

Analogies from 2D to 3D

Exercises in Disciplined Creativity

Carlo H. Séquin

Computer Science Division, EECS Department
University of California, Berkeley, CA 94720
E-mail: sequin@cs.berkeley.edu

Abstract

Human creativity relies to a large part on our ability to recognize and match patterns, to transpose these patterns into different domains, and to find analogies in new domains to known facts in old domains. In the realm of geometrical proofs and geometrical art, such analogies can carry concepts and methods from spaces that are easy to deal with, e.g. drawings in a 2-dimensional plane, into higher dimensions where model making and visualization are much harder to carry out. Students in a graduate course on geometric modeling are challenged with open-ended design exercises that introduce them to this analogical reasoning and, hopefully, enhance their creative thinking abilities. Examples include: constructing a Hilbert curve in 3D-space, finding an analogous constellation to the Borromean rings with four or more loops, or developing 3D shapes that capture the essence of the 2D Yin-Yang figure or of a logarithmic spiral. The proffered solutions lead to interesting discussions of fundamental issues concerning acceptable analogies, the role of symmetry, degrees of freedom, and evaluation criteria to compare the relative merits of the different proposals. In many cases, the solutions can also be developed into attractive geometrical sculptures.




Figure 1: a) Borromean Rings, b) Hilbert Curve, c) Yin-Yang.

1. Introduction

Where does creativity come from? What is “good design”? How do we know that we have a “solution” to a design problem? These are some questions behind the design exercises that I will discuss in this paper.

The human mind is good at discovering patterns. One might argue that human intelligence and creativity relies to a large extent on our ability to recognize patterns and match them with previously stored patterns. We easily see animal shapes and faces in clouds, or goblins and ghosts in tree trunks at night in the forest. We are intrigued by star constellations if they approximate regular triangles or quadrilaterals, or if they lie on a roughly circular arc. We also recognize and enjoy the regularity in tiling patterns. Often we try

to explain the patterns in one domain with patterns from another domain. For example, Kepler tried to explain the relative sizes of the planetary orbits with the suitably nested circumspheres of the Platonic solids (Fig.2), and there were attempts to draw the periodic table onto simple geometrical objects such as cylinders. Often we use analogies to try to explain a new and unfamiliar domain with a model from a well-understood and intuitively plausible domain; as an example, the model of water flowing over a dam of adjustable height has been used to explain the operation of an MOS field effect transistor [8].




Figure 2: *Kepler's attempt to explain the ratios among planetary orbits*

I believe that one's creativity and imagination can be improved by training one's ability to find such analogies. This is why I include such exercises in almost every course I teach at U.C. Berkeley. I try to improve the students' design skills by raising to a conscious level sequences of thoughts and associations that appear in the search for a solution. We ponder questions such as: What is going on in the design process? Where do good ideas come from? What can we do to enhance the flow of good ideas? When do we know that the design is done?

In my graduate courses on geometrical modeling and computer-aided design, these kinds of questions and the corresponding exercises that I discuss in this paper often involve an inductive step going from a 2-dimensional form to a related 3-dimensional shape, or from a configuration of n elements in a highly symmetrical arrangement to a constellation of $n+1$ and more such elements. Each problem was first solved to my own satisfaction to gain an idea of its possibilities and difficulty. I then clarified the design task and curtailed the solution space so that a structured exercise resulted. Typically, these tasks are attacked by the students in two waves: First they are simply given the short, open-ended problem statement and are asked to think about it and bring their ideas and questions to the next class. In a joint discussion we then agree on the salient features that the solution should have and identify some criteria by which we could judge the quality of the various designs. In spite of the stated constraints and of the focussing effect of our discussions, I normally have the pleasure to obtain new and unexpected solutions that add to the richness of the problem and transcend the previously known solutions. Often I find out later that others had pondered the same questions, and sometimes even wrote a whole book about related issues — as in the case of the delightful book “Orderly Tangles” by Alan Holden [6]. Since the exercises often lead to artistically interesting and pleasing results, they should be particularly well suited for this conference, bridging the gap between logic and the arts.

2. Hilbert Curve

An exercise that dates back to 1983 [9,10] asks the students to develop a 3-dimensional, recursive, self-similar, space-filling, piece-wise linear path inspired by the 2D Hilbert curve [5] shown in Figure 1b. This Peano curve, which in the limit fills the unit square, was discovered in 1891. It can be nicely described by a recursive procedure. The problem statement makes it implicitly clear that we want the new curve to visit all the grid points of a cubic array with $2^n \times 2^n \times 2^n$ points so that no grid point has more than two line segments attached to it. The basic approach is fairly obvious: The overall cube is split into eight equal octants which are visited in a particular sequence. These octants are split into octants again, which are then visited in the same geometric order. All the interesting issues lie in the details!

After the students have pondered the question for a while, we start discussing desirable criteria that one might use to rank-order different designs. For instance, the higher-order Hilbert curves should maintain the symmetry of the starting frame, and that symmetry should be as high as possible. The 3D solution may well have more symmetry than the 2D curve. To achieve that goal, we might want to modify the top level of the recursion in the original Hilbert curve so that it forms a closed loop, thus introducing a second axis of mirror symmetry and leading to a prominent “H” in the second generation (Fig.3a). Other aesthetic considerations may suggest that the extra-long line segment of three unit lengths occurring near the center of the 3rd-order planar curve should be avoided. Ideally one might want to avoid even two subsequent colinear line segments. Similarly, one might want to minimize the number of subsequent elements that lie in the same plane. While three subsequent coplanar segments cannot be avoided, if we want to visit all eight corners of a unit cube, sequences of more than three segments may be avoidable. Ideally we would like to find a simple recursive formulation for such a structure.




Figure 3: a) Closed 2D Hilbert curve, b) 3D Starting frame, c) corner element for 2nd generation curve.

A plausible starting frame consists of the closed loop shown in Figure 3b. To use this frame as the basic corner element at the next level, one of its eight segments, i.e., the dark one in the lower part of Figure 3b, must be opened up, and two new connections to adjacent, identical corner elements must be created by suitably twisting the L-elements that were formerly attached to it (Fig.3c). To construct the second generation of the 3D Hilbert curve (Fig.4a), we place eight such modules, reduced to half-size, into the eight octants of the original cube. They have to be properly oriented, and some units have to be mirrored (shown darker in Fig.4a) so that they can be connected readily with their neighbors. This should be done in such a way that a closed path results that basically follows the path of the original frame (Fig.3b). I have built a physical implementation of such a second generation structure from 64 plastic 3/4-inch pipe corner pieces. The problem with constructing larger physical structures lies in the fact that, regardless of the level of recursion, two halves of the sculpture are connected only with either two or four pipe segments, which renders the physical structures rather weak. On the other hand, in the virtual space of computer modeling, the process can be continued without limits using ever smaller copies of the original module. The basic design can then be turned into an impressive virtual sculpture with a suitable choice of pipe dimensions, texture, and coloring (Fig.4c).

For Hilbert curves of order three and higher, some interesting choices have to be made. I was able to design a 512-segment 3D pipe in which there are never more than three coplanar line segments. However, I had to start with a different corner element for the 2nd-generation curve. Rather than removing the dark segment in the base frame (Fig.3b), I chose to remove one of the elements adjacent to it (of which there are four to start with) and to orient the connecting L-pieces so that they turn away from the plane last visited inside the base-frame. These corner elements now show C2 symmetry around the mid-point of their middle segment, and they can now be properly oriented in all eight corner positions to form a 2nd-generation Hilbert path (Fig.4b). To produce the corner element for the 3rd generation, we again break open a segment near one of the corners and suitably twist the adjacent L-elements outwards. This unit can then be assembled into a symmetrical closed loop with carefully chosen orientations and mirroring operations. The drawback of this solution is that the connection operation needs to modify one of the lowest level corner elements, thus a simple recursive composition of eight identical corner elements is not possible.




Figure 4: (a), b) 2nd- and c) 3rd-generation virtual 3D Hilbert pipes, with 64 and 512 pieces respectively.

The solution presented in Figures 4a,c, has the same basic symmetry, but it exhibits sequences of four subsequent coplanar pipe segments; on the other hand, I was able to describe it with a nice recursive formulation. If a closed curve is desired, then one has to change the orientations of the first and last corner elements at the top level of the recursive procedure. The approach generalizes to higher dimensional cubes [3]. The initial starting frame can always be seen as an n -bit reflected Gray code which runs through all permutations of an n -bit string in such a way that each string differs from its predecessor in only a single bit. Scaled-down versions of this traversal of the starting frame are then placed — with suitable orientation — into each “corner” of the original hyper-cube.

3. Symmetrical Constellations of Interlocking Loops

Two tightly intertwined rings form a simple yet intriguing configuration that seems to have symbolic meaning in several cultures (Fig.5a). An attempt to place three loops in space as compactly and as symmetrically as possible, leads to an arrangement known as the Borromean rings (Fig.1a, Fig.5b). Individual pairs of rings are not actually interlocked; the configuration only holds together when all three rings are present. However, when we try to place three perfectly toroidal rings into a constellation of high symmetry, the result is a pairwise interlocking configuration with three-fold symmetry (Fig.5c).




Figure 5: a) Two interlocking rings, b) tight Borromean configuration, c) three interlocking ring pairs.

Next, we aim to cluster more than three rings around the origin, without much concern whether individual pairs of rings mutually interlock. The task given to the students was to find a constellation of four loops with the highest possible symmetry. Every loop should be in an identical position within the constellation, so that the basic symmetry operations transpose any one loop into any other loop.

If one has not seen the solution beforehand, this problem turns out to be surprisingly difficult. Two approaches have proven helpful to guide the students to finding a solution. The first one is to ask what symmetry groups one might possibly expect. After some contemplation, one finds that with four rings it has to be the tetrahedral or the octahedral group. The second approach starts with the notion that one might want to interlock four **triangles**. The choice of triangles to represent the loops seems natural, because each loop has to interact with three other loops, and if we want to do this in a symmetrical manner, we should choose a loop with 3-fold symmetry. Given that we want to place four triangles symmetrically in 3D space, we need to define the positions for twelve vertices — in a symmetrical manner. So the question then turns to how one can place twelve vertices uniformly and symmetrically onto the surface of a sphere. The insight to this secondary problem might come from thinking about the densest sphere packing, or from contemplating the Platonic and Archimedean solids and looking for the occurrence of the number 12 — preferably in an object with tetrahedral or octahedral symmetry. When I initially contemplated this problem, I first thought of the twelve edges on a cube. So I placed the vertices at the midpoints of these edges and connected them into 4 triangles — and the solution emerged almost immediately (Fig.6a).

I implemented this configuration as a physical sculpture from 4-inch cardboard tubes [9], spray-painted with copper enamel on the outside, and with a touch of fluorescent yellow near its center— which made the sculpture glow on the inside when hit with indirect sunlight (Fig.6b). It should be pointed out, that in this arrangement every pair of triangles mutually interlocks; cutting away one triangle would still leave the other three entangled. Also, the topology of Figure 6b is the mirror image of that of Figure 6a.

I continued my quest by looking for constellations of five and more intertwined loops. For five loops, an extension of the process that had helped me find the 4-triangle structure was employed — i.e., I tried an inductive approach. Each of the five loops would have to interact with four others, thus using squares seemed like a reasonable start. This then required twenty vertices positioned symmetrically in space. The

twenty vertices of the pentagon-dodecahedron offered themselves conveniently, and it did not take long to find a grouping of the vertices into four planar polygons — which however were rectangular rather than square (Fig.6c). Also the resulting constellation does not carry the full symmetry of the Platonic solid from which it was derived, it just has one axis of 5-fold symmetry (C_5) and five axes of C_2 symmetry.




Figure 6: a) Four symmetrically clustered loops, b) physical realization, c) five “Borromean” rectangles.

It is interesting to study the interlocking pattern of this structure. No two rectangles interlock! It looks like a more complicated Borromean arrangement in which each loop surrounds exactly two other ones in a cyclical relationship. Encouraged with this result, we might try again to look for a Borromean arrangement with four loops. However, a little conceptual reasoning will soon let us see the difficulty of this quest. Let’s use the notation $A \rightarrow B$ to indicate that loop A encircles loop B on the outside. Thus the Borromean rings have the cyclic relationship indicated in Figure 7a. If we try to draw a similar diagram for the 4-ring constellation, then we run into the difficulty that in a complete graph with four vertices, there are three edges joining at each vertex; thus the number of incoming and outgoing arrows cannot be made the same everywhere. To draw a symmetrical diagram in which all vertices are identical, we would have to use double-headed arrows, which we can interpret as an indication that the two rings (vertices) connected by such a double arrow are mutually interlocking (Fig.7b). On the other hand, the complete graph with five vertices has four edges joining at every vertex, and we can readily draw such a graph with two incoming and two outgoing arrows at each node (Fig.7c). This corresponds to the configuration of five loops discussed above and shown in Figure 6c.




Figure 7: Interlocking schemes: a) Borromean rings, b) 4 rings, c) 5 rings, d) 4 “Borromean” rings

Before I had a chance to push my quest much beyond the constellation with five rings, I stumbled onto the delightful book “Orderly Tangles” by Alan Holden [6]. This is an invaluable resource containing dozens of such symmetrical, interlocking constellations with as many as twenty loops. Figure 8b shows a computer simulation of a tangle with ten triangles inspired by a model built by Holden. In this case, the vertices

of the triangles lie on the midpoints of the thirty edges of the dodecahedron. The inductive reasoning outlined above for the steps from three to four and then to five loops can be continued. Once one understands the search procedure, it is not too difficult to find the more complicated tangles.




Figure 8: a) Four “Borromean” triangles, b) ten interlocking triangles, c) tangle of five tetrahedra (SLS).

Holden’s book also contains a Borromean configuration formed by four triangles; its linking logic corresponds to Figure 7d, and it is realized by a 3-ring Borromean configuration that rigidly holds in place a fourth, non-entangled triangle (Fig.8a). Holden also shows how such symmetrical tangles can be carried beyond just simple loops; he shows models of interlocking tubular tetrahedral frames. Another realization of the classical tangle of five tetrahedra with all 20 vertices lying at the corners of a dodecahedron is shown in Figure 8c. This part has been constructed with Selective Laser Sintering (SLS), one of the emerging layered Solid Free-Form (SFF) fabrication technologies [7].

4. Yin-Yang

Yin-Yang symbolizes the two complementary forces that comprise the Tao, the eternal dynamic way of the universe: Yin is the earthly, dark, passive, or female principle. Yang is the heavenly, light, active, or male force. Geometrically, the Yin-Yang symbol divides a circle into two complementary halves that in some sense are “opposites” of one another. The task given to the students was to find a corresponding partitioning of a sphere in 3-space. The richness of the solutions proposed by the students in the Fall 1997 course CS 285, Solid Free-Form Modeling and Fabrication, was unusually rewarding (Fig.9,10).

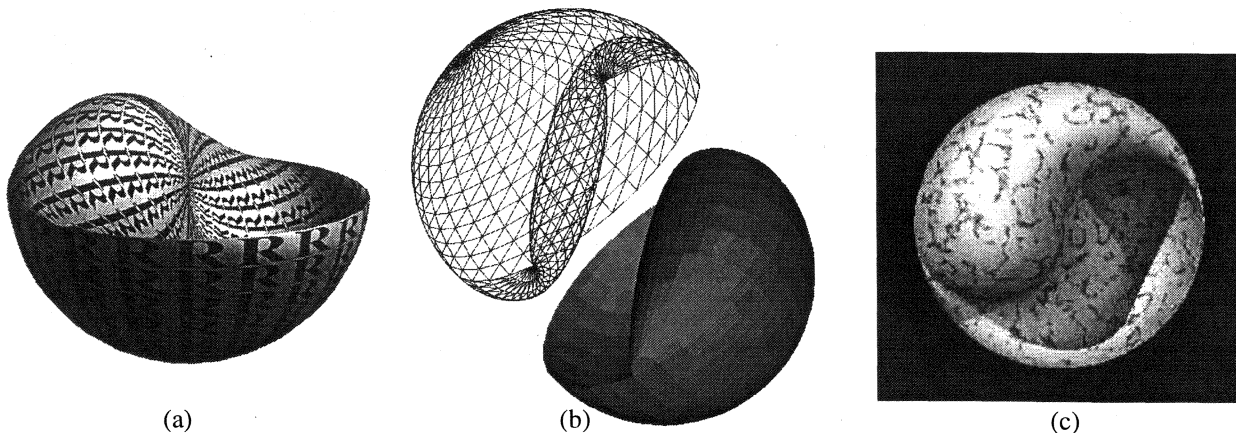


Figure 9: Various solutions to constructing a 3D Yin-Yang.

The most often proposed solution was a sphere cut into two halves with a “band-saw” following the path of the 2D Yin-Yang dividing line. Some solutions were offered as clay models (Fig.10a), others as machined parts (Fig.10b), or as sophisticated computer renderings (Fig.10c). While I feel that this is not the best solution, since it is just an extruded extension of the 2D figure, this is also a shape celebrated by great artists such as Max Bill (Fig.10d).




Figure 10: *The most pervasive solution for a 3D Yin-Yang: A cut with a S-shaped developable surface.*

Still, I was more intrigued with attempts to cut the sphere with a surface that curves in both directions. A key question then arises: should the two halves be identical or mirror images? The most innovative proposal came from a couple of students who reasoned, that the true analogy of going from 2D to 3D demanded that the sphere be cut into **three** identical parts — possibly colored with the three primary colors red, green, blue.




Figure 11: *Yin-Yang symmetry groups: a) mirroring, b) cyclic rotational, c) glide plane reflection.*

The crucial characteristics to order and classify this plethora of shapes is the symmetry of the surface that divides the sphere. The following fundamental possibilities exist: The trivial solution cuts the sphere with a plane through its center; but this does not exhibit any features of the Yin-Yang icon. A more interesting class of dividing surfaces has C2 symmetry with respect to some axis through the sphere center; this results in two congruent halves (Fig.9a,b). The “band-saw-cut” solutions (Fig.10) also fall into this class. A generalization allows C3 symmetry around this axis and would thus cover the case of three identical sub-components. The third and, to me, most interesting class, has glide symmetry, which brings the dividing surface back onto itself when it is rotated 180 degrees and then mirrored on a plane perpendicular to that axis; this leads to complementary mirror-parts (Fig.9c). This shape is most defensible on philosophical grounds; we want to create two halves that are not identical but rather complements of one another. The most beautiful formulation of such a shape (Fig.12a) is composed of three spherical surface pieces and two cyclides [2]. This shape was also discovered by C. E. Peck in 1992.

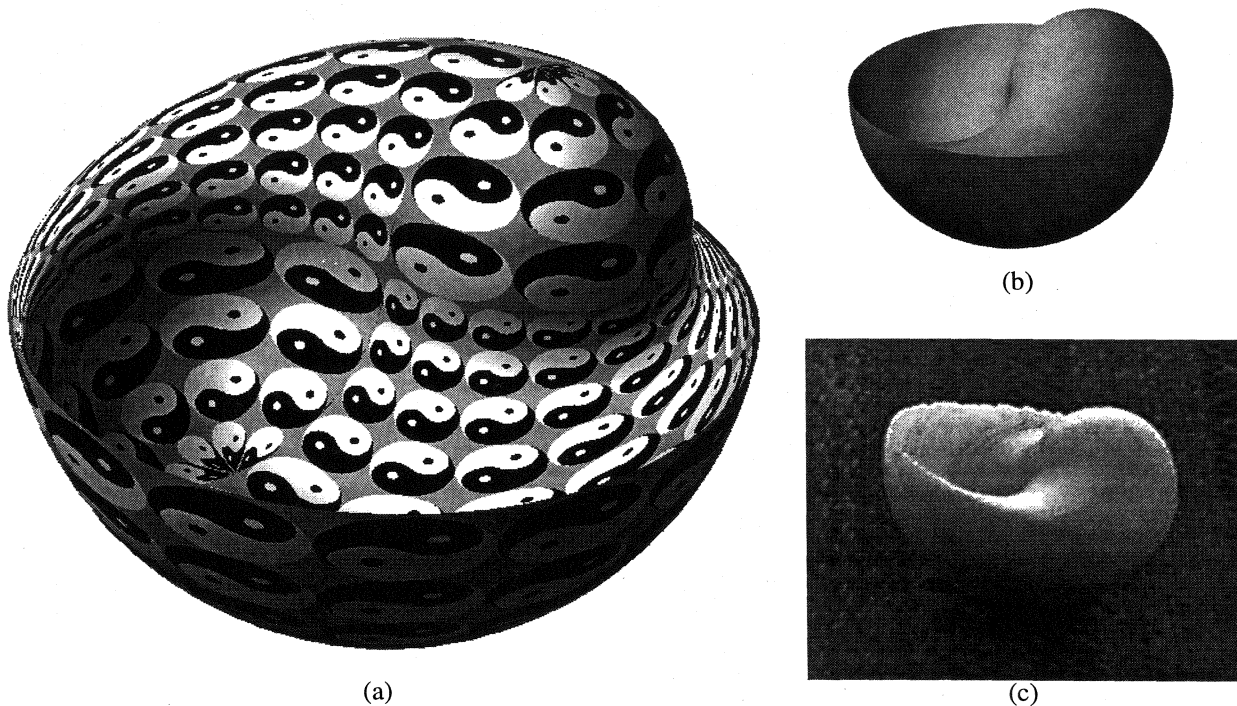


Figure 12: 3D Yin-Yang models with mirror complements: a) cyclid-based, b) torus-based.

Attempts at machining such a shape on a milling machine run into the problem, that the part shows a concave groove that leads into a point with infinite curvature, which cannot be cut with any tool of finite dimension. The 2D Yin-Yang has constant curvature and uses only circles of two radii that differ by a factor of two. I found a corresponding solution in 3D that replaces the two cyclid surfaces with two tori in which the major radius R is twice the size of its minor radius r . This shape can readily be described as a Boolean expression of its five curved shapes and a few half-planes. The resulting shape is shown in Figure 12b, and an early attempt at machining it on a milling machine in Figure 12c.

5. Spiral Surface

The logarithmic spiral (Fig.13a) is a fascinating curve, and may be considered the ultimate solution to self-similarity at arbitrary scales. As in the first problem dealing with the Hilbert curve, we might ask what an analogous curve through 3D-space might look like. One might argue that a 3D spiral curve should (eventually) pass through all possible directions emerging from the origin of the coordinate system, and, at the

same time, gradually move outwards at an exponential rate. The problem of visiting all points on a sphere with a continuous smooth path has been addressed by Dan Asimov's "Grand Tour" [1], which, given a long enough time span, will approach every point on the surface of a sphere with arbitrary closeness. All we have to do, is to let the radius grow exponentially as a function of time.

However, the task we want to focus on here, is to find a **surface** that captures the spiral properties. Ideally, we would like to obtain a spiral intersection curve whenever the surface is cut with an arbitrary plane through the origin — however, this would be asking for too much! But can we get spirals in at least three cutting planes that are mutually perpendicular to one another?




Figure 13: a) *Logarithmic Spiral*, b, c) *emerging pipe-cleaner skeleton*.

This problem has not been presented to any student yet. The way I approached it myself was to make a wire skeleton from pipe cleaners (Fig.13b) that contained the three coordinate axes (black) and three spirals in the three main coordinate planes (white). Additional pipe cleaners (grey) produced some connectivity among these spirals and formed partially spherical shells (Fig.13c). It became clear that it was not possible to connect all spiral branches smoothly with one another; some jumps from one "onion shell" to another had to occur. Thus the surface cannot be totally closed. It turns out that this is a useful feature, since it would be rather dull, if we could not view the internal structure of this surface. Later it became apparent that the edges that had to be introduced to avoid the discontinuous jumps from one shell to the next inner/outer one could themselves take on the shape of spirals. What serendipity!




Figure 14: a) *Spider pattern with patch subdivision lines*, b) *paper model of Spiral Surface*.

One emerging solution exhibited D3-symmetry along the $\{111\}$ space diagonal in the coordinate system in which I had placed the original three spirals; and it showed six openings with helical edges leading inwards toward this axis. At this stage, I started to build paper models to get a better feeling for the surface topology itself. I designed the spider pattern shown in Figure 14a and made several suitably scaled copies. These were joined together in a nested manner, where the long arms of one spider join the short arms of the next larger spider. The result is shown in Figure 14b.

Subsequently this surface was modeled on the computer. Using three Bézier patches to form the shape of one “L” with a 60-degree corner, it takes 18 patches to compose one complete spider (Fig.14a). The boundary constraints to guarantee smooth continuation of the patches are not hard to derive; and the inner control points of the cubic patches are adjusted to minimize any apparent bumps. Finally, Jane Yen added highlights and shadows and rendered that surface (Fig.15a) with the Blue-Moon Rendering Tools [4]. The next step is to make the surface thicker and to derive a solid description from which a 3D model can then be built with one of the layered Solid Free-Form (SFF) fabrication techniques [7].




Figure 15: a) Virtual model of Spiral Surface, b) half that surface showing spiral curve in cross section.

6. Discussion and Conclusions

The design exercises discussed in this paper are a good example of an activity that bridges the realms of logical reasoning on the one hand, and of intuitive and even artistic contemplation on the other. The region between art and mathematics is particularly suitable to study the creative process on (somewhat) open-ended design problems. The problem statements are loose enough to allow the mind to roam free and to come up with potentially wild and unorthodox solutions. At the same time, these geometrical puzzles possess enough structure and quantifiable properties so that one can apply an acceptable metric to the results and rank-order different solutions. This process brings wild, far-flung, non-sequitur creativity into a more disciplined mode where there are design solutions of defensible quality.

The presented problems start with the construction of a simple one-manifold, the 3D Hilbert curve, where connectivity, symmetry, and a recursive formulation are the dominant concerns. The subsequent tasks increase in complexity, adding topological considerations of linking behavior, and evolving from one-manifolds (curves) to two-manifolds (surfaces).

The methods employed to tackle these problems vary, but a dominant role is played by inductive reasoning and a judicious use of symmetry. The key challenge of the 3D Hilbert curve is to find a recursive formulation to build a generic corner element with the desired connection properties at the corners so that the connectivity among the eight octant cuboids stays the same in each generation. In searching for orderly tangles of interlocking loops, it was most productive to determine the expected symmetry group, and then place a suitable number of vertices evenly onto the surface of a sphere so as to span the polygonal implementations of the loops. For the two surface-related problems, paper and pencil, or even computer drawings did not seem adequate to explore the potentially very large solution space. The use of clay, wire-mesh, and/or pipe-cleaners seemed to help a lot in the visualization of the problem and its possible solutions.

What all the solutions have in common is that the “best” solutions in terms of maximum analogy with the starting shape also have a high aesthetic appeal by themselves — which adds a special bonus to these exercise tasks. This bonus becomes even greater with the advent of solid free-form (SFF) manufacturing technologies that allow to turn these virtual design artifacts into nice physical sculptures.

7. Acknowledgments

Some aspects of this work were partially supported by the National Science Foundation under a research grant to study the designer/fabricator interface for rapid prototyping of mechanical parts. The help of Jane Yen and Jordan Smith with the rendering of some of the virtual sculptures is also gratefully acknowledged.

References

1. D. Asimov, “The Grand Tour: A Tool For Viewing Multidimensional Data.” *SIAM Journal of Science & Stat. Comp.*, Vol. 6, pp. 128-143, (1985).
2. W. Boehm, “On Cyclides in Geometric Modeling.” *CAGD*, Vol 7, pp. 243-255, North Holland, (1990).
3. E. N. Gilbert, “Gray Codes and the Paths on the N-Cube.” *Bell System Tech. Jour.* 37, pp. 815-826, (1958).
4. L. Gritz, “Blue Moon Rendering Tools Home Page,” <http://www.bmrt.org/>
5. D. Hilbert and S. Cohn-Vossen, “Geometry and the Imagination.” Chelsea, New York, 1952.
6. A. Holden, “Orderly Tangles.” Columbia University Press, New York 1983.
7. D. Kochan, “Solid Freeform Manufacturing: Advanced Rapid Prototyping.” *Manufacturing Research and Technology 19*, Elsevier, Amsterdam, New York, (1993).
8. C. Mead and L. Conway, “Introduction to VLSI Systems.” Section 1.15, Addison Wesley, Reading MA, (1980).
9. C.H. Séquin. “Creative Geometric Modeling with UniGrafix.” Technical Report UCB/CSD 83/162, U.C. Berkeley, CA (Dec. 1983).
10. C.H. Séquin. “More... Creative Geometric Modeling.” Technical Report UCB/CSD 86/278, U.C. Berkeley, CA (Dec. 1985).