Intuitive Design and Meshing of Non-Developable Ruled Surfaces

Daniel Lordick

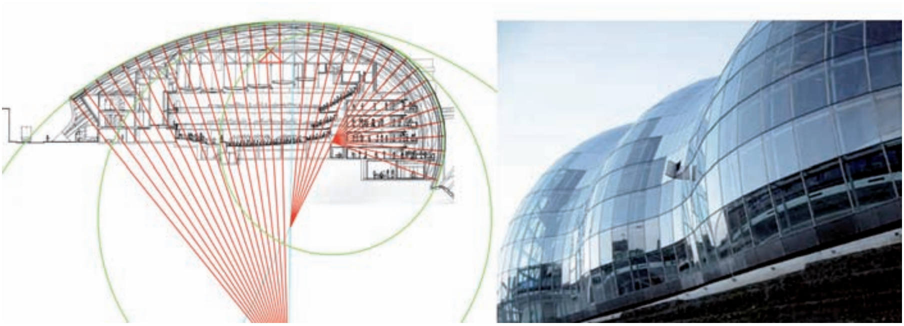

Figure 1: Section of the Sage Gateshaead building, façade to the waterfront (from: Foster 2005).

Introduction

The application of *freeform surfaces* in architectural design has to consider both the static requirements and the effort of fabrication. As a result most surfaces used are not as free as they appear to be. In fact, architects and engineers usually try to take advantage of special geometric properties, such as stability, dependent alone on the shape, and a feasible construction system that relies on a relatively simple concept of generation. For instance, the roof of the *Sage Gateshead* building by *Foster and Partners* (Foster 2005, pp. 168-173) consists of three surfaces of revolution with horizontal axes. The parts are joined seamlessly along the congruent meridian curves and yield the appearance of one large freeform (Figure 1). But because of the properties of rotational surfaces, it was possible to use planar quadrilateral panels for the roof tiling instead of a triangular mesh.

The features mentioned above, shape stability and a simple generation, apply particularly to the class of non-developable ruled surfaces (skew surfaces). They can be generated by constantly moving a straight line, for example by moving a hot wire through foam, or a diamond wire through stone. Skew surfaces offer a great freedom of figuration. Antonio Gaudí was one pioneer in the use of ruled surfaces. He extensively used one-sheet hyperboloids and hyperbolic paraboloids for the design of the Sagrada Família (Burri 2007, pp. 102-107). These surfaces were the answer to Gaudí's question: how can complex shapes be crafted at the construction site? (Figure 2). Another famous example of the use of hyperbolic paraboloids is the *Philips Pavilion* designed by Le Corbusier and Iannis Xenakis for the 1958 Brussels World Fair.

Figure 2: One-sheet hyperboloid crafted with a ruler (from: Burry 2007).

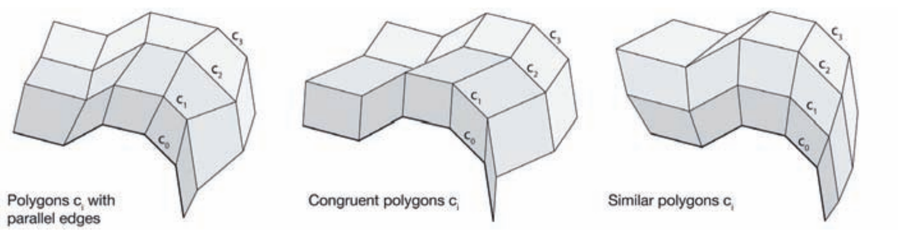

Figure 3: Examples of quadrilateral meshes.

This paper explores the implementation of ruled surfaces into the designing environment. Although ruled surfaces can be seen as a subset of NURBS surfaces spanning two curves, standard software provides no intuitive geometric approach. For instance, the existing tools lack control over the distribution and direction of the generators. One major goal of this paper is to transform non-developable ruled surfaces into planar quadrilateral meshes (called PQ meshes hereafter). At first this paper will present the smooth surfaces that can easily be transformed into a discrete representation with a PQ mesh. Then examples of skew surfaces that fit into the according classes will be investigated and, lastly, insight will be extended into more general cases.

Planar Quadrilateral Meshes

Mesh generation that guarantees planarity
There are several options to proof if a quadrangle is planar: the sum of the corner angles has to be 360°; the diagonals have to intersect; the straight lines carrying facing edges have to intersect or to be parallel. From these observations strategies for mesh constructions can be derived. For example, if two polygons have a sequence of pairwise parallel edges they span a strip of planar quads. This is particularly the case if the polygons are congruent or similar (Figure 3). In the latter two cases the strips are either prisms or pyramids (See 3.1).

Planarity is also guaranteed if the four vertexes of a quad lie on a circle. PQ meshes whose quads possess a circum-circle (circular meshes) have many useful properties. Details on this, and on more sophisticated mesh generation, are given in Pottmann et al. 2007, 677-699.

Surfaces with organic PQ mesh analogues
With respect to the previous paragraph, some smooth surfaces have more or less obvious PQ mesh analogues. Examples of these surfaces include translational and rotational surfaces. Translational surfaces carry congruent and parallel curves and result in meshes with only parallelograms, while rotational surfaces, for reasons of symmetry, provide circular meshes. We will use translational and rotational surfaces as blueprints for the handling of ruled surfaces (see 4 and 5). Two special cases are the following: if a translational surface has a straight generator, it is a cylinder; if the meridian of a rotational surface is a line parallel to the axis, the result is a rotational cylinder. These two cases are developable and covered by 3.1.

By the generalisation of rotational surfaces a *moulding surface* can be obtained. A moulding surface is a sweeping surface generated by a curve on a plane that rolls on a cylinder. For the polyhedral version, this leads to a local rotational axis for each pair of consecutive polygons. All axes are parallel. It is possible to generalize this surface again when the condition of parallelism is dropped (Figure 4 (a)). In both states of generalisation, every strip between consecutive polygons is a polygonal version

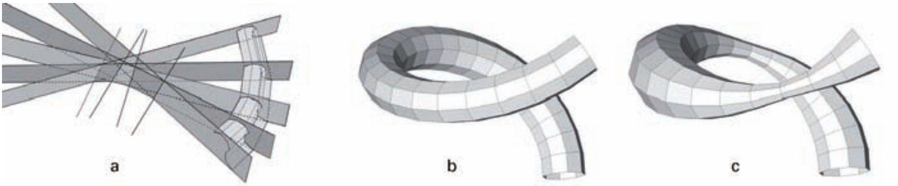

Figure 4: discrete analogues of a generalised moulding (a), pipe (b), and canal (c) surface.

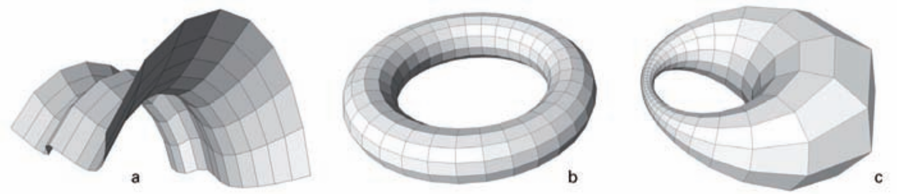

Figure 5: transformed rotational mesh, discrete torus, discrete Dupin cyclide.

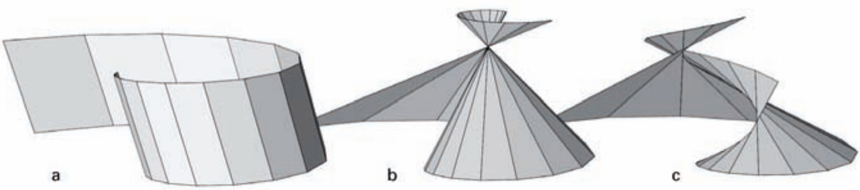

Figure 6: the three kinds of developable surfaces in their discrete analogues: prism, pyramid, and torsal polyhedron.

of a rotational surface and therefore provides a PQ mesh. For more generalisations using affine profiles see Sauer 1970, pp. 132-135.

Another interesting class of surfaces for PQ meshes are the canal surfaces (Figure 4 (c)). A canal surface is the envelope of a family of spheres. The midpoints of the spheres form a curve and the radius may change continuously. If the curve is a straight line we obtain a rotational surface. Each of the spheres touches the canal surface along a circle, which is a principal curvature line of the surface. One class of canal surfaces are the *Dupin cyclides* (Figure 4 (c)). If the radius of the spheres is constant, the result is a *pipe surface* like the *torus* (Figure 4 (b) and 5 (b)).

Lastly, it is important to mention the algebraic surfaces of degree two (quadrics). Any planar intersection of a quadric is a conical section. Moreover, any two parallel intersections of quadrics are similar curves (Lordick 2001, p. 42). Thus, every quadric can be discretized with an endless set of PQ meshes.

Transformations

As soon as a PQ mesh is established we can generate new PQ meshes by certain spatial transformations that preserve planarity. Those are the linear transformations, namely the affine and projective transformations. Scaling in one, two or three axes cannot destroy a PQ mesh. This is an important observation for the construction of PQ meshes in CAD software. For example, the mesh in figure 5 (a) may look complex but was built from a rotational surface. Furthermore, there is an interesting transformation called *inversion*, also referred to as *reflection at a sphere* (Pottmann et al. 2007, pp. 475-478). The inversion can transform the vertices of a PQ mesh into the vertices of a new one. The only condition is that the PQ mesh must be circular. In a circular mesh each quad possesses a circum-circle. Because inversions map circles to circles, the new mesh is also circular. Any PQ mesh that is generated by rotating a polygon around an axis has to be a circular mesh for symmetry reasons and is a perfect candidate for inversions (Figure 5 (b) and (c)). Unfortunately, inversions map lines to circles and therefore do not preserve the straight lines on a ruled surface.

Ruled Surfaces With Organic Pq Meshes

Developable surfaces

The developable surfaces are *cylinders, cones,* and *tangent surfaces of space curves.* Cylinders are translational surfaces with a linear profile. One example of cones is the rotational cone. Surfaces of constant slope belong to the family of the tangent surfaces of space curves. Any developable surface contains a continuous family of straight lines and therefore is a ruled surface. Each of the surfaces has a well-known discrete analogue that consists of planar strips (Figure 6). Obviously, endless variations of PQ meshes can be established on those strips. This paper will not focus on this class, but concentrate on skew surfaces.

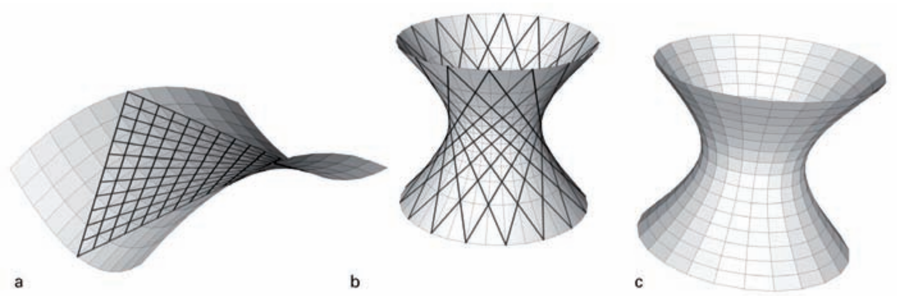

Figure 7: hyperbolic paraboloid with translational parabolas (discretized), one-sheet rotational hperboloid.

Ruled surfaces of degree two, reguli

The ruled surfaces of degree two are the circular cones, the quadratic cylinders, the one-sheet hyperboloids, and the hyperbolic paraboloids. While the first and the second are developable (see last paragraph) they are not in the scope of the further investigations. But one-sheet hyperboloids and the hyperbolic paraboloids are of highest interest. The key property is that both of these shapes contain two families of straight lines. Such a family on a quadric is called a *regulus*. In every point of the surfaces two generators, one of each regulus, intersect. Numerous applications in buildings can be found in the work of Felix Candela, Vladimir Shukhov (Shabolovka radio tower), Le Corbusier (Philips Pavilion), and, as previously mentioned, Antonio Gaudí. One-sheet rotational hyperboloids are also commonly utilised in cooling towers of power stations. Shapes from reguli are especially suitable for shell structures.

Any one-sheet hyperboloid can be obtained from a one-sheet rotational hyperboloid by a linear transformation. Thus it is sufficient to base studies on the latter. Because it is a rotational surface one possible PQ mesh is given by the meridians and the parallel circles (Figure 7(b) and (c)). It is a circular mesh. One strategy for a local mesh construction may use a discrete set of generators of one regulus. Those generators undergo a reflection with respect to a meridian plane. The grid of the generators provides the diagonals in a circular mesh.

Any intersection of a hyperbolic paraboloid (HP surface) with a plane parallel to the axis of the HP surface is a parabola. Moreover, the parabolas in parallel planes are congruent. Thus every HP surface is a translational surface with infinite pairs of characteristics (Wunderlich 1967, pp. 31-32). Any such pair of parabolas provides a PQ mesh (Figure 7 (a)). While in every point of the HP surface the generators correlate with the asymptotic directions, according to the translation, any two characteristic parabolas intersect in conjugate directions.

A good method for the construction of a mesh can be derived as follows: when a HP surface is projected normally onto the tangent plane in its vertex, the projections of the two reguli can be chosen to intersect orthogonally. (Any other angle can be obtained by an affine transformation.) Thus a regular grid of squares or rectangles can be generated. The grid represents the diagonals in a mesh of parallelograms on the HP surface. Obviously the vertices of the mesh coincide with one half of the intersections of the generators (Figure 7 (b) and (c)). In a building construction the generators can represent the cables that are often used for the stabilization of quadrangular glass tiles.

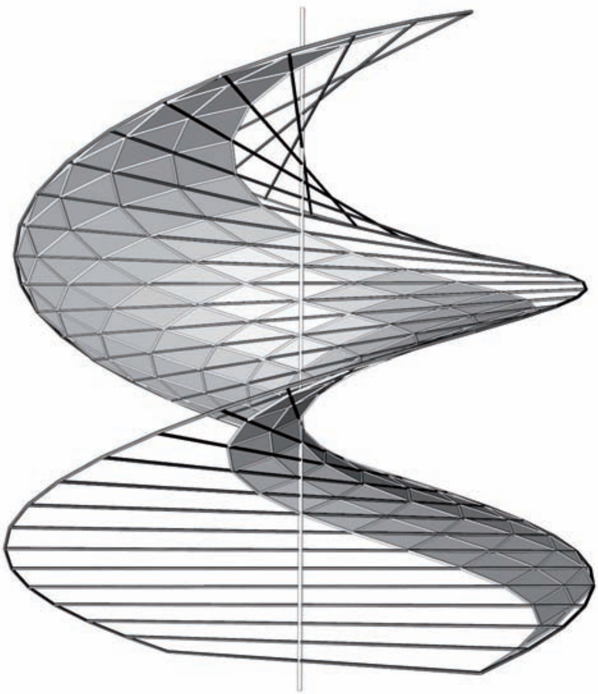

Figure 8: a helicoid interpreted as translational surface and meshed like this.

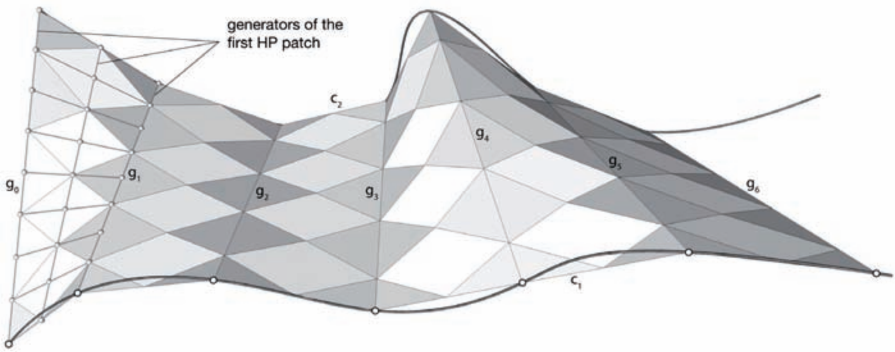

Figure 9: discrete HP patches between consecutive generators.

Helicoid
By applying a helical motion to a straight line that intersects the axis of the motion orthogonally we obtain a *helicoid*. The helicoid is a surface with remarkable features. It belongs to several classes of surfaces. It is a skew surface, a helical surface, a minimal surface, and a translational surface. The latter property is not well known, but can actually be proved by means of descriptive geometry (Wunderlich 1967, p. 178). If a helix is chosen as the generator of a translational surface and is then translated along itself, it sweeps a helicoid within a rotational cylinder. The helix intersects the axis of the helicoid and has half the height of the helicoid. After these observations it is easy to approximate any helicoid with a fair PQ mesh (Figure 8). The mesh on the helicoid can be varied by the use of different helices for the underlying translational surface.

Discretisation of General Skew Surfaces

Local approximation
For every regular generator g of a skew surface exists a right hyperbolic paraboloid that in every point of g has the same tangent planes as the skew surface (Wunderlich 1967, pp. 35-36). Furthermore, the Gaussian curvatures match. That holds for any touching HP surface along g. But that, in general, does not mean they have the same principal curvatures (Müller; Krames 1931, p. 94). Obviously, it is possible to locally approximate any skew surface with an HP surface.

The according method of meshing a skew surface is to add up discrete HP surfaces into the warped strips that span two consecutive generators. An example is given in figure 9. Two polygons c_1 and c_2 with n vertices each represent the boundary curves of a surface patch. Straight lines g_i span corresponding vertices i of c_1 and c_2. Thus quadrilaterals are obtained that are generally not planar and shall be filled with discrete HP patches. By inserting midpoints to the segments of the polygons new generators are obtained (Figure XX (a)). Lines of the other regulus can be generated by regularly dividing the generators g_i. Here a division by 8 is chosen. In general, consecutive HP patches do not have the same tangent planes along the generators g_i. For that reason triangles have to be inserted at the edges of each HP patch. The result is a hybrid mesh of planar quadrilaterals and triangles (Figure XX (b)). Obviously, such a mesh can be achieved from any set of straight lines and therefore can deal even with the worst cases and singularities such as the equivalents of torsal generators. An interesting creative option is to flip the edges of the filled-in triangles. The problem with this version is: the discrete generators g_i in general no longer coincide with edges of the mesh. A related hybrid tessellation on a rotational surface can be seen at the Swiss Re Headquarters built by Foster and Partners (Foster 2005, pp. 272-277).

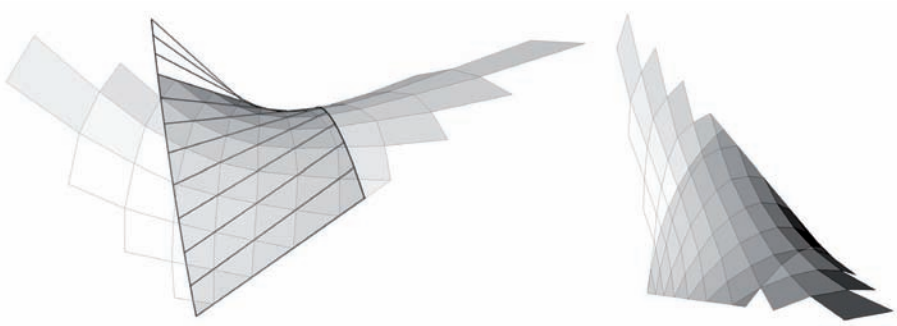

Figure 10: Two views of an arbitrary skew surface with a PQ mesh.

Compound surfaces from HP patches
If the consecutive HP patches along the generators g_i have connections of at least class C^1, the overall tessellation can be a PQ mesh. This obviously excludes the edges of the whole. HP patches with C^1 connections can also be arranged in two directions, because they have two families of generators. Unfortunately compound surfaces derived from this algorithm suffer from massive restrictions concerning creativity.

General algorithm for PQ meshes on skew surfaces
The next strategy to establish a PQ mesh on a set of generators g_i of a skew surface works as follows: select one generator gi and divide it uniformly and accordingly to the design task. The points of the division span planes with the next generator g_{i+1}. Let those planes intersect the next but one generator g_{i+2}. Now the points of intersection on g_{i+2} and the primary chosen points of division can be connected by generators h_j of a regulus that belong to the same surface of order two as g_i, g_{i+1}, and gi+2, in general a one-sheet hyperboloid. Now the intersections of the generators h_j with g_{i+1} are taken as points of division like in the last step and new points are obtained on the generator g_{i+3}. This can be continued until the last generator is reached. Finally a PQ mesh according to 3.2 can be selected (Figure 10). There are two options. In either case the quads have intersecting diagonals – the generators of both reguli.

One problem is the handling of the borders of the skew surface. A simple workaround is to extend the generators and thus to produce many more quads than necessary. At the end of the meshing the superfluous quads are then trimmed. Another question is how to deal with the equivalents of the torsal generators? In this case a useful intersection with the next but one generator cannot be found. The algorithm shall detect the problem and fill the remaining holes with triangles.

Application of Pq Mesh Algorithms to Selected Skew Surfaces

Conoidal surfaces, conoids
Many possible strategies for the generation of ruled surfaces are known (Müller; Krames 1931, pp. 20-60). In general a ruled surface can be given by three curves, called *directrices* or *director curves.* Any straight line intersecting each directrix is a generator of the surface. One directrix can be replaced by a *director plane*, which equals an *ideal line*. The outcome is a *conoidal surface*. Let one of the remaining director curves be a straight line (*linear directrix*). Then the resulting surface is named *conoid*. One example for a conoid is the hyperbolic paraboloid, where both director curves are straight lines. Another example is the helicoid, where the director plane is orthogonal to the linear directrix, which is the axis of the helical motion. Conoids with this property are called *right conoids*. Another kind of conoids has its application in the vaults of cathedrals, in particular in the choirs (Scheffers 1927, pp. 385-389).

The algorithm of 4.3 shall be tested with a right conoid that has a quadratic directrix, in particular an arc. With these specifications the result in general is an algebraic surface of degree four. But if a straight line and an arc are taken as boundary curves of a NURBS surface in CAD software, a different surface occurs. This can be seen with a quick glance at the parameter lines that intersect the directrices but are not parallel to a director plane. To yield the desired conoid another generation has to be used. A good method is to uniformly divide the arc c and to use planes parallel to the director plane to get corresponding points on the linear directrix d. Thus a discrete set of generators is established which serve as the profiles g_i of a loft surface. This loft surface is a fair approximation of the conoid and at the same time provides the discrete basis for the PQ mesh. The mesh can be refined by changing both the subdivision of the arc and the subdivision on the first generator g_0.

Skew surfaces with one linear directrix and other methods
Given are two splines that are arbitrarily parametrized and serve as director curves c1 and c2. We can generate a distribution of generators by several methods. Option (a) shows an equidistant division on c_1 and c_2, option (b) uses the parametrisation of the splines, and option (c) is a conoidal surface. Another option to handle the distribution of the generators on a skew surface with two director curves is to add a linear directrix d. By moving d, the mesh design can be changed dramatically.

Conclusions

The presented methods do not focus on numeric approximation. Rather, the strategy was to adapt geometric properties for the constructions. One aim was to investigate how flexible this approach is. On the other hand it was interesting to observe what possible designs could be gained from the underlying geometric structure. While the PQ mesh is designated to define the panelling of the surface, the generators can serve as additional structural elements such as beams or cables.

Ruled surfaces can be controlled by director curves, by an optional director plane, or a linear directrix. The mesh depends on the discrete set of generators, and the division on a primary selected generator. Although these are sufficient elements to obtain a usable mesh, it still needs some knowledge about ruled surfaces in general to obtain good results. Therefore further possibilities of control have to be explored in communication with users of the algorithms.

Ruled surfaces in renderings often do not appear to be smooth. This is because the underlying standard triangulation used for rendering, if not corrected manually, consists of splinters, i.e. triangles with one or two very small angles. Thus adjacent triangles sometimes have heavily differing normals. In rapid prototyping this leads to insufficient results. The algorithms presented in this paper provide very good approximations with meshes and therefore can lead to much better results.

References

Burry, M. (ed.), 2007: Gaudí unseen, die Vollendung der Sagrada Família. Berlin: Jovis.
Foster, Sir Norman; and Partners, 2005: Catalogue. Munich; Berlin; London; New York: Prestel.
Lordick, D., 2001: Konstruktion der Schattengrenzen krummer Flächen mithilfe von Begleitflächen. Aachen, Karlsruhe: Shaker.
Müller, E.; Krames. J., 1931: Vorlesungen über Darstellende Geometrie. Band III. Leipzig; Wien: Deuticke.
Pottmann, H.; Wallner, J., 2001: Computational line geometry. Berlin; Heidelberg; New York et al.: Springer.
Pottmann, H.; Asperl, A.; Hofer, M.; Kilian, A., 2007: Architectural Geometry. Exton: Bentley Institute Press.
Sauer, R., 1970: *Differenzengeometrie*. Berlin/Heidelberg: Springer Verlag.
Scheffers, G., 1927: Lehrbuch der Darstellenden Geometrie. In zwei Bänden, Zweiter Band. Berlin, Springer.
Wunderlich, W., 1967: *Darstellende Geometrie II*. Mannheim: Bibliographisches Institut.