

TOPOLOGICAL CRYSTALS

JOHN C. BAEZ


ABSTRACT. Sunada’s work on topological crystallography emphasizes the role of the ‘maximal abelian cover’ of a graph X . This is a covering space of X for which the group of deck transformations is the first homology group $H_1(X, \mathbb{Z})$. An embedding of the maximal abelian cover in a vector space can serve as the pattern for a crystal: atoms are located at the vertices, while bonds lie on the edges. We prove that for any connected graph X without bridges, there is a canonical way to embed the maximal abelian cover of X into the vector space $H_1(X, \mathbb{R})$. We call this a ‘topological crystal’. Crystals of graphene and diamond are examples of this construction. We prove that any symmetry of a graph lifts to a symmetry of its topological crystal. We also compute the density of atoms in this topological crystal. We give special attention to the topological crystals coming from Platonic solids. The key technical tools are a way of decomposing the 1-chain coming from a path in X into manageable pieces, and the work of Bacher, de la Harpe and Nagnibeda on integral cycles and integral cuts.

1. INTRODUCTION


The ‘maximal abelian cover’ of a graph plays a key role in Sunada’s work on topological crystallography [12]. Just as the universal cover of a connected graph X has the fundamental group $\pi_1(X)$ as its group of deck transformations, the maximal abelian cover, denoted \overline{X} , has the abelianization of $\pi_1(X)$ as its group of deck transformations. It thus covers every other connected cover of X whose group of deck transformations is abelian. Since the abelianization of $\pi_1(X)$ is the first homology group $H_1(X, \mathbb{Z})$, there is a close connection between the maximal abelian cover and homology theory.

In this paper we prove that for a large class of graphs, the maximal abelian cover can naturally be embedded in the vector space $H_1(X, \mathbb{R})$. We call this embedded copy of \overline{X} a ‘topological crystal’. The symmetries of the original graph can be lifted to symmetries of its topological crystal, but the topological crystal also has an n -dimensional lattice of translational symmetries. In the 3-dimensional case, the topological crystal can serve as the blueprint for an actual crystal, with atoms at the vertices and bonds along the edges.

The most famous example arises from the tetrahedron. Here X is the complete graph on 4 vertices:




Since $H_1(X, \mathbb{R})$ is 3-dimensional, the corresponding topological crystal is a graph embedded in 3-dimensional Euclidean space. It looks like this:




Up to translation this crystal has 4 kinds of atoms, corresponding to the 4 vertices of the tetrahedron. Each atom has 3 equally distant neighbors lying in a plane at 120° angles from each other. These planes lie in 4 families, each parallel to one face of a regular tetrahedron. This structure was discovered by the crystallographer Laves [9] and dubbed the 'Laves graph' by Coxeter [4]. Later Sunada called it the ' K_4 lattice', studied its energy minimization properties, and raised the question of whether it could serve as the pattern for a crystal form of carbon [11]. This form of carbon has not yet been seen, but the Laves graph plays a role in the structure of certain butterfly wings [7].

The Laves graph is exciting because it was studied mathematically before being found in nature. A more familiar example arises from this graph:



The corresponding topological crystal provides the pattern for a diamond:



Up to translation this crystal has two kinds of atoms, corresponding to the two vertices of the original graph.

The general construction of topological crystals was developed by Kotani and Sunada [8], who used the term ‘standard realization’, and later by Eon [6], who used the term ‘archetypical representation’. Sunada uses ‘topological crystal’ for an even more general concept [12], but we only need the following special case.

We start with a graph X . This has a space $C_0(X, \mathbb{R})$ of 0-chains, which are formal linear combinations of vertices, and a space $C_1(X, \mathbb{R})$ of 1-chains, which are formal linear combinations of edges. There is a boundary operator

$$\partial: C_1(X, \mathbb{R}) \rightarrow C_0(X, \mathbb{R}),$$

the linear operator sending any edge to the formal difference of its two endpoints. The kernel of this operator is the space of 1-cycles, $Z_1(X, \mathbb{R})$. There is an inner product on the space of 1-chains such that edges form an orthonormal basis. This determines an orthogonal projection


$$\pi: C_1(X, \mathbb{R}) \rightarrow Z_1(X, \mathbb{R}).$$

For a graph, $Z_1(X, \mathbb{R})$ is isomorphic to the first homology group $H_1(X, \mathbb{R})$. So, to obtain the topological crystal of X , we need only embed its maximal abelian cover \overline{X} in $Z_1(X, \mathbb{R})$. We do this by embedding \overline{X} in $C_1(X, \mathbb{R})$ and then projecting it down via π .


To accomplish this, we need to fix a basepoint for X . Each path γ in X starting at this basepoint determines a 1-chain c_γ . It is easy to show that these 1-chains correspond to the vertices of \overline{X} . Furthermore, the graph \overline{X} has an edge from c_γ to $c_{\gamma'}$ whenever the path γ' is obtained by adding an extra edge to γ . We can think of this edge as a straight line segment from c_γ to $c_{\gamma'}$.

The hard part is checking that the projection π maps this copy of \overline{X} into $Z_1(X, \mathbb{R})$ in a one-to-one manner. In Theorem 7 we prove that this happens precisely when the graph X has no ‘bridges’: that is, edges whose removal would disconnect X . Kotani and Sunada [8] noted that this condition is necessary; with considerably more work we show it is also sufficient. The main technical tool is Lemma 5, which for any path γ decomposes the 1-chain c_γ into manageable pieces.


We call the resulting copy of \overline{X} embedded in $Z_1(X, \mathbb{R})$ a ‘topological crystal’. For example, let X be this graph:



Since X has 3 edges, the space of 1-chains is 3-dimensional. Since $H_1(X, \mathbb{R})$ is 2-dimensional, the space of 1-cycles is a plane in this 3-dimensional space. If we consider paths γ in X starting at the red vertex, form the 1-chains c_γ , and project them down to this plane, we obtain the following picture:



Here the 1-chains c_γ are the white and red dots. These are the vertices of \overline{X} , while the line segments between them are the edges of \overline{X} . Projecting these vertices and edges onto the plane of 1-cycles, we obtain the topological crystal for X . The blue dots come from projecting the white dots onto the plane of 1-cycles, while the red dots already lie on this plane. The resulting topological crystal provides the pattern for graphene, a 2-dimensional form of carbon:



For any connected graph X , there is a covering map

$$q: \overline{X} \rightarrow X.$$

The vertices of \overline{X} come in different kinds, or ‘colors’, depending on which vertex of X they map to. We study the group of ‘covering symmetries’, $\text{Cov}(X)$, consisting of all symmetries of \overline{X} that map vertices of the same color to vertices of the same color. In Lemma 8 we prove that this group fits into a short exact sequence

$$1 \longrightarrow H_1(X, \mathbb{Z}) \longrightarrow \text{Cov}(X) \longrightarrow \text{Aut}(X) \longrightarrow 1$$

where $\text{Aut}(X)$ is the group of all symmetries of X . Thus, every symmetry of X is covered by some symmetry of \overline{X} , while $H_1(X, \mathbb{Z})$ acts on \overline{X} in a way that preserves the color of every vertex. In Theorem 11 we prove that $\text{Cov}(X)$ also acts as symmetries of the topological crystal associated to X . More precisely, there is an action of $\text{Cov}(X)$ as affine isometries of $H_1(X, \mathbb{R})$ for which the embedding of \overline{X} in this space is equivariant. In this action, the subgroup $H_1(X, \mathbb{Z})$ acts as translations.

These results are most interesting in cases where X is highly symmetrical. For example, suppose that $\text{Aut}(X)$ acts transitively on ‘arcs’, meaning pairs consisting of a vertex and an edge incident to that vertex. In this case $\text{Cov}(X)$ acts transitively on the arcs of the corresponding topological crystal. One example of this phenomena is the Laves graph, coming from the tetrahedron. The symmetry group of the tetrahedron is the Coxeter group

$$A_3 = \langle s_1, s_2, s_3 \mid (s_1 s_2)^3 = (s_2 s_3)^3 = s_1^2 = s_2^2 = s_3^2 = 1 \rangle.$$

Thus, the group of covering symmetries of the Laves graph is an extension of A_3 by \mathbb{Z}^3 .

Indeed, the vertices and edges of any Platonic solid form a graph whose symmetry group acts transitively on arcs. For example, the symmetry group of the cube and octahedron is the Coxeter group

$$B_3 = \langle s_1, s_2, s_3 \mid (s_1 s_2)^3 = (s_2 s_3)^4 = s_1^2 = s_2^2 = s_3^2 = 1 \rangle.$$

Since the cube has 6 faces, $H_1(X, \mathbb{R})$ is 5-dimensional for the cube graph. The corresponding topological crystal is thus 5-dimensional, and its group of covering symmetries, an extension of B_3 by \mathbb{Z}^5 , acts transitively on arcs. Similarly, the octahedron gives a 7-dimensional topological crystal whose group of covering symmetries, an extension of B_3 by \mathbb{Z}^7 , acts transitively on arcs. The cuboctahedron can be seen either as a truncated cube or a truncated octahedron, so it too has B_3 as its symmetry group, and in fact this group acts transitively on arcs. Since the cuboctahedron has 14 faces, we obtain a 13-dimensional crystal whose covering symmetries act transitively on arcs.


The symmetry group of the dodecahedron and icosahedron is

$$H_3 = \langle s_1, s_2, s_3 \mid (s_1 s_2)^3 = (s_2 s_3)^5 = s_1^2 = s_2^2 = s_3^2 = 1 \rangle,$$

and these solids give crystals of dimensions 11 and 19. The icosidodecahedron can be seen either as a truncated dodecahedron or a truncated icosahedron; it has H_3 as its symmetry group, and this

group acts transitively on arcs. It has 32 faces, so it gives a 31-dimensional crystal whose covering symmetries act transitively on arcs.

There is also an infinite family of degenerate Platonic solids called ‘hosohedra’ with two vertices, n edges and n faces. These faces cannot be made flat, since each face has just 2 edges, but that is not relevant to our construction: the vertices and edges still give a graph. For example, when $n = 6$, we have the ‘hexagonal hosohedron’:




The symmetry group of the n -gonal hosohedron is $I_n \times \mathbb{Z}/2$, where the Coxeter group

$$I_n = \langle s_1, s_2 \mid (s_1 s_2)^n = s_1^2 = s_2^2 = 1 \rangle$$


is a dihedral group, and $\mathbb{Z}/2$ acts to interchange the two vertices while fixing the edges. The automorphism group of its underlying graph is $S_n \times \mathbb{Z}/2$. The corresponding crystal has dimension $n - 1$, and its group of covering symmetries is an extension of $S_n \times \mathbb{Z}/2$ by \mathbb{Z}^{n-1} . The case $n = 3$ gives the graphene crystal, while $n = 4$ gives the diamond.

There are many other examples of graphs whose symmetry group acts transitively on arcs. One is the Petersen graph:



The symmetry group of the Petersen graph is S_5 . This graph has 10 vertices and 15 edges, so its Euler characteristic is -5 , which implies that its space of 1-cycles is 6-dimensional. It thus gives a 6-dimensional crystal whose group of covering symmetries, an extension of S_5 by \mathbb{Z}^6 , acts transitively on arcs.

Two more nice examples come from Klein’s quartic curve, a Riemann surface of genus three on which the 336-element group $\text{PGL}(2, \mathbb{F}_7)$ acts as isometries. These isometries preserve a tiling of Klein’s quartic curve by 56 triangles, with 7 meeting at each vertex. This picture is topologically correct, though not geometrically:



From this tiling we obtain a graph X embedded in Klein’s quartic curve. This graph has $56 \times 3/2 = 84$ edges and $56 \times 3/7 = 24$ vertices, so it has Euler characteristic -60 . It thus gives a 61-dimensional

topological crystal whose group of covering symmetries, an extension of $\mathrm{PGL}(2, \mathbb{F}_7)$ by \mathbb{Z}^{61} , acts transitively on arcs. There is also a dual tiling of Klein’s curve by 24 heptagons, 3 meeting at each vertex. This gives a graph with 84 edges and 56 vertices, hence Euler characteristic -28 . From this we obtain a 29-dimensional topological crystal whose group of covering symmetries is an extension of $\mathrm{PGL}(2, \mathbb{F}_7)$ by \mathbb{Z}^{29} . Again, this group acts transitively on arcs.


Another interesting property of a topological crystal is its ‘packing fraction’. The set A of atoms is contained in the lattice L obtained by projecting the integral 1-chains down to the space of 1-cycles:

$$L = \{\pi(c) : c \in C_1(X, \mathbb{Z})\}.$$


We can ask what fraction of the points in this lattice are actually atoms. We call this the ‘packing fraction’, and since $Z_1(X, \mathbb{Z})$ acts as translations on both A and L , we define it to be

$$\frac{|A/Z_1(X, \mathbb{Z})|}{|L/Z_1(X, \mathbb{Z})|}.$$

For example, suppose X is this graph:



Then the packing fraction is $\frac{2}{3}$, as can be seen here:



For any bridgeless connected graph X , we prove in Theorem 13 that

$$\frac{|A/Z_1(X, \mathbb{Z})|}{|L/Z_1(X, \mathbb{Z})|} = \frac{|V|}{|T|}$$

where V is the set of vertices and T is the set of spanning trees. The main tool here is Bacher, de la Harpe and Nagnibeda’s work on integral cycles and integral cuts [1].

Here is some data about the examples discussed above. We used Mathematica to count the spanning trees. In this table, $\mathbb{Z}^k . G$ denotes an extension of the group G by \mathbb{Z}^k .

Polyhedron	Covering Symmetries	Spanning Trees	Packing Fraction
tetrahedron	$\mathbb{Z}^3 \cdot A_3$	16	1/4
cube	$\mathbb{Z}^5 \cdot B_3$	384	1/48
octahedron	$\mathbb{Z}^7 \cdot B_3$	384	1/64
cuboctahedron	$\mathbb{Z}^{13} \cdot B_3$	331,776	1/27,648
dodecahedron	$\mathbb{Z}^{11} \cdot H_3$	5,184,000	1/259,200
icosahedron	$\mathbb{Z}^{19} \cdot H_3$	5,184,000	1/432,000
icosidodecahedron	$\mathbb{Z}^{31} \cdot H_3$	208,971,104,256,000	1/6,965,703,475,200
n -gonal hosohedron	$\mathbb{Z}^{n-1} \cdot (S_n \times \mathbb{Z}/2)$	n	$2/n$
Petersen graph	$\mathbb{Z}^6 \cdot S_5$	2,000	1/200
Klein quartic heptagonal tiling	$\mathbb{Z}^{29} \cdot \text{PGL}(2, F_7)$	38,542,412,611,584,000,000	1/688,257,368,064,000,000
Klein quartic triangular tiling	$\mathbb{Z}^{61} \cdot \text{PGL}(2, F_7)$	846,083,041,649,491,968	1/35,253,460,068,728,832

2. MAXIMAL ABELIAN COVERS

We start by reviewing some concepts from Sunada's *Topological Crystallography* [12]. It is convenient to work with graphs having two copies of each edge, one pointing in each direction. Thus, he defines a **graph** $X = (E, V, s, t, i)$ to consist of a set V of **vertices**, a set E of **edges**, maps $s, t: E \rightarrow V$ assigning to each edge its **source** and **target**, and a map $i: E \rightarrow E$ sending each edge to its **inverse**, obeying

$$s(i(e)) = t(e), \quad t(i(e)) = s(e), \quad i(i(e)) = e$$

and

$$i(e) \neq e$$

for all $e \in E$.

If $s(e) = v$ and $t(e) = w$ we write $e: v \rightarrow w$, and draw e as an interval with an arrow on it pointing from v to w . We write $i(e)$ as e^{-1} , and draw e^{-1} as the same interval as e , but with its arrow reversed. The equations obeyed by i say that taking the inverse of $e: v \rightarrow w$ gives an edge $e^{-1}: w \rightarrow v$ and that $(e^{-1})^{-1} = e$. No edge can be its own inverse.

A **map of graphs**, say $f: X \rightarrow X'$, is a pair of functions, one sending vertices to vertices and one sending edges to edges, that preserve the source, target and inverse maps. By abuse of notation we call both of these functions f .

From a graph X we can build a topological space $|X|$ called its **geometric realization**. We do this by taking one point for each vertex and gluing on one copy of $[0, 1]$ for each edge $e: v \rightarrow w$, gluing the point 0 to v and the point 1 to w , and then identifying the interval for each edge e with the interval for its inverse by means of the map $t \mapsto 1 - t$. Any map of graphs gives rise to a continuous map between their geometric realizations, and we say a map of graphs is a **cover** if this continuous map is a covering map. For simplicity we denote the fundamental group of $|X|$ by $\pi_1(X)$, and similarly for other topological invariants of $|X|$. However, there are contexts in which we need to distinguish between a graph X and its geometric realization $|X|$.

Any connected graph X has a **universal cover**, meaning a connected cover

$$p: \tilde{X} \rightarrow X$$

that covers every connected cover of X . The geometric realization of \tilde{X} is connected and simply connected. The fundamental group $\pi_1(X)$ acts as **deck transformations** of \tilde{X} , meaning invertible maps $g: \tilde{X} \rightarrow \tilde{X}$ such that $p \circ g = p$. We can take the quotient of \tilde{X} by the action of any subgroup $G \subseteq \pi_1(X)$ and get a cover $q: \tilde{X}/G \rightarrow X$.

In particular, if we take G to be the commutator subgroup of $\pi_1(X)$, we call the graph \tilde{X}/G the **maximal abelian cover** of the graph X , and denote it by \bar{X} . We obtain a cover

$$q: \bar{X} \rightarrow X$$

whose group of deck transformations is the abelianization of $\pi_1(X)$. This is just the first homology group $H_1(X, \mathbb{Z})$, which is a free abelian group.

We need a concrete description of the maximal abelian cover. We start with the universal cover, but first we need some preliminaries on paths in graphs.

Given vertices x, y in X , define a **path** from x to y to be a word of edges $\gamma = e_1 \cdots e_\ell$ with $e_i: v_{i-1} \rightarrow v_i$ for some vertices v_0, \dots, v_ℓ with $v_0 = x$ and $v_\ell = y$. We allow the word to be empty if and only if $x = y$; this gives the **trivial path** from x to itself. Given a path γ from x to y we write $\gamma: x \rightarrow y$, and we write the trivial path from x to itself as $1_x: x \rightarrow x$. We define the composite of paths $\gamma: x \rightarrow y$ and $\delta: y \rightarrow z$ via concatenation of words, obtaining a path we call $\gamma\delta: x \rightarrow z$. We call a path from a vertex x to itself a **loop** based at x .

We say two paths from x to y are **homotopic** if one can be obtained from the other by repeatedly introducing or deleting subwords of the form $e_i e_{i+1}$ where $e_{i+1} = e_i^{-1}$. If $[\gamma]$ is a homotopy class of paths from x to y , we write $[\gamma]: x \rightarrow y$. We can compose homotopy classes $[\gamma]: x \rightarrow y$ and $[\delta]: y \rightarrow z$ by setting $[\gamma][\delta] = [\gamma\delta]$.

If X is a connected graph, we can describe the universal cover \tilde{X} as follows. Fix a vertex x_0 of X , which we call the **basepoint**. The vertices of \tilde{X} are defined to be the homotopy classes of paths $[\gamma]: x_0 \rightarrow x$ where x is arbitrary. The edges in \tilde{X} from the vertex $[\gamma]: x_0 \rightarrow x$ to the vertex $[\delta]: x_0 \rightarrow y$ are defined to be the edges $e \in E$ with $[\gamma e] = [\delta]$. In fact, there is always at most one such edge. There is an obvious map of graphs

$$p: \tilde{X} \rightarrow X$$

sending each vertex $[\gamma]: x_0 \rightarrow x$ of \tilde{X} to the vertex x of X . This map is a cover.

Now we are ready to construct the maximal abelian cover \bar{X} . For this, we impose a further equivalence relation on paths, which is designed to make composition commutative whenever possible. However, we need to be careful. If $\gamma: x \rightarrow y$ and $\delta: x' \rightarrow y'$, the composites $\gamma\delta$ and $\delta\gamma$ are both well-defined if and only if $x' = y$ and $y' = x$. In this case, $\gamma\delta$ and $\delta\gamma$ share the same starting point and share the same ending point if and only if $x = x'$ and $y = y'$. If all four of these equations hold, both γ and δ are loops based at x . So, we shall impose the relation $\gamma\delta = \delta\gamma$ only in this case.

We say two paths are **homologous** if one can be obtained from another by:

- repeatedly introducing or deleting subwords $e_i e_{i+1}$ where $e_{i+1} = e_i^{-1}$, and/or
- repeatedly replacing subwords of the form $e_i \cdots e_j e_{j+1} \cdots e_k$ by those of the form $e_{j+1} \cdots e_k e_i \cdots e_j$, where $e_i \cdots e_j$ and $e_{j+1} \cdots e_k$ are loops based at the same vertex.

Our use of the term ‘homologous’ is a bit nonstandard, but we shall see some nice relations to homology theory, which we hope justifies this terminology.

We denote the homology class of a path γ by $[[\gamma]]$. Note that if two paths $\gamma: x \rightarrow y$, $\delta: x' \rightarrow y'$ are homologous then $x = x'$ and $y = y'$. Thus, the starting and ending points of a homology class of paths are well-defined, and given any path $\gamma: x \rightarrow y$ we write $[[\gamma]]: x \rightarrow y$. The composite of homology classes is also well-defined if we set $[[\gamma]][[\delta]] = [[\gamma\delta]]$.

We construct the maximal abelian cover of a connected graph X just as we constructed its universal cover, but using homology classes rather than homotopy classes of paths. Fix a basepoint x_0 for X . The vertices of \bar{X} , or **atoms**, are defined to be the homology classes of paths $[[\gamma]]: x_0 \rightarrow x$ where x is arbitrary. Any edge of \bar{X} , or **bond**, goes from some atom $[[\gamma]]: x_0 \rightarrow x$ to some atom $[[\delta]]: x_0 \rightarrow y$. The bonds from $[[\gamma]]$ to $[[\delta]]$ are defined to be the edges $e \in E$ with $[[\gamma e]] = [[\delta]]$. There is at most one bond between any two atoms. Again we have a covering map

$$q: \bar{X} \rightarrow X.$$

The homotopy classes of loops based at x_0 form a group, with composition as the group operation. This is the **fundamental group** $\pi_1(X)$ of the graph X . This is isomorphic to the fundamental group of the space associated to X . By our construction of the universal cover, $\pi_1(X)$ is also the

set of vertices of \tilde{X} that are mapped to x_0 by p . Furthermore, any element $[\gamma] \in \pi_1(X)$ defines a deck transformation of \tilde{X} that sends each vertex $[\delta]: x_0 \rightarrow x$ to the vertex $[\gamma][\delta]: x_0 \rightarrow x$.

Similarly, the homology classes of loops based at x_0 form a group with composition as the group operation. Since the additional relation used to define homology classes is precisely that needed to make composition of homology classes of loops commutative, this group is the abelianization of $\pi_1(X)$. It is therefore isomorphic to the first homology group $H_1(X, \mathbb{Z})$ of the geometric realization of X . By our construction of the maximal abelian cover, $H_1(X, \mathbb{Z})$ is also the set of vertices of \bar{X} that are mapped to x_0 by q . Furthermore, any element $[[\gamma]] \in H_1(X, \mathbb{Z})$ defines a deck transformation of \bar{X} that sends each vertex $[[\delta]]: x_0 \rightarrow x$ to the vertex $[[\gamma]][[\delta]]: x_0 \rightarrow x$.

3. ATOMS

Given a connected graph X , we call the vertices of its maximal abelian cover ‘atoms’ because they play the role of atoms in a topological crystal. Now we describe a systematic procedure for mapping these atoms into a vector space with inner product, namely the space of 1-cycles $Z_1(X, \mathbb{R})$. We show that this map is an embedding if and only if the graph has no bridges.

We begin with some standard material. Let X be a graph. The group of **integral 0-chains** on X , $C_0(X, \mathbb{Z})$, is the free abelian group on the set of vertices of X . The group of **integral 1-chains** on X , $C_1(X, \mathbb{Z})$, is the quotient of the free abelian group on the set of edges of X by relations $e^{-1} = -e$ for every edge e . The **boundary map** is the homomorphism

$$\partial: C_1(X, \mathbb{Z}) \rightarrow C_0(X, \mathbb{Z})$$

such that

$$\partial e = t(e) - s(e)$$

for each edge e , and

$$Z_1(X, \mathbb{Z}) = \ker \partial$$

is the group of **integral 1-cycles** on X .

Any path $\gamma = e_1 \cdots e_n$ in X determines an integral 1-chain

$$c_\gamma = e_1 + \cdots + e_n.$$

For any path γ we have

$$c_{\gamma^{-1}} = -c_\gamma,$$

and if γ and δ are composable then

$$c_{\gamma\delta} = c_\gamma + c_\delta.$$

We define vector spaces of **0-chains** and **1-chains** by

$$C_0(X, \mathbb{R}) = C_0(X, \mathbb{Z}) \otimes \mathbb{R}, \quad C_1(X, \mathbb{R}) = C_1(X, \mathbb{Z}) \otimes \mathbb{R},$$

respectively, and extend the boundary map to a linear map

$$\partial: C_1(X, \mathbb{R}) \rightarrow C_0(X, \mathbb{R}).$$

We let $Z_1(X, \mathbb{R})$ be the kernel of this linear map, or equivalently,

$$Z_1(X, \mathbb{R}) = Z_1(X, \mathbb{Z}) \otimes \mathbb{R},$$

and we call elements of this vector space **1-cycles**. The space of 1-cycles is isomorphic to the first cohomology with real coefficients, $H_1(X, \mathbb{R})$. Since $Z_1(X, \mathbb{Z})$ is a free abelian group, it forms a lattice in the space of 1-cycles. Any edge of X can be seen as a 1-chain, and there is a unique inner product on $C_1(X, \mathbb{R})$ such that edges form an orthonormal basis (with each edge e^{-1} counting as the negative of e .) There is thus an orthogonal projection

$$\pi: C_1(X, \mathbb{R}) \rightarrow Z_1(X, \mathbb{R}).$$

We now come to the main construction, first introduced by Kotani and Sunada [8]:

Definition 1. Let X be a connected graph with V as its set of vertices and a chosen basepoint $x_0 \in V$. Let its set of atoms be

$$A = \{[\alpha] : \alpha: x_0 \rightarrow x \text{ for some } x \in V\}.$$

Define the map

$$i: A \rightarrow Z_1(X, \mathbb{R})$$

by

$$i([\alpha]) = \pi(c_\alpha).$$

That i is well-defined follows from Lemma 3 below: homologous paths give the same 1-chain. For a bridgeless graph, i embeds the set of atoms in the space of 1-cycles:

Theorem 2. Let X be a connected graph. Then the following are equivalent:

- (1) X has no bridges.
- (2) The map $i: A \rightarrow Z_1(X, \mathbb{R})$ is one-to-one.

Proof. The map i is one-to-one if and only if for any atoms $[\alpha]$ and $[\beta]$, $i([\alpha]) = i([\beta])$ implies $[\alpha] = [\beta]$. Note that $\gamma = \beta^{-1}\alpha$ is a path in X with $c_\gamma = c_\alpha - c_\beta$, so

$$\pi(c_\gamma) = \pi(c_\alpha - c_\beta) = i([\alpha]) - i([\beta]).$$

Since $\pi(c_\gamma)$ vanishes if and only if c_γ is orthogonal to every 1-cycle, we have

$$c_\gamma \text{ is orthogonal to every 1-cycle} \iff i([\alpha]) = i([\beta]).$$

On the other hand, Lemma 3 implies that

$$c_\gamma = 0 \iff [\alpha] = [\beta].$$

Thus, to prove (1) \iff (2), it suffices to show that X has no bridges if and only if every 1-chain c_γ orthogonal to every 1-cycle has $c_\gamma = 0$. We show this in Lemma 6. \square

Lemma 3. Let X be a connected graph. Two paths $\alpha, \beta: x \rightarrow y$ in X are homologous if and only if $c_\alpha = c_\beta$.

Proof. The ‘only if’ direction is clear, since β is homologous to α precisely when it can be obtained from α by:

- repeatedly introducing or deleting a subword $e_i e_{i+1}$ where $e_{i+1} = e_i^{-1}$, and/or
- repeatedly replacing a subword of the form $e_i \cdots e_j e_{j+1} \cdots e_k$ by a subword of the form $e_{j+1} \cdots e_k e_i \cdots e_j$, where $e_i \cdots e_j$ and $e_{j+1} \cdots e_k$ are loops based at the same vertex.

and these moves do not change the corresponding 1-chain.

Conversely, suppose $c_\alpha = c_\beta$. Then $\gamma = \beta^{-1}\alpha$ is a loop based at x with $c_\gamma = c_\alpha - c_\beta = 0$. In Section 2 we saw that the group of homology classes of loops based at any vertex is the abelianization of $\pi_1(X)$, namely $H_1(X, \mathbb{Z})$. For a graph we have $H_1(X, \mathbb{Z}) \cong Z_1(X, \mathbb{Z})$, and the cycle associated to any loop γ is c_γ . Thus, $c_\gamma = 0$ implies that γ is homologous to the trivial loop. This implies that $\beta\gamma$ is homologous to β . It is also easy to see that $\beta\gamma = \beta\beta^{-1}\alpha$ is homologous to α . Thus α is homologous to β . \square

The following lemmas are the technical heart of Theorem 2. We need to show that any nonzero 1-chain coming from a path in a bridgeless graph has nonzero inner product with some 1-cycle. The following lemmas, inspired by an idea of Ilya Bogdanov [5], yield an algorithm for actually constructing such a 1-cycle. This 1-cycle also has other desirable properties, which will come in handy later.

To state these, let a **simple path** be one in which each vertex appears at most once. Let a **simple loop** be a loop $\gamma: x \rightarrow x$ in which each vertex except x appears at most once, while x appears exactly twice, as the starting point and ending point. Let the **support** of a 1-chain c , denoted $\text{supp}(c)$, be the set of edges e such that $\langle c, e \rangle > 0$. This excludes edges with $\langle c, e \rangle = 0$, but also those with $\langle c, e \rangle < 0$, which are inverses of edges in the support. Note that

$$c = \sum_{e \in \text{supp}(c)} \langle c, e \rangle e.$$

Thus, $\text{supp}(c)$ is the smallest set of edges such that c can be written as a positive linear combination of edges in this set. Given 1-chains c and c' , we have $\text{supp}(c') \subseteq \text{supp}(c)$ if and only if $\langle c, e \rangle = 0$ implies $\langle c', e \rangle = 0$ and $\langle c, e \rangle > 0$ implies $\langle c', e \rangle \geq 0$.

Lemma 4. *Let X be any graph and let c be an integral 1-cycle on X . Then for some n we can write*

$$c = c_{\sigma_1} + \cdots + c_{\sigma_n}$$

where σ_i are simple loops with $\text{supp}(c_{\sigma_i}) \subseteq \text{supp}(c)$.

Proof. If $c = 0$ we have $c = c_\sigma$ for any trivial loop σ . Otherwise $\text{supp}(c)$ is nonempty, and we can construct a simple loop $e_1 e_2 \cdots$ in X as follows.

We start by picking any edge $e_1: v_0 \rightarrow v_1$ in $\text{supp}(c)$ and moving to the vertex v_1 . In general, when we arrive at the vertex v_i , if it is not one we have visited before we pick an edge $e_{i+1}: v_i \rightarrow v_{i+1}$ in $\text{supp}(c)$. To see that such an edge must exist, note that since c is a cycle, the sum of the coefficients $\langle c, e \rangle$ over edges $e \in \text{supp}(c)$ with v_i as source must equal the sum of these coefficients over edges in $\text{supp}(c)$ with v_i as target. Since we are assuming v_i is the target of an edge in $\text{supp}(c)$, namely e_i , it must also be the source of an edge in $\text{supp}(c)$.

Since the support of c is finite, eventually we must reach a vertex we have visited before. So, eventually $v_i = v_j$ for some $j < i$. Then $\sigma_1 = e_{j+1} \cdots e_i$ is a simple loop. By construction we have $\text{supp}(c_{\sigma_1}) \subseteq \text{supp}(c)$ as desired.

Now let $c_2 = c - c_{\sigma_1}$. We claim that also $\text{supp}(c_2) \subseteq \text{supp}(c)$. To see this, note that both $\langle c, e \rangle$ and $\langle c_{\sigma_1}, e \rangle$ vanish except for edges in $\text{supp}(c)$ and their inverses. If $e \in \text{supp}(c)$ then $\langle c, e \rangle \geq 1$, while $\langle c_{\sigma_1}, e \rangle \leq 1$. Thus $\langle c, e \rangle = 0$ implies $\langle c_2, e \rangle = 0$, while $\langle c, e \rangle \geq 0$ implies $\langle c_2, e \rangle \geq 0$. It follows that $\text{supp}(c_2) \subseteq \text{supp}(c)$.

Since c_2 is again an integral 1-cycle, either $c_2 = 0$ or we can repeat the argument just given with c_2 replacing c . Continuing on, we can construct 1-cycles c_2, c_3, \dots , each with support contained in the support of the previous one, and simple loops $\sigma_1, \sigma_2, \dots$ with $c_{k+1} = c_k - c_{\sigma_k}$ and $\text{supp}(c_{\sigma_k}) \subseteq \text{supp}(c_k)$. With each new 1-cycle the coefficient of at least one edge is less than before, so eventually c_k reaches zero and the procedure terminates. We thus obtain simple loops $\sigma_1, \dots, \sigma_n$ with $c = c_{\sigma_1} + \cdots + c_{\sigma_n}$ and $\text{supp}(c_{\sigma_i}) \subseteq \text{supp}(c)$. \square

Lemma 5. *Let $\gamma: x \rightarrow y$ be a path in a graph X . Then for some $n \geq 0$ we can write*

$$c_\gamma = c_\delta + c_{\sigma_1} + \cdots + c_{\sigma_n}$$

where $\delta: x \rightarrow y$ is a simple path and σ_i are simple loops with $\text{supp}(c_\delta), \text{supp}(c_{\sigma_i}) \subseteq \text{supp}(c_\gamma)$.

Proof. Suppose that $\gamma: x \rightarrow y$ is a path in X . If $x = y$ then c_γ is a cycle and the lemma follows from Lemma 4, taking δ to be any trivial path. Thus we assume $x \neq y$.

We claim there is a simple path $\delta: x \rightarrow y$ with $\text{supp}(c_\delta) \subseteq \text{supp}(c_\gamma)$. We can construct such a path, say $\delta = e_1 e_2 \cdots$, as follows. Start by setting $k = 1$ and let $c_k = c_\gamma$ in this case. As we proceed we will repeatedly increment k and subtract a 1-cycle from c_k to define a new 1-chain c_{k+1} with $\text{supp}(c_{k+1}) \subseteq \text{supp}(c_k)$.

We start at the vertex $v_0 = x$. In general, if v_i is neither y nor a vertex we have visited before, we pick an edge $e_{i+1}: v_i \rightarrow v_{i+1}$ in $\text{supp}(c_k)$ and move to the vertex v_{i+1} . Such an edge must exist, because $\partial(c_k) = y - x$, so the sum of the coefficients $\langle c_k, e \rangle$ over edges $e \in \text{supp}(c_k)$ with v_i as source is greater than or equal to the sum of these coefficients over edges in $\text{supp}(c_k)$ with v_i as target.

Since the support of c_k is finite, eventually v_i is either y or a vertex we have visited before. If $v_i = y$ we stop: by construction $\delta = e_1 \cdots e_i$ is a simple path from x to y . Moreover, $\text{supp}(c_\delta) \subseteq \text{supp}(c_k) \subseteq \text{supp}(c_\gamma)$ as desired.

If v_i is a vertex we have visited before, say $v_i = v_j$ for $j < i$, then $\sigma_k = e_{j+1} \cdots e_i$ is a simple loop. By construction we have $\text{supp}(c_{\sigma_k}) \subseteq \text{supp}(c_k)$. The argument given in Lemma 4 shows that $c_{k+1} = c_k - c_{\sigma_k}$ has $\text{supp}(c_{k+1}) \subseteq \text{supp}(c_k)$. Since $\partial(c_{k+1}) = \partial(c_k) = y - x$, the new 1-chain c_{k+1} is still nonzero. We may thus increment k by 1 and restart the process of building a simple path from x to y .

This algorithm must eventually succeed in building a simple path $\delta: x \rightarrow y$, since the sum of the coefficients of c_{k+1} is strictly less than that for c_k , yet c_{k+1} can never vanish, given that $\partial c_{k+1} = y - x \neq 0$.

If we let $c = c_\gamma - c_\delta$ then c is a 1-cycle, and our construction ensures that $\text{supp}(c) \subseteq \text{supp}(c_\gamma)$. Thus, we can use Lemma 4 to write $c = c_{\sigma_1} + \cdots + c_{\sigma_n}$ where σ_i are simple loops with $\text{supp}(c_{\sigma_i}) \subseteq \text{supp}(c_\gamma)$. This implies that

$$c_\gamma = c_\delta + c_{\sigma_1} + \cdots + c_{\sigma_n}$$

where δ and σ_i have the desired properties. \square

Lemma 6. *Let X be a graph. Then the following are equivalent:*

- (1) X has no bridges.
- (2) For any path γ in X , if c_γ is orthogonal to every 1-cycle then $c_\gamma = 0$.

Proof. To prove (2) \implies (1), we assume X has a bridge e and show $\langle c, e \rangle = 0$ for every 1-cycle c . This will give a path $\gamma = e$ for which $c_\gamma \neq 0$ but c_γ is orthogonal to every 1-cycle.

Since every 1-cycle is a linear combination of 1-cycles c_λ coming from loops λ , it suffices to show $\langle c_\lambda, e \rangle = 0$ for every loop λ . Since e is a bridge, removing e breaks X into two connected components. Any loop $\lambda = e_1 \cdots e_n$ must start and end in the same component, so the number of edges e_i that equal e must equal the number of e_i that equal e^{-1} . Thus $\langle c_\lambda, e \rangle = 0$.

To prove (1) \implies (2), suppose X has no bridges. It suffices to show that if $\gamma: x \rightarrow y$ is a path for which c_γ is orthogonal to every 1-cycle, then $c_\gamma = 0$. We use Lemma 5 to write

$$c_\gamma = c_\delta + c_{\sigma_1} + \cdots + c_{\sigma_n}$$

where $\delta: x \rightarrow y$ is a simple path and σ_i are simple loops with $\text{supp}(c_\delta), \text{supp}(c_{\sigma_i}) \subseteq \text{supp}(c_\gamma)$. If any of the 1-chains c_{σ_i} are nonzero we are done, since then

$$\langle c_{\sigma_i}, c_\gamma \rangle \geq \langle c_{\sigma_i}, c_{\sigma_i} \rangle > 0$$

so c_γ is not orthogonal to every 1-cycle. Thus, we assume $c_\gamma = c_\delta$. We assume that $c_\delta \neq 0$, and need to construct a 1-cycle that is not orthogonal to c_δ . We also assume $x \neq y$, since otherwise we can use c_δ itself as the desired 1-cycle.

Write $\delta = e_1 \cdots e_n$ with $e_i: v_{i-1} \rightarrow v_i$. These edges are distinct since δ is a simple path, and there is at least one of them since $x = v_0$ is not equal to $y = v_n$. Since the last edge e_n is not a bridge, there must be a path in X from v_n to v_0 that does not include e_n . If we follow that path only as far as the first vertex $v_i \neq v_n$ that it reaches, we obtain a path $\alpha: v_n \rightarrow v_i$ that includes no edges e_i , nor their inverses. Thus, we have

$$\langle c_\alpha, c_\delta \rangle = 0.$$

If $\beta = e_{i+1} \cdots e_n$ is the portion of δ that goes from v_i to v_n , then

$$\langle c_\beta, c_\delta \rangle = \langle c_\beta, c_\beta \rangle \geq 1.$$

The path $\alpha\beta$ is a loop, so $c_\alpha + c_\beta$ is a 1-cycle, and this 1-cycle is not orthogonal to c_δ , since

$$\langle c_\alpha, c_\delta \rangle + \langle c_\beta, c_\delta \rangle \geq 1. \quad \square$$

4. TOPOLOGICAL CRYSTALS

In the previous section we took a connected bridgeless graph X and embedded its atoms into the space of 1-cycles via a map

$$i: A \rightarrow Z_1(X, \mathbb{R}).$$

These atoms are the vertices of the maximal abelian cover \overline{X} , or equivalently, the vertices of its geometric realization $|\overline{X}|$. Now we extend i to an embedding

$$j: |\overline{X}| \rightarrow Z_1(X, \mathbb{R}).$$

We call the image of this embedding the **topological crystal** associated to X .

The idea is that just as i maps each atom to a point in the vector space $Z_1(X, \mathbb{R})$, j maps each edge of $|\overline{X}|$ to a straight line segment between such points. These line segments serve as the ‘bonds’ of a topological crystal. The only challenge is to show that these bonds do not cross each other.

Let X be a connected bridgeless graph, and fix a vertex x_0 . Technically, an atom is a homology class of paths $[\alpha]$ starting at x_0 and ending at any vertex x . Given an edge $e: x \rightarrow y$ in X , there is

a unique edge in \overline{X} from the atom $[\alpha]: x_0 \rightarrow x$ to the atom $[\alpha e]: x_0 \rightarrow y$. Every edge of \overline{X} arises in this manner in a unique way. So, an edge in \overline{X} amounts to a pair

$$[\alpha]: x_0 \rightarrow x, \quad e: x \rightarrow y.$$

As explained in Section 2, each edge in \overline{X} gives a closed interval in the geometric realization $|\overline{X}|$, with a preferred parametrization $\phi: [0, 1] \rightarrow |\overline{X}|$. The inverse of an edge gives the same closed interval with the reverse parametrization, $1 - \phi$. We call these closed intervals **bonds** to distinguish from the edges of an abstract graph.

Given a finite-dimensional vector space V , we say a map $f: |\overline{X}| \rightarrow V$ is **affine on bonds** if for any given bond the composite map $f \circ \phi: [0, 1] \rightarrow V$ is affine, meaning

$$f(\phi(t)) = a + bt$$

for some $a, b \in V$. It follows that f maps each bond to a straight line segment in V . It also follows that f is continuous. Note also that given f defined on the atoms, there exists a unique map extending f that is affine on bonds. The reason is that an affine map from $[0, 1]$ to V is determined by its values at the endpoints.

Theorem 7. *If X is a connected graph, the map $i: A \rightarrow Z_1(X, \mathbb{R})$ extends uniquely to a map*

$$j: |\overline{X}| \rightarrow Z_1(X, \mathbb{R})$$

that is affine on bonds. If X is also bridgeless, then j is one-to-one.

Proof. As mentioned, the existence and uniqueness of an extension j that is affine on bonds is automatic. The task is to show that j is an embedding when X has no bridges.

Each bond in \overline{X} arises from a pair

$$[\alpha]: x_0 \rightarrow x, \quad e: x \rightarrow y,$$

and the points on this bond are parametrized by numbers $t \in [0, 1]$. Thus, we can describe any point $p \in |\overline{X}|$ as a triple

$$[\alpha]: x_0 \rightarrow x, \quad e: x \rightarrow y, \quad t \in [0, 1].$$

However, there is some redundancy in this description. If $t = 0$ then p is an atom, and in this case e becomes irrelevant: any choice of $e: x \rightarrow y$ gives the same point p . In general, any point described by the triple

$$[\alpha]: x_0 \rightarrow x, \quad e: x \rightarrow y, \quad t \in [0, 1]$$

is also described by the triple

$$[\alpha e]: x_0 \rightarrow y, \quad e^{-1}: y \rightarrow x, \quad 1 - t \in [0, 1].$$

If $t \in (0, 1)$, the point p is not an atom, and the above triples are the only two that describe it.

Suppose we have a point $p \in |\overline{X}|$ corresponding to some triple $[\alpha]: x_0 \rightarrow x, e: x \rightarrow y, t \in [0, 1]$. We define a 1-chain $c_p \in C_1(X, \mathbb{R})$ by

$$c_p = c_\alpha + te.$$

Thanks to Lemma 3, c_p is independent of the choice of path α representing the atom $[\alpha]$. Furthermore, different triples describing the same point p give the same result for c_p . Thus, c_p is well-defined.

Define

$$j: |\overline{X}| \rightarrow Z_1(X, \mathbb{R})$$

by

$$j(p) = \pi(c_p).$$

This equals $i(p)$ when p is an atom, since $c_p = c_\alpha$ when p is the atom $[\alpha]$. The map j is affine on each bond since c_p depends affinely on the parameter t and π is linear. So, j is the unique map equalling i on atoms that is affine on each bond.

To prove the theorem, it suffices to show that if p and q are points in \overline{X} such that $c_p - c_q$ is orthogonal to every 1-cycle, then $p = q$. To do this, first we show that $c_p = c_q$ implies $p = q$. Then we show that if $c_p - c_q$ is orthogonal to every 1-cycle then $c_p = c_q$.

Suppose $c_p = c_q$. These 1-chains are integral if and only if p and q are atoms, in which case $p = q$ by Lemma 3. So, we may assume that neither 1-chain is integral and neither p nor q is an atom. Thus, we may write p as a triple

$$[[\alpha]], \quad e, \quad t \in (0, 1)$$

and q as a triple

$$[[\beta]], \quad f, \quad u \in (0, 1),$$

implying

$$c_\alpha + te = c_\beta + uf.$$

Working modulo integral 1-chains we have

$$te = uf \pmod{C_1(X, \mathbb{Z})},$$

so either $f = e$ and $u = t$ or $f = e^{-1}$ and $u = 1 - t$. In the first case we have

$$c_\alpha = c_\beta$$

so $[[\alpha]] = [[\beta]]$ by Lemma 3, and thus $p = q$. In the second case we have

$$c_\alpha + e = c_\beta$$

so $[[\alpha e]] = [[\beta]]$ by Lemma 3, and again $p = q$.

Next suppose that $c_p - c_q$ is orthogonal to every 1-cycle. Write

$$c = c_p - c_q$$

and note that

$$c = c_\alpha - c_\beta + te - uf.$$

We need to prove that if c is orthogonal to every 1-cycle then $p = q$.

Form a path γ by either composing $\beta^{-1}\alpha$ with the edges f^{-1} and e , or not, in such a way that the coefficient of every edge in c_γ agrees in sign with its coefficient in c . More precisely, let

$$\gamma = \begin{cases} f^{-1}\beta^{-1}\alpha e & \text{if } \langle c, -f \rangle > 0 \text{ and } \langle c, e \rangle > 0 \\ f^{-1}\beta^{-1}\alpha & \text{if } \langle c, -f \rangle > 0 \text{ and } \langle c, e \rangle \leq 0 \\ \beta^{-1}\alpha e & \text{if } \langle c, -f \rangle \leq 0 \text{ and } \langle c, e \rangle > 0 \\ \beta^{-1}\alpha & \text{if } \langle c, -f \rangle \leq 0 \text{ and } \langle c, e \rangle \leq 0. \end{cases}$$

Thus $\text{supp}(c_\gamma) = \text{supp}(c)$. Next, using Lemma 5, write

$$c_\gamma = c_\delta + c_{\sigma_1} + \cdots + c_{\sigma_n}$$

where $\delta: x \rightarrow y$ is a simple path and σ_i are simple loops with $\text{supp}(c_\delta), \text{supp}(c_{\sigma_i}) \subseteq \text{supp}(c_\gamma)$.

If any 1-cycle c_{σ_i} is nonzero, then since $\text{supp}(c_{\sigma_i}) \subseteq \text{supp}(c_\gamma) = \text{supp}(c)$ we have

$$\langle c, c_{\sigma_i} \rangle > 0,$$

contradicting our assumption that c is orthogonal to every 1-cycle. We thus assume there are no nonzero 1-cycles c_{σ_i} . This implies $c_\gamma = c_\delta$. If δ is a path with no edges, then $c_\gamma = 0$. Since $\text{supp}(c_\gamma) = \text{supp}(c)$, this implies $c = 0$ as desired.

We are left with the case where δ has at least one edge. However, we shall see that this leads to a contradiction. Write $\delta = e_1 \cdots e_n$ with $e_i: v_{i-1} \rightarrow v_i$. There is at least one of these edges, and they are distinct since δ is a simple path. Since the last edge e_n is not a bridge, there must be a path in X from v_n to v_0 that does not include e_n . If we follow that path only as far as the first vertex $v_i \neq v_n$ that it reaches, we obtain a path $\alpha: v_n \rightarrow v_i$ that includes no edges e_i , nor their inverses. Since $\text{supp}(c) = \text{supp}(c_\delta)$, this implies

$$\langle c_\alpha, c \rangle = 0.$$

If $\beta = e_{i+1} \cdots e_n$ is the portion of δ that goes from v_i to v_n , then $\langle c_\beta, c_\delta \rangle \geq 1$, and since $\text{supp}(c) = \text{supp}(c_\delta)$, we have

$$\langle c_\beta, c \rangle > 0.$$

Since $\alpha\beta$ is a loop, $c_\alpha + c_\beta$ is a 1-cycle, and

$$\langle c_\alpha + c_\beta, c \rangle > 0.$$

This contradicts our assumption that c is orthogonal to every 1-cycle. \square

5. SYMMETRIES

Every bridgeless graph gives a topological crystal. How are the graph's symmetries related to those of its crystal? To tackle this, we start by asking how the symmetries of a graph X are related to those of its maximal abelian cover

$$q: \overline{X} \rightarrow X.$$

We shall not study all the symmetries of \overline{X} , only those that cover symmetries of X . These form a group $\text{Cov}(X)$. As we shall see, this group contains $H_1(X, \mathbb{Z})$, which acts as deck transformations of \overline{X} . The group $\text{Cov}(X)$ also maps onto the group of symmetries of X : in other words, every symmetry of X can be covered by some symmetry of \overline{X} .

In fact, we shall prove that $\text{Cov}(X)$ is an extension of the symmetry group of X by $H_1(X, \mathbb{Z})$. We also prove that $\text{Cov}(X)$ acts on the vector space $Z_1(X, \mathbb{R})$ as affine isometries in a way that preserves the topological crystal embedded in this space.

Define an **automorphism** of a graph X to be a map of graphs $f: X \rightarrow X$ that has an inverse. The automorphisms of X form a group $\text{Aut}(X)$. We say an automorphism $g \in \overline{X} \rightarrow \overline{X}$ **covers** an automorphism $f: X \rightarrow X$ if

$$q \circ g = f \circ q.$$

There is a group $\text{Cov}(X)$ where an element is an automorphism of \overline{X} that covers some automorphism of X . We call elements of $\text{Cov}(X)$ **covering symmetries**. Since q is onto, any covering symmetry covers at most one automorphism of X . There is thus a map

$$\psi: \text{Cov}(X) \rightarrow \text{Aut}(X)$$

sending any $g \in \text{Cov}(X)$ to the automorphism $f \in \text{Aut}(X)$ that it covers. It is easy to check that ψ is a group homomorphism.

Lemma 8. *Let X be a connected graph with basepoint. The homomorphism ψ is onto and its kernel is $H_1(X, \mathbb{Z})$, so we have a short exact sequence*

$$1 \longrightarrow H_1(X, \mathbb{Z}) \longrightarrow \text{Cov}(X) \xrightarrow{\psi} \text{Aut}(X) \longrightarrow 1.$$

In other words, $\text{Cov}(X)$ is an extension of $\text{Aut}(X)$ by $H_1(X, \mathbb{Z})$.

Proof. An element $g \in \text{Cov}(X)$ is in the kernel of ψ if and only if it covers the identity map, which means that it is a deck transformation of the maximal abelian cover of X . In Section 2 we saw that the group of these deck transformations is $H_1(X, \mathbb{Z})$.

To show that ψ is onto, we choose any automorphism $f: X \rightarrow X$ and find an automorphism $g: \overline{X} \rightarrow \overline{X}$ that covers f . We need to describe how g acts on vertices and edges. A vertex of \overline{X} is an atom, that is, a homology class of paths in X starting from the basepoint, say $[[\alpha]]: x_0 \rightarrow x$. The automorphism f sends any path $\alpha: x_0 \rightarrow x$ to a path we call $f(\alpha): f(x_0) \rightarrow f(x)$. Unfortunately this path no longer defines an atom unless $f(x_0) = x_0$. To deal with this, arbitrarily choose a path $\beta: x_0 \rightarrow f(x_0)$ and let

$$g([[\alpha]]) = [[\beta f(\alpha)]]: x_0 \rightarrow f(x).$$

An edge \overline{X} is a pair

$$[[\alpha]]: x_0 \rightarrow x, \quad e: x \rightarrow y.$$

We let g map this edge to the edge

$$g([[\alpha]]): x_0 \rightarrow f(x), \quad f(e): f(x) \rightarrow f(y).$$

One can check that g is a map of graphs. Since it is one-to-one and onto on both vertices and edges, it is an automorphism. One can also check that g covers f , since the definition of the projection $q: \overline{X} \rightarrow X$ implies that

$$q([[\alpha]]) = x$$

and

$$q(g([[\alpha]])) = f(x). \quad \square$$

The proof of the above lemma gives an explicit description of the group $\text{Cov}(X)$ and its action on \overline{X} . In the proof, two choices of $\beta: x_0 \rightarrow f(x_0)$ define the same automorphism of \overline{X} if and only if they are homologous. A homology class of paths $[\beta]: x_0 \rightarrow f(x_0)$ is just an atom mapping to $f(x_0)$ via the projection $q: \overline{X} \rightarrow X$. Thus, we have:

Lemma 9. *Elements of $\text{Cov}(X)$ are in one-to-one correspondence with pairs $(f, [\beta])$ where $f \in \text{Aut}(X)$ and $[\beta]: x_0 \rightarrow f(x_0)$. Such a pair acts on \overline{X} by mapping each atom*

$$[\alpha]: x_0 \rightarrow x$$

to the atom

$$[\beta f(\alpha)]: x_0 \rightarrow f(x).$$

and each edge

$$[\alpha]: x_0 \rightarrow x, \quad e: x \rightarrow y$$

to the edge

$$[\beta f(\alpha)]: x_0 \rightarrow f(x), \quad f(e): f(x) \rightarrow f(y).$$

There may in principle be other automorphisms of \overline{X} not lying in the group $\text{Cov}(X)$; we have not found examples. Automorphisms in $\text{Cov}(X)$ have the advantage that they also act naturally as affine transformations of the vector space $Z_1(X, \mathbb{R})$. Furthermore, this space has an inner product coming from the inner product on 1-chains, which in turn defines a metric. The group $\text{Cov}(X)$ acts as **affine isometries** of $Z_1(X, \mathbb{R})$, that is, affine transformations that preserve this metric:

Theorem 10. *There exists a unique action ρ of $\text{Cov}(X)$ as affine transformations of $Z_1(X, \mathbb{R})$ for which the embedding $i: A \rightarrow Z_1(X, \mathbb{R})$ is equivariant, meaning that*

$$\rho(g)i([\alpha]) = i(g[\alpha])$$

for all $g \in \text{Cov}(X)$ and $[\alpha] \in A$. Moreover, the transformations $\rho(g)$ are affine isometries.

Proof. Since $Z_1(X, \mathbb{Z}) \subseteq A$, every point in $Z_1(X, \mathbb{R})$ is an affine combination of points $i([\alpha])$ where $[\alpha]$ is an atom, so an affine transformation of $Z_1(X, \mathbb{R})$ is uniquely determined by its action on such points. This proves the uniqueness of ρ .

For existence, take any $g \in \text{Cov}(X)$ and use Lemma 9 to write it as a pair $(f, [\beta])$ where $f \in \text{Aut}(X)$ and $[\beta]: x_0 \rightarrow f(x_0)$. For any $c \in Z_1(X, \mathbb{R})$, define

$$\rho(g)(c) = i([\beta]) + f_*(c)$$

where

$$f_*: C_1(X, \mathbb{R}) \rightarrow C_1(X, \mathbb{R})$$

is the linear transformation with $f_*(e) = f(e)$ for any edge e of X . Since f_* preserves the inner product on 1-chains and the subspace of 1-cycles we have

$$f_*\pi = \pi f_*$$

where π is the projection from 1-chains to 1-cycles. Recall from Definition 1 that $i([\alpha]) = \pi(c_\alpha)$. Thus we have

$$\begin{aligned} \rho(g)(i([\alpha])) &= i([\beta]) + f_*(i([\alpha])) \\ &= \pi(c_\beta) + f_*(\pi(c_\alpha)) \\ &= \pi(c_\beta) + \pi(f_*(c_\alpha)) \\ &= \pi(c_\beta + c_{f(\alpha)}) \\ &= \pi(c_{\beta f(\alpha)}) \\ &= i([\beta f(\alpha)]) \\ &= i(g[\alpha]) \end{aligned}$$

where in the fourth step we used the fact that $f_*(c_\alpha) = c_{f(\alpha)}$ for any path α .

To check that ρ is an action we need to show $\rho(gh) = \rho(g)\rho(h)$ for any $g, h \in \text{Cov}(X)$. This can be done by brute force, but it suffices to check it on points of the form $i(\llbracket\alpha\rrbracket)$, since every point of $Z_1(X, \mathbb{R})$ is an affine combination of these. For this, note that:

$$\begin{aligned}\rho(gh)(i(\llbracket\alpha\rrbracket)) &= i((gh)(\llbracket\alpha\rrbracket)) \\ &= i(g(h\llbracket\alpha\rrbracket)) \\ &= \rho(g)(i(h\llbracket\alpha\rrbracket)) \\ &= \rho(g)(\rho(h)(i(\llbracket\alpha\rrbracket))).\end{aligned}$$

Finally, it is clear that $\rho(g)$ is an affine isometry, since

$$\rho(g)(c) = i(\llbracket\beta\rrbracket) + f_*(c)$$

implies that $\rho(g)$ is a linear isometry followed by a translation. \square

We conclude by showing that $\text{Cov}(X)$ acts on the space of 1-cycles in a way that preserves the topological crystal. In general, any automorphism f of any graph functorially determines an automorphism of its geometric realization, which we call $|f|$. Thus, $\text{Cov}(X)$ acts as automorphisms of $|\overline{X}|$. In Theorem 7 we saw that the map $i: A \rightarrow Z_1(X, \mathbb{R})$ extends uniquely to an embedding $j: |\overline{X}| \rightarrow Z_1(X, \mathbb{R})$ that is affine on bonds. To show that $\text{Cov}(X)$ preserves the topological crystal, that is, the image of this embedding, we prove that j is equivariant:

Theorem 11. *The embedding $j: |\overline{X}| \rightarrow Z_1(X, \mathbb{R})$ is equivariant with respect to the action of $\text{Cov}(X)$, meaning that*

$$\rho(g)j(p) = j(|g|(p))$$

for all $g \in \text{Cov}(X)$ and $p \in |\overline{X}|$.

Proof. As noted in the proof of Theorem 7, every point $p \in |\overline{X}|$ corresponds to a triple

$$\llbracket\alpha\rrbracket: x_0 \rightarrow x, \quad e: x \rightarrow y, \quad t \in [0, 1],$$

and then we have

$$j(p) = \pi(c_\alpha + te).$$

If we let $\beta = \alpha e$ we thus have

$$j(p) = (1-t)\pi(c_\alpha) + t\pi(c_\beta)$$

or in other words

$$j(p) = (1-t)i(\llbracket\alpha\rrbracket) + ti(\llbracket\beta\rrbracket).$$

By Lemma 9 any element $g \in \text{Cov}(X)$ corresponds to some pair $f \in \text{Aut}(X)$, $\llbracket\beta\rrbracket: x_0 \rightarrow f(x_0)$. The point $|g|(p)$ then corresponds to the triple

$$g(\llbracket\alpha\rrbracket): x_0 \rightarrow f(x), \quad g(e): f(x) \rightarrow f(y), \quad t \in [0, 1].$$

Since $\rho(g)$ is affine, it follows that

$$\begin{aligned}\rho(g)j(p) &= (1-t)\rho(g)(i(\llbracket\alpha\rrbracket)) + t\rho(g)i(\llbracket\beta\rrbracket) \\ &= (1-t)i(g(\llbracket\alpha\rrbracket)) + ti(g(\llbracket\beta\rrbracket)) \\ &= j(|g|(p)).\end{aligned} \quad \square$$

6. THE PACKING FRACTION

In Section 3 we took a connected bridgeless graph X with a basepoint x_0 and embedded its set of atoms, A , into its space of 1-cycles, $Z_1(X, \mathbb{R})$. Now we shall use this embedding to reinterpret A as a subset of $Z_1(X, \mathbb{R})$, as follows:

$$A = \{\pi(c_\alpha) : \alpha: x_0 \rightarrow x \text{ for some } x\}.$$

It is interesting to ask how densely packed these atoms are. Computing the density of the corresponding sphere packing seems hard, but we can also ask what fraction of ‘potential locations for atoms’ are actually filled by atoms. To make sense of this, note that A is contained in the lattice L obtained by projecting the integral 1-chains down to the space of 1-cycles:

$$L = \{\pi(c) : c \in C_1(X, \mathbb{Z})\}$$

where π is the projection of the space of 1-chains onto the space of 1-cycles. We can think of points in L as potential locations for atoms.

We cannot define the fraction of potential atoms that are actual atoms to be $|A|/|L|$, because both the numerator and denominator in this fraction are infinite. To deal with this, note that

$$Z_1(X, \mathbb{Z}) \subseteq A \subseteq L$$

and $Z_1(X, \mathbb{Z})$ acts as translations on both A and L . For L this is obvious, because $Z_1(X, \mathbb{Z})$ is a sublattice of L . For A it follows from

$$Z_1(X, \mathbb{Z}) = \{\pi(c_\gamma) : \gamma: x_0 \rightarrow x_0\}$$

and the fact that composing paths has the effect of adding their 1-chains: we can compose any loop $\gamma: x_0 \rightarrow x_0$ with any path $\alpha: x_0 \rightarrow x$ and get an atom $\pi(c_{\gamma\alpha}) = \pi(c_\gamma) + \pi(c_\alpha)$ which is the result of translating the atom $\pi(c_\alpha)$ by $\pi(c_\gamma) \in Z_1(X, \mathbb{Z})$.

When X is a finite graph, $A/Z_1(X, \mathbb{Z})$ and $L/Z_1(X, \mathbb{Z})$ are finite abelian groups. In this case we define the **packing fraction** to be the ratio

$$\frac{|A/Z_1(X, \mathbb{Z})|}{|L/Z_1(X, \mathbb{Z})|}.$$

This is a regularized version of the meaningless ratio $|A|/|L|$.

To compute the packing fraction, we need another description of the lattice L . Given any lattice Λ in a finite-dimensional real inner product space V , the **dual lattice** Λ^* is

$$\Lambda^* = \{v \in V : \langle v, w \rangle \in \mathbb{Z} \text{ for all } w \in \Lambda\}.$$

We say Λ is **integral** if $\Lambda \subseteq \Lambda^*$, or in other words, if the inner product of any two vectors in Λ is an integer. We say Λ is **self-dual** if $\Lambda = \Lambda^*$.

It is easy to see that $C_1(X, \mathbb{Z})$ is a self-dual lattice in $C_1(X, \mathbb{R})$, since edges form an orthonormal basis. It follows that $Z_1(X, \mathbb{Z})$ is an integral lattice in $Z_1(X, \mathbb{R})$, where the latter space inherits its inner product from $C_1(X, \mathbb{R})$. However, $Z_1(X, \mathbb{Z})$ is not in general self-dual, and this gives another description of the lattice L :

Lemma 12. *If X is a finite graph then $L = Z_1(X, \mathbb{Z})^*$.*

Proof. This was shown by Bacher, de la Harpe and Nagnibeda [1, Lemma 1]. □

Using this we can prove:

Theorem 13. *If X is a finite graph without bridges, then its packing fraction is*

$$\frac{|A/Z_1(X, \mathbb{Z})|}{|L/Z_1(X, \mathbb{Z})|} = \frac{|V|}{|T|}$$

where V is the set of vertices of X and T is the set of spanning trees in X . In fact

$$|A/Z_1(X, \mathbb{Z})| = |V|$$

and

$$|L/Z_1(X, \mathbb{Z})| = |T|.$$

Proof. To show that $|A/Z_1(X, \mathbb{Z})| = |V|$, we use several facts. First, A is the set of vertices of the maximal abelian cover of X . Second, $Z_1(X, \mathbb{Z}) \cong H_1(X, \mathbb{Z})$. Third, via this isomorphism, the action of $H_1(X, \mathbb{Z})$ as deck transformations on vertices of the maximal abelian cover is equivalent to the action of $Z_1(X, \mathbb{Z})$ by translations on A . Thus, the quotient $A/Z_1(X, \mathbb{Z})$ is isomorphic to the set of vertices of X .

To show that $|L/Z_1(X, \mathbb{Z})| = |T|$ we use the work of Bacher, de la Harpe and Nagnibeda on integral cuts [1]. Besides the already mentioned inner product on 1-chains, there is an inner product on 0-chains for which the vertices of X form an orthonormal basis. This lets us define the adjoint

$$\partial^*: C_0(X, \mathbb{R}) \rightarrow C_1(X, \mathbb{R}).$$

Concretely, for any vertex x we have

$$\partial^* x = \sum_{e: y \rightarrow x} e$$

where we sum over all edges with target x . It is well known [2] that there is an orthogonal direct sum decomposition

$$C_1(X, \mathbb{R}) = Z_1(X, \mathbb{R}) \oplus N_1(X, \mathbb{R})$$

where

$$N_1(X, \mathbb{R}) = \text{im } \partial^*$$

is called the space of **cuts**. Inside the space of cuts there is a lattice

$$N_1(X, \mathbb{Z}) = N_1(X, \mathbb{R}) \cap C_1(X, \mathbb{Z})$$

called the lattice of **integral cuts**. However, the decomposition of 1-chains into 1-cycles and cuts fails at the integral level:

$$C_1(X, \mathbb{Z}) \neq Z_1(X, \mathbb{Z}) \oplus N_1(X, \mathbb{Z}).$$

This is the key to understanding the quotient $L/Z_1(X, \mathbb{Z})$.

Since $C_1(X, \mathbb{Z})$ is an integral lattice, so are $N_1(X, \mathbb{Z})$ and $Z_1(X, \mathbb{Z})$. If we had $C_1(X, \mathbb{Z}) = Z_1(X, \mathbb{Z}) \oplus N_1(X, \mathbb{Z})$ then $N_1(X, \mathbb{Z})$ and $Z_1(X, \mathbb{Z})$ would be self-dual, but in fact they are usually not. With the help of Kirchhoff's spanning tree theorem, Bacher, de la Harpe and Nagnibeda [1, Prop. 2] prove that

$$|N_1(X, \mathbb{Z})^*/N_1(X, \mathbb{Z})| = |T|.$$

Then they use a very general fact about lattices [1, Lemma 2] to construct isomorphisms

$$\frac{Z_1(X, \mathbb{Z})^*}{Z_1(X, \mathbb{Z})} \cong \frac{C_1(X, \mathbb{Z})}{Z_1(X, \mathbb{Z}) \oplus N_1(X, \mathbb{Z})} \cong \frac{N_1(X, \mathbb{Z})^*}{N_1(X, \mathbb{Z})}.$$

This implies

$$|Z_1(X, \mathbb{Z})^*/Z_1(X, \mathbb{Z})| = |T|.$$

But $Z_1(X, \mathbb{Z})^* = L$ by Lemma 12, so

$$|L/Z_1(X, \mathbb{Z})| = |T|$$

as desired, and thus

$$\frac{|A/Z_1(X, \mathbb{Z})|}{|L/Z_1(X, \mathbb{Z})|} = \frac{|V|}{|T|}. \quad \square$$

The concept of 'cut' has other applications to topological crystals. In the proof above we saw that for any vertex x ,

$$\partial^* x = \sum_{e: y \rightarrow x} e$$

is a cut, so it projects to zero in $Z_1(X, \mathbb{R})$:

$$\sum_{e: y \rightarrow x} \pi(e) = 0.$$

Multiplying by -1 , we obtain

$$\sum_{e: x \rightarrow y} \pi(e) = 0$$

where we sum over all edges with x as source. For any atom $\pi(c_\alpha)$ arising from a path $\alpha: x_0 \rightarrow x$, the bonds coming out of this atom correspond to edges $e: x \rightarrow y$, and they connect it to atoms $\pi(c_{\alpha e}) = \pi(c_\alpha) + \pi(e)$. The above equation thus says that the bonds coming out of any atom give vectors in $Z_1(X, \mathbb{R})$ summing to zero.

This also follows from Kotani and Sunada's characterization of topological crystals in terms of energy minimization [8, 12]. If we think of the bonds as springs, all obeying Hooke's law with the same spring constant, the total force on each atom must vanish in equilibrium, so the bonds must give vectors that sum to zero.


7. CONCLUSIONS

We leave the reader with two questions. First, when are all symmetries of the maximal abelian cover of X covering symmetries? That is, when is $\text{Cov}(X) = \text{Aut}(\overline{X})$?


Second, when is $\text{Cov}(X)$ actually a semidirect product of $\text{Aut}(X)$ and $H_1(X, \mathbb{Z})$? The short exact sequence

$$1 \longrightarrow H_1(X, \mathbb{Z}) \longrightarrow \text{Cov}(X) \xrightarrow{\psi} \text{Aut}(X) \longrightarrow 1$$

splits, making $\text{Cov}(X)$ into a semidirect product, whenever every automorphism of X fixes the basepoint x_0 . It does not split when X is this graph:



So, the short exact sequence does not describe $\text{Cov}(X)$ as a semidirect product in this case. The reason is that while $\text{Aut}(X)$ is the symmetry group of a regular hexagon, this group does not act as symmetries fixing some atom in the corresponding crystal:



Nonetheless, $\text{Cov}(X)$ can still be seen as a semidirect product of $\text{Aut}(X)$ and $H_1(X, \mathbb{Z})$ in this case, because this crystal has hexagonal symmetries about a point that is not an atom.

On the other hand, suppose X is the graph giving the diamond crystal:



In this case $\text{Aut}(X)$ is the symmetry group of a cube, and there is no way to get this group to act as symmetries of the diamond, so there is no way to express $\text{Cov}(X)$ as a semidirect product of $\text{Aut}(X)$ and $H_1(X, \mathbb{Z})$. What general phenomenon is at work here?

Acknowledgments. I thank Greg Egan, without whose help this project could never have been done: many of the key ideas are due to him. I thank David Speyer and Ilya Bogdanov for their help in proving Lemma 6. In particular, Speyer found a proof of Lemma 6 based on the lattice of integral cuts [10], while Bogdanov noticed the importance of the ‘no bridges’ condition and laid the groundwork for the proof here [5]. I thank Pbroks13 and Tom Ruen for the picture of the hexagonal hosohedron, which is available on [Wikicommons](#) under a Creative Commons Attribution-Share Alike 3.0 Unported license. I thank Leshabirukov for the picture of the Petersen graph, which is available on [Wikicommon](#) under the same license. All the other pictures were created by Egan.

REFERENCES

- [1] R. Bacher, P. de la Harpe and T. Nagnibeda, The lattice of integral flows and the lattice of integral cuts on a finite graph, *Bull. Soc. Math. France* **125** (1997), 167–198. Also available at <http://citeseerx.ist.psu.edu/viewdoc/summary?doi=10.1.1.50.6020>.
- [2] N. Biggs, *Algebraic Graph Theory*, Cambridge U. Press, Cambridge, 1993.
- [3] J. H. Conway and N. J. A. Sloane, *Sphere Packings, Lattices and Groups*, Springer, Berlin, 1991.

- [4] H. S. M. Coxeter, On Laves' graph of girth ten, *Canad. Jour. Math.* **7** (1955), 18–23. Also available at <http://cms.math.ca/10.4153/CJM-1955-003-7>.
- [5] I. Bogdanov, Which paths in a graph are orthogonal to all cycles?, *MathOverflow*, 2015. Available at <http://mathoverflow.net/a/238167/2893>.
- [6] J.G. Eon, Euclidian embeddings of periodic nets: definition of a topologically induced complete set of geometric descriptors for crystal structures, *Acta Crystallographica Section A* **67** (2011), 68–86.
- [7] S. T. Hyde, M. O'Keeffe, and D. M. Proserpio, A short history of an elusive yet ubiquitous structure in chemistry, materials, and mathematics, *Angew. Chem. Int. Ed.* **47** (2008), 7996–8000.
- [8] M. Kotani and T. Sunada, Standard realizations of crystal lattices via harmonic maps, *Trans. Amer. Math. Soc.*, **353** (2000), 1–20. Also available at <http://www.ams.org/journals/tran/2001-353-01/S0002-9947-00-02632-5/>.
- [9] F. Laves, Zur Klassifikation der Silikate. Geometrische Untersuchungen möglicher Silicium-Sauerstoff-Verbände als Verknüpfungsmöglichkeiten regulärer Tetraeder, *Zeitschrift für Kristallographie* **82** (1932), 1–14.
- [10] D. Speyer, Which paths in a graph are orthogonal to all cycles?, *MathOverflow*. Available at <http://mathoverflow.net/a/238391/2893>.
- [11] T. Sunada, Crystals that nature might miss creating, *Notices of the American Mathematical Society* **55** (2008), 208–215. Also available at <http://www.ams.org/notices/200802/tx080200208p.pdf>.
- [12] T. Sunada, *Topological Crystallography*, Springer, Berlin, 2012.

DEPARTMENT OF MATHEMATICS, UNIVERSITY OF CALIFORNIA, RIVERSIDE CA 92521, USA, AND CENTRE FOR QUANTUM TECHNOLOGIES, NATIONAL UNIVERSITY OF SINGAPORE, SINGAPORE 117543
E-mail address: `baez@math.ucr.edu`