Modular Toroids Constructed from Nonahedra

Yifat Amir and Carlo H. Séquin

CS Division, U.C. Berkeley; yifatamir@berkeley.edu; sequin@cs.berkeley.edu

Abstract

Inspired by a piece of artwork by Bente Simonsen at Bridges 2016, we explore various ways of making modular polyhedral toroids from the same basic nonahedron building block. Since the chosen building block cannot be an ideal, regular Johnson solid, we use greedy optimization methods, such as gradient descent, to minimize the deviations from regular n-gons for the facets of the nonahedron, while maintaining the planarity of the connecting faces for each respective toroidal configuration. In general, we formulate our objective to maximize overall symmetry and create a visually aesthetic result. Furthermore, we have constructed polyhedra with a genus as high as 11 using a 3D printer.

An Intriguing Puzzle

Figure 1a shows a sculpture by Bente Simonsen that was exhibited at the Bridges 2016 Art Exhibit [3]. At first glance, it seems to be nothing very special. But it gets much more interesting after one takes a closer look and reads the accompanying cryptic description:

"The sculpture is created from five nine-sided polygons [sic], each with 3 square sides and 6 pentangle [sic] sides, all with the same side length, and fitted together to build a pentagonal torus."

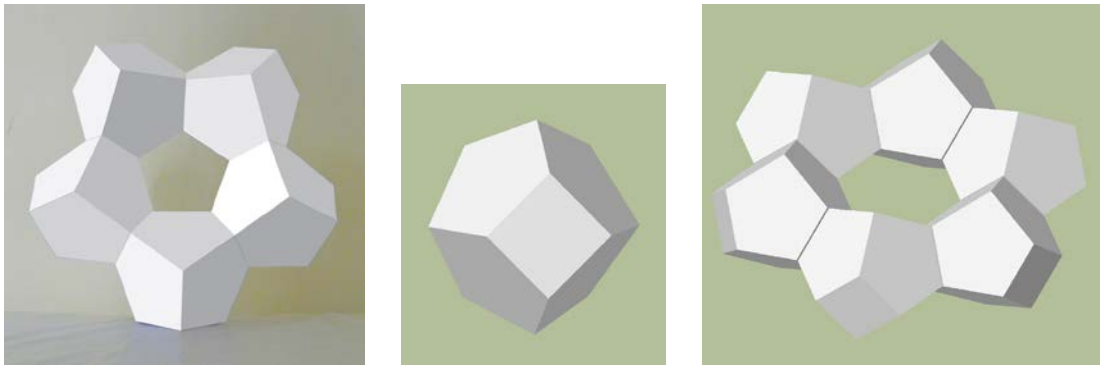

(a) (b) (c)

Figure 1: *(a) Bente Simonsen:"Torus" [3]; (b) nonahedron building block; (c) 6-module toroid.*

To a geometrically astute observer, a couple of intriguing puzzles present themselves. Figure 1b shows what the implied nonahedron module might look like. A check in Wikipedia [4] reveals that this *truncated triangular bipyramid* is a "near-miss" Johnson solid. Johnson solids are convex polyhedra made from all completely regular n-gons. But the shape in Figure 1b cannot be constructed from regular pentagons and square faces. Either the edges need to be of somewhat different lengths, or some faces need to be non-planar and/or have a variation of interior angles. Moreover, the building block seems to have 3-fold rotational symmetry, and this implies that the angle between the face-planes of the square faces must be 60°. This further implies that *six*, rather than five, of those modules would readily form a modular, 6-fold symmetrical, toroidal ring (Figure 1c). Therefore, we conclude, that some "cheating" is going on in Figure 1a. The question then arises, whether all faces in this toroidal shape can be made planar, when all the edges are constrained to be of equal length. If it is not possible, then how can deviations from regularity be hidden most effectively? To get some understanding of these issues, we analyze the degrees of freedom (DOF) of the shapes in Figure 1 and investigate how various imperfections can be traded off against one another.

Variations of the Nonahedron (Enneahedron) Module

Let's first explore some design options for the 9-faced, near-miss Johnson nonahedron in isolation.

A: We could start with the (yellow) corner of the pentagonal dodecahedron consisting of three regular pentagons, and put two of them (yellow and green) together so that three horizontal pairs of edges are fused (Figure 2a); but this results in non-planar quad openings.

B: Next, we start with 3 perfect unit-square faces and place them symmetrically around the *z*-axis, so that the 3 pairs of closest corners are a unit length apart. Subsequently we choose a point on the *z*-axis that is exactly one unit-length away from the upward-pointing corners of the square faces, and we use this vertex to complete the three yellow pentagons. We also apply the same construction at the green bottom pentagons. This construction leads to non-planar pentagons (Figure 2b).

C: A third option is to enforce planarity of the yellow pentagons by raising the vertex on the *z*-axis, forming the apex of a 3-sided pyramid, to an appropriate height – and also applying the same construction to the green bottom pentagons. But this lengthens these pyramid edges considerably (Figure 2c).

D: Since the pentagons are no longer regular, what can we also profit by allowing the quads to become rhombi? Can we now make all faces planar *and* keep all edges the same length? (No, we cannot!)

E: Perhaps the best compromise lies in a mixture of all these approaches, where we allow some violation of *all* the constraints of a perfect Johnson solid. Hopefully, all these violations can then be kept small enough, so that they are not immediately visible in any display of this polyhedron (Figure 2d).

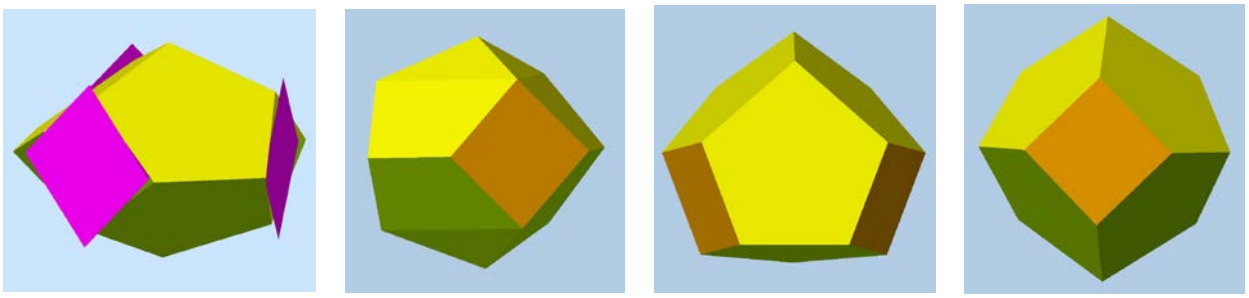

(a) (b) (c) (d)

Figure 2: *Various nonahedral modules: (a) starting with 6 regular pentagons, (b) starting with 3 planar squares and all equal edge lengths, (c) all planar faces, (d) a visually "optimized" module.*

To study what is possible and to find a combination of small violations that gives the "best" overall result, we need to introduce an efficient parameterization for the geometry of this solid and devise a tunable "cost" (or "penalty") function that will allow us to measure the deviation from perfection.

Parameterization

Our building block has 14 vertices, and so we have to set the values of 42 coordinate components. But doing the needed optimization in 42-dimensional space would be inefficient. We can get a complete geometry description with much fewer parameters, if we maintain strict 12-fold "D_{3h} symmetry" (Schönflies notation), also known as "*223" symmetry (Conway notation), composed of a 3-fold rotation axis (the *z*-axis), and 1 horizontal and 3 vertical mirror planes. Maintaining D_{3h} symmetry is good for obtaining visually pleasing results, and it reduces the number of parameters necessary for representing the nonahedron from 42 to 4. Since our optimization is scale-independent, let's start by making the 3 equatorial edges of length 1.0 and placing them into the *x*-*y*-plane at a distance, *ex*, from the origin. The top and bottom pyramid tips lie on the *z*-axis, and they will be placed at a height of ± *tz*. The "corner vertices" at the "shoulders" need two parameters: a radial distance, *rc*, and a height, *hc*. These four parameters are sufficient to capture all variations of our 12-fold symmetrical building block; and any optimization can also be carried out in this 4-dimensional search space.

Cost Function

Our cost function to be optimized will have three major components, measuring separately the differences in edge length (*ELD*), the deviations from the internal angles of a regular polygon (*PAD*), and the degree of non-planarity for all faces (*FNP*).

Edge-Length Differences (ELD): We determine the average length, *avl*, of all 21 edges, and then calculate the sum of squares of all the deviations from this average length for all edges.

$$avl = \sum (edge\text{-}length_i) / 21; \quad ELD = \sum (length_i - avl)^2 .$$

Polygonal Angle Deviations (PAD): We measure the internal angles at all corners in all faces, and then calculate the sum of squares of all the deviations from the corresponding angles in a regular *n*-gon.

$$PAD = \sum (angle_{s4} - 90)^2 + \sum (angle_{p5} - 108)^2 .$$

Face Non-Planarity (FNP): For each face we find the best-fitting plane [1], then calculate the offset, $dist_i$, of the face-contour vertices from this plane, and sum the square of these distances:

$$FNP = \sum (dist_i)^2 .$$

Weighted Combination: The above three terms can be combined into a single cost function, where each component is given some weighting coefficient that expresses the "visibility" of the corresponding type of deviation from the ideal and accounts for natural differences in magnitude among the three terms:

$$COST = \alpha * ELD + \beta * PAD + \gamma * FNP .$$

For instance, if we do not care about the regularity of our faces, we can simply set $\beta = 0$.

Gradient Descent

With a particular cost function in place, we can tackle the task of minimizing this cost function via gradient descent. For each independent parameter, we look at the effect that a small change in this parameter value has on the overall cost function. We combine these parameter sensitivities into a resultant vector that includes each parameter with a weight that is proportional to the contribution that this parameter can make to reduce the overall cost function. This defines the direction of the local gradient, and we now let our system take a small step in this gradient direction.

We still have a choice for the values of α, β, and γ. Originally these values may be set to yield about equal sensitivities with respect to the three different components of the cost function. But in the end, they will have to be fine-tuned by the user to yield the most aesthetic result. The tastes of different users may vary, and the sensitivities for different realizations of the nonahedron may be different. For instance, if the nonahedron has a mirror-like finish, then any non-planarity of its faces would be much more apparent.

Of course, we could readily use more than just three weighting functions, and introduce separate weights for the planarity of the quadrilaterals and for the pentagons. We can also let some components of the cost-function dominate the others, e.g., demand perfect planarity for the quadrilaterals. We may then end up with a fully constrained system with no remaining degrees of freedom. One such example is shown in Figure 1b where the quadrilaterals were perfect squares and all edges were forced to be of equal lengths.

Some Optimal Parameter Settings and Corresponding Results

Table 1 shows the geometric parameter values for various "optimized" nonahedra. The first one (by Séquin) was optimized based on visual inspection of an interactive display, while adjusting the four defining parameters manually. Others were optimized by gradient descent with different settings for α, β, and γ. The values of these parameters cannot be compared directly, since students wrote their own optimization programs with somewhat different ways to capture regularity and planarity of faces. The main consequence for the resulting geometric parameters seems to be that in the visually more rounded shapes, the value *tz* is considerably larger. This improves the planarity of the pentagons, but increases the variations in the edge lengths.

Table 1: *Geometrical Parameter Values for the Isolated Nonahedron:*

param. (a)	Séquin (b)	Jeon (c)	Taori (d)	Jawale (e)	Amir (f)
ex	1.111	1.146	1.104	1.070	1.045
rc	0.990	0.988	0.931	0.999	0.895
hc	0.750	0.741	0.707	0.679	0.642
tz	1.333	1.252	1.072	1.030	1.012

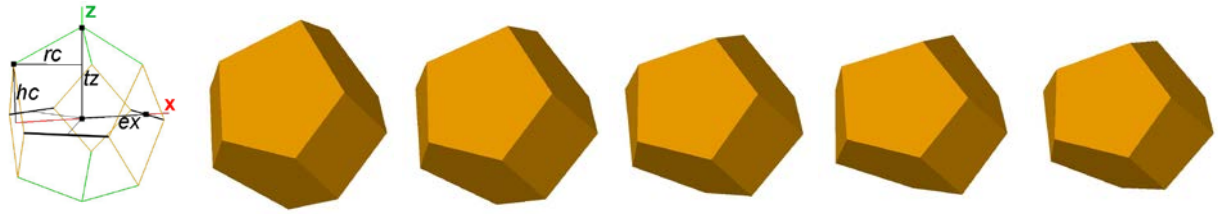

(a) param. (b) Séquin (c) Jeon (d) Taori (e) Jawale (f) Amir

Figure 3: *Various "optimized" nonahedra, displayed from identical viewing directions.*

Constructing Toroidal Rings

Next, we will use this parameterized nonahedron as a building block for the construction of various toroidal configurations. We make small changes to the shape of the nonahedron, so that several identical copies can be assembled seamlessly into different symmetrical toroidal rings. The symmetry of the anticipated structures will define new constraints for the individual nonahedron building block.

6-Ring

Because of the 3-fold rotational symmetry of our building block, pairs of quad faces occur with a dihedral angle of 60° between them, and a ring composed of six nonahedra is the most natural assembly (Figure 1c). In order to glue these building blocks together with their quad faces, we make these faces planar. This can readily be achieved by calculating *rc* to lie directly above the middle of the diagonal of the quad. However, *hc* could still be retained as an independent parameter, if we allow the quad faces to assume a slightly rhombic shape. This extra DOF can then be used to even out edge-length variation or to improve the planarity of the pentagons.

5-Ring and 7-Ring

Next, we try to arrange five of these nonahedra into a 5-fold symmetrical toroidal ring to emulate Bente Simonsen's design [3]. This is not possible without introducing some further distortions that break the 3-fold symmetry of the individual modules. The angle between the face-planes of the quadrilateral contact faces has to be changed from 60° to 72°. In doing this, we can follow the implied recipe given by the original artist, (Section 1), and keep all the quad faces as perfect squares and then aim to make the length of all the other edges the same. There are just enough DOFs to do this. We start by forming the tunnel of the toroid as a regular pentagon with unit-length edges. To this, we attach five vertical squares radiating outwards with angles of 72° between them. Then we add another five vertical squares in a tangential direction to the emerging torus and move them radially outwards until the outer edges between them, which lie in the equatorial plane, also assume unit length. This now defines the overall geometry, except for the ten top and bottom pyramid vertices. The position of these vertices is found by putting them at a unit distance from the three vertices to which they connect. The result is shown in Figure 4a. The drawback is that the pentagons are noticeably non-planar.

Just as in the case of the individual nonahedron, a tradeoff can be made between face planarity and edge-length uniformity. Since the slope of the inner (cyan) pentagons is determined by the inner edges of the five (magenta) quad contact faces, a planarity constraint defines the ratio between the height and the radial distance of the pyramid vertices as measured from the inner equatorial edges. There are now only

three DOFs left: the length, *ie*, of the inner equatorial edges; the radial distance, *fx*, of the outer quad face, which implicitly defines the length, *oe*, of the outer equatorial edges; and the height/distance of the pyramid vertices. Rendering the outer pentagons also planar, uses up two of these DOFs (the first three vertices of a pentagon define a plane that places one positional constraint each on the other two vertices). The last remaining DOF has been used to minimize the variations in the edge lengths (Figure 4b).

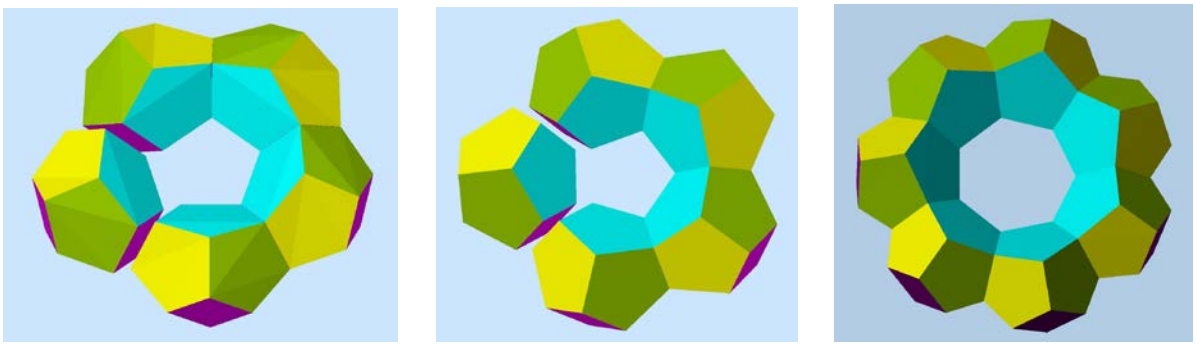

(a) (b) (c)

Figure 4: *(a) Ring of five nonahedron building blocks with perfect squares and all equal edge lengths; (b) keeping all pentagons planar; (c) heptagonal toroid with planar pentagons.*

Given the objective of finding the visually most pleasing toroidal structure, we can take into account that not all deviations from regularity are equally visible. While a deviation from a perfect square on the outer quad faces would quickly become noticeable, the exact shape of the contact quads is much less critical. In particular, any deviation of the corner angles of these quads is very hard to judge, and even deviations in edge length in these quads are less noticeable. Thus we can forego the original constraints placed on these quads, and use the additional DOFs gained to further increase the edge-length uniformity of the other, more visible edges. Programming this is geometrically more challenging, but results in a slightly more improved 5-unit toroid (Figure 5a). The benefit is, that once the necessary dependencies have been put in place, they can also be used with only slight modifications to design an optimal 7-module toroid.

The approach for the 7-unit toroid is the same as for the 5-ring, except that we start with a regular tunnel-*heptagon* of side length 1.0 in the *x*-*y*-plane, symmetrically placed around the origin. We then attach *seven* vertical quadrilaterals angled outward with 360°/7 between them. The quad faces are again defined by their horizontal diagonal, *hd*, and their vertical half-diagonal, *hc*. The parameters *fx*, *tz*, and *tr* are defined in the same way as for the 5-ring, and the available DOFs can be used in the same way (Figure 4c).

(a) (b)

Figure 5: *Toroidal rings: (a) with five modules; (b) with seven modules.*

Figure 5 shows toroidal rings composed of five and seven nonahedral units, respectively. In both cases, those units have been optimized to have all planar faces. The exposed quad faces are perfect squares, while the quad contact faces were allowed to take on irregular shapes to free up additional DOFs that allow to minimize edge-length variations.

An 8-module Toroid

We can also glue together nonahedra by their pentagonal faces. Figure 6a shows a sketch of a possible construction of an 8-module toroid. In this case, the tunnel of the toroid is formed by a regular octagonal prism surface. To this, we attach eight (green) pentagons radiating outward with angles of 45° between them (Figure 6b). Since these contact faces are not directly visible, they can be somewhat irregular, as long as they are planar and mirror symmetric in the vertical direction; they provide three DOFs. The outer tips of the yellow pentagons add two more DOFs. The outer openings of the eight cells are now closed off with two (magenta) quads and two (cyan) pentagons, maintaining overall dihedral D_8 symmetry (Figure 6c). This introduces another three DOFs associated with the two outermost vertices, which are placed with C_2 symmetry around a horizontal symmetry axis. We now have eight DOFs total, because the nonahedral building block here is distorted in a way that destroys its mirror symmetry. On the other hand, enforcing planarity constraints on all faces uses up four of these DOFs.

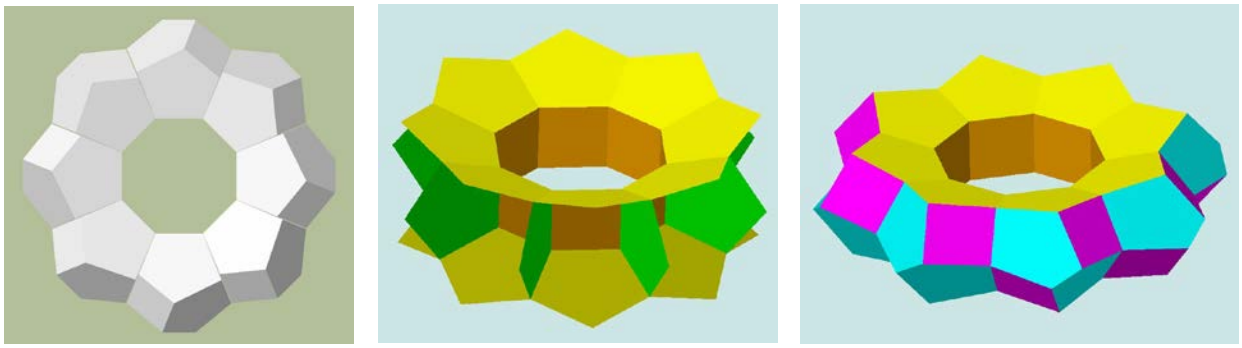

(a) (b) (c)

Figure 6: *(a) Sketch of an 8-module toroid; (b) inner geometry of a symmetrical construction; (c) a pleasing solution with near-planar faces.*

Symmetrical Structures of Higher Genus

Drawing inspiration from [2], the nonahedral building blocks can also be used to make "handle-bodies" of higher genus, where the "handles" are formed from partial toroidal rings of the types discussed above.

Planar Networks

First, we can readily assemble multiple instances of the hexagonal toroid shown in Figure 1c, placing them into a single plane. Figure 7 shows planar structures of genus 2 and genus 3.

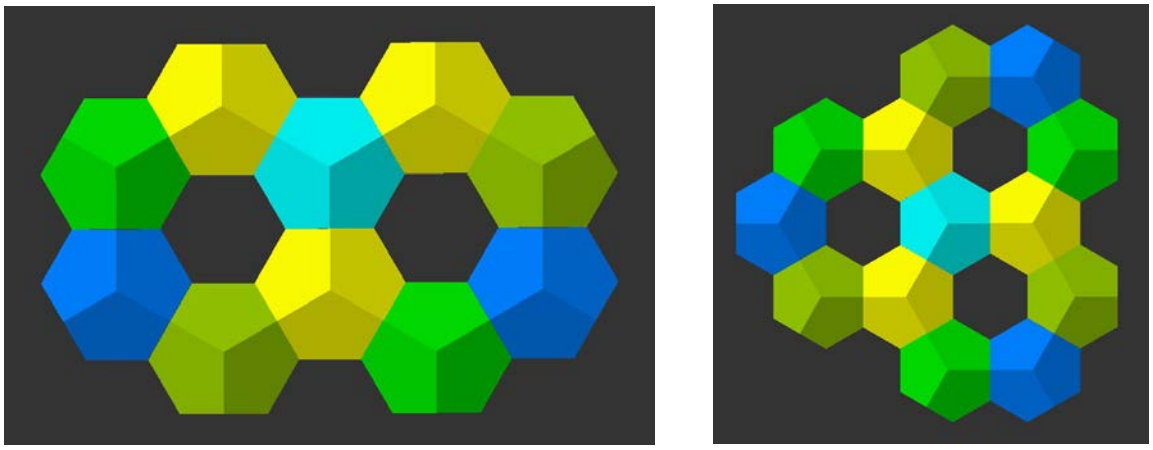

(a) (b)

Figure 7: *Higher-genus objects constructed from hexagonal toroids: (a) genus 2, (b) genus 3.*

3D Cage Structures

Higher-genus modular assemblies become more interesting and more compact, if we branch out into the third dimension. Figure 8 shows structures of genus 2 that have 12-fold D_{3h} symmetry. Figure 8a is a very compact configuration. The vertical module pairs are connected to the (cyan) pyramid tips by their pentagonal faces, angled at 45°, so that all the modules have the same shape. One arch from one of the pyramid tips to the other one constitutes half of another possible 8-module toroid, in which the contact faces alternate between squares and pentagons.

Attaching additional modules in the hexahedral pattern of Figure 7, does not change the slope of the pentagonal faces. Thus, the same vertical module pairs can be used to make connections in different locations between two copies of such a planar network. We can thus form a looser genus-2 cage with larger loops between top and bottom (cyan) pyramid modules (Figure 8b). We can also combine the inner, tight connectors (Figure 8a) and the looser, outer loops (Figure 8b) and create a cage of genus 5 (Figure 8c).

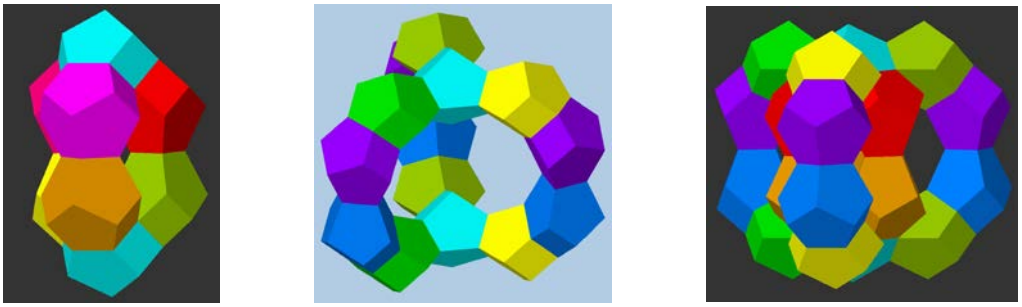

(a) (b) (c)

Figure 8: *3-dimensional cage structures of higher genus: (a) a compact genus-2 cage, (b) a looser genus-2 cage, (c) a combination resulting in a genus-5 cage.*

We can combine the techniques of Figure 7 and Figure 8 in a modular way and connect two planes of planar, hexagon-based networks with the connector types introduced in Figure 8 into cage structures of arbitrarily large genus.

Structures Based on Platonic Solids

If we are willing to accept more strongly deformed nonahedral building blocks, we can construct even more compact polyhedral surfaces of higher genus. Some of these may be derived from Platonic solids with valence-3 vertices (because of the 3-fold symmetry of the nonahedron). For instance, we can place eight nonahedra at the corners of a cube and let them connect to form a genus-5 cage (Figure 9).

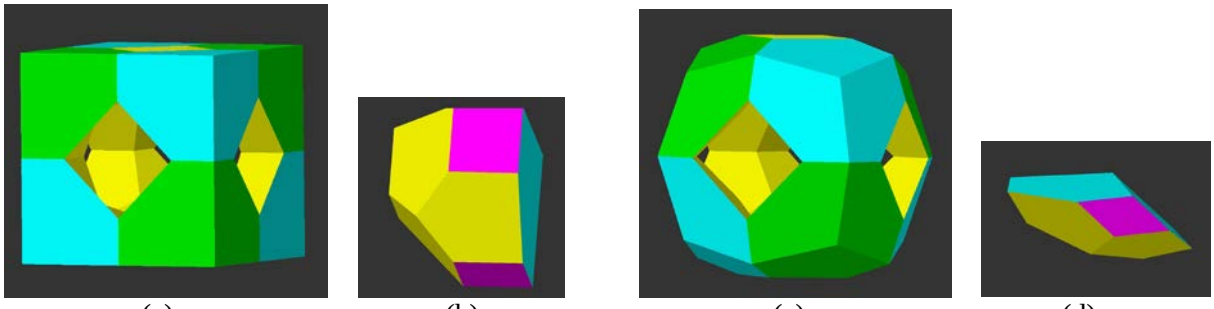

(a) (b) (c) (d)

Figure 9: *Genus-5 structures: (a) based on a cube and (c) a cuboid; (b,d) corresponding modules.*

Figure 9a shows a simple solution with all planar faces and perfectly square contact faces, where all the outer faces lie in the surface of a cube. The nonahedron geometry is easily defined by placing the red contact squares (Figure 9b) perpendicular to the cube edges at their midpoints. This leaves two DOFs: the size of the red squares, and the position of the yellow inner pyramid vertex; the latter can be chosen to

render the inner (yellow) pentagons planar. Figure 9c shows a more rounded solution with rhombic contact faces, where all faces again have been adjusted to be planar. This leaves two DOFs: the lengths of the two diagonals of the red rhombic contact faces. In Figure 9c, they were adjusted to give the resulting polyhedron an overall pleasing look. The planarity constraints for the pentagons determines the positions of the inner and outer corner vertices.

The same basic approach can be taken to make a highly symmetrical structure of even higher genus by placing 20 nonahedral modules at the corners of a dodecahedron. If the overall shape has a convex hull in the form of a perfect dodecahedron, then the angles in the red rhombic contact faces (Figure 10b) are defined by the dihedral angle of 116.5° of the dodecahedron; their size can still be chosen to set the diameter of the pentagonal tunnels. The positions of the inner pyramid vertices are again determined by the planarity constraints of the yellow pentagons. Figure 10c shows the physical realization of the dodecahedral model with two identical 3D-prints composed of ten nonahedral modules each.

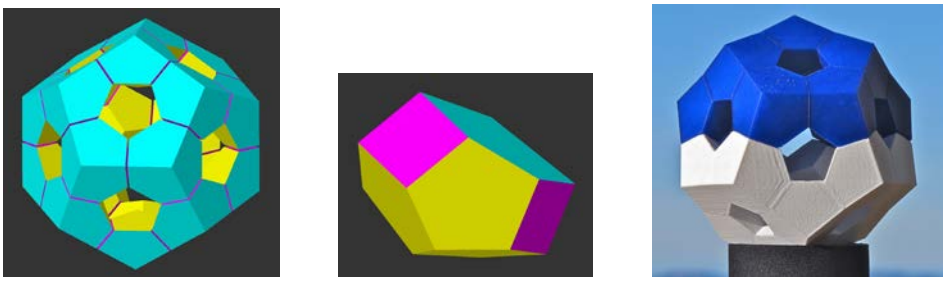

(a) (b) (c)

Figure 10: *Genus-11 structure based on a dodecahedral assembly of twenty nonahedral units: (a) overall assembly, (b) one module showing the (red) rhombic contact faces; (c) 3D-print.*

Summary and Conclusions

We have investigated the challenges and trade-offs in making many "miracle" sculptures composed of a nine-sided, near-miss Johnson solid. First, one should establish a list of the fixed, non-negotiable demands, such as the overall symmetry of the sculpture or the regularity of certain polygonal faces. Then one needs to figure out how these constraints can be incorporated most easily into the design, and how many degrees of freedom (DOF) remain after these constraints have been satisfied. Next, one has to make some choice about the relative weights given to different kinds of imperfections, e.g., non-planar faces or variations in edge-length; the remaining DOFs are then used to optimize the design under these considerations. This can be done approximately with an interactive computer-aided modeling environment, or it can be done more rigorously with some procedural optimization algorithm. If the initial design is close enough to the final goal, simple greedy gradient descent is often a practical method. This design process has been demonstrated with compositions ranging from a simple toroid of genus one to a genus-11 dodecahedral cluster.

Acknowledgments

We would like to express our thanks to Ruta Jawale, Hong Jeon, Alex Romano, and Rohan Taori, who helped in the search for the "optimal" nonahedral building block and to the staff of Jacobs Hall for their help with the construction of the various models on their 3D printers.

References

[1] D. Gordon. "Computing the Plane Equation of a Polygon." http://cs.haifa.ac.il/~gordon/plane.pdf

[2] B. M. Stewart. "Adventures among the Toroids." *(Second Edition 1964)*

[3] B. Simonsen. "Torus." *Bridges 2016 Art Exhibit.* http://gallery.bridgesmathart.org/exhibitions/2016-bridges-conference/bente-simonsen

[4] Wikipedia. "Near-miss Johnson solid." https://en.wikipedia.org/wiki/Near-miss_Johnson_solid