Cyclides in Surface and Solid Modeling

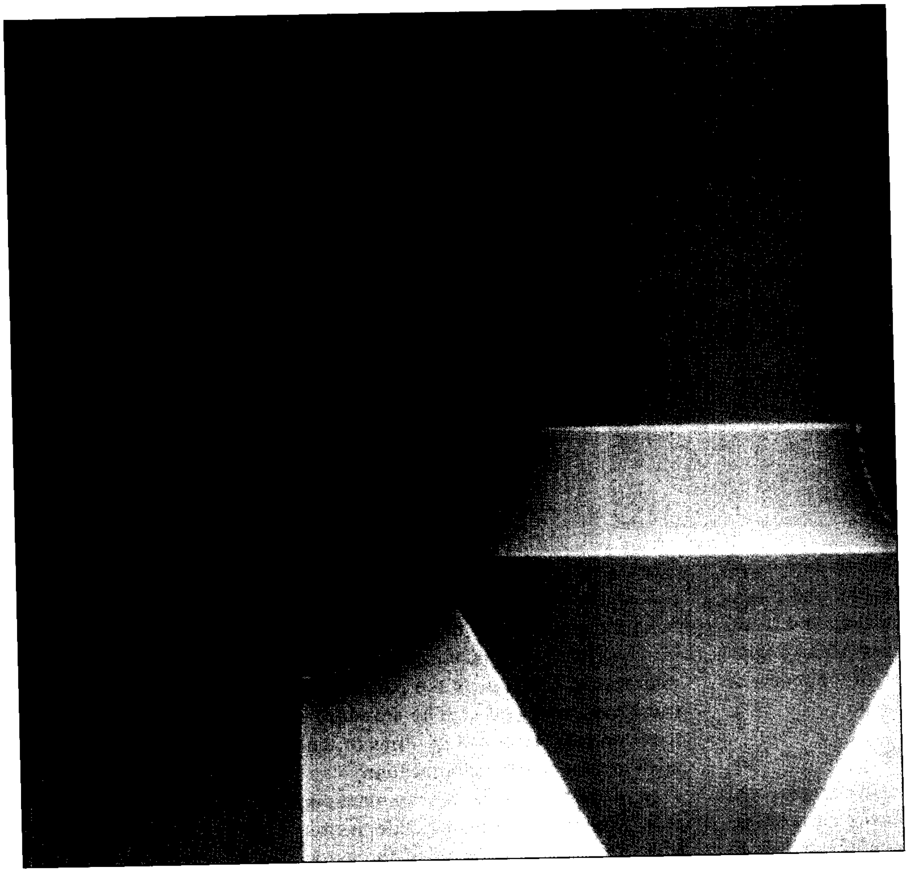

Debasish Dutta
University of Michigan
Ralph R. Martin
University of Wales College of Cardiff
Michael J. Pratt
LMR Systems, United Kingdom

As surfaces in CAGD, Dupin cyclides have low algebraic degree, rational parametric forms, and an easily comprehensible geometric representation in terms of three shape parameters.

Parametric and implicit representations of surfaces are most common in computer-aided geometric design (CAGD). Parametric forms [$x = f(s, t)$, $y = g(s, t)$, $z = h(s, t)$] are particularly well suited for tasks such as surface editing and rendering, but their conversion to implicit [$F(x, y, z) = 0$] form leads to a degree explosion. For example, a four-sided polynomial patch of degree n has algebraic degree $2n^2$. Implicit forms are useful for tasks such as point membership classification. Low-degree algebraic surfaces such as cone, cylinder, sphere, and plane (collectively referred to as *natural quadrics* in CAGD) also enable robust and relatively fast algorithms for geometric computations. Restricting to natural quadrics leads to a limited geometric coverage, but researchers have yet to find intuitive shape parameters for designing with higher degree algebraic surfaces.

Therefore, in the context of CAGD, a surface that preserves all the advantages and avoids most of the disadvantages of parametric and implicit representations would have to satisfy the following characteristics:

- Low-degree algebraic representation (to facilitate robust geometric algorithms).
- Low-degree parametric representation (to ensure compatibility with existing CAD systems and data-exchange standards).
- Simple geometric representation (to provide the designer with intuitive shape parameters for easy surface creation and manipulation).
- Flexibility for free-form surface design when used in a piecewise manner.

In this article, we consider Dupin cyclides, surfaces that fully satisfy the first three criteria and partially satisfy the fourth. Even the widely used parametric bicubic surface fails

0272-17-16/93/0100-0053$03.00

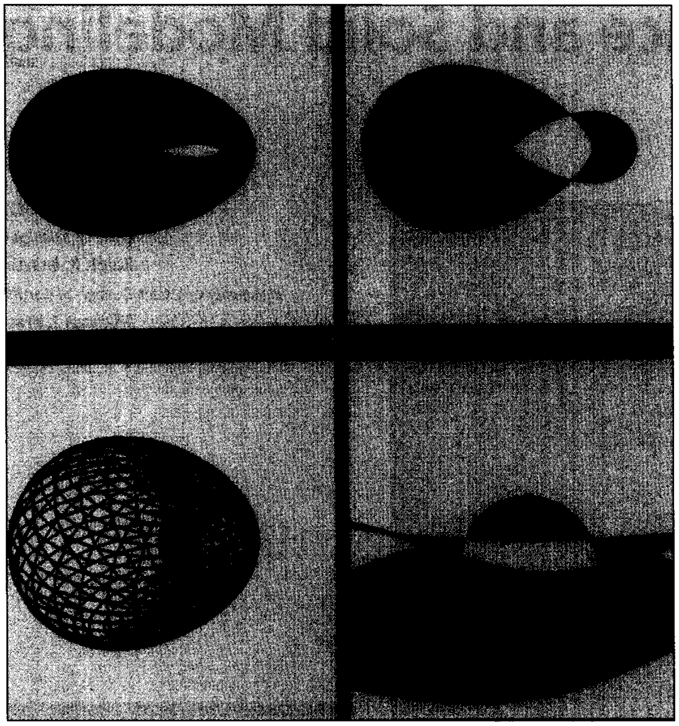

Figure 1. Clockwise from the bottom left: spindle, ring, horn, and parabolic cyclides.

on the first and fourth criteria. Space restrictions permit only a cursory treatment of the geometry, properties, and uses of the Dupin cyclide in free-form surface modeling and blending. However, the references we cite point to detailed original work and, more importantly, to a larger set of references on related topics that we have been forced to omit.

Dupin's cyclides

Early in the last century, the French mathematician and naval architect Charles Pierre Dupin defined a class of surfaces he named *cyclides*.[1] Dupin's definition of the cyclide included all surfaces that were envelopes of spheres of possibly varying radii, which were tangent to three given fixed spheres. Dupin's cyclides have since been investigated by James Clerk Maxwell in the context of an optics problem, and by Arthur Cayley and several other mathematicians. In this article, the term *cyclide* always refers to Dupin's cyclides unless otherwise noted. (In earlier works, the term *cyclide* sometimes refers to a wider class of surfaces.)

In the early 1980s, Martin and Nutbourne[2,3] at Cambridge University and McLean[4] at Chrysler rekindled interest in cyclides, this time in the context of CAGD. In both efforts, cyclide patches were researched for use in free-form surface modeling. The Cambridge work also included cyclide surface intersections.[5] More recently, the use of cyclides as blending surfaces has been investigated.[6-8] We discuss some methods here.

Geometry and properties of cyclides

In general, the cyclide is a quartic surface. The loci of the centers of two families of variable spheres associated with a cyclide define a pair of spine curves for the cyclide. These spine curves are special: They are a pair of conics (an ellipse and a hyperbola, in general) on mutually perpendicular planes, such that the foci of one are coincident with the vertices of the other. Cyclides have an algebraic equation[8]

$$(x^2 + y^2 + z^2 - \mu^2 + b^2)^2 = 4(ax - c\mu)^2 + 4b^2y^2$$

and parametric equations (for $0 \le \theta < 2\pi$, $0 \le \psi < 2\pi$)

$$x(\theta,\psi) = \frac{\mu(c - a\cos\theta\cos\psi) + b^2\cos\theta}{a - c\cos\theta\cos\psi},$$

$$y(\theta,\psi) = \frac{b\sin\theta(a - \mu\cos\psi)}{a - c\cos\theta\cos\psi},$$

$$z(\theta,\psi) = \frac{b\sin\psi(c\cos\theta - \mu)}{a - c\cos\theta\cos\psi}$$

Here, x, y, and z are Cartesian coordinates. Of the constants a, b, and c associated with the spine curve conics, only two are independent, since for the special conics described, $c^2 = a^2 - b^2$. Therefore, we can uniquely define the cyclide by three parameters: a, c, and μ. In the parametric description, the parameters are θ and ψ. The standard substitution of the trigonometric functions in terms of tangents of half angles leads to a rational biquadratic form.[9]

Here, we summarize the important points of cyclide geometry.[10] The morphology of cyclides permits a bilevel classification. The first is in terms of the associated spine curves. Cyclides belong to one of four families:

1. *central cyclides* with central conics as spine curves,
2. *parabolic cyclides* (cubic surfaces) with parabolic spine curves,
3. *revolute cyclides* with straight lines and circles as spine curves, or
4. *degenerate cyclides* with degenerate conics such as points and lines as spine curves.

Within a family, we can further classify a cyclide on the basis of the relative magnitude of the spine conic parameters a, c, and μ, leading to three subforms—namely, *horn*, *ring*, and *spindle* cyclides. Figure 1 shows spindle, ring, horn, and parabolic cyclides, clockwise from the bottom left.

An easy-to-visualize cyclide is the *ring central cyclide*, which resembles a squashed torus. The planes containing the spine curves are the cyclide's planes of symmetry. On those planes, the cross sections consist of two bounding circles whose radii are simple functions of a, c, and μ, as shown in

Figure 2. In fact, by comparison with the torus, a is analogous to the major radius and μ to the minor radius, and c is a measure of the asymmetry of the surface about the third coordinate plane ($c = 0$ gives a standard torus). The horn and spindle forms of the central cyclides arise when the bounding circles in the x-y plane and the x-z plane intersect. For parabolic cyclides, one bounding circle degenerates to a line, and the surface extends to infinity.

The three constants defining a particular cyclide have a straightforward geometric interpretation easily grasped by a designer. Also, if we know the radii of the cross-sectional circles in the two planes of symmetry, the cyclide is completely determined. On the basis of this fact, Dutta and Srinivas at the University of Michigan have developed a cyclide visualization routine. A Silicon Graphics Personal Iris renders all forms of the cyclide in real time, as the user interactively repositions the bounding circles or changes their radii.

In CAGD, cyclides have some notable properties:

1. Both families of lines of curvature on a cyclide are families of circles.
2. Cyclides are closed under normal offsetting. Varying the parameter μ yields the constant radius offsets of a given cyclide.
3. The planes of the circular lines of curvature of either family all intersect in a common line.
4. Along any line of curvature on a cyclide, the angle ϕ between the surface normal and the principal normal to the curve is constant—a consequence of Joachimstal's theorem. It follows that there is a right circular cone, with semivertex angle ϕ, tangent to a cyclide around any of its lines of curvature.
5. With regard to inversion in a sphere, the inverse of a Dupin cyclide is also a Dupin cyclide, provided the center of inversion does not lie on the original surface.

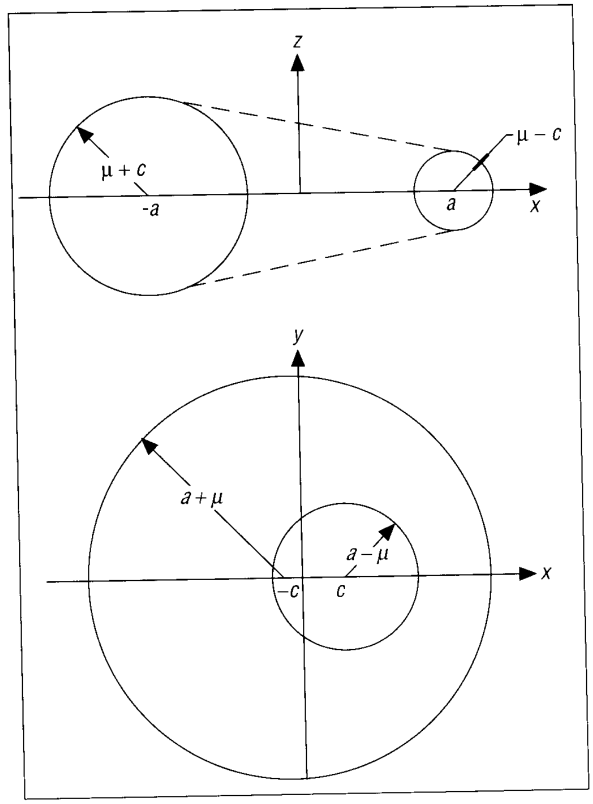

Figure 2. Ring central cyclide: two bounding circles whose radii are simple functions of *a*, *c* and μ.

Patching with cyclides

We will now consider the construction of composite surfaces made from pieces of several cyclides placed together with appropriate continuity.

Principal patches

In the 1980s, researchers at Cambridge University investigated the possibility of designing patched surfaces whose fairness could be controlled. The approach they chose was to create four-sided patches whose sides lie along *lines of curvature*. They called such a patch a *principal patch.*[5]

We wish to stress that there is only *one* pair of orthogonal families of lines of curvature on a surface, and these families have an intrinsic relationship to the underlying geometry. On the other hand, the directions of boundaries of conventional patches used for surface description are independent of the surface geometry. The same surface can be made up of sets of patches that run in any direction across the surface. However, the orientation of principal patches is uniquely defined for a given surface.

We must impose conditions on patch boundary curves if they are to bound a principal patch. Consider a triad of orthogonal unit vectors, where $\mathbf{t}_1$ is the tangent to some curve on a surface, $\mathbf{N}$ is the surface normal, and $\mathbf{t}_2$ is perpendicular to $\mathbf{t}_1$ and $\mathbf{N}$. As we go across the surface by a distance s_1 along the curve, the triad changes orientation as governed by the differential equation[2,5]

$$\frac{d}{ds_1}\begin{bmatrix}\mathbf{t}_1\\ \mathbf{t}_2\\ \mathbf{N}\end{bmatrix}=\begin{bmatrix}0 & \kappa_g & \kappa_n\\ -\kappa_g & 0 & \tau_g\\ -\kappa_n & -\tau_g & 0\end{bmatrix}\begin{bmatrix}\mathbf{t}_1\\ \mathbf{t}_2\\ \mathbf{N}\end{bmatrix}$$

where κ_n is surface normal curvature along the curve, κ_g is the geodesic curvature of the curve, and τ_g is its geodesic torsion. In particular, for any line of curvature, $\tau_g = 0$.

Consider four arcs A, B, C, and D, in clockwise order, which are to lie along lines of curvature and are to be the boundaries of a principal patch. As well as setting $\tau_g = 0$ for each curve, we must also ensure that these curves meet at

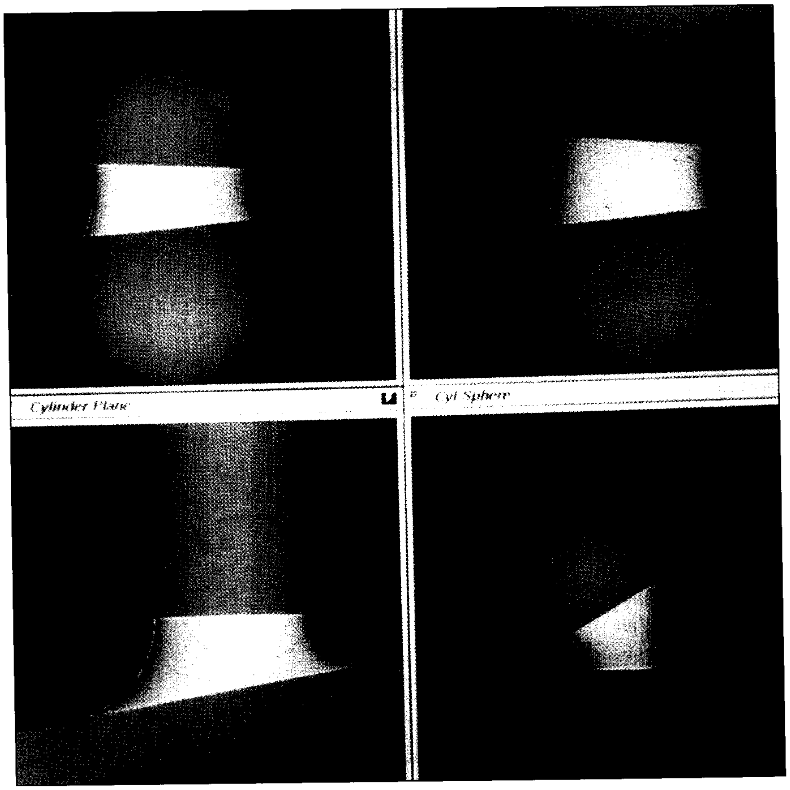

Figure 3. Sphere-cylinder intersection (bottom right), cone-sphere blend (top right), sphere-sphere blend (top left), and cylinder-plane intersection (bottom left).

right angles, as at any point on the surface the principal directions are orthogonal. To ensure that these arcs are the boundary of a principal patch, we must fulfill two conditions, which are also necessary and sufficient:[5]

1. Position matching: Going along side A, then side B, we must reach the same point as going along side C, then side D.
2. Frame matching: Going along side A, then side B must leave the surface triad in the same orientation as going along side C, then side D.

Solutions of the frame- and position-matching equations for general classes of principal patches are available elsewhere.[2,3,5] As **N** is continuous across adjacent patches, principal patches are automatically G^1 continuous.

All of this may seem rather far from our original aim in this article, which is to consider the use of cyclides in CAGD. However, if we ask, "What are the simplest nontrivial curves to use as lines of curvature for construction of principal patches?" the obvious answer is "Circles." Surfaces whose lines of curvature are circles are Dupin's cyclides. Details of how cyclide patches can be considered as principal patches are presented in other works.[2,5,11]

Composite surfaces from cyclide patches

Let the sides of a cyclide patch in clockwise order be A, B, C, D. If we take sides A and C as known—that is, as given by the initial piecewise circular arc curves in our composite design procedure or as sides of patches that have already been computed—the question arises as to how much freedom there is left for sides B and D.

Each side of a cyclide patch is defined by three quantities: the curvature K of the circle to which the side belongs, the angle ϕ between the normal to the surface and the normal to the side (ϕ is constant along the side), and the arc length s of the side. (These quantities are related to the normal and geodesic curvatures mentioned earlier by $\kappa_n = K \cos \phi$ and $\kappa_g = K \sin \phi$.) By the time we impose the frame- and position-matching constraints, only *one* degree of freedom remains for choosing sides B and D.

Martin[2] used a method based on patch subdivision to construct interior lines of curvature of the cyclide patch. Following a suggestion by Malcolm Sabin, he also showed that cyclides are *rational biquadratic* patches. This provides an important link between the principal patch approach to cyclide patches and the rational parametric polynomial methods widely used elsewhere in CAGD. For a Bezier formulation for a cyclide patch, see Pratt's paper.[9]

The basic method of creating composite cyclide surfaces by starting from two piecewise circular curves has the disadvantage that altering an earlier patch in the sequence affects all subsequently created patches. This is obviously undesirable, and de Pont outlined a method[5] to *locally* modify a surface created with cyclide patches. He also proposed that the basic unit for design with cyclide patches not be a single cyclide patch, but rather a 2 × 2 group of cyclide patches assembled into a "super-patch." He calls this a *cyclide q-patch* and treat it as a single unit. In the cyclide q-patch, each subpatch contributes one degree of freedom. Thus, a cyclide q-patch has four degrees of freedom, compared with the single one available for a simple cyclide patch. This means that we can now position the fourth corner arbitrarily in space (within certain limits), and it is not restricted to lying on a circle. We can use the other remaining degree of freedom to alter the shape of the cyclide q-patch interior.

At Chrysler, McLean based his approach to free-form surface modeling with cyclides on *Gaussian images*.[4] With this approach, we choose four cocircular vertices on the Gaussian sphere. Choosing a boundary circle between two consecutive vertices uses up the remaining freedom and fixes a cyclide patch. We choose an initial cyclide patch; others are then built on the sides of the existing patches. McLean concluded that after four or five patches had been added, controlling the overall shape was difficult.

The lack of freedom in shape control when using cyclide patches prompted other authors to consider possible extensions of the cyclides for shape design. Recently, Degen[12] proposed the use of *double Blütel surfaces*. These are projective images of Dupin's cyclides, and hence are also a special case of fourth-order algebraic surfaces with a rational biquadratic parameterization.

Cyclide intersections

In this section, we review algorithms for intersecting cyclides with the natural quadrics. Since parametric and implicit forms of the cyclide exist, by substitution we can readily obtain an expression for the intersection curve in parameter space. We can find successive points along the intersection curve of a cyclide with a plane by solving a biquadratic equation in u, v parameters. Thus an *analytic* solution is possible for these points.[5]

Furthermore, assuming these points are joined by straight line segments, we can bound the resulting errors in a straightforward way. For cyclide-cyclide intersections, we can compute the intersection curve analytically, this time by solving a biquartic equation. Cyclides have advantages over bicubic and rational bicubic patches in this respect. With such patches, the corresponding equations to be solved will be of higher degree and thus will require iterative or other numerical techniques.

Recently, Johnstone and Shene at Johns Hopkins University proposed an alternative method for computing intersections of cyclides.[13] They consider the cyclide as a ring surface, that is, a surface swept out by a variable radius circle moving in space. Circles of one ring surface are inverted into lines, which are then intersected with the other cyclide.

Blending with cyclides

When we construct mechanical objects by combining constructive solid geometry (CSG) primitives such as cones, cylinders, and spheres, we can use cyclides for smoothly blending the intersection curves. We can obtain constant radius blends by using cylinders and tori (special cases of cyclides), and variable radius blends by using nondegenerate cyclides. Furthermore, we can also use cyclides for smoothly connecting various, possibly nonintersecting, CSG primitives, as in piping design. Chandru, Dutta, and Hoffmann[7] analyze the problem of approximate blending using piecewise cyclides, while Böhm[6] and Pratt[8] are concerned with exact blending. All the authors seek tangent-plane or G^1 continuity in constructing the blends.

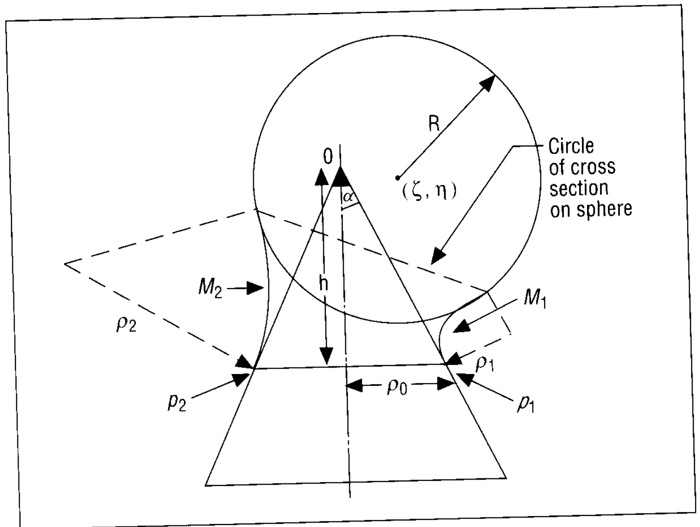

Figure 4. Cone-sphere blend: circles M_1 and M_2 are the cross section of the blending cyclide on plane P.

Cyclide blending is a direct consequence of the property that around any line of curvature on a cyclide there exists a right circular tangent cone. Therefore, a surface that blends smoothly into a cone C around a circular meridian c_1 will also blend smoothly into any cyclide tangent to the cone C around c_1. The problem now reduces to constructing blends to cones, cylinders being regarded as cones with vertices at infinity. In this section, we review two important subcases: the blend between a cone and a sphere, and the blend between two cones.[6-8]

Cone-sphere blend

The cone-sphere blend is easy to construct (see Figure 3, top right). The cone's axis, together with the sphere's center, defines a plane P about which both the cone and the sphere are symmetrical. Clearly, the cyclide blending the intersection is also symmetrical about P. First, we choose a meridian c_1 of the cone to act as one of the blend boundaries. Circle c_1 intersects plane P at two points, p_1 and p_2. Next, in plane P we construct two circles M_1 and M_2, each tangent to the cone at one of these points, say p_1, and also tangent to the third circle, which is the cross section of the sphere (see Figure 4). Circles M_1 and M_2 are the cross section of the blending cyclide on plane P; we can find the three constants a, c, and μ defining it from their centers and radii (see Figure 2). One degree of freedom is available to the designer—the choice of the first blend boundary—but everything else follows automatically.

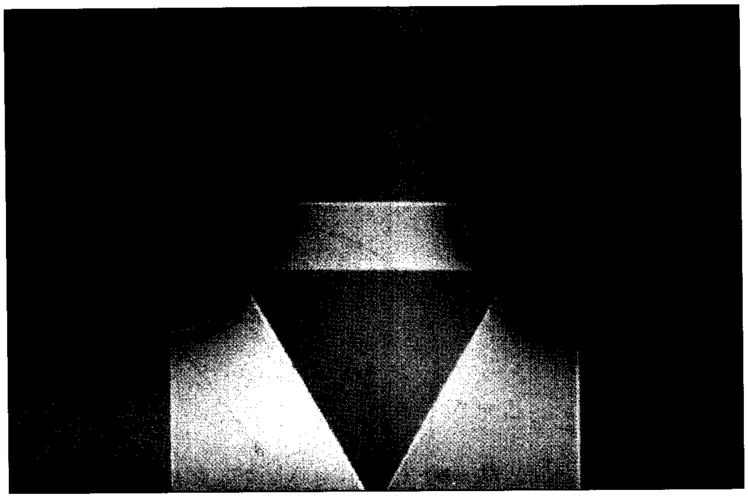

Figure 5. Blend of two intersecting cylinders.

In the top left and bottom left, Figure 3 shows how we can also use cyclides to blend sphere-sphere and cylinder-plane intersections. The salient property shared by spheres and planes is that at points on *any* circle on either surface, there is a constant angle ϕ between the principal curve normal and the normal to the surface.

Parabolic cyclides—cubic surfaces that extend to infinity—are useful in blending, particularly for a sphere-cylinder intersection where the sphere is tangent to a generating line of the cylinder (see Figure 3, bottom right). We can also blend two intersecting cylinders, as shown in Figure 5, if we use a sphere as a transitional surface. We return to this topic later.

Cone-cone blend

Blending between two cones is fundamental because it forms the basis of blending between two general cyclides. Malcolm Sabin's theorem states that two cones possess a family of blending cyclides if and only if the cones contain a common inscribed sphere.[8] This condition requires the axes of both cones to pass through the sphere's center, so they define a plane of symmetry. The generation of the blending cyclide now follows in a similar way to that of the cone-sphere blend. The designer has one degree of freedom. When the cone axes intersect, but the inscribed sphere condition does not hold, the situation is more complicated. We can still construct a cyclide blend, but only from parts of two cyclides.[8]

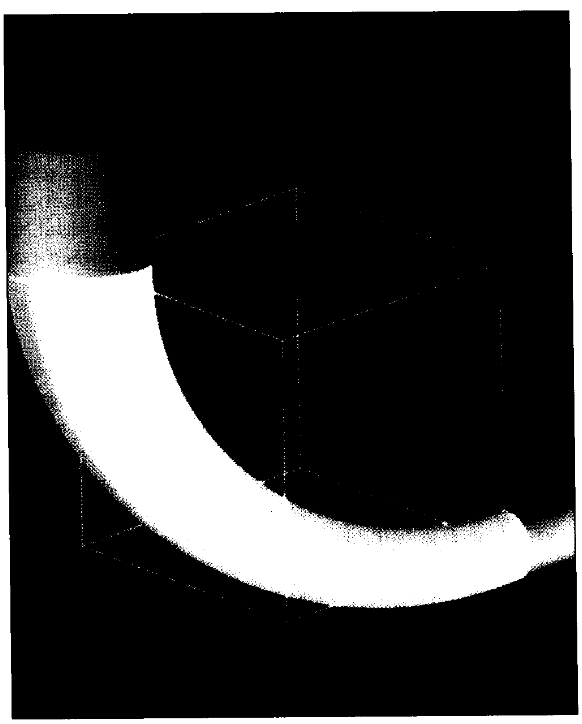

Figure 6. Two cylinders joined using three spheres and four ring cyclides.

Complex cyclide blends

The blends we described thus far require the surfaces, that is, the primitives, to share a common plane of symmetry. Asymmetrical situations are more complex and require multiple cyclide pieces. Therefore, by complex cyclide blends, we refer to exact or approximate blending surfaces constructed by smoothly joining pieces of cyclides.

An example is the double-cyclide blend between two pipes of unequal radii. Briefly, the idea is as follows: The boundaries of the blending surface are circular lines of curvature on the blending cyclide. Now we choose another line of curvature L (circle) between the boundaries of the blending cyclide. Clearly, L divides the blended object into two parts. Either of these could be rotated about the axis of L without losing C^1 continuity. This procedure results in a double-cyclide blend that uses portions of one cyclide (rotated about a line of curvature). We can also construct double-cyclide blends using two different cyclides, but this is more complex.[8]

Another method[7] for constructing approximate rolling-ball blending surfaces involves approximating the rolling-ball spine curve with piecewise (conic) cyclide spines. This is done by constructing cyclide pieces over elements of an approximated spine. While we can accomplish a conic approximation of space curves using the technique of biarcs,[3] construction of smooth cyclide pieces over them imposes additional restrictions. As mentioned earlier, we can smoothly join two cyclides at a circle of curvature if and only if they share a common tangent cone there.

Multisurface blends

Cyclides can also connect multiple, nonintersecting primitives. Suppose we have three cylindrical pipes of distinct diameters, and we wish them to have a smooth junction. We can achieve this by choosing an intermediate sphere of radius larger than any of the pipes and blending each pipe onto the sphere. Intermediate spheres can also act as "control spheres" that cyclide pieces smoothly interpolate. For a designer, these intermediate spheres provide a mechanism for controlling the global shape of the connecting piece. Figure 6 shows two cylinders joined using three spheres and four ring cyclides.[14] Böhm has also developed a method for blending between three cones using spherical surfaces.[6]

Conclusions

As surfaces in CAGD, cyclides have attractive properties such as low algebraic degree, rational parametric forms, and an easily comprehensible geometric representation using simple and intuitive geometric parameters. The alternative representations permit the transition between forms when one or the other is more convenient for some specific purpose. Cyclides provide a very useful extension of *geometric coverage* in solid modeling, primarily as blending surfaces for many commonly occurring situations. Research continues to extend the range of techniques. In the domain of free-form surface modeling, cyclides appear marginally insufficient with respect to the degrees of freedom for shape control. ❑

Acknowledgment

Thanks are due to Y.L. Srinivas of the University of Michigan for the color pictures in Figures 1, 3, 5, and 6. Dutta was supported in part by NSF grant DDM 90-10411 and by AFOSR grant F49620-92-J-0095.

References

1. C.P. Dupin, *Applications de geometrie et de mechanique*, Bachelier, Paris, 1822.
2. R.R. Martin, *Principal Patches—A New Class of Surface Patch Based on Differential Geometry*, Eurographics 83, North-Holland, Amsterdam, 1984, pp. 47-55.
3. A.W. Nutbourne and R.R. Martin, *Differential Geometry Applied to Curve and Surface Design: Vol. 1 Foundations*, Ellis Horwood, Chichester, UK, 1988.
4. D. McLean, "A Method of Generating Surfaces as a Composite of Cyclide Patches," *Computer J.*, Vol. 28, No. 4, 1984, pp. 433-438.
5. R.R. Martin, J. de Pont, and T.J. Sharrock, "Cyclide Surfaces in Computer Aided Design," in *Mathematics of Surfaces*, J.A. Gregory, ed., Oxford Univ. Press, Oxford, UK, 1986, pp. 253-267.
6. W. Böhm, "On Cyclides in Geometric Modelling," *Computer Aided Geometric Design*, Vol. 7, 1990, pp. 243-255.
7. V. Chandru, D. Dutta, and C.M. Hoffmann "Variable Radius Blending Using Dupin Cyclides," in *Geometric Modeling for Product Engineering*, M.J. Wozny, J. Turner, and K. Preiss, eds., North-Holland, Amsterdam, 1990.
8. M.J. Pratt, "Cyclides in Computer Aided Geometric Design," *Computer Aided Geometric Design*, Vol. 7, 1990, pp. 221-242.
9. M.J. Pratt, "Applications of Cyclide Surfaces in Geometric Modelling," in *Mathematics of Surfaces*, Vol. 3, D.C. Handscomb, ed., Oxford Univ. Press, Oxford, UK, 1989.
10. V. Chandru, D. Dutta, and C.M. Hoffmann, "On the Geometry of Dupin Cyclides," *The Visual Computer*, Vol. 5, No. 5, 1989, pp. 277-290.
11. T.J. Sharrock, *Surface Design with Cyclide Patches*, PhD thesis, Cambridge Univ., Cambridge, UK, 1985.
12. W.L.F. Degen, "Generalised Cyclides for Use in CAGD," in *Mathematics of Surfaces*, Vol. 4, A. Bowyer, ed., Oxford Univ. Press, Oxford, UK, to appear.
13. J.K. Johnstone and C.K. Shene, "Computing the Intersection of a Plane and a Natural Quadric," *Computers & Graphics*, Vol. 16, No. 2, June 1992, pp. 179-186.
14. Y.L. Srinivas and D. Dutta, "Blending and Joining Using Cyclides," *Proc. ASME Design Automation Conf.*, ASME, New York, 1992.

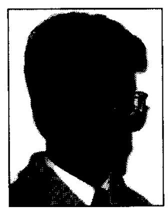

Debasish Dutta is an assistant professor in the Department of Mechanical Engineering and Applied Mechanics at the University of Michigan, Ann Arbor. His research interests include geometric and solid modeling, mechanical design, and manufacturing.

Dutta received his bachelor's degree in mechanical engineering from Jadavpur University, Calcutta, India, and his PhD in industrial engineering from Purdue University. He is a member of ACM, ASME, IIE, and Sigma Xi.

Readers may contact Dutta at Dept. of Mechanical Engineering, University of Michigan, Ann Arbor, MI 48109. His e-mail address is dutta @ caen.engin.umich.edu.

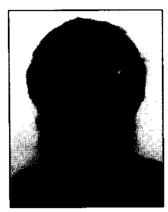

Ralph Martin is a lecturer in computing mathematics at the University of Wales College of Cardiff. He has held various grants, including an SERC ACME grant for research into geometric algorithms for manufacturing engineering. He has published over 30 papers in the field of geometric computing, as well as coauthored two books.

Martin received his BA in physics and theoretical physics and his PhD in engineering, both from Cambridge University, England. He is a fellow of the Institute of Mathematics and Its Applications and a member of the British Computer Society.

Mike Pratt is a partner in the consultancy organization LMR Systems. Until recently, he was head of the Department of Applied Computing and Mathematics at the Cranfield Institute of Technology, England. His research interests include solid and surface modeling, and he also has a strong interest in form features and the CAD/CAM interface.

Pratt received an MA in physics from Oxford University, and an MSc in aeronautical science and a PhD in mechanical engineering from Cranfield. He is a member of the IEEE Computer Society and ASME, as well as other organizations.

Readers may contact Martin at Dept. of Computing Mathematics, Mathematics Institute, Senhhennydd Road, Cardiff, Wales, UK, CF2 4A9, e-mail ralph@cm.cf.ac.uk.