,T_2)$ are *equivalent* if there is a homeomorphism of pairs from $(B_1\,,\,T_1)$ to $(B_2\,,\,T_2)$, and $(B\,,\,T)$ is *untangled* or *trivial* if it is equivalent to a product $(D\,,\,\{x,y\}) \times I$, where D is a disk and x, y are points in its interior.

In the definition of tangle equivalence given above, for a homeomorphism $h : (B\,,\,T_1) \to (B\,,\,T_2)$ there is no restriction on the effect of h on the boundary 2–sphere of B, other than the requirement that it map ∂T_1 to ∂T_2. For example, any tangle represented as a diagram of a *rational tangle* [Conway 1967] is equivalent to a tangle where T consists of two parallel line segments, *i.e.* it is trivial.

A tangle $(B\,,\,T)$ is *locally unknotted* if each 2–sphere in B meeting T transversely in two points bounds in B a ball meeting T in an unknotted spanning arc. Otherwise we say that $(B\,,\,T)$ is *locally knotted*; an example is illustrated in Figure 4(i). Observe that if (i) $(B\,,\,T)$ is locally unknotted, and (ii) there exists a properly embedded disk in B separating the arcs of T, then $(B\,,\,T)$ is trivial.

Given tangles $(B_1\,,\,T_1)$, $(B_2\,,\,T_2)$, we may glue them together by means of some homeomorphism of (2–sphere , four points) pairs to obtain a link L of one or two components in the 3–sphere. Such a pair $(S^3\,,\,L)$ is called a *sum* of the tangles $(B_1\,,\,T_1)$, $(B_2\,,\,T_2)$. Given two tangles drawn in the usual way as diagrams, one way of summing them is to join the diagrams by arcs in the projection plane.

If $(B\,,\,T)$ contains a 2-sphere S exhibiting local knottedness, with knotted arc α in the ball bounded by S, then there is a well-defined non-trivial knot K obtained by joining the ends of α with an arc in S, and this knot K will persist as a connected summand of any knot formed by summing $(B\,,\,T)$ with an arbitrary tangle. Therefore if we can sum $(B\,,\,T)$ with a tangle so as to obtain the unknot, $(B\,,\,T)$ is locally unknotted. In more complicated situations we have the following effective test for local unknottedness:

Proposition 3.1 *Let $(B\,,\,T)$ be a tangle for which there exist tangles $(B_1\,,\,T_1)$, $(B_2\,,\,T_2)$, such that summing $(B\,,\,T)$ with the $(B_i\,,\,T_i)$ in turn produces distinct prime knots K_1 , K_2. Then $(B\,,\,T)$ is locally unknotted.*

Proof If $(B\,,\,T)$ is locally knotted, there exists a non-trivial knot that is a connected summand of each of the distinct prime knots K_1 , K_2, and this contradicts uniqueness of factorization of knots. □

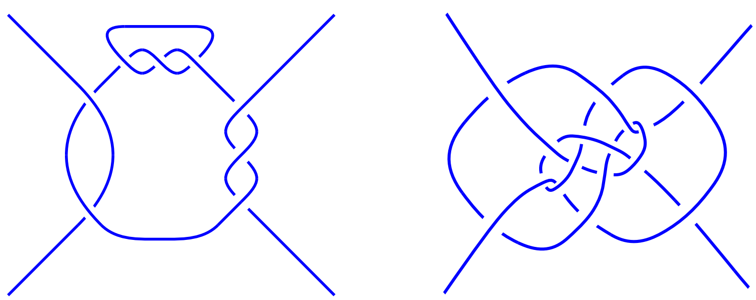

Figure 6: (i) A locally knotted tangle and (ii) a companion tangle

There is a special type of tangle that pertains to satellite knots. We define a *companion tangle* to be a tangle $(B\,,\,T)$ where T consists of two parallel, knotted arcs in B; an example is illustrated in Figure 6(ii). Each companion tangle $(B\,,\,T)$ contains a properly embedded annulus A in $B - T$ that "follows" the two strands of T in tube-like fashion, *i.e.* there is a homeomorphism $h\,:\,S^1 \times I \to A$ such that each section $h(S^1 \times \{t\})$ of A bounds a disk in B meeting T transversely in two points. A companion tangle cannot be locally knotted, as it is always possible to sum it with a trivial tangle so as to obtain the unknot; in Figure 6(ii) this can be seen by taking two arcs in the projection plane, one joining the two left-hand ends and the other the two right-hand ends.

Let us now consider a knot (S^3, K) constructed as a sum of a companion tangle (B_1, T_1) with a locally unknotted tangle (B_2, T_2). We may form a torus F in the complement of K as the union of the "following" annulus A_1 of (B_1, T_1) described above, with a boundary-parallel annulus A_2 in (B_2, T_2) that "swallows" T_2. Let V be the solid torus containing K that is bounded by F; V is the union of two "halves" $V \cap B_1$, $V \cap B_2$ glued together along cross-sectional disks D_1, D_2, both in $\partial B_1 = \partial B_2$, and each meeting K in two points. The core Γ of V is a non-trivial knot in S^3, as it it the union of a knotted arc in B_1 that is the core of A_1 with an unknotted arc in B_2 that is the core of A_2.

Any cross-sectional disk of V not meeting K would have to separate the strands of the second tangle (B_2, T_2), but the local unknottedness of (B_2, T_2) would force that tangle to be trivial, and we would be in the situation described above where K is the unknot. On the other hand, if there is no cross-sectional disk of V separating the strands of (B_2, T_2), K is a satellite of Γ and the torus F is incompressible in $S^3 - K$.

The next theorem provides the method for showing that the 920 outstanding knots are prime. It is closely related to results in [Schubert 1953, Lickorish 1981, Cromwell 2004], as explained below; however the full proof is given here as the hypotheses are an exact fit to our situation, and moreover should be applicable to future tabulations with more than 20 crossings.

Theorem 3.2 *Let K be a knot that is a sum of a companion tangle with a locally unknotted tangle. If K is non-trivial, then K is prime.*

Proof We adopt the notation of the preceding discussion: (S^3, K) is a non-trivial knot that is the sum of a companion tangle (B_1, T_1) with a locally unknotted tangle (B_2, T_2), $F = A_1 \cup A_2$ is the incompressible torus in $S^3 - K$ that "follows" T_1 and "swallows" T_2, and V is the solid torus with boundary F.

Let S be a 2–sphere in S^3 meeting K transversely in two points. Before proceeding further it is useful to observe that each simple closed curve C in $S - K$ is either null-homotopic in $S - K$ (hence also null-homotopic in $S^3 - K$), or else it separates the punctures. In particular, a circle on F that bounds a cross-sectional disk of V cannot lie on S.

We first consider the special case where S is contained in V. Since by hypothesis each of the constituent tangles (B_1, T_1), (B_2, T_2) is locally unknotted, the conclusion of the theorem holds for S contained in either "half" $V \cap B_i$, and we are motivated to consider the transverse intersection of S with the two disks D_1, D_2 along which the halves of V are glued together. We may assume that the two points of $S \cap K$ are away from the D_i. The set $S \cap (D_1 \cup D_2)$ is the union of a disjoint collection of circles on S; let C be a circle from this collection that is innermost on S, say without loss of generality $C \subset S \cap D_1$. Then C bounds a disk $\Delta_1 \subset S$ and a disk $\Delta_2 \subset D_1$. The union of the disks Δ_i is an embedded 2–sphere Σ, bounding a ball B' contained in one of the B_i. The number n of points of $\Delta_2 \cap K$ is $0, 1$ or 2, and we consider each of these cases. We can exclude the possibility $n = 2$ summarily, as in this case C would be a simple closed curve on S homotopic in $S^3 - K$ to a meridional curve of $F = \partial V$, a situation ruled out in the previous paragraph.

Suppose that $n = 0$; then Δ_1 meets K in 0 or 2 points. In the latter case, from the hypothesis of local unknottedness applied to Σ, the ball B' would meet K in an unknotted arc; we deduce from this that one of the components of $S^3 - S$ would meet K in this arc, and the conclusion of the theorem would follow. Otherwise the ball B' does not meet K. The circle C might not be innermost on D_1, but nonetheless Δ_1 can be pushed by an isotopy through B', taking with it all components of $S \cap B'$, reducing the number of components of $S \cap (D_1 \cup D_2)$. If $n = 1$, then from the hypothesis of local unknottedness B' meets K in an unknotted arc, so again there is an isotopy that pushes Δ_1 across B', including if necessary another component of $S \cap B'$ meeting K in one point. Here the isotopy will move points of K along the unknotted arc, but can be assumed to fix K setwise. We conclude that there is an isotopy of S into one of the B_i without affecting transversality of $S \cap K$, whence S bounds a ball on one side meeting K in an unknotted arc, and the conclusion of the theorem follows for this special case.

For the remainder of the proof we assume that S has non-empty transverse intersection with $F = \partial V$; the proof will be completed by showing that there an isotopy of S in S^3, maintaining transversality of S with K, that moves S to a 2–sphere contained in V.

Recall that a simple closed curve in the twice-punctured sphere $S - K$ is either null-homotopic in $S - K$ or is homotopic to a meridian of curve of K. The torus F does not contain any simple closed curve of the second type, so each component C of $S \cap F$ is a simple closed curve bounding a disk in $S - K$, and also bounding a disk in F owing to the incompressibility of F. Let us take a component C of $S \cap F$ bounding a disk $\Delta_1 \subset S - K$ whose interior does not meet F; also let Δ_2 be the disk on F bounded by C. Then $\Delta_1 \cup \Delta_2$ is an embedded 2–sphere in $S^3 - K$, and in a manner similar to that of the special case we can perform an isotopy of S that reduces the number of components of $S \cap F$. Repeating the process will eventually move S into V, and the proof of the theorem is complete. □

There is overlap between Theorem 3.2 and results in the literature, most notably H. Schubert's paper [Schubert 1953] where the notion of companionship tree of a knot is introduced, and where it is shown that doubled knots and cabled knots are prime.

In [Lickorish 1981] a tangle is called *prime* if it is non-trivial and locally unknotted, and it is proved in that paper that a sum of two prime tangles is a prime link. A companion tangle is prime according to this definition, as it is locally unknotted, and cannot be trivial as its individual strands are knotted arcs. Theorem 3.2 shows that, apart from the obvious single exception, a knot formed as a sum of a companion tangle with a trivial tangle is also prime, thus confirming a special case of the conjecture stated in §4 of [Lickorish 1981].

Theorem 4.4.1 of [Cromwell 2004] deals more generally with primality of satellite knots; it includes the hypothesis that the pair $(V\,,\,K)$ is locally unknotted, this being the conclusion of the special case dealt with in the proof of Theorem 3.2.

Recall that there are 920 knots in our tabulation that are under examination for primeness. It was suspected that five of these are satellites of the figure-eight knot, and these were easily found in the list; they are all obtained by summing the companion tangle of Fig. 4(ii) with a 4–crossing rational tangle. Since they are already known to be non-trivial, and rational tangles are certainly locally unknotted, application of Theorem 3.2 shows that they are prime.

Naturally one suspects that the remaining 915 knots are satellites of the trefoil knot. In order to apply 3.2 we need diagrams of these knots that show each as a sum of a companion tangle with a locally unknotted tangle. Undoubtedly it would be possible to find such diagrams directly; however, a different approach was used here. The list of $199,631,989$ prime *alternating* 20–crossing knots provides, up to flype equivalence, all projections of prime non-alternating knots, and an easy search through this list found 434 projections of tangle sums of the required kind. Suitable over- and under-passes were applied to these, resulting in a refined list of 915 knot diagrams that (i) visibly were sums of tangles $(B_1\,,\,T_1)\,,\,(B_2\,,\,T_2)$ with $(B_1\,,\,T_1)$ a companion tangle, and (ii) matched the tabulated 915 knots.

As the knots were already known to be non-trivial it remained to check that in each case the tangle $(B_2\,,\,T_2)$ was locally unknotted. In all but ten cases, verification was immediate, as the diagrams of $(B_2\,,\,T_2)$ are either alternating, in which case they are subject to Theorem 1 of [Menasco 1984], or they are standard diagrams of arborescent tangles [Bonahon and Siebenmann 1979].

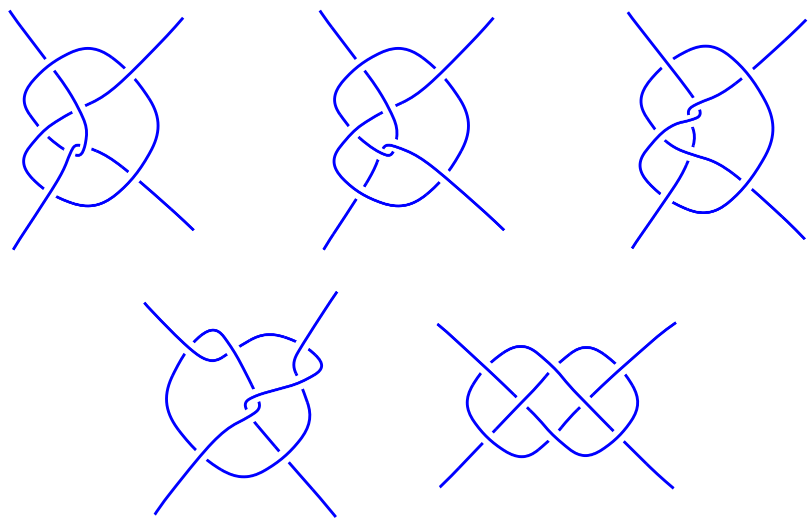

Figure 7

The ten exceptional cases come in five pairs, each pair consisting of a diagram and its reflection in the projection plane. It was only necessary to check one tangle from each pair, and they are illustrated in Figure 7.

Under mild scrutiny the individual strands of the third, fourth and fifth tangles of Figure 7 are all revealed to be unknotted, so local knots for these tangles are ruled out. One can also notice that the first two tangles are equivalent: there is a homeomorphism that interchanges the lower two tangle ends. Therefore, in order to complete the proof that all 920 satellite knots are prime, we just need to check that the first tangle is locally unknotted. This follows quickly from Proposition 3.1: summing (in the "diagrammatic" sense) with a tangle with no crossings produces a prime 8–crossing knot, and one can obtain a prime 9–crossing knot, also a prime 10–crossing knot by summing with a 2–crossing tangle.

4 Acknowledgments

The author would like to thank the referee and the editors for useful suggestions that resulted in a significantly improved exposition.

It is a pleasure to acknowledge the support of Ben Burton in this endeavour; a certain error in a line of the author's earlier code would probably have persisted, had he not

been able to compare his results for 17 crossings with Ben's. Historically tables of knots have been error-prone, but the probability of errors is greatly reduced by having independent tabulations to hand.

5 Appendix

5.1 The rate of growth of the number of knots

It is natural to ask whether one can estimate the number of prime knots with a given number of crossings without an actual tabulation. There are very few results in existence on this topic, but the following is known. It is stated for links rather than knots, but it suggests that the number of non-alternating knots grows exponentially faster than that of alternating knots.

Theorem 5.1 **(i)** *[Sundberg and Thistlethwaite 1998] Let A_n denote the number of prime alternating link types with n crossings. Then*

$$\lim_{n\to\infty} A_n^{1/n} \;=\; \frac{101+\sqrt{21001}}{40} \;\approx\; 6.1479$$

(ii) *[Thistlethwaite 1998] Let λ be the limit stated in (i). There exists a set $\mathcal{B}$ of prime links, strictly containing the set of prime alternating links, such that if B_n is the number of links in $\mathcal{B}$ with n crossings, then $\lim_{n\to\infty} B_n^{1/n}$ exists and is strictly greater than λ.*

5.2 The number of prime knot types with n crossings, $3 \leq n \leq 20$

The first correct tabulations of knots of $17, 18$ and 19 crossings were produced by Ben Burton [Burton 2020].

When reading Table 3, it is worth noting that the only prime, alternating, non-hyperbolic knots are the $(2,n)$–torus knots (with n necessarily odd) [Menasco 1984]. Thus for even n all n–crossing prime alternating knots are hyperbolic, and for odd n there is a single non-hyperbolic prime alternating knot, namely the $(2,n)$–torus knot. Also it follows that all prime satellite knots are non-alternating.

# crossings	# alt. knots	# nonalt. knots	# hyperbolic	# torus	# satellites
3	1	0	0	1	0
4	1	0	1	0	0
5	2	0	1	1	0
6	3	0	3	0	0
7	7	0	6	1	0
8	18	3	20	1	0
9	41	8	48	1	0
10	123	42	164	1	0
11	367	185	551	1	0
12	1288	888	2176	0	0
13	4878	5110	9985	1	2
14	19536	27436	46969	1	2
15	85263	168030	253285	2	6
16	379799	1008906	1388694	1	10
17	1769979	6283414	8053363	1	29
18	8400285	39866181	48266380	0	86
19	40619385	253511073	294130212	1	245
20	199631989	1647687439	1847318507	1	920

Table 3

References

[Bonahon and Siebenmann 1979] Bonahon, F. and Siebenmann, L. (1979) New geometric splittings of classical knots and the classification and symmetries of arborescent knots. *Preprint, University of Southern California Preprints.*

[Burton et al. 2023] Burton, B. and Budney, R. and Pettersson, W. et al. (2023) *Regina, software for low-dimensional topology, http://regina-normal.github.io/* , 1999–2023.

[Burton 2018] Burton, B. (2018) *Private communication.*

[Burton 2020] Burton, B. (2020) The next 350 million knots. 36th International Symposium on Computational Geometry (SoCG 2020), https://doi.org/10.4230/LIPIcs.SoCG.2020.25

[Conway 1967] Conway, J.H. (1967) An enumeration of knots and links, and some of their algebraic properties. *Computational Problems in Abstract Algebra* (Proc. Conf., Oxford, 1967), 329–358, Pergamon, Oxford.

[Cromwell 2004] Cromwell P. (2004). Knots and Links. *Cambridge University Press*, 2004.

[Culler et al. 2007] Culler, M., Dunfield, N. and Goerner, M. (2007). SnapPy. *Computer Software, snappy.computop.org.*

[Dehn 1914] Dehn, M. (1914). Die beiden Kleeblattschlingen. *Math. Ann.* 75 (1914), 402–413.

[Dowker and Thistlethwaite 1983] Dowker, C.H. and Thistlethwaite, M.B. (1983). Classification of knot projections. *Topology Appl.* 16 (1983), no. 1, 19–31.

[Epstein and Penner 1988] Epstein, D. and Penner, R. (1988) Euclidean decompositions of noncompact hyperbolic manifolds. *J. Differential Geom.* 27 (1988), no. 1, 67–80.

[Ewing and Millett 1991] Ewing, B. and Millett, K. (1991) A load balanced algorithm for the calculation of the polynomial knot and link invariants. *The Mathematical Heritage of C.F. Gauss*, 225–266, World Sci. Publ., River edge, NJ, 1991.

[Fox 1962] Fox, R. (1962) A quick trip through knot theory. *Topology of 3-Manifolds and Related Topics (Proc. The Univ. of Georgia Institute, 1961)* 120–167, Prentice-Hall, Englewood Cliffs, NJ.

[Gordon and Luecke 1989] Gordon, C. and Luecke, J. (1989) Knots are determined by their complements. *J. Amer. Math. Soc.* 2 (1989), no. 2, 371–415.

[Hartley 1983] Hartley, R. (1983) Identifying non-invertible knots. *Topology* 22 (1983), no.2, 137–145.

[Hoste et al. 1998] Hoste, J. Thistlethwaite, M., Weeks, J. (1998) The first 1,701,936 knots. *Math. Intelligencer* 20 (1998), no. 4, 33–48.

[Kirkman 1885] Kirkman, T.P. (1885) The enumeration, description and construction of knots of fewer than ten crossings. *Trans. Roy. Soc. Edinburgh* 32 (1885), 483–506.

[Lickorish 1981] Lickorish, W.B.R. (1981) Prime Knots and Tangles. *Trans. Amer. Math. Soc.* 267 (1981), no.1, 321–332.

[Little 1885] Little, C.N. (1885) On knots, with a census of order ten. *Trans. Connecticut Acad. Sci.* 18 (1885), 374–378.

[Menasco 1984] Menasco, W. (1984) Closed incompressible surfaces in alternating knot and link complements. *Topology* 23 (1984), no.1, 37–44.

[Perko 1974] Perko, K. (1974) On the classification of knots. *Proc. American Math. Soc.* 45 no. 2 (1974), 262–266.

[Perko 1976] Perko, K. (1976) On dihedral covering spaces of knots. *Inventiones Math.* 34, 77–82 (1976).

[Poincaré 1895] Poincaré, H.. (1895) Analysis situs. *Journal de l'Ecole Polytechnique* (2) 1, 1–123 (1895).

[Rankin et al. 2004] Rankin, S., Flint, O., Schermann, J. (2004) Enumerating the prime alternating knots, I. *J. Knot Theory Ram.* 13 (2004), no. 1, 57–100.

[Riley 1975] Riley, R. (1975) Discrete parabolic representations of link groups. *Mathematika* 22 no. 2 (1975), 141–150.

[Sakuma and Weeks 1995] Sakuma, M., Weeks, J. (1995) The generalized tilt formula. *Geom. Dedicata* 55 (1995), no. 2, 115–123.

[Schreier 1924] Schreier, O. (1924) Über die Gruppen $A^a B^b = 1$. *Abh. Math. Sem. Univ. Hamburg* 3 (1924), 167–169.

[Schubert 1953] Schubert, H. (1953) Knoten und Vollringen. *Acta Math.* 90 (1953), 131–286.

[Sundberg and Thistlethwaite 1998] Sundberg, C., Thistlethwaite, M. (1998) The rate of growth of the number of prime alternating links and tangles. *Pacific J. Math.* 182 (1998),no. 2, 329–358.

[Tait 1896] Tait, P.G. (1896) On knots I, II, III. *Scientific Papers, Vol. 1, Cambridge University Press* (1896), 273–347.

[Thistlethwaite 1998] Thistlethwaite, M. (1998) On the structure and scarcity of alternating links and tangles. *J. Knot Theory Ram.* 7 (1998),no. 7, 981–1004.

[Trotter 1964] Trotter, H. (1964) Non-invertible knots exist. *Topology* 2 (1964), 275–280.

[Weeks 1989] Weeks, J. (1989) SnapPea. *Computer Software, www.geometrygames.org.*

Department of Mathematics, University of Tennessee, Knoxville, TN 37996, USA

morwen@math.utk.edu

www.math.utk.edu/~morwen