Introduction

The plaid model grew out of my attempt in [**S1**] to understand outer billiards on kites. A *kite* is a convex quadrilateral having a line of symmetry that is also a diagonal. In particular, let K_A be the kite with vertices

$$(-1,0), \quad (0,1), \quad (0,-1), \quad (A,0). \tag{1}$$

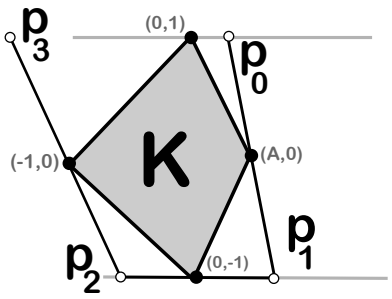

Figure 0.1: Outer billiards on the kite K_A.

Figure 0.1 shows outer billiards on K_A for $A = 4/9$. Given $p_0 \in \boldsymbol{R}^2 - K_A$, we define a map $p_0 \to p_1$ by the rule that the line segment $\overline{p_0 p_1}$ is tangent to K_A at its midpoint, and K_A is on the right-hand side as one walks along the segment from p_0 to p_1. We then consider the orbit $p_0 \to p_1 \to p_2 ....$ See [**S1**] for an extensive discussion of outer billiards and a long bibliography.

We call an outer billiards orbit on K_A *special* if it lies in the union

$$\boldsymbol{R} \times \{\pm 1, \pm 3, \pm 5, ...\} \tag{2}$$

of odd-integer horizontal lines. The orbit shown in Figure 0.1 is special. In [**S1**] I proved the following result.

Theorem 0.1 *When A is irrational, outer billiards on K_A has an unbounded special orbit.*

Theorem 0.1 is an affirmative answer to the *Moser-Neumann problem*, from 1960, which asks whether an outer billiards system can have an unbounded orbit. The orbits in Theorem 0.1 are quite complicated. They return infinitely often to every neighborhood of the vertices of K_A. I called such orbits *erratic*.

The key step in understanding the special orbits on K_A is to associate an embedded lattice polygonal path to each special orbit. This path encodes

the symbolic dynamics associated to the second return map to the union $\boldsymbol{R} \times \{-1, 1\}$ of lines. These lines are partially shown in Figure 0.1. When $A = p/q$ is rational, it is possible to consider the union of all these lattice paths at once. I call this union the *arithmetic graph* and denote it by Γ_A. When pq is even, every component of Γ_A is an embedded lattice polygon. Part 3 of the monograph has a detailed description.

One of the key results in [**S1**] is the *Hexagrid Theorem*. This result gives large-scale structural information about Γ_A. Basically, it says that Γ_A must intersect certain lines in certain places, and must avoid certain lines in certain places. Some years later I discovered that the Hexagrid Theorem is just the first in a series of results which allowed this large-scale structure to extend down to increasingly fine scales. When all these results are assembled into one package, the result is the plaid model.

We will formally define the plaid model in the next chapter. The plaid model is a rule for assigning a square tiling of the plane to each parameter $A = p/q \in (0, 1)$ with pq even. We call such parameters *even rational*. There is a similar construction when pq is odd, but the details are sufficiently different that we do not treat it here. Here we give a rough feel for the plaid model. Based on the parameter A we assign even integers to the lines of the usual infinite grid of integer-spaced vertical and horizontal lines. We call these integers *capacities*. At the same time, we define a second grid of slanting lines and we assign odd integers to these lines. We call these odd integers *masses*. We then place a *light point* at every intersection of the form $\sigma \cap \tau$ where

- σ is a slanting line.
- τ is a horizontal or vertical line.
- The mass of σ has the same sign as the capacity of τ and smaller absolute value.

Figure 0.2 illustrates the rule on a made-up example. The pictures in the next chapter show the real rules.

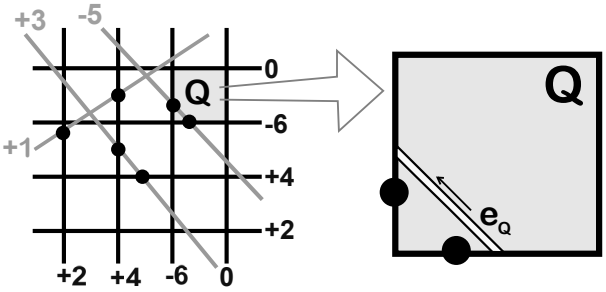

Figure 0.2: A caricature of the plaid model.

The horizontal and vertical lines divide the plane into unit integer squares. It turns out that each such square Q has either 0 or 2 edges containing an

odd number of light points. In the former case we associate the empty set to Q. In the latter case, we associate to Q a directed edge e_Q which joins the centers of the two sides having an odd number of light points. (We will explain later how the edge direction is determined.) Figure 0.2 shows a made-up example of the assignment $Q \to e_Q$.

The edges fit together to form an infinite family PL_A of polygons in the plane which we call the *plaid polygons*. Again $A = p/q \in (0,1)$ and pq is even. The lines of capacity 0 divide the plane into larger squares of side length $p+q$, which we call *blocks*. No polygon crosses the boundaries of these blocks. Figure 0.3 shows two of the blocks associated to $PL_{4/9}$. We do not show the orientations of the edges but they are consistently oriented around each polygon, one way or the other.

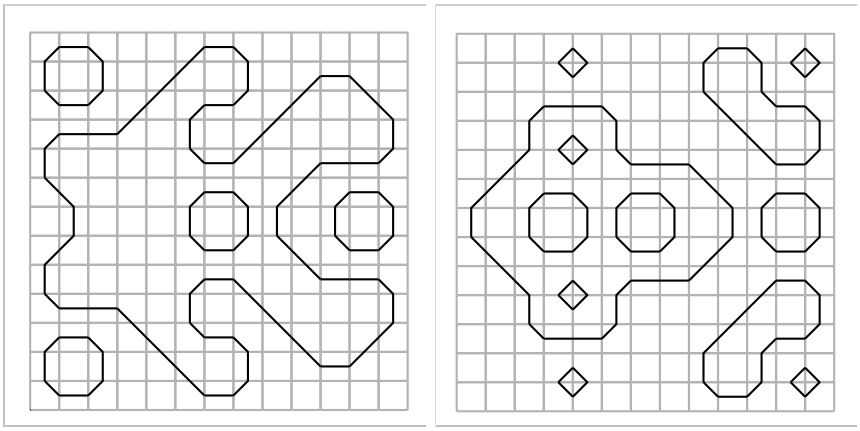

Figure 0.3: Two blocks of $PL_{4/9}$.

Here is a concrete connection between the plaid model and outer billiards.

Theorem 0.2 (Projection) *Let $A \in (0,1)$ be an even rational parameter. Modulo the vertical translations which preserve PL_A, there is a bijection between the polygons in PL_A that lie in the right half-plane and the special outer billiards orbits relative to the kite K_A. Moreover, the plaid polygon π may be (monotonically) parameterized as $\pi = \{(x_t, y_t) |\ t \in [0,N]\}$ in such a way that the point $2x_k$ lies within 3 units of the kth point of*

$$S_\pi = O_\pi \cap (\boldsymbol{R}_+ \times \{-1,1\})$$

for all $k \in \{1,...,N\}$. Here O_π is the special orbit associated to π and N is the number of points S_π.

In other words, if you put your finger on one of the polygons π of PL_A (that lies to the right of the Y-axis) and trace around it at the correct speed, the horizontal motion of your hand will track the first return map of the corresponding special outer billiards orbit up to a factor of 2 and an error of at most 3 units. Note, however, that π typically does not have N vertices and x_k need not be a vertex of π.

The Projection Theorem is a consequence of the Quasi-Isomorphism Theorem below, which gives a more precise result about the connection between PL_A and Γ_A. The Quasi-Isomorphism Theorem is the multiscale extension of the Hexagrid Theorem from [**S1**].

The plaid model has a three-dimensional interpretation which reveals connections to Pat Hooper's Truchet tile system [**H**]. When we forget the orientations on the polygons, it turns out that there are $p + q$ distinct blocks modulo translation symmetry of the tiling. (When we remember the orientations there are twice as many.) We will take $p + q$ such blocks, one representative from each translation equivalence class, and stack them on top of each other in a special order. We will then canonically interpolate between the polygons at consecutive heights in the stack to form polyhedral surfaces. The result is a cubical array of $(p + q)^3$ unit integer cubes that is filled with pairwise disjoint embedded polyhedral surfaces.

By construction, the slices of the 3D polyhedral surfaces at integer heights in the XY direction are the plaid model polygons discussed above. When the surfaces in the plaid model are sliced in the other coordinate directions, namely the XZ and YZ directions, what emerges (at least for some slices) is a pattern of curves that is combinatorially isomorphic to the curves produced by Pat Hooper's system. Figure 0.4 gives an example. The plaid parameter is $4/9$ and the Truchet parameter is $\alpha = \beta = 3/8$. Theorem 7.2, the Truchet Comparison Theorem, establishes a combinatorial isomorphism like this for all even rational parameters. Thus, the plaid model is a kind of marriage between outer billiards on kites and the Truchet tile system.

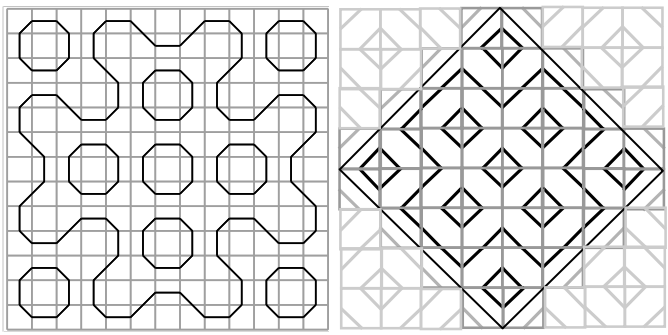

Figure 0.4: A YZ slice for $4/9$ compared to a Truchet tile.

Another curious connection between the plaid model and the Truchet tile system is that the left edge of the figure on the left side of Figure 0.3, which shows the union of plaid polygons in a certain block for the parameter $4/9$, exactly matches the left edge of the left side of Figure 0.4, which shows a special YZ slice for the same parameter $4/9$. This seems to happen for every parameter. We will give several plausible explanations in §23.8, but we will stop short of giving a proof.

The monograph has 5 parts. The rest of this introduction is a detailed description of these parts.

0.1 PART 1: THE PLAID MODEL AND ITS PROPERTIES

In Part 1, I will define the plaid model and study its properties. One technical claim, Theorem 1.4, will not be apparent from any of the definitions. This result basically says that each unit integer square has 0 or 2 sides containing exactly one light point. I will assume Theorem 1.4 in Part 1 and will deduce it in Part 2 as an immediate corollary of Theorem 8.2.

After studying the basic properties of the model, I will explain how one can use the hierarchical nature to get information about the large-scale structure of the tilings in an algorithmic way. In particular, I will give a heuristic explanation of why the model exhibits coarse self-similarity and rescaling phenomena.

After giving the basic definitions, I will explain how to assemble the two-dimensional blocks into embedded polyhedral surfaces. Finally, I will establish the connection between the XZ and YZ slices of these surfaces and the Truchet tilings. Again, the main result is the Truchet Comparison Theorem from §7.3.

Part 1 is rather long and involved, but most of the material is not needed for Parts 2-4. The (dis)interested reader can skip to Part 2 after reading §1, §2.2, §2.3, §4.2, §4.3, §4.4, and §5.2.

0.2 PART 2: THE PLAID PET

Let Π denote the set of unit integer squares. For each parameter $A = p/q$, the union of plaid polygons PL_A defines a dynamical system on Π. We simply follow the directed edge in each tile and move to the tile into which the edge points. When the tile is empty, we do not move at all. We call this dynamical system the *PL_A-dynamics*. This system is similar to the *curve-following dynamics* defined in [**H**].

In Part 2, I will connect this dynamical system to higher dimensional polytope exchange transformations. The parity claim from Part 1 will follow from this. Let X be a flat torus. A *polytope exchange transformation* (or *PET*) on X is given by a partition of X into polytopes

$$X = \bigcup A_i = \bigcup B_i, \tag{3}$$

so that there are translations T_i such that $T_i(A_i) = B_i$ for all i. Such a system gives rise to a global and almost everywhere defined map $T : X \to X$ defined so that $T|_{A_i} = T_i$. This map is not defined on the boundaries of the polytopes of the partition. However, it is an invertible piecewise defined translation.

Now we describe what we mean by a *fibered integral affine PET*. Let

$$\widehat{X} = \boldsymbol{R}^3 \times (0, 1).$$

We will work with a quotient of the form $X = \widehat{X}/\Lambda$, where Λ is a discrete group of affine transformations acting on $\widehat{X}$. The quotient X is topologically

the product of a 3-torus and $(0,1)$. The group Λ preserves each slice $\boldsymbol{R}^3 \times \{P\}$ and acts there as a group of translations. The quotient $X_P = (\boldsymbol{R}^3 \times \{P\})/\Lambda$ is a flat torus whose isometry type depends on P.

In a fibered integral affine PET, we have the same partitions as above, except that each map T_i is a locally affine map, and we have the following additional features:

- The linear part of T_i is independent of i.
- T_i preserves each slice X_P.
- The restriction of each T_i to X_P is a translation.
- All vertices of all lifts of all polytopes in the partitions have integer coordinates.

Theorem 0.3 (Plaid Master Picture) *There is a 4-dimensional fibered integral affine PET X with the following property: When A is even rational and $P = 2A/(1+A)$, there is a locally affine map $\Phi_A : \Pi \to X_P$ which conjugates the PL_A-dynamics on Π to the PET dynamics on X_P.*

Remarks: (i) Since Π is a discrete set of points, we have to say what we mean by a *locally affine map* from Π into X_P. We mean a restriction of a planar affine map to Π.
(ii) The Plaid Master Picture Theorem says that the PL_A-dynamics encodes the symbolic dynamics associated to a certain 3-dimensional PET, X_P. Here $P = 2A/(1+A)$. It might seem a bit funny to use the parameter P instead of A but this change of coordinates turns out to be useful and natural.
(iii) The Plaid Master Picture Theorem also says that these individual slices $\{X_P\}$ fit together into a 4-dimensional fibered integral affine PET. This shows a kind of coherence between the plaid model at one parameter and the plaid model at a different parameter, even though the plaid model polygons themselves vary wildly from parameter to parameter.

0.3 PART 3: THE GRAPH PET

In Part 3 we do for the arithmetic graph what we did for the plaid polygons in Part 2. When $A = p/q$ is rational, the arithmetic graph Γ_A defines a dynamical system on $\boldsymbol{Z}^2$. We just move from vertex to vertex according to the oriented polygons. We call this system the Γ_A-*dynamics*. Here is our main result.

Theorem 0.4 (Graph Master Picture) *There is a 4-dimensional fibered integral affine PET Y with the following property. When $A = (0,1)$ is rational, there is a locally affine map $\Psi'_A : \boldsymbol{Z}^2 \to Y_A$ which conjugates the Γ_A-dynamics on $\boldsymbol{Z}^2$ to the PET dynamics on Y_A.*

Remarks: (i) In this result, A need not be *even* rational.
(ii) We will see that there is a change of coordinates that conjugates Ψ'_A to a map Ψ_A with a nicer formula. That is the reason for our notation.

I will explain Theorem 0.4 in two ways. First, I will deduce Theorem 0.4 from [**S1**, Master Picture Theorem], which is presented here (with a few cosmetic changes) as Theorem 13.2.

Second, I will follow [**S2**] and my unpublished preprint [**S3**], and prove a very general version of Theorem 0.4, namely Theorem 16.9, which works for any polygon without parallel sides. (The no-parallel-sides condition is not really essential, but it makes the argument easier.) Since Theorem 0.4 already follows from Theorem 13.2, I will not give a formal proof that the PET produced by Theorem 16.9, in the case of special orbits on kites, is identical to the PET from Theorem 0.4. However, in §16.7 I will explain precisely how to match up the two PETs, and I will explain how to see a computer demonstration of the match-up.

I'd like to mention that I had many helpful and interesting discussions with John Smillie about this general topic, and he has a different way to view these kinds of compactifications in terms of Dehn-invariant like constructions involving the tensor product.

0.4 PART 4: PLAID-GRAPH CORRESPONDENCE

In this part of the monograph we relate the plaid model to outer billiards. Our next result says that the combinatorially defined plaid model gives a uniformly accurate model of the dynamically defined outer billiards special orbits. When the time comes, we will deduce the Projection Theorem above as a consequence.

Theorem 0.5 (Quasi-Isomorphism) *For each even rational parameter A, there exists an affine transformation $T_A : \boldsymbol{R}^2 \to \boldsymbol{R}^2$, and a bijection between the components of $T_A(\Gamma_A)$ and the components of PL_A with the following property. If $\gamma \in T_A(\Gamma_A)$ and $\pi \in PL_A$ are corresponding components then there is a homeomorphism $h : \gamma \to \pi$ which moves points by no more than 2 units.*

Figure 0.5 below shows the Quasi-Isomorphism Theorem in action for the parameter $A = 3/8$. The grey polygons are the components of $T_A(\Gamma_A)$ in $[0, 11]^2$ and the black polygons are components of PL_A in $[0, 11]^2$.

One thing that is surprising about the Quasi-Isomorphism Theorem is that the grid Π of half-integer points associated to the plaid model sits in a funny way with respect to the grid $G_A = T_A(\boldsymbol{Z}^2)$ which contains the vertices of $T_A(\Gamma_A)$. The grid G_A has co-area $1 + A$ whereas Π has co-area 1. When A is irrational, these grids are incommensurable. The homeomorphism from the Quasi-Isomorphism Theorem does not map vertices to vertices but nonetheless the two sets of polygons are close to each other.

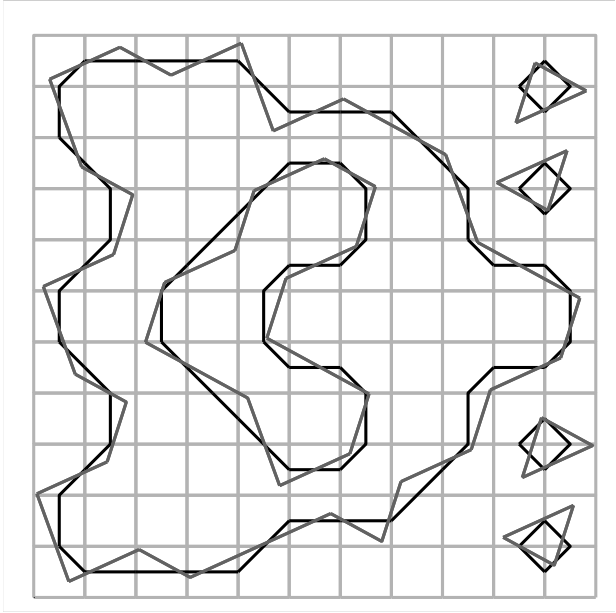

Figure 0.5: The Quasi-Isomorphism Theorem in action.

The Quasi-Isomorphism Theorem is a consequence of an orbit equivalence between our two PETs. Call a subset $Z \subset X$ *dynamically large* if every nontrivial orbit on X intersects Z.

Theorem 0.6 (Orbit Equivalence) *There is a dynamically large subset $Z \subset X$ and a map $\Omega : Z \to Y$ with the following property. For any $\zeta \in Z$ with a well-defined orbit the following three statements hold:*

1. *There is some $k = k(\zeta) \in \{1, 2\}$ such that*
$$\Omega \circ F_X^k(\zeta) = F_Y \circ \Omega(\zeta). \tag{4}$$
2. *If $F_X(\zeta) \in Z$ then there is some $\ell = \ell(\zeta) \in \{0, 1\}$ such that*
$$\Omega \circ F_X(\zeta) = F_Y^\ell \circ \Omega(\zeta). \tag{5}$$
3. *ζ is a fixed point of F_X if and only if $\Omega(\zeta)$ is a fixed point of F_Y.*

The set Z is a union of 2 open convex integral prism quotients and the restriction of Ω to each one is an integral projective transformation that maps the slice Z_P into the slice Y_A, where $P = 2A/(1 + A)$. Finally, Ω is at most 2-to-1 and $\Omega(Z)$ is open dense in Y and contains all the well-defined F_Y-orbits.

A *prism quotient* is the quotient of a set of the form $H \times \boldsymbol{R} \times [0, 1]$ under the map $(x, y, z, P) \to (x, y, z + 2, P)$, where H is a convex polygon. The prism quotient is *convex integral* if H has integer vertices. The Orbit Equivalence Theorem is saying roughly that the graph PET is a *renormalization* of the plaid PET. We will explain this point of view in §18.8.

0.5 PART 5: THE DISTRIBUTION OF ORBITS

Part 5 concerns the distribution and geometry of orbits in the plaid model.

We will explain how to assign a polygonal path π to an orbit O of the plaid PET which lies in a slice X_A where A is not even rational. The path π is unique up to translation. We call π the *plaid path* associated to O. One could view π as the geometric (i.e., Hausdorff topology) limit of plaid polygons associated to a sequence of orbits $\{O_n\}$ in even rational slices.

We say that a path π in the plane is *thin* if it is contained in an infinite strip. Otherwise we say that π is *fat*. For instance, a parabola is fat. The bulk of Part 5 is devoted to proving the following result.

Theorem 0.7 *For every irrational parameter A, the slice X_A of the plaid PET has an infinite orbit whose associated plaid path π is fat. More precisely, the projection of π onto the X-axis is the set $[1/2, \infty)$ and the union of vertices of π having first coordinate $1/2$ contains a large-scale Cantor set.*

By a *large-scale Cantor set*, we mean a set of the form

$$\{1/2\} \times \bigcup_{\beta} \left(\sum_{j} \beta_j Y_j \right). \tag{6}$$

The union takes place over all finite binary sequences β, the sequence $\{Y_j\}$ is unbounded, and the ordering on the points in the union coincides with the lexicographic ordering $(0, 1, 10, 11, 100...)$ on the finite binary strings.

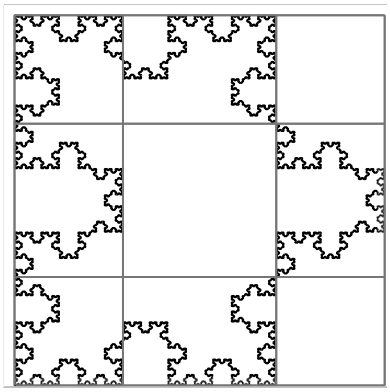

Figure 0.6: Cantor-set like properties of the left edge.

Figure 0.6 shows the picture of a particular plaid polygon for the parameter 169/408. The square here has side length 577. The extra lines (making a tic-tac-toe pattern) have significance we will explain in the proof of Theorem 3.1. One can see a kind of finite version of a large-scale Cantor set going up along the left edge of the figure. The curve π from Theorem 0.7 is a geometric limit of these finite curves taken with respect to a continued-fraction-like approximating sequence of rational numbers.

Theorem 0.7 combines with the Projection Theorem to give another proof of Theorem 0.1. At the same time, one can extract Theorem 0.7 from the main results in [**S1**] and from the Quasi-Isomorphism Theorem. See §21.6. The ideas underlying our proof here are similar to what we do in [**S1**]. However, by going through the details here we can give a proof that is independent from [**S1**].

We will combine the proof of Theorem 0.7 with some results we prove in Part 1 to get some additional information about the distribution of orbits that goes beyond what we could do in [**S1**]. Say that a polygon is *N-fat* if it is not contained in a strip of width N.

Theorem 0.8 *Let $\{p_k/q_k\} \subset (0,1)$ be any sequence of even rational numbers with a limit that is neither rational nor quadratic irrational. Let $\{B_k\}$ be any sequence of associated blocks. Let N be any fixed integer. Then the number of N-fat plaid polygons in B_k is greater than N provided that k is sufficiently large.*

Remark: The restriction that the limit not be quadratic irrational is probably just an artifact of the proof rather than a necessary hypothesis.

The Projection Theorem lets us translate Theorem 0.8 into the language of outer billiards. Say that the *essential diameter* of a special outer billiards orbit on K_A is the diameter of its intersection with the set

$$[0,\infty) \times \{-1,1\}.$$

A special orbit of diameter D essentially fills out an octagon-shaped "annulus" of width D.

Corollary 0.9 *Let $\{p_k/q_k\} \subset (0,1)$ be any sequence of even rational numbers with a limit that is neither rational nor quadratic irrational. Let $\omega_k = p_k + q_k$. Let $\{I_k\}$ be any sequence of intervals of the form $[n\omega_k, n\omega_k + \omega_k]$. Let N be any fixed integer. Then there are more than N distinct orbits in the interval $I_k \times \{1\}$ which have essential diameter at least N provided that k is sufficiently large.*

0.6 COMPANION PROGRAM

The monograph comes with several companion computer programs which illustrate most of the results. You can download these from the following location.

https://press.princeton.edu/titles/13339.html

At least in the short run, the same programs can also be found at the following location.

http://www.math.brown.edu/~res/Java/PLAID.tar

Once you download the tarred directory, you untar it. The unpacked files live in a directory called *Plaid*. This directory has several sub-directories and a *README* file. The *README* file has further instructions. One of the programs can also be run directly on the web. At least in the short run, the web version can also be accessed from the following location.

http://www.math.brown.edu/~res/Javascript/Plaid/Main.html

This web program gives a quick view of the 3D plaid model and the plaid surfaces.

I discovered all the results in this monograph using the program, and I have extensively checked my proofs against the output of the program. While this monograph mostly stands on its own, the reader will get much more out of it by using the program while reading. I would say that the program relates to the material here the way a cooked meal relates to a recipe. Throughout the text, I have indicated computer tie-ins which give instructions for operating the computer program so that it illustrates the relevant phenomena. I consider these computer tie-ins to be a vital component of the monograph.