Construction of periodic approximants for the canonical-cell model of a quasicrystal

M. E. J. Newman and C. L. Henley

Laboratory of Atomic and Solid-State Physics, Cornell University, Ithaca. NY 14853–2501.

M. Oxborrow

Niels Bohr Institute, Blegdamsvej 17, DK–2100 Copenhagen, Denmark.

'Canonical-cell' packings are (possibly random) three-dimensional tilings related to atomic structure models of icosahedral quasicrystals. Here we describe an algorithm for constructing a list of all canonical-cell tilings with a given periodic structure. The algorithm is more general than its application to the canonical-cell model and could be used to find periodic structures in other tiling models with other symmetries. It involves considering all possible vertical towers of canonical cells with periodic boundary conditions in the two horizontal directions, and constructing a subset of the surfaces that cut through the tower such that every possible tower of cells can be specified by the sequence of transitions from one surface to another up the tower. In this way, we have produced an exhaustive list of the 32 distinct tilings (316 if we include those generated by symmetry operations) of a cube with edge $\tau^3 = 2 + \sqrt{5}$ times the smallest possible cube edge. Using averages over this list of finite structures as an approximation to an average over the possible states of an infinite random tiling of canonical cells, we calculate diffraction patterns which exhibit diffuse arcs similar to those seen in diffraction experiments on icosahedral TiMnSi and other quasicrystalline alloys.

I. INTRODUCTION

Quasicrystals are materials which possess a point symmetry, normally 5-fold, 8-fold, 10-fold or 12-fold, which is incompatible with a periodic Bravais lattice, and yet which have sufficient translational order that they exhibit sharp diffraction peaks. In some alloys, such as icosahedral AlCuFe (Bancel 1991) and icosahedral AlPdMn (Boudard *et al.* 1992), the translational order is good enough that the diffraction peaks are resolution-limited. Many aperiodic structure models have been suggested to explain these properties. Some of these models describe the origin of the symmetry and long-range order, others specify the atomic positions, and a few address both questions; our concern here is with the symmetry and long-range order.

The models for the origins of quasicrystalline symmetry fall into two general classes. One class consists of perfectly quasiperiodic structures in three dimensions, often based on quasiperiodic tilings which are generalizations of the two-dimensional Penrose tiling (Penrose 1974, Gardner 1977). These models are analogous to the description of crystalline solids in terms of Bravais lattices. Quasiperiodic lattices are constructed in a deterministic way using any of a variety of methods such as projection from higher dimensions (Duneau and Katz 1985, Elser 1986, Bak 1986), dual grid methods (Socolar *et al.* 1985, Gähler and Rhyner 1986), or inflation rules (Penrose 1974, Gardner 1977, Smith 1993). Alternatively, lattices may be specified by local 'matching rules' (Levine and Steinhardt 1984, Katz 1988, Levitov 1988, Socolar 1989, Ingersent 1991) which force quasiperiodicity, suggesting a mechanism by which it might arise during growth. However, to permit growth without defects, the matching rules must be supplemented with additional 'growth rules' (Onoda *et al.* 1988), making this scenario less convincing.

There is a second class of 'random', i.e., nondeterministic, models in which one specifies how clusters of atoms or tiles are to be fitted together locally, without restricting the arrangements of the structure at long distances. These models further divide into 'icosahedral glasses' (Stephens and Goldman 1986, Stephens 1989, Elser 1989, Robertson and Moss 1991, 1993, Gibbons and Kelton 1993), which consist of networks in which clusters of atoms are joined together by 'bonds' or 'linkages' running in certain specified directions, and 'random tilings' (Elser 1987, Henley 1991) which are similar to icosahedral glasses, but include additional constrains to ensure that the network is built from only a small finite number of tiles. Random-tiling models have been studied in two dimensions with 8-fold (Li *et al.* 1992), 10-fold and 12-fold (Oxborrow and Henley 1993) symmetries, and in three dimensions with icosahedral symmetry.

Random tiling models have a nonzero system entropy arising from the many ways in which their elements may be packed. Ideally there is a one-to-one mapping of tiling configurations onto real atomic configurations of alloys by means of rules specifying how the tiles are to be decorated with atoms. This means that the tiling entropy is identical to that of the atomic configurations, provided the configurations are nearly degenerate. This entropy can reduce the free energy of the random-tiling quasicrystal with respect to a similar deterministic crystal, offering an explanation of how the quasicrystal state might be stabilized. It has been shown that random tiling models, like the ideal quasiperiodic models, are consistent with the long-range order seen experimentally.

Given a random tiling model, several questions naturally arise:

(i) What is the value of the entropy density? This affects the entropic stabilization. In the case of the twelve-fold 'square-triangle' tiling (Oxborrow and Henley 1993), an exact solution has been found (Widom 1993) to this question, and in other cases approximate numerical values are available (Widom *et al.* 1989, Strandburg 1991, Shaw *et al.* 1991, Li *et al.* 1992).

(ii) What are the values of the elastic constants which determine the 'phason elasticity'? Icosahedral tilings can be represented as the projection into real space of a subset of the vertices of a six-dimensional hypercubic lattice. This projection from six into three dimensions leaves us with three other degrees of freedom in the structure, which form a three-dimensional perpendicular or 'perp' space orthogonal to real space. It turns out that the entropy of a random tiling is related to the mean slope, or 'phason strain', of the plane that the tiling forms in perp space with respect to its position in real space (Elser 1985, Jarić and Nelson 1988). The entropy has a maximum at some equilibrium strain $\mathbf{E}_0$. The coefficients of the quadratic variation about this maximum are called the 'phason elasticity' (Henley 1991), and are a measure of the strength of the long-range order in the structure. They affect such properties as the size of diffuse wings around the Bragg diffraction peaks and possible instabilities of the quasicrystal phase with respect to related crystalline phases (Widom 1991).

(iii) 'Phason fluctuations' around the average phason strain of the tiling determine other quantities related to intermediate-range order. Calculating these quantities involves averaging over many competing tiling configurations, and determines such measurable quantities as phason Debye-Waller factors (needed for structure refinements from diffraction data) and the diffuse arc features seen in diffraction patterns (see Sec. V below).

(iv) Another goal, important from a practical point of view is the generation of an ideal quasiperiodic version of the random tiling, if this is not already known.

In this and a subsequent paper (Newman and Henley, to be published) we study a particular random-tiling model, the 'canonical-cell model' (Henley 1991), and address the four questions just outlined. A separate problem is to find rules for decorating the tiles with atoms so as to reproduce the correct stoichiometry, density, diffraction pattern, and cohesive energy; we do not address those questions at present.

A. Random-aggregation models, *bc*-networks, and canonical-cell tilings

Several atomic structure models for icosahedral quasicrystals are formulated as decorations of a network of nodes. In the real quasicrystal these nodes would represent icosahedrally-symmetric clusters of atoms and the spaces between the clusters would be filled with 'glue atoms'. In the models with which we shall be dealing, the nearest-neighbor nodes are joined by 'linkages' of two types, which we refer to as type b and type c. Type b linkages run parallel to the axes of two-fold symmetry of the reference icosahedron and are of length b; type c linkages run parallel to the axes of three-fold symmetry, and are of length $c = b\sqrt{3}/2$. We call such networks bc-networks. They can be periodic, quasiperiodic, or random, and they can be constructed with a wide range of densities. The shortest distance we allow between two nodes in these networks is c, so it is convenient to measure the network's density in terms of the packing fraction when spheres of diameter c are placed on the nodes. If icosahedral symmetry is the ruling principle in our material, it is natural to hypothesize that the volume occupied by icosahedral clusters should be maximized, which implies that the packing fraction of the nodes ought to be maximized. The maximum packing fraction is 0.68, which is attained by a simple periodic b.c.c. packing corresponding to relatively simple periodic alloy structures found in nature and closely related to quasicrystals (Elser and Henley 1985, Guyot and Audier 1985, Henley and Elser 1986).

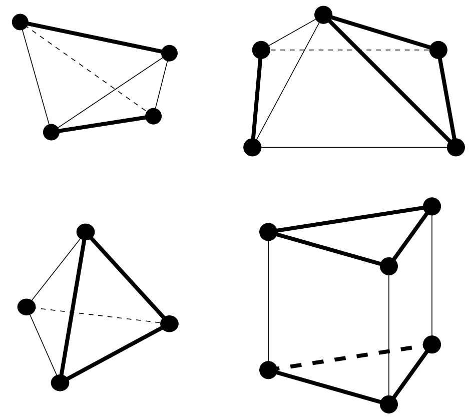

FIG. 1: The four canonical cells. There are two irregular tetrahedrons, an irregular square pyramid, and an equilateral triangular prism. The single lines represent the b-type linkages, which run parallel to the two-fold axes of the icosahedron; the double lines represent the c-type linkages, which run parallel to the three fold axes.

Several models based on bc-networks have been suggested for icosahedral alloys, but they suffer from rather low packing fractions, on the order of 0.55. 'Canonical-cell' tilings are a class of bc-networks with much higher packing fractions, in the range 0.60–0.62. They achieve this by restricting the allowed shapes of the spaces between the nodes. The linkages in the network outline various irregular polyhedra, which have linkages for their edges and nodes for their vertices. In a canonical-cell tiling each such polyhedron is required to be one of the four canonical cells depicted in Fig. 1.

An important consequence of the restriction to the four canonical cells is that there are only 32 possible ways that corners of the cells may be arranged around a node to completely fill its 4π of solid angle, and correspondingly 32 allowed ways in which linkages may meet at a node. These 'node environments' were enumerated by Henley.[50] For a *bc*-network to be a canonical-cell tiling it is necessary but unfortunately not sufficient for every node in the network to be surrounded by one of the 32 allowed node environments. It is a pity this is not a sufficient condition, since it would give us a simple method of checking whether a tiling was made up of canonical cells, if all we had to do was check the types of node environments appearing in the network. Luckily, there turns out to be a straightforward augmentation of the idea of the node environment which *does* give us a simple check like this. The solution is to add to the model, simply for the purposes of checking the structure, a new type of linkage, the *d*-linkage, corresponding to the third-nearest neighbor distance $d = b\sqrt{5}/2$ (the two shortest distances being the *b*- and *c*-type linkages) which is the length of the diagonal of the rectangular face of the two rightmost cells in Fig. 1. We then retabulate the allowed node environments, including the *d*-type linkages. Now if we ensure that each environment in a network of nodes is of an allowed type, including the *d*-type linkages, we can demonstrate rigorously that the network will be entirely composed of correct canonical cells.[51]

By restricting our structures to the four canonical cell shapes we ensure that our model forms a true three-dimensional tiling. By contrast, there are a number of icosahedral glass models (Elser 1989, Robertson and Moss 1991, 1993, Gibbons and Kelton 1993), which take the form of *bc*-networks with packing fractions comparable to those of the canonical-cell model, but which do not have restrictions on the number of different cell shapes in the structure. These icosahedral glasses are much less like three-dimensional tilings of space than they are like random-aggregations of nodes in a network. There will usually be a finite number of cell types in a given icosahedral glass model, though this depends on the exact rules used in the construction of the glass. But the number of types may be large, and in any case it is not going to be easy to break the network down into a tiling of polyhedra. In addition, it can be shown (Henley 1991) that the canonical cells are the four best-packed polyhedra, so that the presence of any other polyhedra, as in an icosahedral glass, necessarily lowers the packing fraction. From the similarity of packing fractions between the canonical-cell tilings and the icosahedral glasses, we infer that the best-packed icosahedral glasses consist almost entirely of canonical cells.

The tiling nature of the model is important when one considers how it is to be related to a model for a real physical quasicrystal. Specifying the cluster of atoms with which the nodes are to be decorated captures the most important features of the structure, but further work needs to be done in the interstitial spaces between nodes to complete the picture: either the clusters leave gaps in these spaces, in which case additional rules must be specified for filling the gaps with atoms (often called 'glue'), or else the clusters overlap, in which case additional rules must be specified for modifying the clusters in the overlap region so that the atom placements will be mutually consistent. In either case, the extra decoration rules must depend on the local pattern of nodes around the space in question. If there is a finite set of these patterns, as there is in a tiling, it becomes relatively simple for us to specify the extra rules in a deterministic way, and thereby define a one-to-one mapping of tilings onto atomic structures. This gives us a convenient way to model configurational randomness, for example in the calculation of diffraction patterns. Furthermore, since the decoration is deterministic, the structural entropy of the real material is that of the tiling, so that thermodynamic calculations on the tilings (Newman and Henley 1993) are relevant to the real system. The decoration of canonical-cell tilings to produce real atomic models of the icosahedral phases is to be discussed in a later paper (Henley, Mihalkovič, Newman and Oxborrow, to be published).

B. Previous methods for building periodic structures

We are motivated by the desire to construct periodic, space-filling, canonical-cell structures: i.e., packings of canonical cells which obey periodic boundary conditions in three independent spatial directions. Of course, quasiperiodic structures are incompatible with periodic boundary conditions. However, given a well-defined quasiperiodic structure one can construct a series of periodic 'approximants', which approximate to the quasiperiodic structure and converge towards it in the same way that rational approximants of higher and higher order converge to an irrational number. (For example, the series $\frac{1}{1}$, $\frac{2}{1}$, $\frac{3}{2}$, $\frac{5}{3}$, $\frac{8}{5}\ldots$ converges to $\tau = (1+\sqrt{5})/2$.) Each successively larger unit cell is indistinguishable from a fragment cut out of the infinite quasiperiodic structure, and forms a better approximation to it.

A number of different periodic structures which can be built out of canonical cells have been described. These range from simple structures entirely composed out of one type of cell (the b.c.c. packing of *A*-type cells, for instance (Henley 1991)), to the large cubic structures of Mihalkovič and Mrafko (1993). The principal present shortcoming of the canonical-cell model however, is the absence of any technique for generating arbitrarily large periodic approximants to some infinite quasiperiodic packing of cells, in the way that one can for the Ammann tiling of rhombohedra, for example. Ideally, we would like either some projection method (Duneau and Katz 1985, Elser 1986) which would create a canonical-cell tiling by projection from a six-dimensional hypercubic lattice, or we would like a set of inflation rules (Penrose 1974, Gardner 1977, Smith 1993) that would produce from some given starting structure a hierarchy of successively larger and larger approximants to an infinite tiling. Neither of these goals has yet been achieved, and in their absence other methods have evolved for finding large periodic structures. One of these is the Monte Carlo technique given by Mihalkovič and Mrafko (1993), which has found the largest structures so far.

In this paper we present a new technique, which takes a different approach. Our technique gives us a way to search exhaustively for all canonical-cell structures with a given set of boundary conditions, thus finding, for example, all possible periodic structures with cubic boundary conditions of a given size. A computational implementation of the technique has given us periodic approximants considerably larger than those previously found by hand. Although these are still smaller than those found using Monte Carlo methods by Mihalkovič and Mrafko (1993), our program produces an exhaustive list of such structures. The Monte Carlo technique by contrast, finds only a small subset of the structures of a given size, and tends on average to find structures of particularly high packing fraction.

C. Outline of the paper

The paper is organized as follows. In Section II we describe the construction of the 'stacking graph' which codifies a subset of the possible structures that can be built out of canonical cells. In Section III we give examples of the construction of these graphs for three different sizes of system. Section IV explains how the stacking graph can then be used to generate periodic structures made of canonical cells, and gives the results of these calculations for the three systems examined in Section III. Amongst other things, we have produced a list of structures obeying cubic boundary conditions of the so-called '$\frac{5}{3}$' size. In Section V we make use of these new structures to produce an approximate diffraction pattern for the random tiling of canonical cells.

Apart from the construction of periodic structures, the stacking graph can be used as the basis for a transfer matrix calculation of the thermodynamic properties of the ensemble of random tilings of canonical cells, including the random tiling entropy and the phason elasticity. We will present details of these calculations in a later paper (Newman and Henley, to be published). We have also used the periodic structures discussed in Section IV as the basis for realistic atomic models of icosahedral alloys, by specifying atomic decorations for the canonical cells. We expect to publish these studies separately as well (Henley, Mihalkovič, Newman and Oxborrow, to be published).

II. GROWTH AND THE STACKING GRAPH

The essential idea behind our way of constructing new periodic canonical-cell structures, is the development of a method for describing and cataloging all the possible structures that can be built from canonical cells, whether they are periodic or not. Given a way to catalog such structures, and a method for generating the complete catalog, we will be able to search amongst the structures in it for those which satisfy a certain set of boundary conditions. Unfortunately, the task of creating the catalog is a hard one, because the structures are three-dimensional, and necessarily do not fall easily onto any lattice. The approach we have developed turns the three-dimensional problem into a one-dimensional one by imposing periodic boundary conditions on our structure along two independent directions. The structures in our catalog are then all possible towers of canonical cells having a particular, usually square, base, and obeying periodic boundary conditions in, say, the x and y directions, but not necessarily obeying them in the z direction.

This formulation is similar, but not identical, to the transfer-matrix approach previously developed for two-dimensional tilings (Widom *et al.* 1989, Shaw and Henley 1991, Henley 1991, Li *et al.* 1992). In those approaches, one slices the structure up into layers of cells. In general, these layers cover the entire cross-section of the system, so that the average number of cells in a layer scales with the cross-sectional area. In the present approach (see Sec. II A) we similarly divide our tower of cells into layers, but the cells in a layer are now just those which are forced by the packing rules, and their number is expected to be finite rather than scaling with the cross-section of the tower.[52] The method is analogous in some ways to the method used in constructing transfer-matrices for Ising models on lattices, where one takes spiral boundary conditions and adds just one spin at a time (Nightingale 1990), and its advantages are similar.

The two ways of describing a packing by a sequence of layers can both be used, in principle, for any kind of random-tiling structure. However, when applied to the canonical-cell problem, the enumeration of possible new layers using the earlier method is similar to the task of constructing a set of 2D tilings of the cross-section, under constraints imposed by the layer below. The magnitude of this task is expected to scale exponentially with the cross-sectional area, and so becomes very large for large systems. The new method on the other hand involves only a finite amount of work for each layer, tending to a constant as the cross-section becomes large.

A. Dead surfaces

Imagine the process of building a tower of canonical cells, starting from some initial layer of canonical cells which spans the entire cross-section. We build the tower up by adding cells to the open top surface, which we will refer to as the 'growing surface'. This surface can be specified by stating the types and orientations of the cell faces which make it up, or alternatively, the positions of the nodes at the corners of the faces, and the bonds which join them to one another. It turns out that for the majority of the possible growing surfaces satisfying the boundary conditions, there is at least one move that is forced; there is some crevice on the surface which, because of its shape, can only be plugged with one particular cell or combination of cells. And if we make this forced move, we are very likely to find that as a result some other move is also forced. We thus have no choices to make about how the construction is to proceed until we finally arrive at a growing surface in which every crevice can be filled with cells in at least two different ways. We call such surfaces 'dead surfaces'. Only when we are at a dead surface do we have to make a choice about how the building is to proceed.[53]

Now if we are given a particular tower of canonical cells to catalog, we know that we must be able to represent it by the sequence of dead surfaces passed through as it is constructed from some initial layer of cells at the bottom. Since by definition no choices about how the construction is to proceed are possible between the dead surfaces, this sequence of surfaces specifies the structure, and this is the basis for our cataloging system. The number of dead surfaces needed to describe every possible structure within a tower of finite cross-section is itself finite, so we can label each surface with a number, and each structure is then associated by a sequence of these numbers.

These sequences are in fact not *quite* unique. It is rare but possible for one to make two different choices about how some crevice on the growing surface is to be plugged, and find them both ending up with the same new dead surface. The cells added in the two cases must be different, but the eventual outcome can, occasionally, be the same. Thus, to completely catalog all structures, we must catalog not the sequence of dead surfaces, but the sequence of moves from one dead surface to the next. For example: $(1 \to 2)$, $(2 \to 5)$, $(5 \to 3)'\ldots$ Here the prime on the last transition indicates that there are at least two ways of adding canonical cells to surface 5 to reach surface 3, no other choices being made in between, and that the choice we have made in this case is the one denoted $(5 \to 3)'$, not $(5 \to 3)$. Given a catalog of moves like this, we have a unique, one-dimensional description of every possible tower of canonical cells with the given (normally square) cross-section and periodic boundary conditions.

The name 'dead surface' was originally introduced by Onoda and co-workers (1988) to describe a similar concept that arises in their growth algorithm for two-dimensional Penrose tilings. However, there are differences between our technique and theirs. Unlike their algorithm, we are not trying to generate a perfect quasiperiodic tiling in any sense. Also there are no matching rules in force between our cells; we require only that our structures fit together with no gaps between cells. This allows the entropy of the ensemble of all possible such tilings to be finite. A related difference is that the number of cells between one dead surface and the next is finite in our case, and does not increase indefinitely with the cross-sectional area. (We find that about ten are added on average.) This means that the number of dead surfaces in a structure scales like the volume, and this makes the entropy extensive.

We can estimate the number of different dead surfaces satisfying a given set of boundary conditions by regarding the dead surfaces as a random tessellation of irregular triangles and quadrilaterals. Such a tessellation will have an extensive entropy, so the number N of possible dead surfaces should scale exponentially with the area A of the base of the system:

$$N = N_0 \mathrm{e}^{\lambda A}. \tag{1}$$

The numbers of surfaces found for the $\frac{2}{1}$, $\frac{3}{2}$, and $\frac{5}{3}$ systems fit this dependence well, with coefficients $N_0 = 0.52$, $\lambda = 0.45/b^2$. In Section III C we use this relation to predict the number of dead surfaces in a large system, prior to their explicit enumeration.

B. Moves and forcing

Complete information about all the canonical-cell structures of a given cross-section is contained in a list of all the possible transitions from one dead surface to another. This information is best represented in the form of a 'stacking graph', whose construction we now explain. Starting with any one dead surface, we consider all the possible ways in which we can add a single cell, or a cluster of cells, to fill the unoccupied solid angle surrounding one of the nodes on the surface. The choices we can make are most conveniently specified in terms of the linkages surrounding the node, which will form an incomplete version of one of our 32 allowed node environments. Since no moves are forced on a dead surface, we know that each such incomplete environment must be compatible with at least two out of the 32 possibilities.

Since we are primarily interested in using the dead surfaces as a description of our structure, and not as a record of the order in which the structure was built, it is not necessary to investigate the ways in which every node in the surface can be completed. It is only necessary to consider the possible completions of one single node in the surface. This node may be chosen in any way we like. We could choose it at random from the nodes in the surface if we wanted, but it is desirable that our methods be deterministic, so that the description of a particular structure in terms of dead surfaces should always be the same. So we invent a rule for choosing one node from the surface, and in the present calculation the rule was that we go for the one node which has the fewest possibilities for completion. When several nodes tie for first place, we narrow things down by considering only the ones that are lowest-lying in the tower (that is, the ones with the smallest z-coordinate), and of these choose the one nearest the origin. This specifies the node uniquely. The vast majority of the time, we find that the node we choose has only two possibilities for completion. Once we have made our choice and completed the node, we also add any other nodes and linkages that are forced as a result, until we reach another dead surface. Then the process repeats. It is clear that the string of choices made and the dead surfaces that result from them are sufficient to reconstruct the structure, given the original starting surface.[54]

We now construct a sort of decision tree, which we call a 'stacking graph'. The stacking graph codifies all the possible structures we can build out of canonical cells with periodic boundary conditions of the type we have been enforcing. To construct the stacking graph we take our starting surface, which we will assume to be a dead surface. (If it isn't, we simply add cells to it until it is, and then use that as our starting surface.) We label this as surface number 1. Now we choose the node on that surface which we are going to complete, and in turn we try out each possible completion of that node, and find the new dead surface to which each completion leads us. Since there are typically two possible completions of our node, we will probably get two new dead surfaces, which we label '2' and '3'. Occasionally, there will be three or more possible completions and three or more new dead surfaces. Other times forced moves which follow upon the completion of the starting node will not fit together properly, and one or more lines of investigation can come to a dead end, as described in the footnote at the beginning of Section II A, so that a dead-surface leads to only one further surface, or none.

When the completion of our chosen node does lead to a new surface, we then repeat the growth process for that dead surface, giving unique labels to the new dead surfaces *it* produces, and so on. This generates a tree-like graph, in which the numbers representing the dead surfaces are joined by arrows representing the possible transitions from one dead-surface to the next. Sometimes dead surface n will have a possible successor surface that has appeared already in the list as dead surface m; in that case we don't give the successor a new label, but merely draw an arrow from n to m. Also there is no need to examine the possible ways of adding cells to this repeated surface, since we will have already done that the previous time this surface appeared. In order to recognize repeated surfaces properly, we have to allow for the fact that when they reappear, they may be translated from the position in which they previously appeared. (One can also envisage a form of the algorithm in which rotations and inversions of a surface would be treated as identical with it. We did not make this identification in our calculations, but such an algorithm would work just as well as the one described here.)

As we continue to complete nodes in dead surfaces, we find our algorithm leading to more and more surfaces that we have seen before, and fewer and fewer new ones. Eventually, the whole process closes in on itself as all remaining lines of investigation lead to old surfaces and we run out of new ones to examine. We believe that this will always happen, since the number of dead surfaces needed to describe all possible structures is, as we have said, finite. It is possible for the stacking graph to close in on itself before the complete set of dead surfaces has been found,

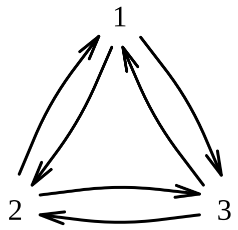

FIG. 2: The decision tree for the system with square base of length τb on each side. There are only three dead surfaces, labeled 1, 2, and 3.

in which case the remaining dead surfaces fall on one or more other disjoint pieces of the stacking graph, with no transitions leading from the surfaces of these pieces to the surfaces of the main stacking graph. This happens in the case of the $\frac{5}{3}$ system (see Section III C), for which the stacking graph has three disjoint pieces. When we refer to the stacking graph, we must include all disjoint pieces, if we wish to include all possible structures built of canonical cells, and when we construct such a stacking graph we have to go through the procedure described above for each disjoint piece.[55] Any possible structure that can be built from canonical cells, with these periodic boundary conditions, is then completely specified by a particular path through the stacking graph, since this tells us the sequence of dead surfaces that comprises the structure, and the moves made to get from each one to the next.

Our method is applied to the construction of periodic approximants in Section III. In the majority of the cases we will consider, we will be applying square periodic boundary conditions to our system in the x and y directions, building a square tower of cells. However, the square prism is not the only alternative. For instance, we have also investigated a system with a rhombic base, to which periodic boundary conditions can also be easily applied.

III. EXAMPLES OF CONSTRUCTION OF STACKING GRAPHS

Using the methods described in Section II B, we have constructed stacking graphs for a number of different systems. Here we give results for three systems with square base and square periodic boundary conditions.

A. Base of side τb

As a first example we take a system with square base of side τb, where the axes are directed along any three mutually perpendicular two-fold axes of our choice. This system is so small and so constrained that there turn out to be only three dead surfaces. Each is accessible in one move from every other, which is not the case for

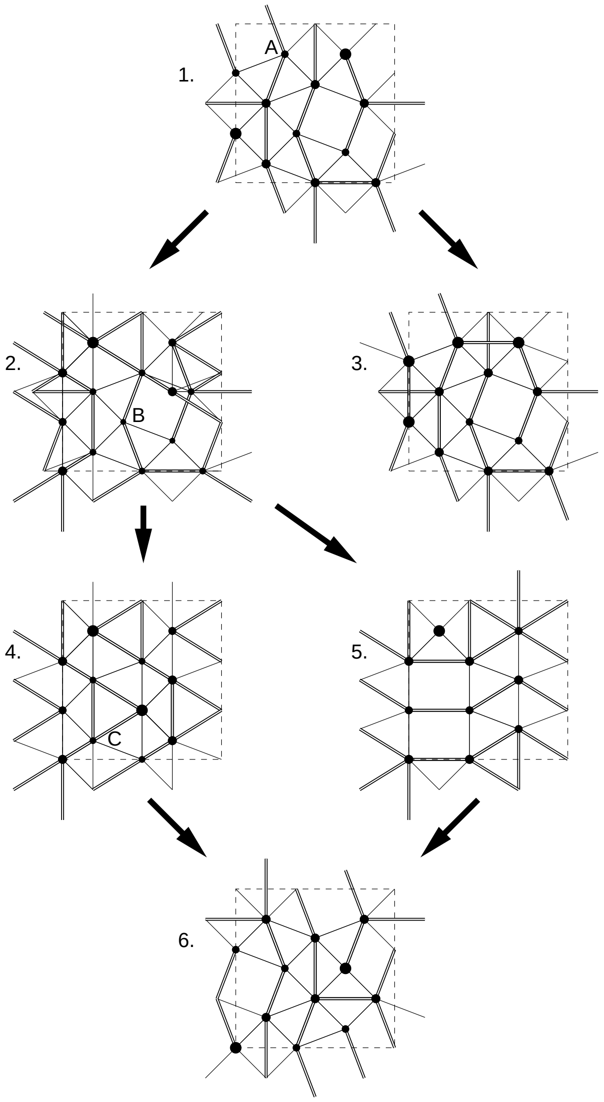

FIG. 3: A sequence of dead surfaces for the system with square base of length $\tau^2 b$ on each side, viewed from above. (See Section III B.) As in Fig. 1, the single and double lines represent the b- and c-type linkages respectively. The arrows indicate how each surface can be grown from a previous one by choosing a single canonical cell to add, and then completing the surface by adding all others which are forced.

larger less-constrained systems. The stacking graph for this system therefore looks like Fig. 2.

B. Base of side $\tau^2 b$

A less trivial example is afforded us by the system with square base $\tau^2 b$ on each side, which is the base of the '$\frac{3}{2}$' cubic approximant. Fig. 3 illustrates a part of the process of constructing the stacking graph for this system. Each part shows a dead surface viewed from above, so that the edges of the cells lying in the surface are visible as lines; some of these edges are continued outside the dotted square representing the square base, to illustrate how the surface wraps around on itself.

The first part of the figure shows a suitable starting surface, which we have labeled '1'. It turns out that the node marked 'A' has just two possibilities for completion which, given that this is a dead surface, is the minimum number possible, so we choose this node as the one we will grow from. In turn we try each of the two possibilities for the completion of node A, adding as well any linkages which are forced as a result of the additions which complete the node. The two different outcomes are illustrated in the next part of the figure, and labeled '2' and '3'. Although the area around node A has now been covered up, the node itself no longer being visible at all beneath the cells we have placed on top of it, there are other regions of the surface which are untouched from the previous stage. This is typical of the growth process in which a finite number of new cells are added between each dead surface and the next, even in the limit where the area of the surfaces becomes large.

Examination of surface 2 now reveals that the node labeled 'B', has the minimum possible number of completions (two). Constructing these two possibilities and filling in the cells forced by these moves, we produce another two new surfaces, which we number '4' and '5'. (The same is done for surface '3', but the results are not shown in the figure). Repeating the process yet again with the node labeled 'C' on surface 4, we arrive at the final surface shown in the figure, which we have labeled '6'. (Node C actually has two possibilities for completion, but one of them turns out to lead to an unlawful node environment, as discussed earlier, so we discard this possibility.) Careful inspection now reveals that surface 6 is just the same as surface 1, with a reflection in a plane which cuts the surface left-to-right. We have thus found a closed cycle in the stacking graph, and this amounts to finding our first periodic structure with this base, since repeating the cycle twice returns us to the original orientation of surface 1. The resulting structure is the known '$\frac{3}{2}$' cubic approximant structure given by Henley (1991).

Continuing this process to completion produces the complete stacking graph for this system, shown in Fig. 4. There are 17 dead surfaces in all. Clearly many distinct closed loops can be constructed on this diagram, corresponding to all sorts of periodic structures. Among them is one other, previously unknown cubic structure with the same unit cell as the cubic $\frac{3}{2}$ approximant, as well as many structures much larger in the stacking direction. We consider these structures further in the next section.

The $\frac{3}{2}$ system is the largest for which the stacking graph is still small enough to be constructed by hand using models, as well as on the computer. As a check on our algorithm and on the computer implementation, we have performed the calculation both ways, and found agreement.

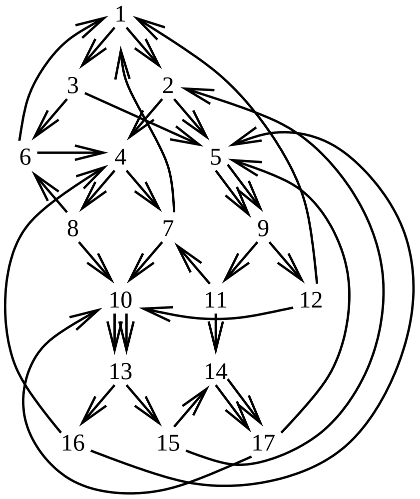

FIG. 4: The complete decision tree for the system with square base of length $\tau^2 b$ on each side. Moves from one surface to another which are indicated with double arrows are possible in two different ways. That is, one can get from the initial surface to the final one by adding two different sets of cells to the initial surface. It is important to keep track of these different possible transitions, as they give rise to different final structures.

C. Base of side $\tau^3 b$

Our most important new results are those for a system with a square base $\tau^3 b$ on each side. Since no structures of this size were known when we began this work (but see Mihalkovič and Mrafko (1993)), we had first to construct an initial dead surface to start the algorithm going. This we did by hand using models made of wood and plastic. We anticipated however, that there were going to be a very large number of dead surfaces for this system. Using Eq. 1 of Section II, we estimated the number of dead-surfaces to be 6900 ± 1500, and so we turned to a computational method to enumerate them exhaustively.

Our computer algorithm is a straightforward implementation of the method described in Section II. It entails taking a dead surface, choosing one node on that surface, and completing it in every possible way. Almost every completion leads us to another dead surface, which we check against all those previously generated, to see if we have seen it before. If not, we add it to our growing list of surfaces. The program starts with a list composed of just the one initial surface, and ends when every surface in the list has been investigated. For the system of square base $\tau^3 b$ we find a stacking graph composed of 7417 dead surfaces, which decomposes into three disjoint sections having 5327, 1296 and 794 surfaces respectively.

IV. PERIODIC STRUCTURES

The first use to which we have put our stacking graphs, and the one which is the main concern of this paper, is the construction of periodic structures. Periodic structures are useful technically as the starting point for constructing atomic models using decoration rules, or as the initial conditions for Monte Carlo simulations. In addition, in the present instance, where we have not yet discovered a rule for the quasiperiodic structure, the approximants could offer a clue to that structure.

Different structures are represented by different paths around the stacking graph. All such structures obey the periodic boundary conditions in the x and y directions, and they will be periodic in the z direction too if they begin and end with the same surface. Thus closed loops in the graph correspond to three-dimensional space-filling periodic structures. These structures are not guaranteed to be cubic since the vertical distance between the initial and final surfaces need not be the same as the base length, and even if it is, the final surface may not be vertically above the initial one, but may be translated. However, it is a simple matter to check whether this is the case, and thereby extract cubic or any other structures from the stacking graph. Knowing no better technique, we have simply reconstructed *every* possible structure in turn, by enumerating all possible closed loops of dead surfaces from our stacking graph, and then checked each one to see if it satisfies cubic or other periodic boundary conditions. The time taken to find all possible approximants with a given number of dead surfaces increases exponentially with the number of surfaces.[56]

A. Structures with base of side τb

The method was first applied to the system with base τb on a side, demonstrating that the known eight-node $\frac{2}{1}$ approximant is the only cubic structure with this base. This was previously known (Henley 1991). However, a whole new class of tetragonal periodic structures also comes out of the analysis, corresponding to sequences of dead surfaces of the form

$$1 \to 2 \to 1 \to 3 \to 1 \ldots$$
$$1 \to 2 \to 1 \to 2 \to 1 \to 3 \to 1 \ldots$$
$$1 \to 2 \to 1 \to 2 \to 1 \to 2 \to 1 \to 3 \to 1 \ldots$$

and so forth. There are several other classes of periodic structures generated in a similar fashion. Also there are oblate structures corresponding to sequences such as $1 \to 2 \to 1 \ldots$, which repeat every b units in the z direction.

B. Structures with base of side $\tau^2 b$

For the system with base $\tau^2 b \times \tau^2 b$ we have, as mentioned above, found one new cubic periodic structure in addition to the '$\frac{3}{2}$ approximant' described by Henley (1991). All three cubic structures contain 32 nodes. An infinite number of tetragonal structures can also be built in a manner similar to the system of base τb. In addition we see for the first time in this system, periodic structures in which surfaces repeat with a sideways translation so that they are not vertically above the previous occurrence of the same surface. This gives rise to a large number of structures which may be pseudo-tetragonal, monoclinic or orthorhombic.

C. Structures with base of side $\tau^3 b$

For the system with base $\tau^3 b \times \tau^3 b$, the list of 7417 surfaces that we find makes a large stacking graph. There are many cycles of dead-surfaces in it, and almost all of them do not correspond to cubic structures. However, again with the aid of the computer, we have searched for loops of up to 50 dead surfaces and reconstructed the structures to find their shape. In this way we have found 316 new cubic periodic structures of side $\tau^3 b$. They have varying degrees of symmetry, the highest symmetry group being the cubic subgroup $\mathrm{Pa}2_1 3$. There are 32 different structures which are inequivalent by symmetry. Some properties of the periodic structures are summarized in Table I. The "perp space variance" column indicates the variance of the magnitudes of the perp space positions of the nodes of each structure. To derive these figures, all factors of τ in the projection of the structure into perp space have been replaced by the 'Fibonacci ratios' 1/1, 2/1, 3/2, and 5/3 for the respective approximants. This ensures that the perp-space displacement corresponding to one period of the periodic boundary conditions is zero. The variance figures are a useful guide for separating structures which are related by symmetry from those which are genuinely different from one another. The column marked "multiplicity" gives the number of elements in the point group for each structure.

We believe this list of periodic structures to be an exhaustive list of the periodic structures of this size. This means, for instance, that if there exists in this list a structure which is related to one of the smaller cubic structures by some inflation rule—a $\frac{5}{3}$ approximant—then that structure must be present in our list. Work is under way to try and find a set of inflation rules by comparing small and large approximants in this way.

The number of nodes in our $\frac{5}{3}$ cubic structures ranges from 134 to 138, with the packing fraction varying between 0.600 and 0.617. Amongst these structures we find the two generated by Mihalkovič and Mrafko (1993) using Monte Carlo methods.

D. Discussion and possible extensions

It is natural to ask whether we could extend this calculation to a larger system. The next step would be to try a system of base $\tau^4 b$ on a side. However, this is not likely to be tractable—using Eq. 1 of Section II we estimate that the stacking graph would consist of about a million dead surfaces, putting the calculation well beyond our present computational resources. A more reasonable thing to try would be a square system $2\tau^2 b$ on each side, i.e., four times the area of the base in the $\frac{3}{2}$ system described above. We have not yet explored this possibility.

One other calculation we have performed is that for a system with a 60° rhombic base with sides of length $\tau^3 b$. This system grows along an axis of three-fold symmetry of the reference icosahedron, rather than a two-fold axis as in the case of the square boundary conditions. Despite the large area of its base, this appears to be a very constrained system with only 7 different dead surfaces. If Eq. 1 still holds, clearly the values of at least one of the constants N_0, λ must be different for systems which grow along a three-fold direction in this way. We were interested to see if we could find any periodic structures in the shape of a rhombic prism. None were found, though 3 other periodic structures were found.

V. DIFFRACTION PATTERNS

We have used our set of large $\frac{5}{3}$ periodic structures to investigate the diffraction pattern of a random tiling made of canonical cells. The canonical-cell model was originally suggested as a random-tiling model, in which many tiling configurations are possible with internal energies differing by amounts not greatly exceeding kT. For such a system, the entropy associated with the many configurations could make a significant contribution to the stability of the icosahedral structure. It has been suggested that this is the mechanism responsible for the stabilization of the *i*-phase in alloys such as AlCuFe and AlPdMn. It is interesting to investigate the effect of the tiling disorder on the diffraction pattern; comparison of the results with real scattering experiments may give us a clue as to whether real quasicrystals are also random tilings.

In this paper we perform only the simplest diffraction calculation, in which the structure factor is evaluated for a system of point scatterers centered on the node positions in the lattice. In order to turn this into a realistic simulation of scattering from an alloy, we would have to convolve the results with the scattering amplitude of a single icosahedral cluster of atoms, plus a contribution from the glue atoms. We should also take account of the fact that we have real atoms by including form factors and Debye-Waller factors for the atoms we wish to simulate. We are at present working on such a calculation which is to be published later. However, earlier studies on icosahedral glasses (Robertson and Moss 1991, 1993, Gibbons and Kelton 1993) have shown that the pattern of diffraction from the bare nodes contains important information on the degree of long-range order and already exhibits the striking features, such as arcs of diffuse scattering, which are characteristic of the short-range order at the scale of the separation of the nodes.

To simulate the effect of the disorder on the diffraction

Fibonacci ratio	number of nodes	packing fraction	volume of primitive unit cell	perp-space variance	Bravais lattice	space group	multiplicity
1/1	2	0.680	1/2	0	bcc	$Im\bar{3}$	1
2/1	8	0.642	τ^3	3/20	cubic	$Pa\bar{3}$	1
3/2	32	0.606	τ^6	9/52	cubic	$Pa\bar{3}$	1
3/2	32	0.606	τ^6	21/104	rhomb	$R\bar{3}$	4
5/3	134	0.600	$\tau^9/2$	246/1139	rhomb	$R\bar{3}$	4
5/3	135	0.604	τ^9	874/3825	rhomb	$R3$	8
5/3	135	0.604	τ^9	1114/3825	rhomb	$R3$	8
5/3	136	0.608	τ^9	249/1156	rhomb	$R\bar{3}$	4
5/3	136	0.608	τ^9	68495/314432	tricl	$P1$	24
5/3 †	136	0.608	τ^9	267/1156	rhomb	$R\bar{3}$	4
5/3 †	136	0.608	τ^9	267/1156	rhomb	$R\bar{3}$	4
5/3	136	0.608	τ^9	539/2312	ortho	$P2_12_12_1$	6
5/3	136	0.608	τ^9	273/1156	rhomb	$R\bar{3}$	4
5/3	136	0.608	τ^9	68495/314432	tricl	$P\bar{1}$	12
5/3	136	0.608	τ^9	75/289	rhomb	$R\bar{3}$	4
5/3	136	0.608	τ^9	303/1156	rhomb	$R\bar{3}$	4
5/3	136	0.608	τ^9	41631/157216	tricl	$P1$	24
5/3	136	0.608	τ^9	165/578	rhomb	$R\bar{3}$	4
5/3	136	0.608	τ^9	11283/39304	rhomb	$R3$	8
5/3	136	0.608	τ^9	90303/314432	tricl	$P1$	24
5/3	136	0.608	τ^9	339/1156	rhomb	$R\bar{3}$	4
5/3	136	0.608	τ^9	23067/78608	tricl	$P1$	24
5/3	136	0.608	τ^9	46239/157216	tricl	$P1$	24
5/3	136	0.608	τ^9	1455/4913	rhomb	$R3$	8
5/3	136	0.608	τ^9	183/578	rhomb	$R3$	8
5/3	136	0.608	τ^9	25371/7860	tricl	$P1$	24
5/3	136	0.608	τ^9	759/2312	cubic	$P2_13$	2
5/3	136	0.608	τ^9	105/289	rhomb	$R\bar{3}$	4
5/3	137	0.613	τ^9	90084/319073	rhomb	$R3$	8
5/3	137	0.613	τ^9	96999/319073	rhomb	$R3$	8
5/3	138	0.617	$\tau^9/2$	86/391	rhomb	$R\bar{3}$	4
5/3	138	0.617	τ^9	195/782	mono	Pb	12
5/3	138	0.617	τ^9	44957/161874	mono	Pb	12
5/3	138	0.617	τ^9	109/391	mono	$P2_1$	12
5/3	138	0.617	τ^9	24859/80937	mono	$P2_1$	12
5/3	138	0.617	τ^9	241/782	rhomb	$R\bar{3}$	4

TABLE I: The cubic canonical-cell approximants, up to the $\frac{5}{3}$ size. In the Bravais lattice column, "rhomb" stands for "rhombohedral", "ortho" stands for "orthorhombic", "tricl" stands for "triclinic", and "mono" stands for "monoclinic". The two rows marked with a † have identical entries, but they correspond to different approximants. The second and fourth perp-space moments of these two approximants are identical, but their sixth perp-space moments are different.

pattern, we take all the structures with a given unit cell and calculate, one by one, the diffraction pattern from each structure. Then we average the intensities of all these patterns, to get a mean result for the entire ensemble. In the limit as the size of the system becomes large, the result will tend to the correct one for the infinite system. In calculating the average over the many structures of one size, we should weight each structure according to how likely it is to appear, which in turn will depend on its Boltzmann weight in the ensemble. In our calculations, we have made the high-temperature assumption that the energy of interaction of the canonical cells in our structures is negligible by comparison with kT, so that all the Boltzmann weights are close to unity, and can be treated as equal.

We cannot in reality take the limit as the size of our system diverges. The best we can do is to approximate the infinite system with the largest of the finite ones we have investigated—the $\frac{5}{3}$ cubic system. For this finite system there is a sparse reciprocal lattice of possible scattering vectors, by contrast with the infinite system, which would have a dense set of possible vectors, as is typical of quasiperiodic structures. Since the periodic structures we have found satisfy cubic periodic boundary conditions, this sparse reciprocal lattice is also cubic, with one reciprocal lattice vector per $(2\pi/\tau^3 b)^3$ of k-space.

Fig. 5 shows the structure factor in a plane perpendicular to one of the axes of the cube (i.e., an icosahedral pseudo-two-fold axis), for one particular $\frac{5}{3}$ periodic structure taken from the 316 that were found. This picture was produced by calculating the strengths of the δ-function peaks on each reciprocal lattice point in scat-

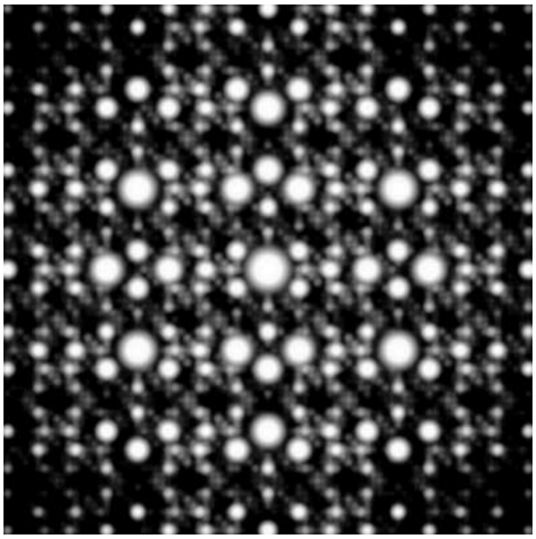

FIG. 5: A logarithmic intensity plot of the structure factor of a single cubic periodic structure composed of canonical cells, having a period in each direction of $\tau^3 b$. The picture is a slice though the reciprocal lattice along a plane perpendicular one of the cubic axes. Notice that for this single structure, there are asymmetries in the pattern about the central peak.

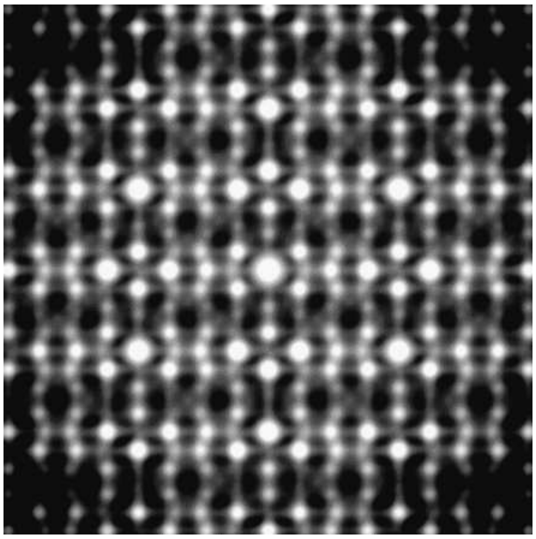

FIG. 6: A logarithmic intensity plot of the structure factor averaged over all the cubic periodic structures of side $\tau^3 b$. The areas marked '1' and '2' correspond to the arcs of diffuse scattering seen in Fig. 7.

tering from point scatterers at the nodes of the infinite lattice formed by the periodic repetition of the structure in question over all space. The δ-functions were then convolved with a Gaussian to broaden the peaks so that they show up. In Fig. 6 we have repeated the calculation for each of the 316 periodic structures of this size that were found, and then averaged the intensities. Notice that while Fig. 5 has clear asymmetries about the central peak, Fig. 6 has none, and displays the complete symmetries of the icosahedral group within this plane.

Two points are worth mentioning about Fig. 6. First, the diffraction pattern is very similar to published diffraction patterns from icosahedral glass models (Robertson and Moss 1991, 1993, Gibbons and Kelton 1993). These glassy models show statistical noise in their scattering patterns which is not present in our results, but otherwise the comparison is good. The similarity in the diffraction patterns is not especially surprising since it is known that the best icosahedral-glasses (Robertson and Moss 1991, 1993) consist almost entirely of packings of canonical cells.

Second, there are dark areas clearly visible in Fig. 6, such as those labeled '1' and '2', in which the intensity of scattering is very low. They are surrounded by a haze, or diffuse background, produced by averaging over a number of bright peaks appearing in different places in different periodic structures. The experimental diffraction patterns of a number of different icosahedral phases, particularly titanium alloys, exhibit similar dark areas and bright arcs at the same places in reciprocal space. A particularly clear photograph of these arcs, observed in experiments on TiMnSi by Viano *et al.* (1993), is reproduced in Fig. 7 for comparison. It seems possible then that, as suggested by Robertson and Moss (1991, 1993), and also by Gibbons and Kelton (1993), these diffuse features are a result of the random tiling nature of the alloys' structure. In respect of these features our tile-based model seems to be at least as successful as the glassy growth models suggested by these authors, and indicates that the diffuse arcs are not related to the large phason fluctuations found in icosahedral glasses, since such fluctuations are not also found in our random tilings. Rather, the arcs must be related to correlations at the intermediate scale of one or two tile edges, a scale at which the distribution of local environments is nearly the same in the equilibrium random tiling and the growth models.

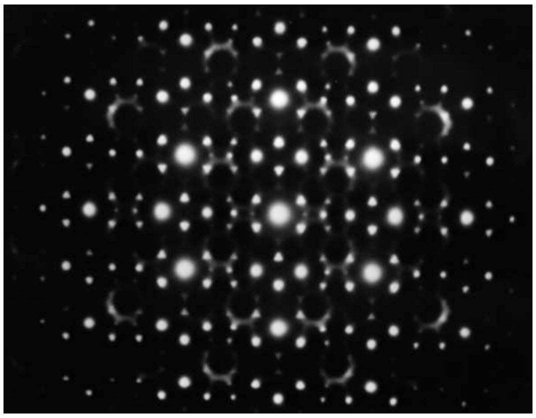

FIG. 7: An x-ray diffraction pattern for TiMnSi taken close to one of the two-fold axes of the icosahedron. The arcs marked '1' and '2' correspond to the similarly marked features in Fig. 6. Reproduced from Viano *et al.* (1993) by permission.

VI. CONCLUSIONS AND DISCUSSION

In this paper we have presented a number of new results concerning the canonical-cell model of an icosahedral quasicrystal. First, we have constructed all possible periodic structures made of canonical cells which fit within three different cubic unit cells. In the case of the so-called '$\frac{5}{3}$' unit cell, for which no periodic structures were known when we began this study, we have found 316 ways of tiling the cube of side $\tau^3 b$ (32 ways, if we count only those which are inequivalent by rotational symmetry). The method involves construction of a 'stacking graph' which in principle provides a description of all possible tilings. The stacking graph has other uses in the study of phason fluctuations and bulk properties of the random tiling. We hope to present our work on these aspects of the problem in a later paper.

Second, as described in Sec. V, we have used the ensemble of $\frac{5}{3}$ periodic structures to compute diffraction patterns (see Figs. 5 and 6), which exhibit the diffuse arcs similar to those appearing in real specimens of certain icosahedral alloys (see Fig. 7). As pointed out in Sec. V, we have calculated only the pattern of point-like scattering from the nodes in the canonical-cell tiling; we have not decorated our periodic structures with actual atoms. This demonstrates that diffuse arcs can arise from correlations in the configurational disorder at the tile or cluster packing level, and do not require the presence of correlations at the atomic level (Robertson and Moss 1991, 1993, Gibbons and Kelton 1993). The arcs are clearest (Fig. 6) when the diffraction patterns are averaged over the ensemble of tilings. This averaging suppresses the relative magnitude of the 'speckle noise' which is inherent to Fourier transforms of any single finite sample of random tiling (Garg and Levine 1988), and which after averaging becomes the uniform diffuse background.

A. Discussion

The approach used here to find periodic structures made of canonical cells is applicable to *any* tiling model, in any number of dimensions. The procedure for generating the stacking graph is simpler (and more reliable) than an alternative approach used in older transfer-matrix calculations (Widom *et al.* 1989, Shaw and Henley 1991, Li *et al.* 1992), where an entire layer of cells is added at a time. Unfortunately, the size of the stacking graph grows exponentially with the cross-sectional area of our system, and there is, consequently, a rather strict limit on the size of the system we can study in this way. In the case of the canonical-cell tiling, the $\frac{5}{3}$ system that we have studied here is probably the largest cubic system that can be tackled using present-day computers.

The list of periodic structures is useful in itself. In the absence of a Monte Carlo algorithm for sampling the random-tiling ensemble, our enumeration (see the table) of the possible arrangements of tiles in a relatively small system is a useful substitute.

1. The range of packing fractions seen in the table is probably a good indication of the range of values taken by the packing fraction in infinite tilings with overall icosahedral symmetry. Our exhaustive enumeration is much better for estimating this range than Monte Carlo approaches (Mihalkovič and Mrafko 1993), which succeed for the larger packing fractions but tend to converge on illegal packings when the packing fraction is in the lower part of the allowed range.

2. Presumably one or two of the structures from our list in the table are the finite approximants of an (as yet hypothetical) ideal *quasiperiodic* canonical-cell packing. Study of these structures may provide clues about how to build the latter. We can identify candidates to be such a special approximant by (i) high symmetry (ii) special ('magic') values of the packing fraction (Henley 1991) or (iii) small size of phason fluctuations.

3. We can also use our periodic structures to create realistic atomic structure models for icosahedral alloys by decorating the tilings (Henley, Mihalkovič, Newman and Oxborrow, to be published). These can then be relaxed under pair potentials as tests of their stability, or their diffraction patterns can be calculated in order to fit the decorations to experimental data.

Apart from the construction of periodic structures, our stacking graph has other applications.

1. The algorithm of Sec. II B, which adds forced tiles to a surface until a new dead surface is reached, could be used unchanged in *growth* models, similar to icosahedral glass models (Robertson and Moss 1991, 1993), by simply choosing at random one way to complete the environment a node, rather than exploring all options for completion. In the language of the stacking graph, one would be generating a random path through the graph rather than accumulating a complete description of it. Since, 'dead ends' (see Section II B) are quite rare in the stacking graph, it should be easy to grow large structures.

2. With slight modification, so as to accommodate chemical potentials that favor nonzero phason strain, the stacking graph becomes a 'transfer matrix' describing an ensemble of random tilings which are infinite in the stacking direction. This transfer matrix may be used for numerically exact calculation of 'thermodynamic' quantities (entropy density, phason elastic constants) for ensembles of tilings with a given finite cross-section. We have carried out such a calculation, and extracted the

first estimates of these quantities for the canonical-cell random tiling; the results will appear in a subsequent paper.

3. By biasing the growth algorithm so as to limit phason fluctuations, or to favor nodes at vertices of an inflated rhombohedral or canonical-cell tiling, it might conceivably be possible to grow periodic structures on a larger base.

VII. ACKNOWLEDGMENTS

We would like to thank Pat Gibbons for his detailed criticism of the manuscript of this article, Mike Widom for helpful discussions concerning the use of the transfer matrix method in quasicrystal structure calculations, Jason Ho for his donation of some model parts, and Pat Gibbons and Ann Viano for supplying us with Fig. 7. This work was supported in part by the Science and Engineering Research Council of Great Britain, and by the DOE under grant number DE–FG02–89ER–45405.

---

[1] Bak, P., 1986, Scr. Metall. **20**, 1199.
[2] Bancel, P. A., 1991, in *Quasicrystals: the State of the Art,* P. J. Steinhardt and D. P. DiVincenzo (eds.), World Scientific, Singapore.
[3] Boudard, M., de Boissieu, M., Janot, C., Heger, G., Beeli, C., Nissen, H.-U., Vincent, H., Ibberson, R., Audier, M., and Dubois, J. M., 1992, J. Phys. Condens. Matter **4**, 10149.
[4] Duneau, M., and Katz, A., 1985, Phys. Rev. Lett. **54**, 2688.
[5] Elser, V., 1986, Acta. Cryst. A **42**, 36.
[6] Elser, V., 1987, in *Proceedings of the XVth International Colloquium on Group Theoretical Methods in Physics,* R. Gilmore (ed.), World Scientific, Singapore.
[7] Elser, V., 1989, in *Aperiodicity and order*, vol. 3, M. V. Jarić and D. Gratias (eds.), Academic Press, London.
[8] Elser, V., and Henley, C. L., 1986, Phys. Rev. Lett. **55**, 2883.
[9] Gähler, F., and Rhyner, J., 1986, J. Phys. A **19**, 267.
[10] Gardner, M., 1977, Sci. Am. **236**, No. 1, 110.
[11] Garg, A., and Levine, D., 1988, Phys. Rev. Lett. **60**, 2160.
[12] Gibbons, P. C., and Kelton, K. F., 1993, J. Non-cryst. Solids **153**, 165.
[13] Guyot, P., and Audier, M., 1985, Phil. Mag. B **52**, L15.
[14] Henley, C. L., 1991, in *Quasicrystals: the State of the Art,* P. J. Steinhardt and D. P. DiVincenzo (eds.), World Scientific, Singapore.
[15] Henley, C. L., and Elser, V., 1986, Phil. Mag. B **53**, L59.
[16] Henley, C. L., 1986, Phys. Rev. B **34**, 797.
[17] Henley, C. L., 1991, Phys. Rev. B **43**, 993.
[18] Ingersent, K., 1991, in *Quasicrystals: the State of the Art,* P. J. Steinhardt and D. P. DiVincenzo (eds.), World Scientific, Singapore.
[19] Jarić, M. V., and Nelson, D. R., 1988, Phys. Rev. B **37**, 4458.
[20] Kalugin, P. A., Kitaev, A. Yu., and Levitov, L. S., 1985, JETP Lett. **41**, 145.
[21] Katz, A., 1988, Comm. Math. Phys. **118**, 263.
[22] Lançon, F., Billard, L., and Chaudhari, P., 1986, Europhys. Lett. **2**, 625.
[23] Lei, T., and Henley, C. L., 1991, Phil. Mag. B **63**, 677.
[24] Levine, D., and Steinhardt, P. J., 1984, Phys. Rev. Lett. **53**, 2477.
[25] Levitov, L. S., 1988, Comm. Math. Phys. **119**, 627.
[26] Li, W., Park, H., and Widom, M., 1992, J. Stat. Phys. **66**, 1.
[27] Mihalkovič, M., and Mrafko, P., 1993, Europhys. Lett. **21**, 463.
[28] Newman, M. E. J., and Henley, C. L., 1993, J. Non-cryst. Solids **153**, 205.
[29] Nightingale, M. P., 1990, in *Finite-Size Scaling and Numerical Simulation of Statistical Systems,* V. Privman (ed.), World Scientific, Singapore.
[30] Onoda, G. Y., Steinhardt, P. J., DiVincenzo, D. P. and Socolar, J. E. S., 1988, Phys. Rev. Lett. **60**, 2653.
[31] Oxborrow, M., and Henley, C. L., 1993, J. Non-Cryst. Solids **153**, 210.
[32] Penrose, R., 1974, Bull. Inst. Math. and its Appl. **10**, 2660.
[33] Robertson, J. L., and Moss, S. C., 1991, Z. Phys B **83**, 391.
[34] Robertson, J. L., and Moss, S. C., 1993, J. Non-Cryst. Solids **153**, 77.
[35] Shaw, L. J., Elser, V., and Henley, C. L., 1991, Phys. Rev. B **43**, 3423.
[36] Shaw, L. J., and Henley, C. L., 1991, J. Phys. A **24**, 4129.
[37] Smith, A. P., 1993, J. Non-Cryst. Solids **153**, 258.
[38] Socolar, J. E. S., 1989, in *Quasicrystals* M. V. Jarić and S. Lundqvist (eds.), World Scientific, Singapore.
[39] Socolar, J. E. S., 1991, in *Quasicrystals: the State of the Art,* P. J. Steinhardt and D. P. DiVincenzo (eds.), World Scientific, Singapore.
[40] Socolar, J. E. S., Steinhardt, P. J., and Levine, D., 1985, Phys. Rev. B **32**, 5547.
[41] Stephens, P. W., 1989, in sl Extended Icosahedral Structures, M. V. Jarić and D. Gratias (eds.), Academic Press, New York.
[42] Stephens, P. W., and Goldman, A. I., 1986, Phys. Rev. Lett. **56**, 1168.
[43] Strandburg, K. J., 1991, Phys. Rev. B **44**, 4644.
[44] Strandburg, K. J., Tang, L.-H.,, and Jarić, M. V., 1989, Phys. Rev. Lett. **63**, 314.
[45] Viano, A. M., Gibbons, P. C., and Kelton, K. F., 1993, Bull. Am. Phys. Soc. **38**, 681.
[46] Widom, M., 1991, Philos. Mag. Lett. **64**, 297.
[47] Widom, M., 1993, Phys. Rev. Lett. **70**, 2094.

[48] Widom, M., Deng, D. P., and Henley, C. L., 1989, Phys. Rev. Lett. **63**, 310.

[49] Widom, M., Strandburg, K. J. and Swendsen, R. H., 1987, Phys. Rev. Lett. **58**, 706.

[50] Descriptions of these environments and other technical information about the model has been deposited with the American Institute of Physics Auxiliary Publications Service. See Henley (1991) for further details.

[51] The demonstration is very straight-forward, since the inclusion of three types of linkage forces a complete set of cells filling the 4π of solid angle around each different node environment. Clearly if every node is surrounded by correct canonical cells, then the entire structure is composed of canonical cells.

[52] It was noted in an earlier paper (Lei and Henley 1991) that the canonical cell tiling can be cut most neatly into slabs transverse to a 5-fold axis. This however is the direction least favorable to periodicity in the transverse directions, so this arrangement is not especially sensible in the present case.

[53] There is a slight subtlety which arises here. It occasionally happens that we make some legitimate addition to a dead surface, but that the extra tiles we are forced to add as a result eventually lead us to construct a disallowed node environment somewhere else on the lattice—an environment in which some part of the solid angle surrounding the node cannot be plugged with any combination of canonical cells. In this case the surface which we started with, which we labeled as 'dead' because it had two or more possible completions for each node, may not have been dead after all. There may only have been one way to continue building on that surface so that we end up with a true canonical-cell structure. The surface was, if you like, only 'locally dead'. However, this local deadness does not affect the answers we get from our algorithm, and it is much easier to spot than the longer-range true deadness, so we will continue to call surfaces dead when they are merely locally dead, and the algorithm we use will be exactly as described in the text.

[54] However, it may not be immediately clear why all possible structures can be described by considering the choices to be made at one particular node on each dead surface. One might argue that if we had chosen some other node on the surface and completed that, we might have been led to some completely different structure which we would not otherwise have discovered. This however cannot be the case. If we choose just one node to complete in the surface and build from there, we will eventually, quite possibly after passing through a number of other dead surfaces, arrive at this other node that we are worried about, and we will then make a choice about how to complete it. The possible completions will not have changed just because we have added some cells elsewhere in the lattice, and the only difference between completing the node now and completing it several moves ago, will be the order in which new nodes and linkages are added to the lattice. The possible final structures will be the same in either case.

[55] In the case of a disjoint stacking graph we require separate starting surfaces for each separate piece of the graph. In the calculation described here, we found starting surfaces for the smaller subgraphs by observing that the list of structures generated by the largest subgraph was clearly incomplete, lacking some members which were related to others appearing in the list by symmetry operations of the superlattice. By generating some of these missing members using the symmetry operations and then taking them apart to extract a dead surface, we were able to produce a suitable starting surface for each subgraph. The fact that the subgraphs found in this manner produce no more structures which lack symmetric counterparts indicates to us that in doing this we have found all the parts of the $\frac{5}{3}$ stacking graph.

[56] We should point out that this method doesn't actually *guarantee* that we will find all the cubic periodic structures of the desired size, since it would be possible for such structures to be represented in the stacking graph by two or more periods of the same structure identified with different sequences of dead surfaces. However, we believe that, because we complete the lowest-lying nodes in the structure first (which makes the dead surfaces flatter), our algorithm will eventually be driven towards the representation of a periodic structure by unique repeating sequence of dead surfaces, although the first few surfaces may be different owing to some odd choice of starting surface. Two other pieces of evidence support this conclusion: (i) We have constructed the stacking graph for the $\frac{5}{3}$ system in a slightly different way by changing the rule for choosing which node to complete. This produces a completely different graph with different dead surfaces, but still we find the exact same set of periodic structures. (ii) If a structure is a valid periodic packing of canonical cells, then we can generate up to 24 additional valid packings by acting on that structure with the 24 operations of the icosahedral group that preserve the cubic Bravais lattice. If a small number of structures were missing from our enumeration we would still see their symmetry-related counterparts appearing, and this would indicate that something was wrong. Nothing of that kind happens here however.