Growth Rules For Quasicrystals

Joshua E.S. Socolar

IBM Research Division, T. J. Watson Research Center,
P.O. Box 218, Yorktown Heights, NY 10598

1 Introduction

All known alloys that form quasicrystals do so when the liquid is subjected to a rapid quench, during which the quasicrystalline phase grows rather quickly. (Representative cases include AlMn [1], AlLiCu [2], and AlFeCu [3].) While the quasicrystals formed upon quenching are not as well-ordered as the best annealed samples [4, 5], they exhibit translational correlations that cannot be explained by simple random accretion models. [2, 6] In addition, equilibration times for quasicrystals are quite long, owing to the slow processes involved in the relaxation of phason strains. [7, 8] In order to understand the origins of long-range quasicrystalline order in quenched samples, it is necessary to investigate mechanisms for establishing quasicrystalline order in the absence of the slow bulk relaxation processes required by equilibrium quasicrystal models. We must determine how, or perhaps whether, quasicrystalline order can be engendered by nonequilibrium growth kinetics alone.

The problem of the growth of quasicrystals differs from the analogous problem for crystals in one important respect. In both cases the details of the surface chemistry of the growing sample — sticking probabilities, diffusion rates, etc. — determine the growth kinetics and can be quite complicated. For crystals, however, the generation of long-range order poses no intrinsic conceptual difficulty. One assumes that the dynamics favor the formation of rigid unit cells and that the kinetics can be treated at the level of the aggregation of unit cells. For quasicrystals, on the other hand, it is quite difficult to see how long-range order can be established even in the most simplistic abstract models. The manner in which the growth is to be continued at certain points may be determined by arbitrarily distant regions, so it would appear that nonlocal interactions have to be invoked if the growth kinetics are to maintain the translational order. Thus while we might aspire to a detailed understanding of the solid-liquid interface that would explain such properties as the macroscopic shape of the sample and the existence of nonequilibrium facets of a given orientation, we will content ourselves for now with an understanding of how quasiperiodic order can be established at all.

The existence of matching rules for quasicrystals, constraints on local configurations that completely determine the structure of the tiling [9], shows that certain quasicrystalline structures are ground states of some Hamiltonians. It is clear that these quasicrystal states can be reached from arbitrary initial conditions if the appropriate relaxation and annealing processes are allowed to occur. In most real quasicrystals, though, the evidence suggests that phason strains, which are variations in the relative phases of incommensurate Fourier components of the density [10, 11], do not occur on the time scales of rapid quenches. [7, 8] Though it is possible that some phason annealing occurs near the surface of the growing cluster, we will ignore such processes here in an effort to isolate the role that growth kinetics can play.

We consider growth rules that require no relaxation processes whatsoever; i.e., accretion rules which do not incorporate any mechanism for repositioning or removing an atom once it has been attached to the growing cluster. In addition, we imagine that the growth can be described by models governing the aggregation of geometric units that represent clusters of atoms and do not worry about the details of the formation of the units. Recognizing that the interactions between units could be extremely complex, we allow rules of arbitrary complexity governing the placement of new units on a growing cluster. The only constraint we require in the interests of physical plausibility is that the rules must be *local*: the probability of adding a given unit at a given surface site must be determined by the local environment of that site. Though it is not *a priori* obvious, it turns out that rules of this type can indeed result in structures with long-range quasicrystalline order.

2 Phason disorder and random accretion models

The observation that some crystalline phases closely related in composition to quasicrystals can be described as packings of icosahedral units suggest that we begin by investigating simple rules for the accretion of icosahedra (or decagons, in 2D). [6, 12] In the simplest case a structure is grown via the attachment of icosahedra to a cluster according to four rules: (1) The orientations of neighboring units (and hence all units) are identical. (2) All bonds between nearest neighbor units have the same length and one of a symmetric set of directions which are defined relative to the orientation of the units. (For example, icosahedra may be joined along common fivefold axes.) (3) Units are not allowed to overlap; i.e., there is a minimum distance between near neighbors. (4) Sites for the addition of new units are chosen randomly from the set of possible sites.

Though these rules do produce longer translational correlation lengths than might have been expected (as measured by the widths of diffraction peaks), they do not generate sufficient correlations to account for experimental observations. [6] The disorder appearing in growth models of this type is best described in terms of the configuration of the phason field in the sample. (Because the model assumes rigid units and bonds, no conventional strain, or phonon strain, is possible.) A feature of the simple algorithm just described is that it produces discontinuities, or tears, in the phason field which are effective in destroying the long-range order. [6]

A useful geometric picture for analyzing phason strain is obtained by mapping the positions of the units onto a 6D hypercubic lattice (or 5D for the decagonal case). [13] For simplicity, consider the case of icosahedral units connected by bonds along their five-fold axes. The position of any icosahedron can then be written as $\sum_{n=1}^{6} k_n \mathbf{e}_n$, where k_n are integers and $\mathbf{e}_n$ are the six vertex vectors of the icosahedron. (Each bond lies along one of the $\mathbf{e}_n$'s.) The $\mathbf{e}_n$'s are then mapped onto the basis vectors of the hypercubic lattice. Note that the mapping of the positions of the units onto the 6D lattice is unique up to a choice of origin since there is no way to write one vertex vector as an integer linear combination of the others.† If units are connected by bonds in symmetry directions other than vertex directions, the bonds can always be written as unique sums of vertex vectors and the same mapping applied. The only difference is that neighboring units in physical space are mapped onto 6D lattice sites separated by an appropriate sum of basis vectors, not to nearest neighbor sites.

Given a network of units placed on the 6D lattice, their positions in real space are given by projecting occupied 6D lattice sites onto the appropriate 3D subspace. The orthogonal complement of this subspace is called "perp-space". For any perfect quasicrystalline configuration of units, the perp-space projections of all of the occupied 6D lattice sites lie within a bounded region. Conversely, it can be shown that the diffraction pattern generated by any infinite structure that is bounded in perp-space contains Bragg peaks at the quasicrystalline reciprocal space positions.

The destruction of quasicrystalline order by static phason fluctuations occurs when fluctuations in the phason variable (the projection onto perp-space) become unbounded, or grow as some strictly positive power of the sample size. [14] Note that small fluctuations about a uniform phason strain would still produce Bragg peaks and hence translational order, but would destroy icosahedral symmetry and hence will be thought of as disordered for present purposes. A tear in the phason field begins to

†For some symmetries the natural choice of bonding vectors includes subsets for which $\sum \mathbf{e}_n = 0$. The mapping of the structure into hyperspace can still be performed, but points in hyperspace that differ by bonds corresponding to one of the degenerate sums must be identified.

form when neighboring units in real space find themselves separated in perp-space by an amount larger than a bond length. When this occurs, the parts of the sample on opposite sides of the tear evolve essentially independently and are unlikely to fluctuate in just the right way so as to close off the tear later. At large length scales the structure tends to look like a tree of strips separated by tears, with each strip executing a random walk in perp-space. Thus large phason fluctuations are generally produced.

To improve agreement with experiment the rules can be revised in a variety of ways. Approaches that have been considered include allowing some annealing near the surface of the growing cluster in a way that favors a higher density of nearest-neighbor bonds [12], disallowing any configurations containing certain near-neighbor distances [6, 15], and limiting the probabilities for selecting sites on the basis of their distance from the origin [15]. These approaches have achieved some success, particularly when the second and third are combined, and it is likely that further refinements can be made. There is no general understanding, however, of the origin of the phason fluctuations or uniform phason strains that arise in these models.

Rather than review the results obtained for random growth models here, we will address the issue from a different point of view. In an effort to understand the intrinsic barriers to perfect quasicrystal formation, if there are any, we will study the problem of whether physically plausible growth rules can eliminate phason disorder altogether. Surprisingly, it turns out that this is possible even without resorting to surface annealing or to rules that refer to distance from the origin. At least for the canonical example of the Penrose tilings, certain *local* rules for tile aggregation have been shown to ensure perfect quasicrystalline order out to arbitrarily large distances. These rules show that the intrinsic limits imposed by the locality requirement are far less severe than may have been anticipated. Their structure and the effects of deviations from them provide a logical starting point for the study of quasicrystal growth in general.

3 Local rules for growing Penrose tilings

This section is concerned with a set of local rules that result in the growth of defect-free 2D Penrose tilings. The Penrose tilings [16] are selected because they exhibit several properties in a particularly simple way while retaining all of the essential complexities of quasicrystal structure. Complications specific to various icosahedral structures or tilings of other symmetries are better addressed once the Penrose case is understood.

Before considering the growth algorithm for Penrose tilings in detail, a few words should be said about some subtleties of the problem and the precise sense in which perfect growth is achieved. Note that there may be a distinction between the results that are sufficient from the point of view of the physics of quasicrystal growth and those required by certain mathematical criteria.

First, the meaning of the locality constraint must be made explicit. In general, the addition of a unit to the growing cluster involves three steps: random selection of a surface site, classification of the site, and selection of the appropriate type and position of a unit to be added at the site. Locality requires that the classification of the selected site require only information about some finite environment around the site and that there be an upper bound on the size of the environment needed. In addition, the action to be taken at the selected site should be determined by the classification of that site alone. (Note, however, that it is perfectly acceptable for the specified action to be "Do nothing at this type of site." It is easy to prove that Penrose tilings cannot be grown by local rules if one insists that some tile must be added at any surface site that happens to be chosen. [17]) The growth rules discussed below are prescriptions for assigning probabilities (possibly 0) to the various actions that might be taken at a site. The rules are local in the sense that the classification of surface sites and consequent assignment of probabilities is determined completely by the local environment of each site.

Second, the extent to which the growth rules result in perfect Penrose tilings must be made clear. The rules discussed in this section are based on the classification of surface sites into three types: forced, unforced, and marginal. (Marginal sites were referred to as "corner sites" in Ref. [18].) At forced sites, the probability of adding a particular tile in a specific orientation is unity ($p_f = 1$). At unforced sites, the probability of adding anything at all is zero ($p_u = 0$). At marginal sites, a specified tile is added with a small probability ($p_m = \epsilon$) — the most probable action taken at a marginal site is to add nothing. The precise sense in which perfect growth is achieved is as follows: Given any arbitrarily large distance, R, and any probability, P, arbitrarily close to unity, p_m can be chosen small enough that the rules will produce with probability greater than P a portion of a perfect Penrose tiling that covers a disk of radius R. In some sense, the limit $p_m \to 0$ produces a perfect, infinite Penrose tiling, although the rate of growth of the tiling goes to zero with p_m. In any case, the relevant question for the physics of quasicrystals is whether or not there is an intrinsic limit on the size of defect-free quasicrystals generated with high probability by local rules and this question is clearly resolved by the growth algorithm below.

Now for the purposes of analysis, it is useful to take the $p_m = 0$ limit obtained by specifying that marginal sites are selected only when no forced sites exist anywhere on the cluster. Strictly speaking, this is a nonlocal rule since it requires a search of the perimeter of the cluster for forced sites. It is appropriate to analyze this limit, though, since in a system in which atoms (or tiles) bombard a cluster from all directions, low probabilities for adding at marginal sites allow an effective perimeter search to occur when the time scale for adding to a marginal site is much longer than the time scale for probing all the surface sites.

Third, while one may, for reasons of mathematical interest, insist on an algorithm that is capable of generating any of the uncountably infinite distinct Penrose tilings, all that is really required for quasicrystal growth is an algorithm that generates a single Penrose tiling. The rules described below do permit the construction of any Penrose tiling. By this we mean the following: The different Penrose tilings can be specified by giving the values of two vectors, **u** and **w**. Tilings with the same **w** and different values of **u** are simple translates of each other and tilings with different values of **w** are related by a uniform phason shift. [19] Given any of the possible seed clusters, the growth rules allow tilings with all possible values of **w** and **u** consistent with the placement of the seed. It is conceivable that a simpler set of rules would suffice to grow one particular Penrose tiling with a special values of **w** and **u** around a particular seed.

Finally, the extent to which a real system can be approximated by rules that rely on probabilities that are strictly zero is not clear. One knows from studies of crystal growth that sticking probabilities at various sites can differ by exponential factors, $e^{\Delta E/kT}$, where ΔE is the difference in the binding energies or the energy barriers associated with sticking at different sites. The growth of whisker crystals around a screw dislocation occurs, for example, because there is a large difference between the probability for sticking at a step on the surface and that for sticking on a flat portion of the surface. There is, however, a fundamental difference between the quasicrystal and a simple crystal regarding the effects of a nonzero probability for sticking at unforced or marginal sites, no matter how small the probability is. In a crystal with a simple, perfectly rigid unit cell there is no choice that will disrupt the translational order (although it is possible to generate vacancies or large holes). In a quasicrystal, on the other hand, making unforced choices or adding prematurely at marginal sites inevitably leads to the creation of some defects that could destroy long-range translational order. For this reason, it is necessary to study the effects of nonzero values of p_u and/or p_m on the growth. This issue will be discussed briefly in section 5, though few rigorous results are presently available.

3.1 *The growth rules in detail*

We now consider the growth of Penrose tilings as a purely mathematical problem. The units to be used are the two Penrose rhombs with matching rule decorations (the canonical arrow decorations, for example.) At the risk of being redundant, let us state the problem as succinctly as possible. We desire a set of rules governing the addition of Penrose rhombs to a growing cluster with the following properties. (1) The rules must be local; the specification of what to add at a given surface site must be made on the basis of local information about that site. Note that the specification includes the tile position and the orientation of the matching rule decoration within the tile. (2) Application of the rules must result in the growth of a portion of radius

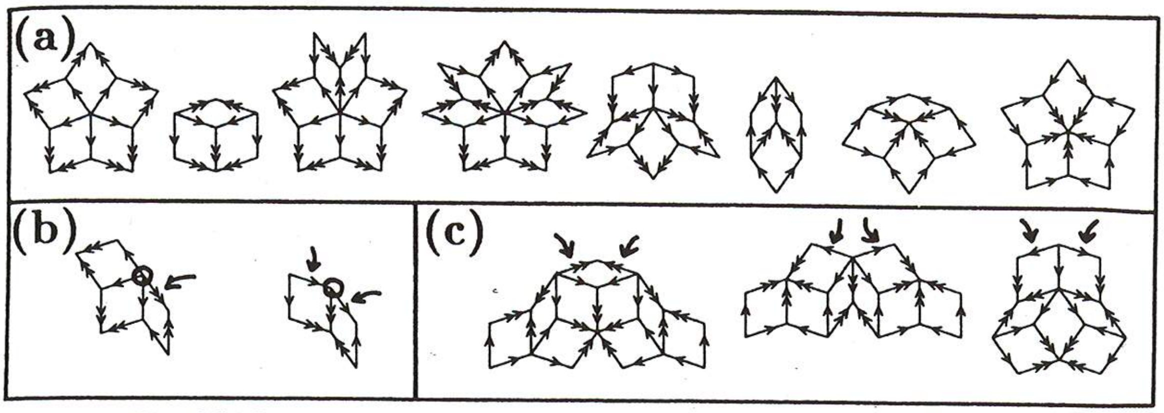

Figure 1: The OSDS classification of sites for Penrose tiling growth. (a) The eight allowed vertex configurations, shown with the edge-arrow decoration of the tiles. (b) Examples of forced edges. There is only one way to add to the edges indicated, owing to the structure around the circled vertex. Note that in the case on the left, the edge is forced even though the entire vertex may be either the third or fourth from the left in (a). (c) The edges indicated are classified as marginal.

R of a perfect Penrose tiling with probability P, where R can be made arbitrarily large and P arbitrarily close to unity.

The following rules, henceforth referred to as the "OSDS rules", are sufficient.[18]

- *Classification of sites:* Begin with a large enough seed cluster that can be found in a Penrose tiling. ("Large enough" means containing enough tiles so that at least one site on the surface must be either forced or marginal, by the definitions below. Eight tiles are sufficient.) Each step in the growth begins with the random selection of a tile edge on the surface of the growing cluster. The arrow decoration of the chosen edge permits only two choices for how a tile can be attached to it while respecting the matching rules. One choice uses a fat tile, the other a skinny. For each of these choices, the two vertices at the ends of the chosen edge are examined. If the choice results in a configuration around either vertex that is *not* consistent with any of the eight vertex configurations found in the Penrose tiling, it is discarded. Edges where only one choice is possible are called "forced". The allowed vertices and some examples of forced edges are shown in Figures 1a and 1b. If both choices produce consistent vertex configurations, the edge is generally classified as "unforced". There are special unforced edges, however, which can be identified by the configuration of all tiles joined to the edge at its endpoints, that are classified as "marginal". Figure 1c shows the only three types of marginal edges. For marginal edges, the choice of adding a skinny tile is discarded, even though it would appear consistent with the local environment. Note that a marginal edge can become forced or unforced when a new tile is added that shares one of its vertices.

- *Assignment of probabilities:* If a forced edge is selected, the unique allowed tile is added with probability 1. If an unforced edge is selected, nothing is added. If a marginal site is chosen, the fat tile is added with probability $p_m << 1$.

To analyze the growth dictated by these rules it is useful to take the $p_m \to 0$ limit by stipulating that growth at a marginal site occurs if and only if there are no forced sites available. When a cluster has no forced sites on its surface the surface is called "dead". When a dead surface is reached, a fat tile is always added at a marginal site. New forced sites are then created and the growth proceeds rapidly until another dead surface is reached. Of course if p_m were really identically zero, the growth would stop at the first dead surface encountered, but the limiting rules just defined are approximated with arbitrary accuracy for clusters up to any given size by sufficiently small, nonzero p_m. In addition, as we will see below, the growth rate need not vanish even for $p_m = 0$ when special seeds are used as nucleation sites.

To prove that the rules work as advertised we must show that:

1. adding a fat tile to a marginal site that is part of a dead surface never introduces a defect;

2. every dead surface contains at least one marginal site;

3. the rule that a fat tile is always added at a marginal site does not preclude the construction of any of the Penrose tilings; i.e., it does not limit the possible values of $\mathbf{u}$ and $\mathbf{w}$.

Note that the addition of forced tiles can never introduce a defect into the tiling.

Proof of 1 and 2: By examining all configurations around an edge for which that edge is unforced and the ways in which such configurations can be patched together to form a dead surface, it is easy to see that all dead surfaces are made up of straight faces and five types of corners. A dead surface and the five possible dead corners are shown in Figure 2. It is well-known that Penrose tilings contain strings of tiles called "worm segments" in which tiles can be rearranged such that no matching rule violations occur in the tiling except at the endpoints of the worm segments and that worm segments of arbitrarily large length occur. [20] The straight faces of dead surfaces are borders of such worm segments. (See Figure 2.) As soon as one tile in the worm segment is placed, forced growth fills out the entire segment and proceeds beyond it until a new long worm segment is reached.

Seeing the structure of the dead surfaces, one is tempted to allow all unforced sites to be marginal, but caution is necessary here. For most dead surfaces, the addition of either possible tile at any of the surface sites is perfectly consistent with the Penrose tilings and will produce no violations of the matching rules as forced growth proceeds. There are, however, certain "dangerous" faces of some dead surfaces for which one

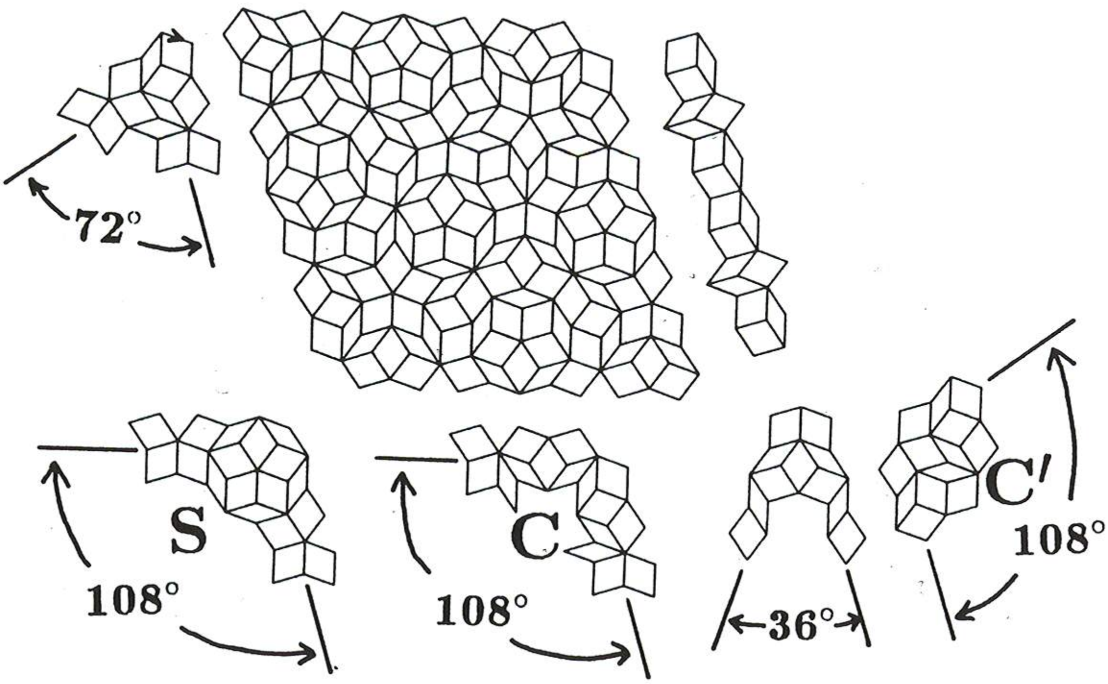

Figure 2: The structure of a dead surface. The large rhombic cluster has a dead surface. The strip of tiles to its right is a worm segment that can be added to the adjacent face. The worm could also be flipped about its long axis and added to the same face of the cluster, indicating that all the edges on that face are unforced. The five possible corners that can join dead faces are shown. The three 108° corners are labelled **S** (for the star of fat tiles), **C** and **C′** (for the chevron of two skinny tiles). The orientation of the 72° corner is defined by direction of the single arrow on the corner edge.

of the choices is inconsistent. To see the inconsistency, one must examine distant regions of the cluster, as will be explained below.

There is no way to avoid the dangerous face problem with a rule that specifies what to add at a generic unforced site. Some dangerous faces require one arrangement of the adjoining worm while others require the other arrangement and there is no way to distinguish them without looking at the corners where the dangerous face terminates, which may be arbitrarily far from the chosen site. These corners, however, which are always either the stars or chevrons of Figure 2, do contain the necessary information, albeit in a highly nontrivial fashion, and so it is the edges at these corners that are designated as marginal.

What makes a face dangerous is best explained in the context of the Ammann line decoration. It is well-known that the Penrose tiles can be decorated with line segments in such a way that in any tiling obeying the matching rules the line segments join to form straight lines that traverse the sample. [21] These Ammann lines form

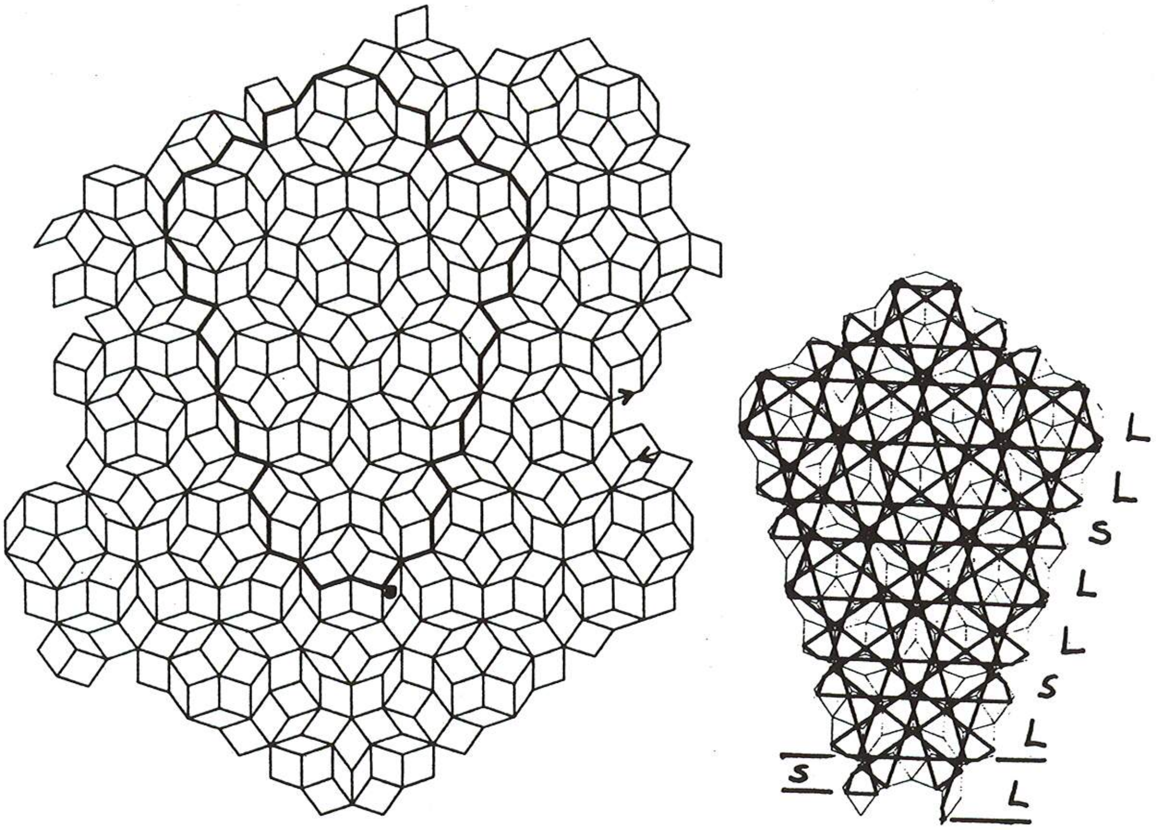

Figure 3: A dead surface with a dangerous face. At left, the heavy line marks a dead surface to which a single (marked) tile was added. Forced growth then proceeded until an edge was found where no consistent addition can be made. At right, the same dead surface is shown with the Ammann line decoration applied to the tiles. The sequence of intervals between horizontal Ammann lines and the effect produced by making the two possible choices at the bottom face are shown.

grids of parallel lines normal to the five directions of the tile edges and the spacings of consecutive lines in each grid form a Fibonacci sequence of long and short intervals. Choosing which tile to add at a dead surface site amounts to choosing the next interval in the Ammann grid with lines parallel to the surface. A dangerous face is one for which the next interval is forced owing to the extended structure of the grid of Ammann lines parallel to the face. In Figure 3, for example, the next interval at the bottom must be an S, but the only way to know this is to examine the sequence of spacings

$$...LLSLLSL... \tag{1}$$

If the first interval on the left were not specified, or if it were an S, then either choice for the next interval on the right would be consistent with the Fibonacci sequence. Similarly, there exist situations in which distant elements dictate that an S be added.

In view of the apparently nonlocal nature of the information that must be used to add the appropriate interval in the 1D sequence, it is remarkable that the 2D structure of the Penrose tiling conspires to make this information available at the corners of the dangerous face. To see that the above rules always add correctly at a dangerous face, one must catalogue all dead surfaces and dangerous faces and explicitly show that no inconsistency is encountered. The complexity of this task is greatly reduced by exploiting the well-known inflation/deflation symmetry of the Penrose tilings. [16]

The inflation of a finite cluster is defined as the cluster of all inflated (larger) tiles that are covered by the original cluster. Inflation of a cluster produces a new cluster with fewer tiles. The deflation of a finite cluster is defined as the cluster of deflated (smaller) tiles that are produced by all of the original tiles plus any small tiles that are forced by that cluster. By these definitions inflation and deflation really are inverse operations except in the case of a shape with a 36° corner, which can never be produced by deflation. One can show that under inflation or deflation dead surfaces remain dead, the macroscopic shape of a dead surface is unchanged, star and chevron corners are interchanged, and the arrow direction at each 72° corner is reversed. Furthermore, if two dead surfaces are related by some number of inflations, then the new dead surfaces reached when appropriate tiles are added at corresponding unforced sites will be related precisely by the same number of inflations. Thus one can effectively examine all dead surfaces and the evolution produced by a given choice for adding to them by exhaustively examining a finite, in fact rather small, subset of them: those that are too small to be inflated. (Note that some boring details traceable to the behavior of 36° corners have been left out here.)

A complete inventory of dead surface shapes, with corner types indicated, is shown in Figure 4. For each shape shown, the dead surfaces are obtained by tracing out any path that forms a convex polygon containing the region marked with a dot. The ellipses are meant to indicate that the number of heavy lines depends upon the absolute size of the cluster, which may be smaller or larger than those shown. The distances between heavy lines form a geometric sequence of powers of the golden mean. Each heavy line is a dangerous face of any dead surface containing it.

One sees immediately that every dead surface contains at least one marginal site. Note that the **C′** corner never appears. Inspection of the extended neighborhood around a **C′** corner shows that it never remains dead unless both faces joined to it are shorter than a few tile edges. In the perfect Penrose tiling, there is only one small dead surface containing such a corner, the surface of the cluster used to illustrate the corner type in Figure 2. In a cluster containing defects, other dead surface with **C′** corners can arise.

A systematic investigation of the growth resulting from all possible additions to each dead surface shows that the only dangerous faces are those marked by heavy lines in Figure 4 and that the addition of a fat tile at the marginal corner site is always the consistent choice.

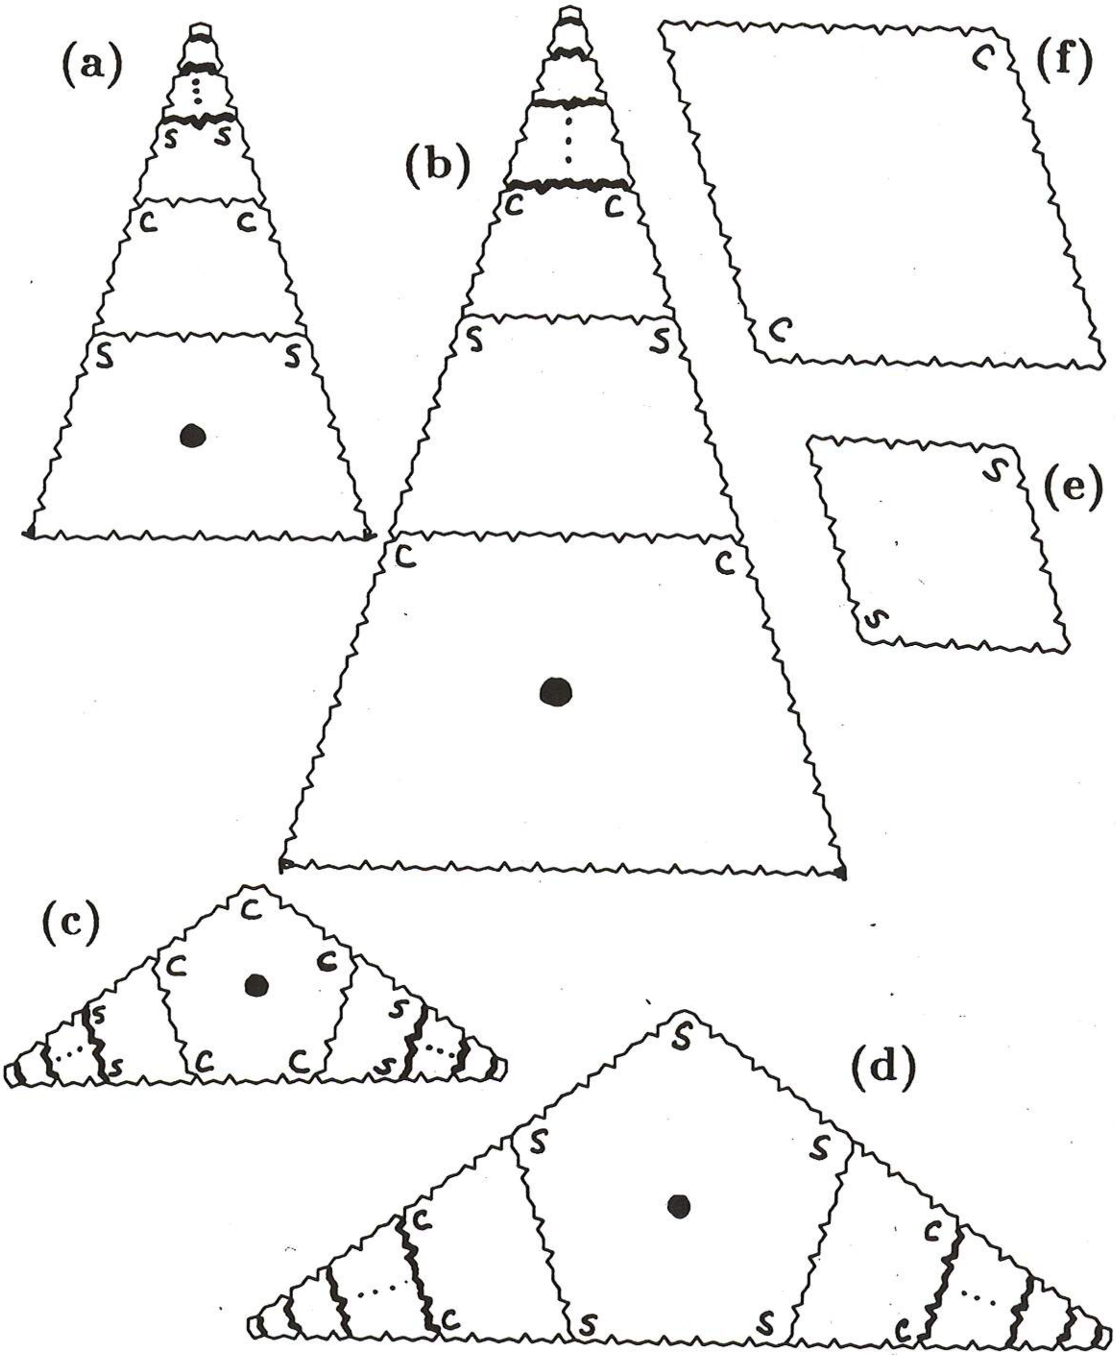

Figure 4: Catalogue of dead surfaces. Each figure represents a series of dead surfaces, all of which include the area marked with a dot. Any path enclosing the dot is a dead surface. The corner types are marked as in Figure 2. The faces marked with heavy lines are dangerous faces, given the presence of the region marked with a dot. Note that in passing from (a) to (b) or from (c) to (d) stars and chevrons are interchanged, the orientation of the 72° corners are reversed, and an additional dangerous face becomes possible. This process can be repeated ad infinitum to produce infinite series of dead surfaces as indicated by the ellipses. It can also be iterated in reverse to produce smaller clusters with dangerous faces.

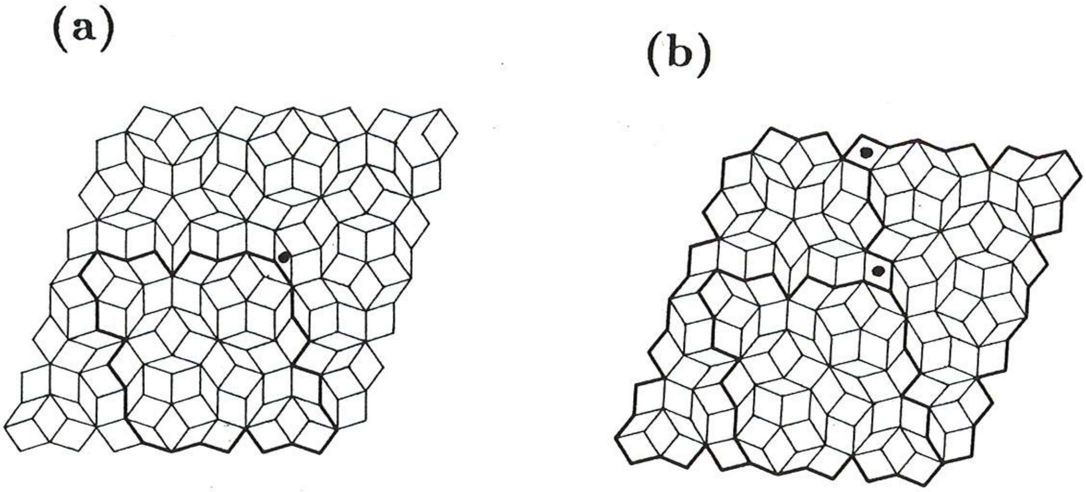

Figure 5: Nonrestrictive nature of the OSDS rules. (a) The forced growth produced by adding a skinny at a marginal site. (b) The final configuration of (a) is reached by following the OSDS rules, making the appropriate choices at two marginal sites. The tiles marked with dots are the choices that were made at dead surfaces.

It may be helpful to consider how the inflation operation enters the argument from a slightly different point of view. The key point is that the operations of filling in all forced tiles and inflating commute (except for a few cases involving 36° corners. Thus the evolution of a given cluster to its next dead surface can be determined by first inflating until there are only a few tiles in the cluster, adding the forced large tiles, then deflating. This is the property that allows a rule for adding at marginal sites to work (or at least to be proven to work).

Proof of 3: To see that no Penrose tilings are excluded by the OSDS rules (i.e., any values of **w** and **u** consistent with the original seed can be produced), one can check by inspection of each dead surface that the evolution forced by the addition of a skinny tile at a marginal site can always be obtained by adding fat tiles to other marginal sites in the appropriate order. Consider, for example, the evolution depicted in Figure 5a, where a skinny tile (marked with a dot) has been added at a marginal site. Figure 5b shows how the same final shape can be obtained through a sequence of additions of fat tiles (marked) to marginal sites of intermediate dead surfaces. The one exception, of course, is that the effect of adding a skinny at a dangerous site cannot be mimicked. Thus we see that any evolution that is allowed by the Penrose matching rules can be obtained via growth according to the OSDS rules.

For completeness we note that the choice of which sites are marginal is not unique. In deriving the OSDS rules, the selection of marginal sites above was motivated by the requirement that no error be made when adding to *any* dead surface. With a suitable choice of marginal site rules, however, one can ensure that the dead surfaces containing dangerous faces never appear. The simplest example is to designate 72° corners as marginal and always add a fat tile to them. Starting from almost any small seed the only dead surfaces generated are rhombi and trapezoids with no dangerous faces.

3.2 *Phenomenology of perfect growth*

The growth of a large perfect cluster according to the OSDS rules is rather irregular. Small values of p_m imply that once a dead surface is reached it will just sit there for a long time. Eventually, a tile will stick at a marginal site and initiate a burst of forced growth out to the next dead surface. During the forced growth, there are times when the majority of perimeter sites are forced and also times when only a few forced sites remain which then generate a new burst of growth.

Figure 6 shows a typical stage in the growth of a cluster according to the OSDS rules. The heavy lines depict the series of dead surfaces that have occurred and the filled tiles are the ones added at marginal sites. If forced growth is allowed to continue, the next dead surface will be a trapezoid with its short base at the top of the figure. Note that three star corners with marginal sites are present on the outer surface. These will be filled by forced growth before the next dead surface is reached.

The catalogue of dead surfaces generated after the addition to a single marginal site of each of the possible dead surfaces indicates that the area of the tiling roughly doubles on average between marginal additions. Assuming that the forced growth occurs at an average rate proportional to the perimeter of the growing cluster, the total time for growing a cluster of N tiles is given by $N^{1/2}/v + t\log_2 N$, where v measures the rate of forced growth and t is the average waiting time for adding at a marginal site. In order to avoid mistakes, we must have t large enough that all forced additions during the last stage of growth occur within one waiting time, implying $t \sim N^{1/2}/v$. We therefore estimate that the total time necessary for growing a perfect cluster of N tiles with high probability is on the order of $N^{1/2}\log N$.

For reasons having to do with the details of the sequences of dead surfaces generated, it is possible to reduce the time required for growth by using an alternative definition of marginal sites. One can reduce t significantly, for example, by taking 72° corners to be the only marginal sites. A better way to speed up the growth, however, is to use a special defective seed that forces growth ad infinitum without ever encountering a dead surface. Such seeds are discussed in the next subsection.

3.3 *Seeds for the growth of the Penrose tiling*

As previously mentioned, the OSDS rules require that an initial seed of eight tiles be present. The reason for this is just that dead surfaces containing seven or fewer tiles

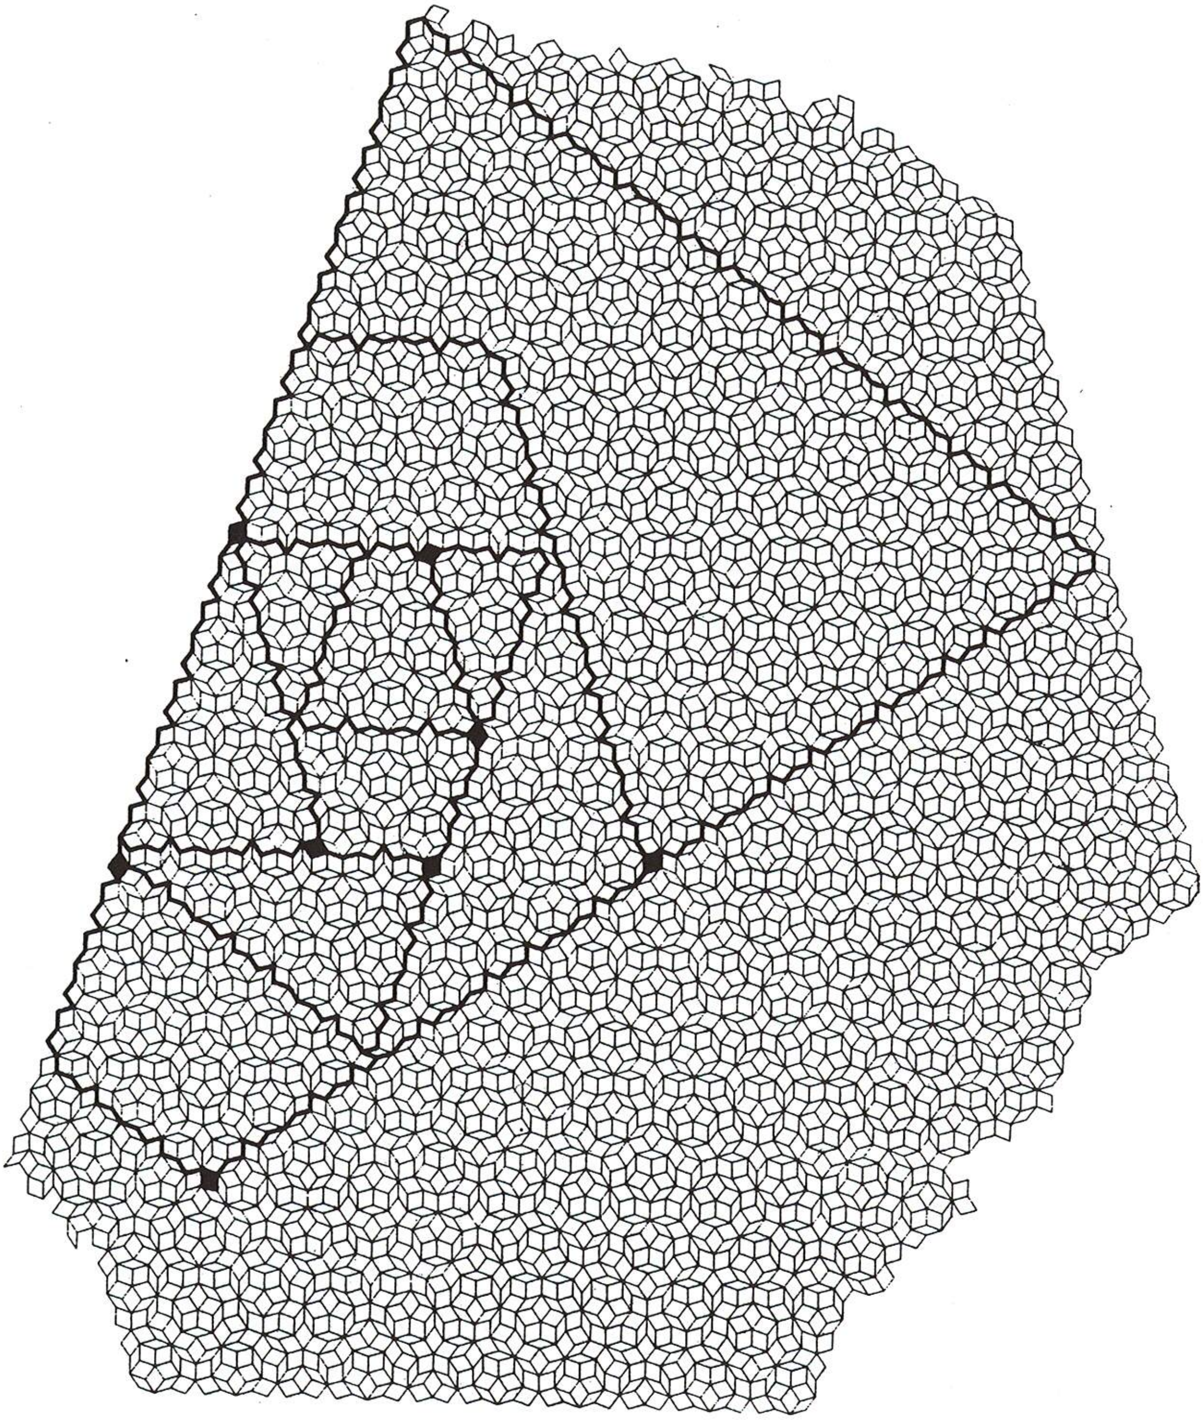

Figure 6: Growth according to the OSDS rules. Heavy lines denote dead surfaces that occurred during growth. Tiles added at marginal sites are colored black. Adding all forced tiles will eventually produce a large trapezoidal dead surface with its short base at the top.

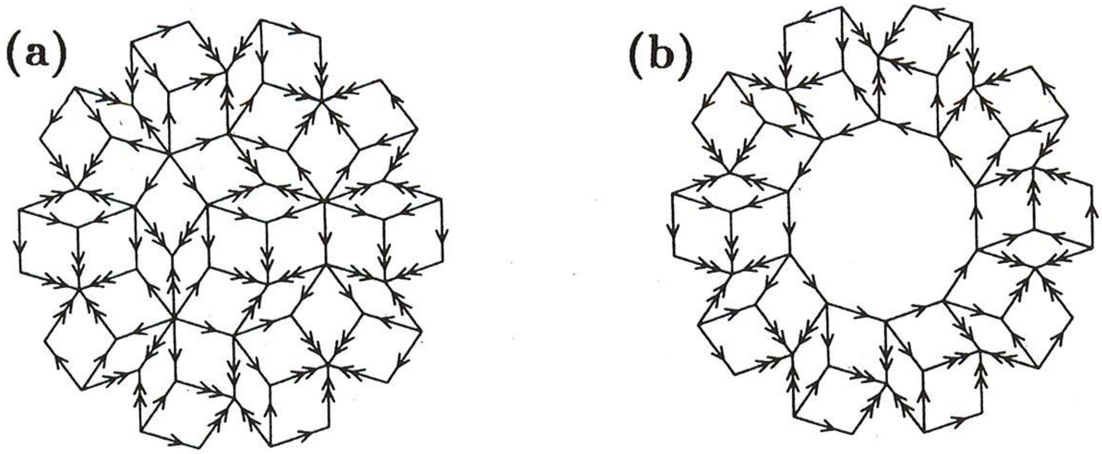

Figure 7: (a) The center of the cartwheel tiling. (b) A decapod. The central decagon cannot be filled in without violating the matching rules. Forced growth around this seed will continue ad infinitum.

may not have any forced or marginal sites. For the sake of completeness, one can simply add to the rules a list of possible starting seeds, which is clearly a local rule since the requisite seeds all contain fewer than eight tiles.

A surprising and perhaps physically significant result is obtained when one considers *imperfect* seeds of a special type. In some of the earliest work on Penrose tiles, it was noted that there exist infinite Penrose tilings that obey the matching rules everywhere except within a single decagon. The decagon lies along an infinite worm, dividing it into two semi-infinite pieces. A defect called a "decapod" occurs within the decagon when one of the pieces is "flipped" with respect to the other [20]; i.e., the two halves of the worm are in opposite orientations and hence cannot be joined without violating the matching rules along some edge within the decagon.

There is a special Penrose tiling, called the "cartwheel" [20], that has ten semi-infinite worms emanating from a central decagon. If the orientations of the worms are consistent as shown in Figure 7a, then the central decagon can be filled in with no matching rule violations. If, however, one flips some of the semi-infinite worms, a decapod can be created in the central decagon.

Consider first the decapod shown in Figure 7b, where the beginnings of the semi-infinite worms have been arranged so as to make a configuration with complete ten-fold symmetry. As we have just noted, this configuration can lie at the center of an otherwise perfect Penrose tiling. Thus one can imagine growing from this seed according to the OSDS rules without ever encountering any inconsistency. In fact, one finds that growth around the seed is greatly improved — forced growth proceeds ad infinitum without ever producing a dead surface! [18] A simple argument shows why. Suppose that a first dead surface were reached. It would have to have ten-fold symmetry, since the forced growth alone can never break the symmetry of the original seed. There is no way, however, to make a dead surface with ten-fold symmetry out of straight faces and corners making angles of 72° or 108°.

A more powerful technique for determining which decapods yield infinite forced growth makes use of the arrow matching rule decorations. A charge can be assigned that measures the circulation of the single arrows around the central decagon. Traversing any closed loop of tile edges that encloses the central decagon, each edge with a single arrow pointing in the counterclockwise direction contributes +1 to the charge and each edge with a single arrow pointing clockwise contributes -1. The fact that the circulation around any single tile is 0 implies that in any Penrose tiling with no defects the charge within any closed loop is zero. The charge of a decapod can be $\pm q$, with $q \leq 10$ an even number since there are exactly 10 single-arrow edges around the central decagon. The decapod of Figure 7b has charge 10.

Examining the arrow decorations along the straight faces of a dead surface and at the different possible corners, one finds that the 72° corner contributes ± 1 to the charge of the dead surface, while all other corners and faces contribute 0. Since the maximum number of 72° corners on a surface with only 72° and 108° corners is 2, the charge of a dead surface can only be ± 2 or 0. Thus any decapod with charge greater than 2 in magnitude acts as a seed for infinite forced growth.

Complete criteria for forcing infinite growth are not yet known, though the 61 possible decapod defects can be checked one by one. The charge alone is not sufficient, as some decapods with charge ± 2 or 0 do *not* yield infinite forced growth while others do.‡ The detailed reasoning required to check for infinite growth will not be reproduced here. It relies on the behavior of decapods under inflation and the fact that any dead surface remains dead under inflation even if it surrounds a decapod. It has been proven that a sufficient criterion for forcing infinite growth is that three consecutive arrows on the central decagon point in the same direction. [22]

‡In Gardner's Scientific American article [20], it is claimed that all but one defective decapod force infinite tilings. This is true, but only in a special sense that is not relevant here. The distinction Gardner makes is between what might be called "versatile" and "nonversatile" decapods. Versatile decapods are those that can exist as the only defect in many distinct tilings. Nonversatile decapods are those for which there is only one particular infinite tiling that obeys the matching rules everywhere outside the decapod. Gardner states that all defective decapods but one are nonversatile. We are interested here, however, in identifying those decapods for which an infinite tiling is obtained by adding only to locally forced sites (the forced sites of the OSDS rules, for example). The important point is that there exist nonversatile decapods that do not yield infinite forced growth. For these decapods, dead surfaces are encountered, but they always have dangerous faces. Thus while there is only one way that they can grow to infinity, the growth cannot be determined by locally forced additions alone.

4 Generalization to other tilings and symmetries

The existence of local growth rules for the Penrose tilings raises the question of whether other tilings support similar rules. It must be emphasized that this is *not* the same question as whether they support matching rules. Matching rules ensure only that infinite tilings are ordered, whereas here we are interested in rules for adding to finite clusters, which therefore have to anticipate matching rule violations that might be forced later by a given addition. It is not even obvious that matching rules are necessary for growth rules to exist, though they are a great help in the analysis since they guarantee that any mistake will eventually show up as an easily recognizable local defect.

The properties of the Penrose tilings that were required *for the above analysis* of local growth rules were the matching rules and the inflation-deflation operations. The matching rules could have been stipulated by specifying all possible configurations of tiles sharing vertices with a given tile rather than using arrow decorations of the tile edges. Their crucial feature is that the configurations that must be specified are bounded in size. Tilings that can be specified by matching rules of this type are called "restorable". The inflation-deflation operation was used to reduce the catalogue of dead surfaces to the point where all evolutions resulting from the addition at a marginal site could be inspected. An essential point was that adding all forced tiles and deflating (or inflating) are commuting operations. Since the determination of whether a tile is forced is made on the basis of its local environment, it is essential that local information be sufficient to determine the inflation of any region. An inflation operation that requires only local information in order to be consistently performed is called a "bounded-context" inflation.

The only quasicrystal tilings known to have both matching rules and bounded-context inflation properties are

- the Penrose tilings and certain other pentagonal tilings produced by the canonical projection procedure [23];
- the octagonal tiling produced by canonical projection;[24, 25]
- the dodecagonal tiling produced by canonical projection;[25]
- two icosahedral tilings – the canonical projection using Ammann rhombohedra [21, 26, 27, 28] and the zonohedral tiling of Ref. [29].

In all cases except the icosahedral ones, results on growth rules are available, though the analysis has not been done in the same detail as in the Penrose case.

Distinct pentagonal projection tilings are produced by different choices for the position of the perp-space acceptance domain, or equivalently by different choices for the phases of periodic grids used in the generalized dual method. [29, 30] Let γ_n

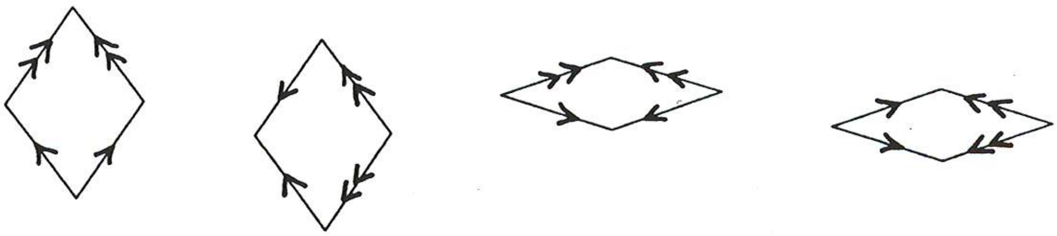

Figure 8: The tile decorations required for producing all of the canonical projection tilings. Note that the left-hand tile of each shape is just a Penrose tile.

denote the phase the n^{th} grid in units where a phase shift of 1 leaves the grid invariant and let $\Gamma = \sum \gamma_n$. Clearly, the integer part of Γ is irrelevant. The fractional part of Γ distinguishes tilings in different local isomorphism classes: tilings that contain different sets of local environments. The Penrose tilings are the duals of pentagrids with $\Gamma = 0$.

It can be shown that the only values of Γ corresponding to restorable tilings are $\Gamma = m\tau$, where m is any integer and τ is the golden mean, $(1+\sqrt{5})/2$. [23] Curiously, the exact same set of Γ's is selected by the criterion that the tilings support a bounded-context inflation. For these values of Γ, growth rules can be developed in the same spirit as the OSDS rules for the Penrose tilings. [23] The only difference is that the size of the environments that must be considered to determine whether a site is forced (or marginal) or not increases linearly with m with a coefficient of order unity.

Infinite forced growth can be obtained, as in the Penrose case, around a point defect seed. It can be shown that the four tiles shown in Figure 8 are sufficient for composing any projection tiling with $\Gamma \neq 0$. [31] The analysis of growth around decapod defects is similar to the Penrose case, though the charge is less effective in selecting decapods that force infinite growth.[22]

The octagonal and dodecagonal cases are quite similar to each other. Analyses similar to the one carried out for the Penrose tilings show that perfect growth can be guaranteed around point defect seeds (the analogues of decapods) without the aid of special marginal sites. [32] In the dodecagonal case a rule has been found that generates perfect growth around a particular *nondefective* seed containing just two tiles with the use of marginal sites, though the rule produces only one particular dodecagonal tiling. In both the octagonal and dodecagonal cases the determination of whether a site is forced must include configurations containing tiles at opposite vertices of a tile edge. It is not sufficient to consider only configurations around a single vertex.

The icosahedral case remains a mystery, mostly for technical reasons. It is much harder to examine three-dimensional configurations by hand or with the aid of a computer than it is to deal with the 2D tilings. The identification of planar dead faces is straightforward since the generalized dual construction can be used to generate the planar analogues of worm segments, but the catalogue of complete dead surfaces is much harder to construct. Although all the intuition built from the comparison of the pentagonal and icosahedral cases with regard to matching rules, inflation, Ammann grids, and projection suggests that growth rules should exist in the icosahedral case, no definitive results are available.

5 Imperfect application of the rules

In physical applications, one might expect that the growth rules would not be perfectly followed. A number of different sorts of imperfections could arise, depending on the manner in which the rules are violated. A full exploration of the possible behaviors, which would require a study of the effect of variations in unit cell shape and conventional elastic distortions on the growth, is beyond the scope of this work. Note that the restriction to rigid units makes it impossible for dislocations to form, which eliminates both the possibility of disrupting the translational order via dislocations and their attendant strain fields and the possibility of "healing" tears in the phason field by allowing them to enclose dislocations. Nevertheless, some nontrivial questions arise even when one considers the simplest deviations from the perfect growth rules for quasicrystals, deviations resulting from slight variations in the sticking probabilities. Most importantly, one would like to know whether tears in the phason field are generically produced in the limit of perfectly rigid units.

Using the OSDS rules as a starting point, a variety of effects can be considered. Perhaps the most obvious issue concerns the limit $p_m \to 0$ required for perfect growth. This issue can be avoided completely, however, if one begins with a decapod seed, where p_m can be set to zero without causing the growth rate to vanish. If one insists that p_m not be identically zero, defects will be generated at some finite length scale. The analysis of the resulting structure is difficult because one does not know how often marginal sites arise during forced stages of the growth, sites which do not remain marginal when the forced growth is continued, so it is hard to determine how many chances there are for mistakenly adding to a marginal site. In any case, the sorts of defects induced when p_m (the probability for sticking at a marginal site) is nonzero are similar to those resulting from nonzero values of p_u (the probability for sticking at an unforced site), though they are more rare.

When p_u is nonzero, there is practically no need to define marginal sites. First, the growth does not stop at a dead surface, so marginal sites are not strictly necessary. Second, for large clusters, the number of unforced sites grows more rapidly than the number of marginal sites, so that any influence of the marginal sites becomes negligible.

The essential issue is then the effect of nonzero values of p_u on the growth. It can be argued that some nonzero value is unavoidable in a physical system, since there is always some probability of adding to neighboring unforced sites simultaneously, at which point each added unit finds itself occupying a forced site and will not detach as expected at an unforced site. It is quite reasonable, however, to consider growth for very small values of p_u.

The rule for adding at an unforced site will be taken to be that one of the two choices is made at random, perhaps with probabilities weighted according to the frequency of occurrence of the different types of tiles. Given a nonzero p_u, there will always be some remote possibility of generating a cluster with arbitrary shape, just as the growth of a crystal surface by any similar rules would eventually produce arbitrarily complicated overhangs. The relevant question becomes one of scale: On what length scale do tears in the phason field typically appear for a given value of p_u?

We first note that the frequency with which tears are generated can depend on additional rules that must be used to specify the growth in the vicinity of a defect. To see what sort of rule is required, we must examine the structure of the defects produced. Fortunately, known properties of the Ammann line decoration make it easy to anticipate the types of defects that have to be handled and the existence of matching rules guarantees that all defects will eventually generate local defective sites.

Two distinct kinds of mistakes can arise when a random choice is made for which tile to add at an unforced site. Recall that the choice determines the next interval in the Ammann grid parallel to the face where the choice is made. "Jags" occur when a given interval has already been explicitly determined to be L or S and the opposite choice is made at an unforced site along the same Ammann line. If forced growth were allowed to run to completion before the unforced choice were made, then the choice would have become forced in a time proportional to the distance between the unforced site and the nearest site along the same Ammann line that has already been determined. Once a single tile is placed that determines the position of a new Ammann line, all surface sites along that line will be forced in sequence. Figure 9 shows an example of an unforced choice that results in a jag. Note that as forced growth proceeds, the jag becomes recognizable as a local violation of the single-arrow matching rule. "Sequence errors" occur when a choice is made that introduces an Ammann line interval inconsistent with the Fibonacci sequence in one direction; e.g., the wrong choice is made at a dangerous face. The matching rules ensure, however, that the sequence error will eventually produce jags in grids in other directions [19] (and nothing more complicated), so the extra rules required to handle defects need only address local configurations around single-arrow mismatches.

Figure 3 shows how a sequence error results in the formation of a jag via forced growth. The sequence error introduces phason strain in the region between the two conflicting Ammann lines. It can be shown that the presence of the phason strain

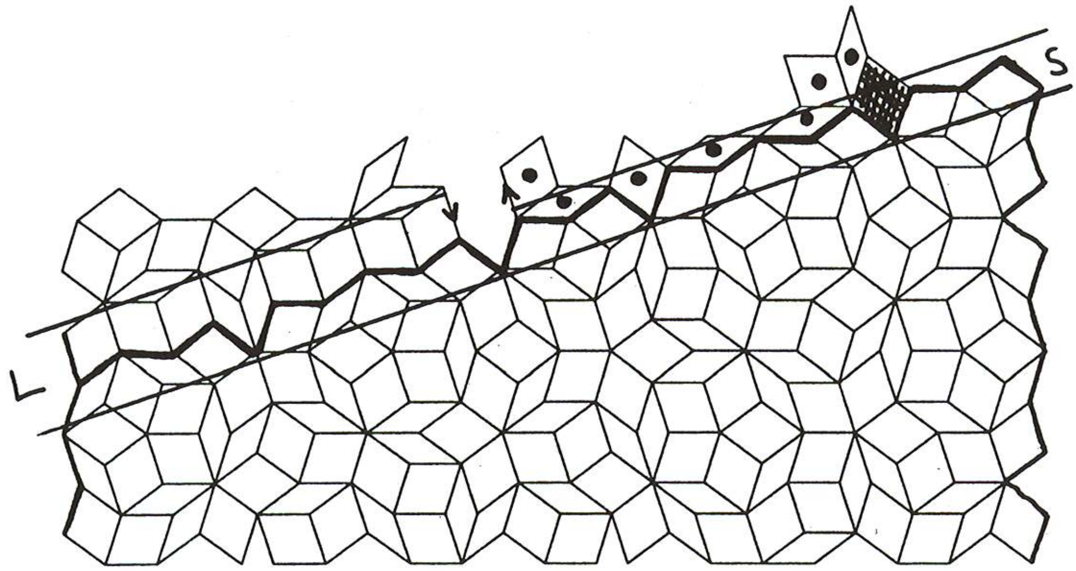

Figure 9: An unforced choice that generates a jag. The shaded tile was added before forced growth was completed. The tiles marked with dots were then forced in sequence. The two edges marked with single arrows make it impossible to continue the forced growth in a consistent way. A few selected Ammann lines are drawn to illustrate the nature of the jag in the Ammann grid context.

must be accompanied by a nonzero density of jags. [19]

Several jags from Ammann lines in different directions can occur in the same neighborhood, but the defects produced can always be enclosed by a decapod. Thus a natural rule suggests itself: Any time a site is encountered where no addition consistent with the matching rules can be made, consider the tiles required to complete a decapod enclosing the matching rule violation as forced. If the rare occasions where the growth sites are selected in an order that produces deep concavities are neglected, then this rule will generate tilings in which the only defects are the single-arrow matching rules violations found in decapods. Tears in the phason field will never develop because they are always prevented at the moment they begin to form.

At present, investigations into the distribution of phason strains generated by such a rule have not been pursued to completion. Note that when more than one decapod is present, their existence can force other decapods to appear. Furthermore, if a sequence error is made at some point, it will force an infinite number of decapods to be inserted in the strip between the two conflicting Ammann lines, unless a decapod of the appropriate type happens to arise in one of the conflicting lines so as to remove the error. The long-range effects engendered by an initial distribution of decapods are therefore difficult to predict.

An alternative to the approach of forcing the decapod surrounding a matching rule violation is to assume that any site where neither tile can be added consistently is simply left empty. [33] It has been shown by computer simulation of the growth that for tilings containing up to 10^5 tiles, values of p_u smaller than 10^{-3} produce phason fluctuations that diverge only logarithmically with sample size. It is not clear, however, whether the fluctuations on larger length scales will remain logarithmic. For $p_u = 10^{-2}$, there is evidence for power-law divergent phason fluctuations.

In large clusters grown using this approach, long lines of correlated mismatches are easily identified. [33] Whether these are properly described as tears is worth considering, for it brings to light some potentially important subtleties associated with the definition of a continuous phason field for a discrete tiling. Consider a portion of a sample in which the phason variable, $\mathbf{w}$ is fixed on the left and right boundaries at different values. A linear variation in $\mathbf{w}$ across the sample would produce defects uniformly distributed in the sample with a density determined by the imposed difference in $\mathbf{w}$. Now if $\mathbf{w}$ is made to jump discontinuously across the midline of the sample, all these defects would pile up on that line. Note, however, that if the jump in $\mathbf{w}$ is small, only a small subset of the Ammann lines will contain a defect, so that there may be large sections of the discontinuity line that do not appear to contain a defect. Thus the determination that a tear exists might be made on the basis of the observed correlations between defects, rather than on the extent of an unbroken line of obvious defects. In addition, one would not expect to observe a perfect alignment of the defects. In general, a tear will appear only as a relatively sharp variation across a region of finite width. It therefore becomes somewhat difficult to distinguish a tear, or the beginnings of one, from a more benign variation. Ultimately, when growth proceeds long enough that the differences in $\mathbf{w}$ across typical tears becomes very large, tears will become recognizable as the familiar "cracks" observed in random growth models, but this takes much longer than one might expect.

It is interesting to consider the growth around a defect under rules that leave defective sites empty. It is observed in simulations that when a defect is produced, two bumps on the surface tend to grow on either side of the defect. Because the bumps grow independently, different random choices made during their growth can result in their evolving rather different values of the phason variable, thereby generating a tear. It is clear that a rule like the one described above for immediately enclosing the defect will not generate separate bumps. One might therefore expect the divergences in the phason variable to be pushed out to much larger length scales. Further work is clearly required on this point.

6 Icosahedral growth: Open questions

The understanding of the growth of Penrose tilings gives us significant insights that can be carried over to the icosahedral case. As mentioned above, there are icosahedral structures that share all the properties of the Penrose tilings that we have used in the analysis. Two different icosahedral tilings provide natural contexts for the investigation of growth rules; the canonical projection tilings composed of Ammann rhombohedra and the "zonohedral tiling" or "Penrose LI class" tiling. [29]. The zonohedral tiling is most useful for aspects of the analysis involving Ammann planes or inflation/deflation. The geometrical simplicity of the Ammann rhombohedra and their matching rules, on the other hand, make them more natural candidates for investigation via computer simulation. The relation between the two structures is that the vertices of the zonohedral tiling are a subset of the vertices of the projection tiling. The situation is a bit complicated, though, as evidenced by the fact that the mapping is two to one: each projection tiling can be formed by decorating either of two distinct zonohedral tilings. [34]

There do not appear to be any fundamental barriers to the development of an algorithm for perfect icosahedral growth along the lines of the OSDS rules. Nevertheless, there may be significant differences that arise in the analysis of defects, including seeds for perfect growth. For icosahedral structures, the analogues of Ammann lines are Ammann planes and a jag becomes a line of mismatches (a step in an Ammann plane) rather than an isolated point mismatch. The analogues of decapods consist of infinite lines of mismatches emanating from a central region which might be called an "icosapod".

A useful relation exists between the matching rules for the Ammann rhombohedra projection tiling and the matching rules for 2D pentagonal projections with arbitrary Γ.[28] One considers a slab of Ammann rhombohedra joined along parallel edges; i.e., all the rhombohedra dual to a single plane in the generalized dual construction. The top surface of this slab can be projected onto a plane perpendicular to the defining edge direction to form a 2D canonical projection tiling. The matching rules for the Ammann rhombohedra are then seen to map precisely onto the rules of Figure 7 and the defect lines emanating from an icosapod correspond to ordinary decapods in certain slabs pierced by the lines. Unfortunately, there is no unique way of reconstructing the 3D slab locally given the 2D projection, so one cannot use the results on decapods to immediately generate infinite 3D growth.

The nucleation and development of tears in the phason field is rather difficult to characterize in the 3D case. The possible topology of steps (2D jags) is not even well-understood at present. Defect lines can form closed loops or infinite curves and it is not clear what sorts of structures will be typically generated during growth. It is conceivable that tears are effectively suppressed even for rather large values of p_u because the constraints imposed by the forced growth in the vicinity of a defect are more severe in three dimensions than in two. The issue, however, is far from being resolved.

The investigation of local growth rules for icosahedral quasicrystals is still in its early stages. Though all indications are that no intrinsic problems prevent their existence, we have yet to arrive at an explicit algorithm for perfect growth in the spirit of the OSDS rules. Such rules would have the status of an existence proof, demonstrating that growth kinetics favoring long-range quasiperiodic translational order are possible in principle. A thorough understanding of them might suggest novel strategies for preparing real quasicrystals with more perfect order, or at least appropriate ways of constructing more realistic models of quasicrystal growth.

Acknowledgements: Conversations with D. DiVincenzo, K. Ingersent, and P. Steinhardt concerning various points discussed in this paper are gratefully acknowledged. The figures, as well as portions of the theory, were generated with the aid of software written by S. Langer.

References

[1] D. Shechtman, I. Blech, D. Gratias, and J. W. Cahn. *Phys. Rev. Lett.*, **53** (1984) 1951.

[2] P. A. Heiney, P. A. Bancel, P. M. Horn, J. L. Jordan, S. LaPlaca, J. Angilello, and F. W. Gayle. *Science*, **238** (1987) 660.

[3] A. Tsai, A. Inoue, and T. Masumoto. *J. Mat. Sci. Lett.*, **6** (1987) 1403.

[4] C. A. Guryan, A. I. Goldman, P. W. Stephens, K. Hiraga, A. P. Tsai, A. Inoue, and T. Masumoto. *Phys. Rev. Lett.*, **62** (1989) 2409.

[5] A. Tsai, A. Inoue, and T. Masumoto. *Jpn. J. Appl. Phys.*, **26** (1987) L1505.

[6] P. W. Stephens. in *Aperiodicity and Order* **3**, ed. M. Jarić and D. Gratias (Academic Press, Boston, 1989).

[7] T. C. Lubensky, S. Ramaswamy, and J. Toner. *Phys. Rev.*, **B32** (1985) 7444.

[8] P. A. Bancel. *Phys. Rev. Lett.*, **63** (1989) 2741.

[9] K. Ingersent, this volume.

[10] P. Bak. *Phys. Rev. Lett.*, **54** (1985) 1517.

[11] D. Levine, T. C. Lubensky, S. Ostlund, S. Ramaswamy, P. J. Steinhardt, and J. Toner. *Phys. Rev. Lett.*, **54** (1985) 1520.

[12] V. Elser. in *Aperiodicity and Order* **3**, ed. M. Jarić and D. Gratias (Academic Press, Boston, 1989).

[13] C. L. Henley. *Comm. Cond. Mat. Phys.*, **13** (1987) 59.

[14] T. C. Lubensky, J. E. S. Socolar, P. J. Steinhardt, P. A. Bancel, and P. A. Heiney. *Phys. Rev. Lett.*, **57** (1986) 1440.

[15] J. L. Robertson and S. C. Moss. University of Houston preprint.

[16] R. Penrose. *Bull. Inst. Math. Appl.*, **10** (1974) 266.

[17] R. Penrose. in *Aperiodicity and Order* **1**, ed. M. Jarić (Academic Press, Boston, 1988).

[18] G. Y. Onoda, P. J. Steinhardt, D. P. DiVincenzo, and J. E. S. Socolar. *Phys. Rev. Lett.*, **60** (1988) 2653.

[19] J. E. S. Socolar, T. C. Lubensky, and P. J. Steinhardt. *Phys. Rev.*, **B34** (1986) 3345.

[20] M. Gardner. *Sci. Am.*, **236** (1977) 110.

[21] D. Levine and P. J. Steinhardt. *Phys. Rev.*, **B34** (1986) 596.

[22] K. Ingersent. PhD thesis, University of Pennsylvania, May 1990.

[23] K. Ingersent and P. J. Steinhardt. *Phys. Rev. Lett.*, **64** (1990) 2034.

[24] B. Grunbaum and G. C. Shephard. *Tilings and Patterns* (Freeman, San Fransisco, 1986).

[25] J. E. S. Socolar. *Phys. Rev.*, **B39** (1989) 10519.

[26] A. Katz and M. Duneau. *J. Phys.*, **47** (1986) 181.

[27] A. Katz. *Commun. Math. Phys.*, **118** (1988) 263.

[28] J. E. S. Socolar. *Commun. Math. Phys.*, **129** (1990) 599.

[29] J. E. S. Socolar and P. J. Steinhardt. *Phys. Rev.*, **B34** (1986) 617.

[30] N. de Bruijn. *Proc. Kon. Nederl. Akad. Wetensch.*, **A84** (1981) 39.

[31] M. Kleman and A. Pavlovitch. *J. Phys.* **C3**, Workshop on Aperiodic Crystals (1986) 229.

[32] S. Langer, D. P. DiVincenzo and J. E. S. Socolar, unpublished.

[33] M. Ronchetti and M. V. Jarić. in *Quasicrystals* ed. M. Jarić and S. Lundqvist (World Scientific, Singapore,1990).

[34] J. E. S. Socolar and K. Benson, unpublished.