Visualization of the Genus of Knots

Jarke J. van Wijk* Arjeh M. Cohen†

Dept. Mathematics and Computer Science
Technische Universiteit Eindhoven

ABSTRACT

The genus of a knot or link can be defined via Seifert surfaces. A Seifert surface of a knot or link is an oriented surface whose boundary coincides with that knot or link. Schematic images of these surfaces are shown in every text book on knot theory, but from these it is hard to understand their shape and structure. In this paper the visualization of such surfaces is discussed. A method is presented to produce different styles of surfaces for knots and links, starting from the so-called braid representation. Also, it is shown how closed oriented surfaces can be generated in which the knot is embedded, such that the knot subdivides the surface into two parts. These closed surfaces provide a direct visualization of the genus of a knot.

CR Categories: H.5.2 [Information Interfaces and Presentation]: User Interfaces; I.3.8 [Computer Graphics]: Applications

Keywords: Visualization, knot theory, topology, genus, Seifert surfaces

1 INTRODUCTION

To introduce the topic discussed in this paper, we start with a puzzle. Consider a trefoil, the simplest knot (fig. 1). It is easy to define a surface that has this knot as its boundary: Take a strip, twist it three times, and glue the ends together. If we try to color the sides of the surface differently, we see that there is something strange. The strip is a kind of Möbius strip, and cannot be oriented, there is only one side. The puzzle now is to define an orientable surface that has the trefoil as its boundary.

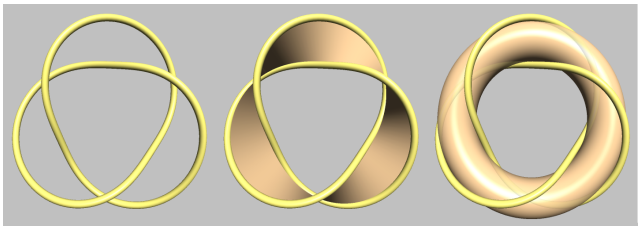

Figure 1: Trefoil

A second puzzle. It is easy to embed a trefoil in a closed surface: A trefoil is a so-called torus knot. However, this knot does not divide the torus into two parts, contrary to what one might expect from local inspection. Can we embed the knot on a closed surface, such that it divides this surface into two parts?

The first puzzle has been solved in 1930 by Frankl and Pontrjagin [6], who showed that such a surface can be found for any knot.

*e-mail: vanwijk@win.tue.nl
†e-mail: A.M.Cohen@tue.nl

Oriented surfaces whose boundaries are a knot K are called Seifert surfaces of K, after Herbert Seifert, who gave an algorithm to construct such a surface from a diagram describing the knot in 1934 [12]. His algorithm is easy to understand, but this does not hold for the geometric shape of the resulting surfaces. Texts on knot theory only contain schematic drawings, from which it is hard to capture what is going on. In the cited paper, Seifert also introduced the notion of the genus of a knot as the minimal genus of a Seifert surface. The present paper is dedicated to the visualization of Seifert surfaces, as well as the direct visualization of the genus of knots.

In section 2 we give a short overview of concepts from topology and knot theory. In section 3 we give a solution for the second puzzle: We show how a closed surface can be constructed in which a knot is embedded, such that it divides the surface in two parts. A Seifert surface consists of disks and bands, such a closed surface consists of spheres and tubes. In section 4 we discuss how these elements can be derived and positioned from an abstract notation of a knot. In section 5 we present how surfaces can be generated. Examples of results are shown in section 6. Finally, in section 7 the results are discussed and suggestions for future work are made.

2 BACKGROUND

In this section we informally enumerate a number of definitions and concepts from topology in general and knot theory in particular. We limit ourselves to those results that are directly relevant for the work presented here. More information can be found in several books, and also on the Web many resources are available. *The Knot Book* [1] of Colin Adams gives a very readable and accessible introduction for non-experts, more depth can be found in [10, 7, 8].

Knot theorists have enumerated knots by means of diagrams or braid words, with invariants like the genus for distinguishing them. Results can be found in the literature and on the Web. The Knot Atlas of Bar-Natan provides many tables of knots and invariants [2]; the KnotInfo table [9] of Livingston and Cha was a very valuable resource for us.

2.1 Topology

Knot theory is a subfield of topology. Topology is the mathematical study of the properties of objects that are preserved through deformations of objects. Two surfaces are *homotopic* if each one can be continuously deformed into the other. If this can be done without passing the object through itself, they are not only homotopic but also *isotopic*. For instance, a torus is isotopic (and hence also homotopic) to a cup with one handle, and homotopic to a tube in the shape of a trefoil.

Two surfaces are homotopic when three conditions are satisfied. First of all, either both should be orientable or neither; secondly, the number of boundary components must be the same; and finally, the Euler characteristic χ must be the same. The *Euler characteristic* χ is an invariant for surfaces. Given an arbitrary (but regular) polygonalization of a surface, $\chi = V - E + F$, with V the number of vertices, E the number of edges, and F the number of faces. Closed oriented surfaces are homotopic to a sphere with g handles

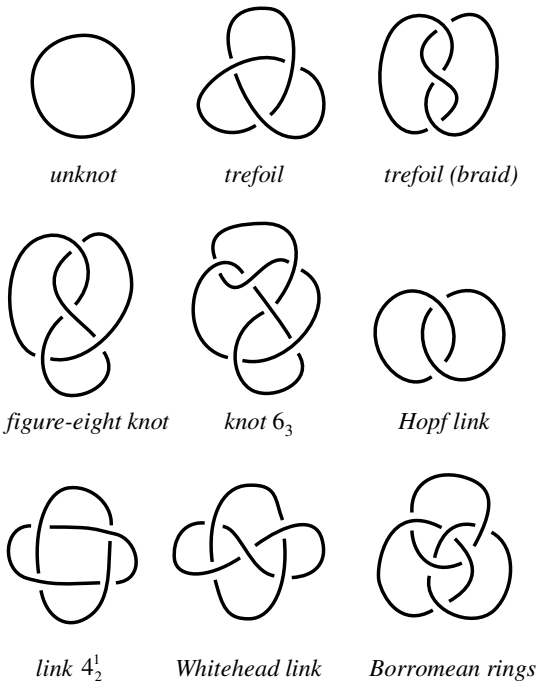

Figure 2: Link diagrams

(or, equivalently, a donut with g holes). The number g is called the genus of the surface. The genus of surfaces with boundaries is defined to be equal to that of the surface that results when all boundaries are capped off with (topological) disks. For surfaces with m boundaries $\chi = 2-2g-m$.

2.2 Knot theory

Knot theory studies the properties of mathematical knots and links. A mathematical knot is a tamely embedded closed curve embedded in $\mathbb{R}^3$. Here an embedding of a closed curve is called tame if it can be extended to an embedding of a tube around the curve. A link consists of multiple components, each of which is a knot. A knot or link can be continuously deformed as long as it does not intersect itself. The result of such a deformation is a knot isotopic to the original one. Up to isotopy, a knot can be represented by a non-intersecting closed polyline (finite sequence of line segments in three-space).

Knots and links are usually studied using *projections* or *diagrams*, such as shown in fig. 2. One knot can be projected in many different ways, as an example two different projections of the trefoil are shown. A diagram consists of edges and crossings. If an orientation is assigned to the knot, we see that two different types of crossings exist: right-hand crossings and left-hand crossings (fig. 3).

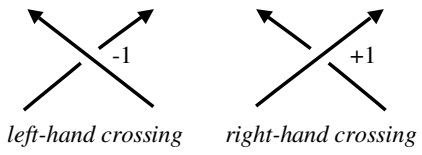

Figure 3: Two different types of crossings

Some main questions in knot theory are whether two knots are the same or not, and especially if a knot is equal to the unknot; how many different knots do exist (given constraints), and how to classify knots. One approach to this is to define invariants of knots. A classic one is the minimum number of crossings in a diagram of a knot; more powerful and distinctive are so-called knot polynomials, such as the Jones polynomial [1].

2.3 Seifert surfaces

The genus of a knot, introduced by Seifert [12], is another classic invariant in knot theory. The Euler characteristic for a 1-dimensional object is 0 when applied to a knot, hence that does not lead to a distinction. Seifert therefore used a connected, oriented, compact surface that has the knot as its boundary to define the genus of a knot. At first sight, it is surprising that such a surface exists for any knot or link. Seifert showed that such a surface can be derived from a knot diagram using a simple algorithm. It consists of four steps (fig. 4). First of all, assign an orientation to the components of the knot or link. Secondly, eliminate all crossings. At each crossing two strands (say, A and B) meet. A crossing is eliminated by cutting the strands, and connecting the incoming strand of A with the outgoing strand of B, and vice versa. This gives a set of non-intersecting (topological) circles, called Seifert circles. Thirdly, if circles are nested in each other, offset them in a direction perpendicular to the diagram. Fill in the circles, giving disks. Finally, connect the disks using twisted bands. Each band corresponds to a crossing, and has one twist, with orientation derived from the crossing type. A twist is a rotation over plus (right-hand) or minus (left-hand) 180 degrees. The resulting surface satisfies the requirements. Different projections of the knot lead to different surfaces, possibly also with a different genus. The genus of a knot is defined as the minimal genus of all surfaces (Seifert surfaces) bounded by the knot.

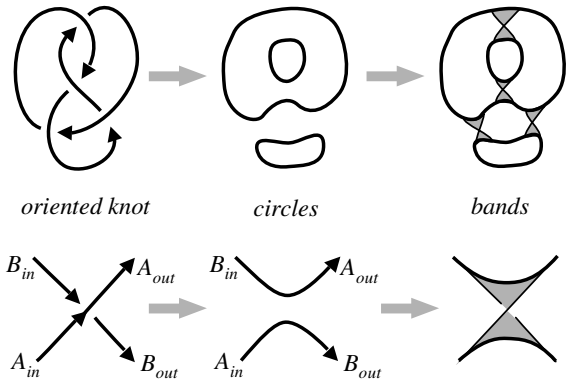

Figure 4: Seifert's algorithm: Assign orientation, eliminate crossings, and add bands; shown for a knot and a crossing

2.4 Challenge

Texts on knot theory show figures similar to fig. 4. From these it is hard to understand the shape of the surface. One reason is that such surfaces are not familiar and are rarely encountered in the real world. We have searched the literature and the Web, but could not find satisfying visualizations of Seifert surfaces. The KnotPlot package of Robert Scharein [11] has a very rich set of features and is a delight to work (and play) with, but even this has no option to show Seifert surfaces. We therefore found it a challenge to develop a method to visualize Seifert surfaces. Specifically, our aim was to enable the viewer to interactively generate and view Seifert surfaces in 3D for arbitrary knots and links in different styles.

One possible route to this is to consider a Seifert surface as a minimal surface (i.e., the surface with zero mean curvature, also known as the soap bubble surface) using the knot as boundary. However, this requires that a three-dimensional knot is available. Also, definition of a suitable initial surface mesh and the iterative calculation of the minimal surface are not easy to implement and are compute intensive. We therefore opted for a different approach. Given an abstract notation of a knot, derive the structure of the Seifert surface and find a smooth geometry in a quick and deterministic way.

2.5 Braid representation

To generate Seifert surfaces for arbitrary knots and links, we need a definition of these knots and links. Many different notations have been developed, such as the Conway notation and the Dowker-Thistlethwaite notation. For our purposes we found the braid representation to be very useful. Using this, several different styles of surfaces can easily be generated, and also, the braid representation lends itself well to experimentation. It does have its limitations though, as we discuss in section 6.

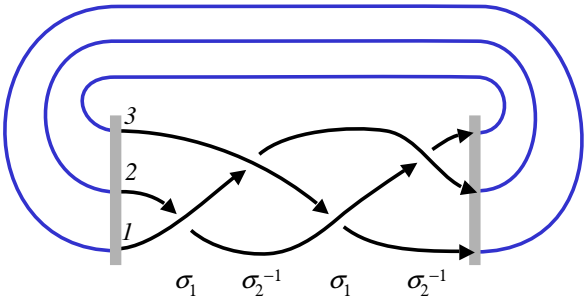

Figure 5: Braid representation

A braid consists of a set of n strings, running (here) from a left bar to a right bar (fig. 5). Strings are allowed to cross, and the pattern can be encoded by enumerating the crossings from left to right. A crossing is denoted by σ_k^j, which means that strings at the k'th and $k+1$'th row are twisted j times, where $j=1$ denotes a right crossing and $j=-1$ a left crossing. The closure of the braid is defined by attaching the left bar to the right bar, such that no further crossings are introduced. In other words, we add n extra strings that connect the beginnings and ends of strings at the same row, without further crossings. Every knot and link can be defined as a braid. A trefoil has the *braid word* $\sigma_1\sigma_1\sigma_1 = \sigma_1^3$, a figure eight knot can be represented as $\sigma_1\sigma_2^{-1}\sigma_1\sigma_2^{-1}$. An alternative notation for braids is to use uppercase letters for right crossings and lowercase letters for left crossings, and where the character denotes the strings effected, according to alphabetic order. Hence, a trefoil is encoded by *AAA*, and a figure eight knot by *AbAb*. Furthermore, every braid word defines a knot or a link, which makes this representation well suited for experimentation.

3 Closed surfaces

Besides visualization of Seifert surfaces, another aim was to make the genus of a knot 'more visible'. A trefoil or a figure eight knot has genus 1, hence the corresponding Seifert surfaces are homotopic to a torus with a hole in the surface. Via a number of steps in which the Seifert surface is deformed, cut, and glued, this equivalence can be shown, but this is not really intuitive. Closed surfaces are easier to understand, hence we studied how a closed surface can be generated that contains the Seifert surface as an embedded subsurface. We call such a surface a closed Seifert surface. The following reasoning is straightforward, but we could not find it in the literature.

The standard approach of topologists is to cap off boundaries (here the m boundaries of the Seifert surface) with (topological) disks. This does lead to a surface that is homotopic to a closed surface, but not isotopic. What we need here to close the surface in a more decent way, is an oriented surface that has the m components of the link as boundary. But this is exactly the definition of a Seifert surface itself, which leads us immediately to a solution. Using a physical analogy, the solution is to take two identical Seifert surfaces, glue them together at the boundaries, and inflate the closed object. This is shown in fig. 6 for a trefoil (which also shows a possible solution to the puzzles posed in the introduction). The Seifert surface consists here of two disks, connected by three bands; the closed Seifert surface consists of two spheres, connected by three tubes. The knot splits the closed surface into two parts.

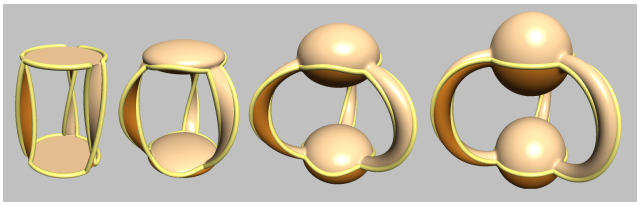

Figure 6: Inflating two Seifert surfaces, glued together at their boundaries

The genus of a closed Seifert surface can be determined as follows. The Euler characteristic of a Seifert surface is $\chi_s = 2 - 2g_s - m$, with g_s the genus and m the number of components of the knot. For the Euler characteristic χ_c of the closed surface we find $\chi_c = 2\chi_s$: The number of vertices, faces and edges doubles, but at the boundaries a certain number of edges and the same number of vertices disappear. However, as V and E have opposite signs in the definition of χ, this does not influence the resulting value. For a closed surface $\chi_c = 2 - 2g_c$, with g_c the genus of the closed surface. This leads to

$$g_c = 2g_s + m - 1.$$

This gives us a direct way of finding oriented closed surfaces in which to embed a knot or link of genus g_s such that the knot divides it into two parts. For instance, for a trefoil or figure eight knot a genus 2 surface can be used (such as a donut with two holes, or two spheres connected by three tubes), and in greater generality, for a knot of one component a donut with twice the number of holes as the genus can be used.

4 Structure

In this section we derive the structure of the Seifert surfaces, starting from the braid word. The aim here is to determine the number of disks (or spheres) and their position in space, and the bands (or tubes), with the number of twists and attachment position to the disks as attributes. The disks are positioned in 3D (x,y,z) space. We take x and y in the plane of the diagram, and z perpendicular to the plane. Disks are parallel to the x,y plane. Each disk has two sides, which we denote with A and B. For each disk a decision must be made if the A or B side is positioned upwards.

Because of the regular structure of braids, various styles of Seifert surfaces can easily be derived from these. Fig. 7 shows four styles for a figure eight knot, using ellipsoids and tubes. First, the *stacked* style. If all closing strings are positioned in the default way,

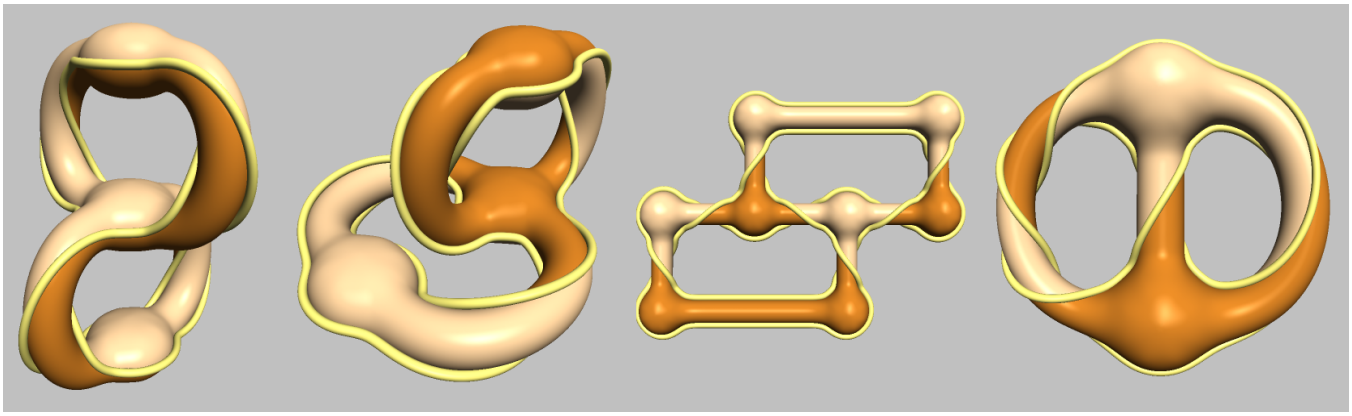

Figure 7: Figure eight knot in stacked, split, flat, and reduced style

it is easy to see that the Seifert circles are all nested. Hence, the corresponding Seifert surface consists of a stack of disks, where each disk is connected with bands to its neighbors (fig. 8). All disks have the A side facing upwards, their position is $(0,0,(i-1)D)$, where i is the index of the row to which the disk corresponds, and D a distance between the disks. A nice geometric representation is obtained by subdividing each disk into k sectors, where k is the total number of crossings, and connecting sectors of neighboring disks with bands when appropriate. Using a suitable setting for the geometry, an object similar to a wedding-cake is generated.

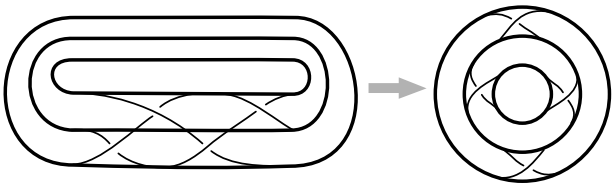

Figure 8: Standard braid representation gives stacked disks

As a variation on this, one set of closing strings can be positioned above, and the remaining set can be positioned below the braid. This gives the *split* style: two sets of stacked disks in wedding-cake style, where the lowest disks of each set are connected by bands in the plane. One set has the A side facing up, the other set has the B side upwards. As an example, in fig. 9 two strings are positioned above and one is positioned below the braid. We introduced this style in order to produce for instance the Seifert surface that results from the standard projection of the figure eight knot.

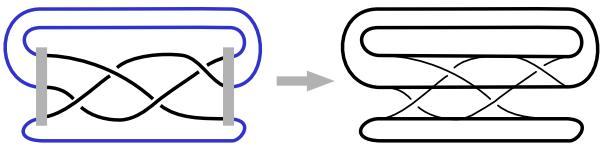

Figure 9: Split style

An alternative style, the *flat* style, is obtained as follows. The upper closing string is positioned above of the braid, the lowest closing string below the braid, and the closing strings in between are put downward, pushed perpendicular to the plane of the braid. Strings of the last kind introduce extra crossings. Their number can be minimized by carefully choosing the path of the string (fig. 10). From this lay-out of the strings, disks and bands can be derived using Seifert's algorithm. Thus a set of non-nested, disjoint Seifert circles will be obtained, so they can be positioned in a plane. The structure can be constructed as follows. Suppose that σ_k^j is the i-th crossing. We add two disks, one with A up (brown) at position $(iD,kD,0)$ and one with B up (yellow) at $(iD,(k+1)D,0)$. In other words, at each upper and lower triangle of an original crossing disks are positioned. Next, vertical bands are added that represent the original crossings, with a twist according to the crossing. Finally, horizontal bands are added between disks on the same row. If both disks have the same side up, no twists are added, otherwise a single positive or negative twist is used, dependent on the order of A and B.

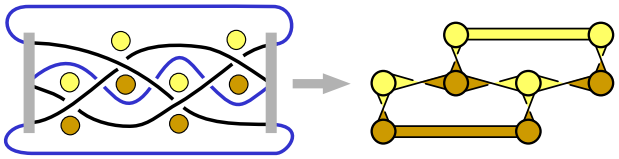

Figure 10: Flat style

The flat style is not particularly interesting, but this planar lay-out can be simplified further, giving the more attractive *reduced* style. Several disks have only two bands attached to them. Such a disk can be removed, and the original two bands can be replaced by a single band, with the number of twists equal to the sum of the number of twists of the original bands. Application of this rule to the figure eight knot leads to a simple structure of two disks, connected by three bands with $1,1$ and -3 twists respectively (fig. 11). Such a knot, with a Seifert surface that consists of two disks, connected by parallel twisted bands, is known as a pretzel knot. The trefoil is a $(1,1,1)$ pretzel knot. The structure of the pattern of

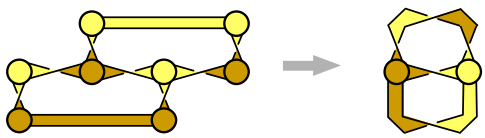

Figure 11: Reduced style

disks and bands can be described as a planar graph, with each disk mapped to a vertex, each band to an edge, and each hole to a face. For the optimal lay-out of such graphs a number of algorithms exist [4]. We implemented a simplistic one (using a trial-and-error approach), which gave satisfactory results for the graphs produced here.

5 GEOMETRY

In the previous section we have discussed how to generate disks and bands from a braid word, and how to position and orient the disks. The next step is to produce a surface to visualize the Seifert surface or the corresponding closed surface. We use ellipsoids as the basic shape for disks and spheres, and curved cylinders with an elliptical cross-section for the bands and tubes. These are approximated with polygons. Furthermore, smoothing can be applied to obtain a more smooth surface.

5.1 Ellipsoids

In standard position, an ellipsoid with two axes of equal length (representing a squeezed sphere) can be described by

$$\mathbf{p}(\mathbf{u}) = (d\cos u\cos v,\ d\sin u\cos v,\ h\sin v)/2,$$

with spherical coordinates $\mathbf{u} = (u,v) \in [-\pi,\pi)\times[-\pi/2,\pi/2]$, and with the diameter d and height h as parameters. Obviously, setting d close to zero gives a disk, setting $d = h$ gives a sphere. The ellipsoid is subdivided into n_s sectors, where each sector has at most one tube attached. Consider one such a sector $(u,v) \in [-U,U]\times[-V,V]$, where $U = \pi/n_s$ and $V = \pi/2$. The top half ($v \in (0,V]$) belongs to either A or B, the bottom part belongs to the other part of the surface. If no band is attached, then this sector can be straightforwardly polygonized with a rectangular mesh. The vertices are $\mathbf{p}_{ij} = \mathbf{p}(\mathbf{u}_R(i,j))$, with

$$\mathbf{u}_R(i,j) = (Ui/I, Vj/J)$$

and $(i,j) \in [-I,I]\times[-J,J]$. Obviously, the vertices at the poles coincide. If a band or tube is attached, a hole must be made in

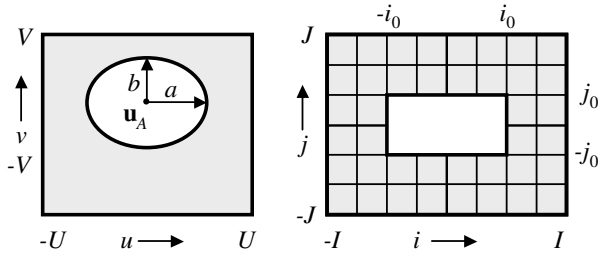

Figure 12: Sector of ellipsoid in (u,v) and (i,j) coordinates

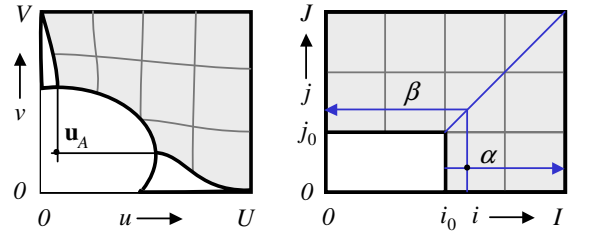

Figure 13: Upper right quadrant in (u,v) and (i,j) coordinates

this mesh, and some care is required to make sure that this hole conforms with the end of the tube. The cross-section of bands and tubes is described as an ellipse, with width w and height d. Obviously, setting w close to 0 gives a band, setting $w = d$ gives a tube. Suppose that the attachment point of the centerline of the tube is $\mathbf{p}_A = \mathbf{p}(\mathbf{u}_A)$. Typically, $u_A = 0$, and v_A is an optional offset in the direction of the poles to move the attachment point closer to the disk to which the other side of the tube points. This was used for instance in fig. 6. We model the boundary of the hole in the ellipsoid in spherical coordinates as

$$\mathbf{u}_B(s) = \mathbf{u}_A + (a\cos s\pi/2,\ b\sin s\pi/2),$$

with $s \in [0,4)$ (fig. 12). The lengths of the semi-axes a and b are chosen such that the distances of $\mathbf{p}(\mathbf{u}_B(0))$ and $\mathbf{p}(\mathbf{u}_B(1))$ to $\mathbf{p}(\mathbf{u}_A)$, measured along the surface of the ellipsoid, match with $w/2$ and $d/2$ respectively. This hole is a perfect ellipse in (u,v) space, and, for our purposes, a good enough approximation of an ellipse in 3D space.

Also, we define a rectangular hole in the mesh space: $(-i_0,i_0)\times(-j_0,j_0)$. The mesh has to be warped such that the inner boundary conforms with the hole in the ellipsoid, while the outer boundary still conforms with the standard boundary of the sector. We have modeled this as follows. Consider the upper-right quadrant of the sector (fig. 13). We measure the position of a mesh-point (i,j) in a kind of polar coordinates (α,β), where $\beta \in [0,1]$ denotes the angle, and $\alpha \in [0,1]$ denotes if we are close to the inner boundary ($\alpha = 0$) or the outer boundary ($\alpha = 1$). Specifically, we use

$$\alpha_{ij} = \max(\alpha_i,\alpha_j) \quad \text{with}$$

$$\alpha_i = \frac{i-i_0}{I-i_0} \quad \text{and} \quad \alpha_j = \frac{j-j_0}{J-j_0};$$

and

$$\beta_{i,j} \quad = \quad \begin{cases} j/L(i,j) & \text{if } \alpha_i > \alpha_j \\ 1-i/L(i,j) & \text{otherwise} \end{cases}$$

with

$$L(i,j) = (1-\alpha_{ij})(i_0+j_0) + \alpha_{ij}(I+J).$$

If only the hole has to be taken care of, mesh points can be found using

$$\mathbf{u}_C(i,j) = \mathbf{u}_A + \begin{array}{l}((\alpha_{ij}a + (1-\alpha_{ij})(U-u_a))\cos\beta_{ij}\pi/2,\\ (\alpha_{ij}b + (1-\alpha_{ij})(V-v_a)))\sin\beta_{ij}\pi/2).\end{array}$$

To obtain a smooth transition from the inner to the outer boundary, we determine the vertices $\mathbf{p}_{ij} = \mathbf{p}(\mathbf{u}_H)$ by blending circular and rectangular coordinates via

$$\mathbf{u}_H(i,j) = (1-h(\alpha_{ij}))\mathbf{u}_C(i,j) + h(\alpha_{ij})\mathbf{u}_R(i,j) \quad \text{with}$$

$$h(t) = -2t^3 + 3t^2.$$

The blending function $h(t)$ gives a smooth transition at the boundaries because $h'(0) = 0$ and $h'(1) = 0$. The other quadrants are dealt with similarly. A result is shown in fig. 14.

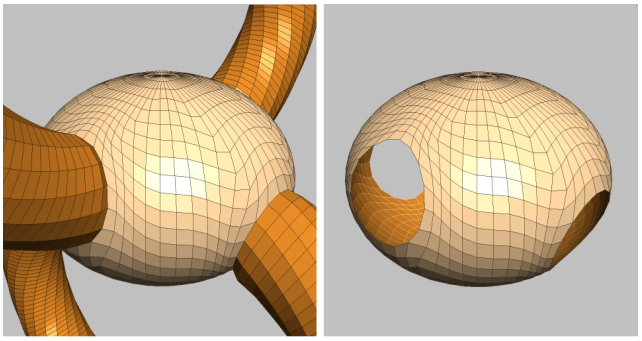

Figure 14: Mesh of ellipsoid

5.2 Tubes

The tubes are also modeled via a rectangular mesh of polygons. We use a mesh $\mathbf{c}_{ij}, i \in [0..P-1], j \in [0..Q]$, where i runs around the cross section of the tube, and j along the centerline. The centerline of a tube is modeled by use of a cubic Bézier curve [5]. Such a curve is given by

$$\mathbf{b}(t) = (1-t)^3\mathbf{b}_0 + 3(1-t)^2 t\mathbf{b}_1 + 3(1-t)t^2\mathbf{b}_2 + t^3\mathbf{b}_3$$

with $t \in [0,1]$. For $\mathbf{b}_0$ and $\mathbf{b}_3$ we use the end points of the tube, i.e., the attachment points $\mathbf{p}_A$. The control point $\mathbf{b}_1$ is derived from the normal $\mathbf{n}_0$ on the surface of the ellipsoid

$$\mathbf{b}_1 = \mathbf{b}_0 + \mu\mathbf{n}_0/3|\mathbf{b}_3 - \mathbf{b}_0|$$

where μ (typically 1) can be tuned to vary the offset of the tubes. The other control point $\mathbf{b_2}$ is defined similarly.

To generate the surface of the tube, contours must be rotated and interpolated. We use a Frenet frame as a natural reference frame along the centerline, given by

$$\mathbf{f}_3(t) = \mathbf{b}'/|\mathbf{b}'|,\ \mathbf{f}_2(t) = \mathbf{f}'_3/|\mathbf{f}'_3|,\ \mathbf{f}_1(t) = \mathbf{f}_3 \times \mathbf{f}_2.$$

A Frenet frame is undefined when the curvature is zero. When the control points are colinear, an arbitrary frame can be chosen instead. When locally the curvature is zero, the frame can rotate over 180 degrees, which has to be checked and corrected for.

Suppose that the start contour consists of a sequence of points $\mathbf{p}_i$, with $i = 0, \cdots, P-1$, such that $\mathbf{p}_0$ is located at the boundary between the A and B part of the surface, and with a counterclockwise orientation when viewed from outside the ellipsoid. Here $P = 4i_0 + 4j_0$. The end contour with points $\mathbf{q}_i$ is defined similarly, also with P points, except that we assume here a clockwise orientation. We use a rotating frame for the rotation of the contour, given by

$$\begin{aligned} \mathbf{g}_1(t) &= \cos\phi\,\mathbf{f}_1 - \sin\phi\,\mathbf{f}_2 \\ \mathbf{g}_2(t) &= \sin\phi\,\mathbf{f}_1 + \cos\phi\,\mathbf{f}_2 \\ \mathbf{g}_3(t) &= \mathbf{f}_3 \end{aligned}$$

with

$$\phi = \phi(t) = (\phi_1 - \phi_0 + T2\pi)t + \phi_0.$$

The offset ϕ_0 is set such that initially $\mathbf{g}_1$ is aligned with $\mathbf{p}_0 - \mathbf{b}_0$. We measure this initial offset relative to the Frenet frame with

$$\phi_0 = \arctan\frac{\mathbf{p}_0^* \cdot \mathbf{f}_2(0)}{\mathbf{p}_0^* \cdot \mathbf{f}_1(0)}$$

where

$$\mathbf{p}_0^* = \mathbf{p}_0 - \mathbf{b}_0 - ((\mathbf{p}_0 - \mathbf{b}_0) \cdot \mathbf{f}_3(0))\mathbf{f}_3.$$

The final offset ϕ_1 is defined similarly. The value of T is chosen such that the total rotation $\phi(1) - \phi(0)$ matches with the desired number of twists R of the tube, e.g.,

$$T = \text{round}\,\frac{\phi_0 - \phi_1 + R\pi}{2\pi}.$$

Contours are interpolated in a local frame, using a cubic Bézier spline again, i.e.,

$$\mathbf{c}_i^*(t) = (1-t)^3\mathbf{c}_{i0}^* + 3(1-t)^2 t\mathbf{c}_{i1}^* + 3(1-t)t^2\mathbf{c}_{i2}^* + t^3\mathbf{c}_{i3}^*.$$

For $\mathbf{c}_{i0}^*$ we use start contour points, transformed by use of the $\mathbf{g}(0)$ frame:

$$\mathbf{c}_{i0}^* = (\mathbf{g}_1(0) \cdot (\mathbf{p}_i - \mathbf{b}_0),\ \mathbf{g}_2(0) \cdot (\mathbf{p}_i - \mathbf{b}_0),\ \mathbf{g}_3(0) \cdot (\mathbf{p}_i - \mathbf{b}_0)).$$

For $\mathbf{c}_{i3}^*$ the end contour points are used:

$$\mathbf{c}_{i3}^* = (\mathbf{g}_1(1) \cdot (\mathbf{q}_i - \mathbf{b}_3),\ \mathbf{g}_2(1) \cdot (\mathbf{q}_i - \mathbf{b}_3),\ \mathbf{g}_3(1) \cdot (\mathbf{q}_i - \mathbf{b}_3)).$$

For the contours in between we use ellipses:

$$\mathbf{c}_{i1}^* = \mathbf{c}_{i2}^* = (w\cos 2\pi i/P,\ d\sin 2\pi i/P,\ 0).$$

The points of the mesh of the tube are now finally given by

$$\mathbf{c}_{ij} = \mathbf{b}(j/Q) + (\mathbf{g}_1(j/Q),\ \mathbf{g}_2(j/Q),\ \mathbf{g}_3(j/Q)) \cdot \mathbf{c}_i^*(j/Q).$$

Figure 15 shows the resulting mesh, for an extreme case where a thin ellipsoid is fitted with thick round tubes.

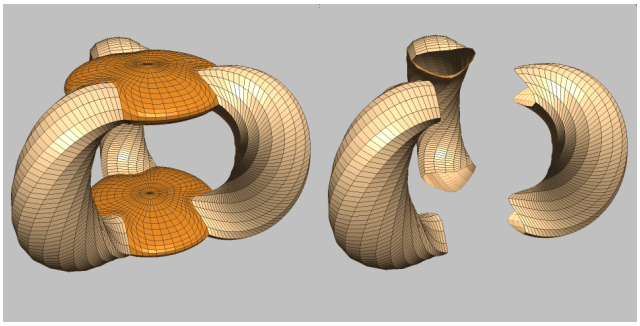

Figure 15: Meshes of tubes

5.3 Smoothing

The preceding approach gives ellipsoids and tubes. To obtain smoother surfaces, especially such that the transitions between tubes and ellipsoids are less pronounced, we use Catmull-Clark subdivision [3]. Here a mesh is recursively subdivided, where in each subdivision a face with n sides is subdivided into n quadrilaterals. In each subdivision step new points are introduced per face and per edge. The face points are the average of the face's original vertices, edge points are the average of the center of the original edge and the two new adjacent face points. The original vertices are replaced with

$$(\mathbf{q} + 2\mathbf{r} + \mathbf{v}(n-3))/n$$

where $\mathbf{q}$ is the average of the new face points around the vertex, $\mathbf{r}$ is the average of the midpoints of the n edges that share the vertex and $\mathbf{v}$ is the original vertex. Furthermore, to obtain an even smoother result, we implemented an option to smooth the surface initially without subdivision. Here, in each step the vertices of the mesh are averaged with their neighboring face and edge points according to the Catmull-Clark scheme. This smoothing process leads to surfaces that approach more or less minimal surfaces. One problem is that the tubes tend to become thin, especially when the original mesh is coarse and multiple twists occur. For a straight circular tube, with P points per cross section, Q cross sections and R twists, we can derive that each smoothing step reduces the radius with a factor

$$\rho = \frac{1}{16}(8 + 3\cos p + 3\cos q + \cos(p+q) + \cos(p-q))$$

with $p = 2\pi/P$ and $q = R\pi/Q$. A solution to the thinning of the tubes is simply to use a larger initial radius for the tubes, for instance by scaling the radius with ρ^{-k} when k smoothing steps are used.

Using these options, a variety of versions, starting from the same geometry can be generated (fig. 16). We implemented these by offering the user a choice between an exact and a smooth surface, and between a coarse, but fast mesh and a fine mesh.

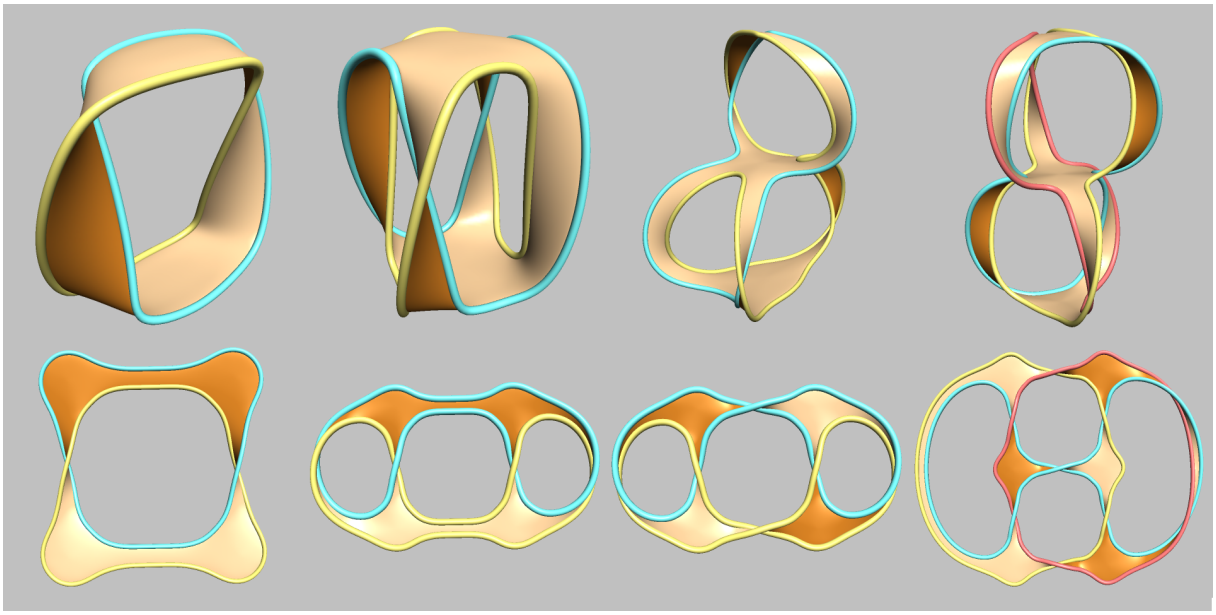

Figure 17: Hopf link (AA), link 4_2^1 ($AAAA$), Whitehead link ($AbAbb$), and Borromean rings ($AbAbAb$), in stacked and reduced style

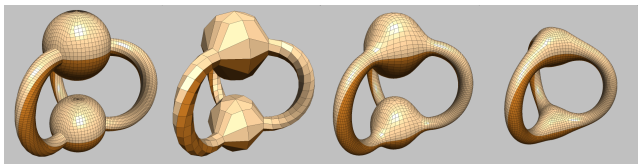

Figure 16: Detailed mesh, coarse mesh, coarse mesh subdivided, coarse mesh smoothed and subdivided (with increased diameter)

5.4 Knot

It is convenient to have an explicit representation of the knot or the components of the link that correspond to the surfaces. Here, the geometry of the knot is derived from the surfaces. Each polygon is assigned to part A or B of the surface, components are found by tracing edges that bound polygons that belong to different parts. The knot is shown as a tube. Optionally, an offset can be specified, such that the knot is shifted perpendicular to the surface in an outward normal direction. In fig. 7 we used an offset of the radius of the tube, such that the knot touches the surface. Also, this is useful for visualizing the linking number of the offset with the original knot, a quantity that plays a role in knot invariants like the Alexander polynomial.

6 Results

The method presented has been implemented in an MS-Windows application. The user is able to define knots and links (via specification of a braid word or by selecting one from a table), set the style, change geometric and graphic parameters, and view the result from all sides. All changes can be viewed in real-time, provided that the resolution of the mesh is chosen not too high (i.e., less than 10,000 polygons). The application is available for download [13]. Interactive viewing provides much better insight in the 3D shape than watching static images. Nevertheless, we show some more examples of results. Figure 17 shows a number of links (generated with the same approach as for knots), fig. 18 shows some more

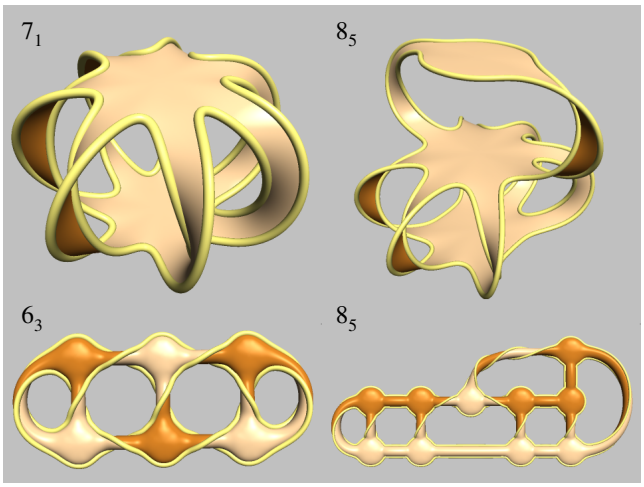

Figure 18: Knot 7_1 ($AAAAAAA$), 8_5 ($AAAbAAAb$), and 6_3 ($AAbAbb$)

knots. It is interesting to see that the Borromean rings in stacked style have similarities with the trefoil (one plateau is added) as well as the figure eight knot (alternating up and down bands). This observation can also be traced back in the braid words. As with the trefoil, one pattern is repeated three times; while the same pattern (Ab) is used as for the figure eight knot.

The 6_3 knot shown in fig. 18 has more crossings than in the standard diagram (9 vs. 6, see figure 2). However, the genus is minimal in the reduced style version, and also, this representation looks more symmetric and aesthetically pleasing.

Unfortunately, the braid representation does not always yield a minimal genus surface. Consider fig. 19, where knot 5_1, also known as the cinquefoil knot, and the almost similar 6_1 knot are compared. The knot 5_1 has the braid word $AAAAA$, the knot 6_1 has the braid word $AABacBc$. If we use these braid words to generate Seifert surfaces, we get a good result for the cinquefoil knot. The

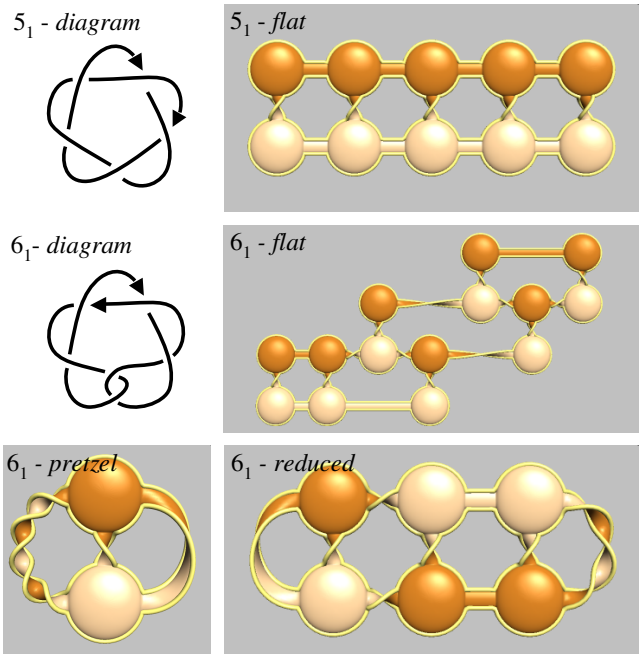

Figure 19: 5_1 cinquefoil knot (*AAAAA*) and 6_1 knot (*AABacBc*)

closed Seifert surface has four holes, which matches with its genus 2. However, the surface for the 6_1 knot also has four holes. The 6_1 has genus 1, and to visualize this, the shape should have two holes, which can be achieved by visualizing the 6_1 knot as a (5,-1,-1) pretzel knot. The flat style, closest to the original braid representation, is messy. Merging bands and eliminating disks gives the more compact reduced representation, but these steps cannot reduce the genus.

This limitation can be explained in various ways. The main difference between the upper parts of 5_1 and 6_1 is that in the former the strands run parallel, while in the latter their direction is opposite. The braid notation excels in representing parallel twisted strands, but cannot compactly represent twisted strands with opposite directions. Knots with many crossings and a low genus typically have twisted strands with opposite directions, pretzel knots are a good example of these.

The braid representation works well for diagrams with few crossings, as well as knots with a (relatively) high genus. For tutorial purposes, the first category is most interesting. Once one is familiar with various instances of the Seifert surfaces of the trefoil, the figure eight, the Hopf link and the Whitehead link, the mental step from diagrams of more complex knots to the corresponding Seifert surfaces will be easy.

7 Discussion

We have presented a method for the visualization of Seifert surfaces, and have introduced closed Seifert surfaces. These surfaces are generated starting from the braid representation, several styles can be produced, and by varying parameters different versions can be produced.

In this field, one answer gives immediately rise to new questions. Some examples. We are interested in producing minimal genus surfaces for knots and links. We are currently working on generating Seifert surfaces directly from the Dowker-Thistlethwaite notation, such that for a larger class of knots and links the Seifert surface with the minimal genus can be obtained. Some first results are shown in fig. 20. When this can be done for all different knots, tables and overviews of Seifert surfaces can be generated automatically.

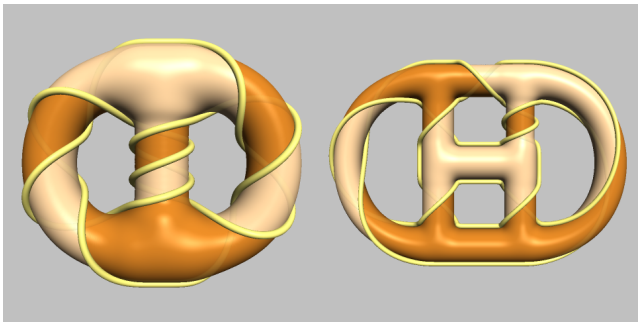

Figure 20: Knot 9_{35} and 9_{38}

We used spheres and tubes here as building blocks, but other primitives are conceivable. In fig. 20 we used for instance T-joints instead of spheres. Direct generation of 'a donut with holes' is on our wish list. Also, implicit surfaces and minimal surfaces are interesting candidates for describing the surfaces.

One remaining puzzle concerns the morphing of shapes. For instance in fig. 7, all shapes are isotopic, but we would like to exhibit this via a smooth animation.

Finally, so far we concentrated on visualizing Seifert surfaces. These surfaces play an important role in computing linking numbers, fluxes, and circulations for space curves. Visualizing these would be helpful in a wide range of applications ranging from knot theory to electromagnetism to fluid dynamics.

One could ask what all this is good for. We admit that our main motivation was simply curiosity about these interesting and hard-to-picture surfaces and the challenge of developing methods to visualize them. Nevertheless, the results are useful for knot theorists to illustrate and explain their work. Our first experience with using the tool in a course on knot theory was very positive in this respect.

Acknowledgements

We thank the reviewers for their valuable and inspiring comments.

References

[1] C. Adams. *The Knot Book: An elementary introduction to the mathematical theory of knots*. W.H. Freeman and Company, 1994.
[2] D. Bar-Natan. www.math.toronto.edu/~drorbn/katlas, 2004.
[3] E. Catmull and J. Clark. Recursively generated B-spline surfaces on arbitrary topological meshes. *Computer-Aided Design*, 10(6):350–355, 1978.
[4] G. di Battista, P. Eades, R. Tamassia, and I.G. Tollis. *Graph Drawing – Algorithms for the visualization of graphs*. Pearson, 1998.
[5] J.D. Foley, A. van Dam, S.K. Feiner, and J.F. Hughes. *Computer Graphics: Principles and Practice in C (2nd Edition)*. Addison-Wesley, 1995.
[6] P. Frankl and L. Pontrjagin. Ein Knotensatz mit Anwendung auf die Dimensionstheorie. *Math. Annalen*, 102:785–789, 1930.
[7] L. Kaufman. *On Knots*. Princeton University Press, 1987.
[8] C. Livingston. *Knot theory*. Math. Assoc. Amer., 1993.
[9] C. Livingston and J.C. Cha. Knotinfo: Table of knot invariants, 2004. www.indiana.edu/~knotinfo.
[10] D. Rolfsen. *Knots and links*. Publish or Perish, 1976.
[11] R.G. Scharein. *Interactive Topological Drawing*. PhD thesis, Department of Computer Science, The University of British Columbia, 1998.
[12] H. Seifert. Über das Geschlecht von Knoten. *Math. Annalen*, 110:571–592, 1934.
[13] J.J. van Wijk. www.win.tue.nl/~vanwijk/seifertview, 2005.