One of the handouts for Sternberg's course "Representation Theory and Applications" discusses the eigenvalues of the Laplacian on the cube. Here I expand on some of Sternberg's points.

Sternberg's "cubic Laplacian" is really a specialized discrete analog to the usual Laplacian. If we divide space up into a cubic lattice with mesh size ε, then the natural discrete analog to $\partial^2 f/\partial x^2$ is

$$\left(\frac{f(x+\varepsilon)-f(x)}{\varepsilon}-\frac{f(x)-f(x-\varepsilon)}{\varepsilon}\right)\frac{1}{\varepsilon}=\frac{f(x+\varepsilon)+f(x-\varepsilon)-2f(x)}{\varepsilon^2}$$

and so the "discrete Laplacian" at a lattice point **v** is $\dfrac{\sum_{\mathbf{w}\in\text{neighbors}(\mathbf{v})} f(\mathbf{w})-6f(\mathbf{v})}{\varepsilon^2}$. To get this into Sternberg's form, we assume that our function f is periodic on the lattice, with period twice the mesh size. That implies that $f(x+\varepsilon)=f(x-\varepsilon)$ (ditto y and z), so we can restrict our attention to a single cube. Our Laplacian becomes $\left(2\Sigma f(\text{neighbors of }\mathbf{v})-6f(\mathbf{v})\right)/\varepsilon^2$, where now **v** has three neighbors instead of six. Sternberg divides by –6 and ignores ε, getting finally:

$$Lf(\mathbf{v})=f(\mathbf{v})-\tfrac{1}{3}\Sigma f(\text{neighbors of }\mathbf{v})$$

Now $L=I-\frac{1}{3}A$, where A is the adjacency operator: $Af(\mathbf{v})=\Sigma f(\text{neighbors of }\mathbf{v})$. We will eventually see that the eigenvalues of A are –3,–1, +1, +3, so the eigenvalues of L are 0, 2/3, 4/3, 2.

The regular "continuous" Laplacian occurs (among other places) in the classical wave equation:

$$\frac{\partial^2 f}{\partial t^2}=\nabla^2 f$$

whose general solution is a superposition of plane wave solutions $\sin(\mathbf{k}\cdot\mathbf{v}-\omega t)$. Here ω is the so-called angular frequency (i.e., the number of radians per second), and $\mathbf{k}=(k_x,k_y,k_z)$, is the so-called wave-number vector (the number of radians per meter along each axis). The period of the wave is 2π/ω; the wavelength is $2\pi/k$ ($k=|\mathbf{k}|$); finally, the speed of the wave is ω/k. (Note that this has the right units, distance/time.)

Plugging $f(\mathbf{v},t)=\sin(\mathbf{k}\cdot\mathbf{v}-\omega t)$ into the classical wave equation, we get $-\omega^2=-k^2$. In other words, the waves must have unit speed. Notice that $f(\mathbf{v},t)$ is an eigenfunction for the operator $\partial^2/\partial t^2$ with eigenvalue $-\omega^2$, and also an eigenfunction for the Laplacian with eigenvalue $-k^2$.

Turning to the cube, let us now restrict attention to plane waves that are periodic along each axis with period twice the lattice mesh. We'll see below how to set things up so that $f(\mathbf{v},0) = \pm 1$ for each vertex **v**. We expect the eigenvalues to be proportional to k^2. It's not difficult to guess 8 such eigenfunctions, falling into 4 orbits under the action of the rotation group G of the cube. Each orbit gives a representation of G. I illustrate and name these below. Have a look, then we'll pick up the discussion afterwards.

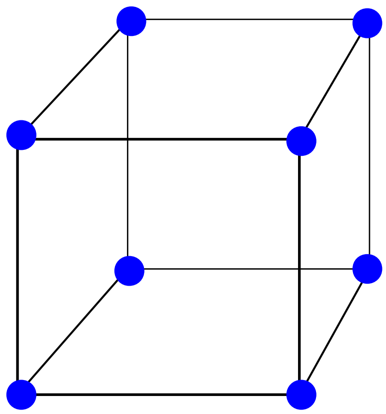

Trivial Representation
$\lambda = 0$

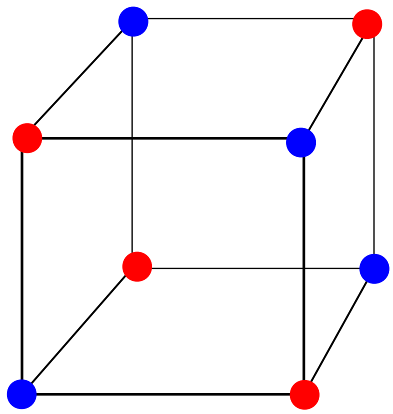

Parity Representation
$\lambda = 2$

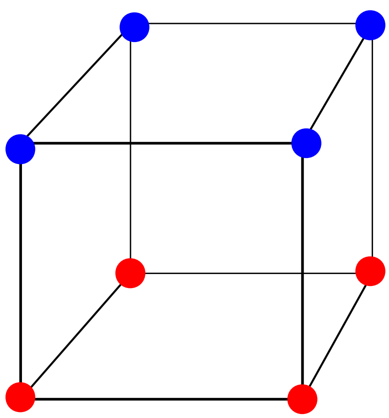

Coordinate Representation
(one of 3 basis elements)
$\lambda = 2/3$

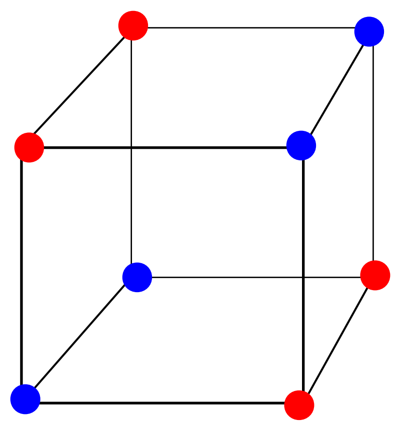

Crosswise Representation
(one of 3 basis elements)
$\lambda = 4/3$

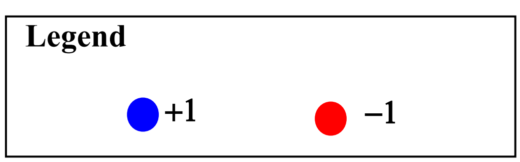

Let's say a word about each of these representations. A couple of notational conventions first, though. I'll drop all references to t from now on, and write just $f(\mathbf{v})$ instead of $f(\mathbf{v},0)$. I'll assume that the center of our cube is at the origin. The cube's edges will have length π, so its vertices are (±π/2, ±π/2, ±π/2). Our waves will be of the form $f(\mathbf{v}) = \sin(\mathbf{k}\cdot\mathbf{v}) = \sin(k_x v_x+k_y v_y+k_z v_z)$, with each k_i either 0 or 1. With this setup, $f(\mathbf{v}) = \pm 1$ at all vertices, as advertised.

Trivial: 1-dimensional, with basis the constant function $f(\mathbf{v}) \equiv 1$. The wavelength is infinite, or in other words $k = 0$. The eigenvalue of Sternberg's L is $\lambda = 0$, as is easily checked.

Coordinate: 3-dimensional. The three basis functions have $\mathbf{k} = (1,0,0)$, $\mathbf{k} = (0,1,0)$, and $\mathbf{k} = (0,0,1)$. We have $k = 1$. We can check that $\lambda = 2/3$, setting the proportionality constant between k^2 and λ: $\lambda = {}^2/_3\, k^2$.

Crosswise: 3-dimensional. The basis functions have $\mathbf{k} = (1,1,0)$, $\mathbf{k} = (1,0,1)$, and $\mathbf{k} = (0,1,1)$. So $k^2 = 2$, and indeed we have $\lambda = {}^2/_3 \cdot 2 = {}^4/_3$.

Parity: 1-dimensional. The basis function has $\mathbf{k} = (1,1,1)$. So $k^2 = 3$, and $\lambda = 2$, as expected.

Now, the rotation group G of the cube is isomorphic to S_4, since the rotations permute the major diagonals, and all 24 permutations are obtained. You can check that the trivial and parity representations, above, correspond to the trivial and parity reps on S_4.

I just made up the names of the other two reps; there don't seem to be standard terms for them. In the coordinate rep, the waves' directions lie along the coordinate axes. In the crosswise rep, the waves' directions lie along the minor diagonals, hence "crosswise".

We now say goodbye to our continuous waves. We might prefer the discrete version as an approximation technique. More likely we would be studying a crystal lattice, where the values of f have physical meaning only at the atoms that occupy the lattice points.

So we keep only the cube and functions defined on the set V of eight vertices. G acts on V, say on the left, and hence also on the 8-dimensional vector space $\mathsf{F}(V)$ of functions on V.

Sternberg starts with this viewpoint. Using power tools, he deduces that the rep of G on $\mathsf{F}(V)$ splits into 4 irreducible representations (irreps for short), with dimensions 1,3,3,1, and that the adjacency operator A: $\mathsf{F}(V) \rightarrow \mathsf{F}(V)$ has eigenvalues –3, –1,1,3 on these irreps. The interest is in showing off the shiny group theory gadgets, of course. In his book, however, he applies the same machinery to the buckyball: the truncated icosahedron or soccer ball, which has 60 vertices and a rotation group isomorphic to A_5. Here it is not so easy to guess the irreps.

Sternberg begins (pp. 6–8) by computing the decomposition of F(*V*) into irreps of S_4. This part is pretty straightforward. On page 9 ("F(*V*) as a subspace of F(*G*)") Sternberg introduces a new idea. Implicitly, he is using an induced representation. Let's see how that works.

In general, if *H* is a subgroup of *G* (both finite), any representation of *H* induces a representation of *G*. Sternberg develops the theory of induced representations in Chapter 3 of his book, and also in handouts 11 and 12 (1260211.pdf and 1260212.pdf).

In general, if the representation of *H* is carried by a space *W*, then the induced representation on *G* is carried by a direct sum of copies of *W*, one for each left coset of *H*. The situation at hand is particularly simple, with the *H* rep being the trivial rep on a 1-dimensional space. So we can identify *W* with the field of scalars (**C**, though it doesn't really matter here what field we use). The induced rep on *G* is then carried by a space of dimension $\#G/\#H$. We can identify this space with the space of functions $G/H \to \mathbf{C}$, i.e., with F(*G*/*H*). We can in turn identify F(*G*/*H*) with a subspace of F(*G*), namely the space of functions which are constant on each left coset of *H*.

H in the case at hand is the isotropy group of some arbitrarily chosen vertex. I will call this vertex the **apex**. So a left coset of *H* is attached to each vertex. If you have a vivid imagination, you can picture the cube with three short whiskers sticking out from each vertex, labeled with the elements of a coset. An element of F(*G*) paints a value (pictured as a color) onto each whisker; an element of F(*V*) paints a value onto each vertex. We map F(*V*) to F(*G*) by taking the value at a vertex and copying it to all three of its whiskers — we do this for all the vertices, of course.

If your imagination could use a little help, just consult the diagram above. (View online, or use a color printer!) Notice the elements of $H = \{1, h, h^2\}$ labeling the whiskers of the apex, and notice how *r*, a vertical 90° rotation, moves these whiskers to another vertex.

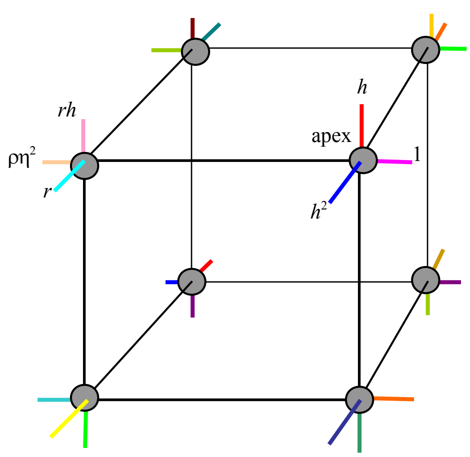

Arbitrary element of F(*G*)

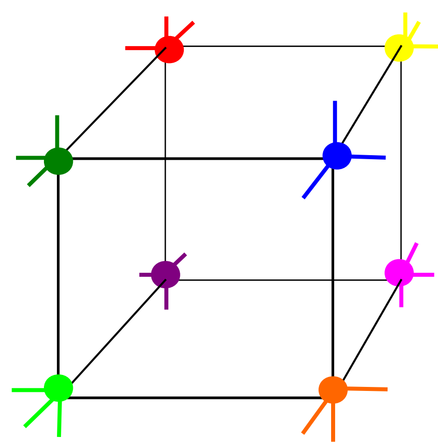

Element from F(V)

$\mathsf{F}(G)$ is really the same as the group algebra $\mathbf{C}[G]$: a function f: $G \to \mathbf{C}$ is the same as the element $\Sigma_g f(g)\, g$. I find this point of view clarifies matters. The left action of G on $\mathsf{F}(G)$ is simply left multiplication $G \times \mathbf{C}[G] \to \mathbf{C}[G]$.

Now suppose $\Sigma_g f(g)\, g$ belongs to $\mathsf{F}(V)$ regarded as a subspace of $\mathbf{C}[G]$. So for any g, $f(gh)$ has the same value for all $h \in H$. Clearly $(\Sigma_g f(g)\, g)\, h = \Sigma_g f(g)\, g$ for any $h \in H$. So H acts trivially on $\mathsf{F}(V)$ when acting on the right.

Also, it is clear that if $\Sigma_g f(g)\, g \in \mathsf{F}(V)$, then so is $g'\Sigma_g f(g)\, g$ for any $g' \in G$. In other words, $\mathsf{F}(V)$ is an invariant subspace of $\mathbf{C}[G]$ under the left action of G on $\mathbf{C}[G]$.

These two facts can be given a geometrical interpretation. The *left* action of G on V rotates the cube; whisker-triplets are carried into whisker-triplets, and this whisker mapping corresponds to the left action of G on G (since there is a 1-1 correspondence between G and the set of all whiskers).

The *right* action of G on G is altogether sneakier, and there is no particularly illuminating geometrical description for it. Since H is not a normal subgroup, this action in general breaks up whisker triplets, and so does not give rise to a well-defined action on V.

However, the right action of H on G can be described geometrically, as was illustrated in the diagram above. We saw the 1 and h whiskers sticking out from the apex. A 120° rotation about the apex carries 1 into h. The r and rh whiskers both stick out of another vertex **v**, and a 120° rotation about **v** carries r into rh. The same is true with any $g \in G$ in place of r. With a sufficiently vivid imagination, you could say that the right action of h "spins each vertex in place".

If we now imagine values painted as colors onto each whisker, we can visualize the left action of G and the right action of H on $\mathbf{C}[G]$.

It's now time to bring the adjacency operator A into the story. Say $f = \Sigma_g f(g)\, g \in \mathsf{F}(V)$. Af replaces the value at a vertex (or at all three of the vertex's whiskers) with the sum of the values at the neighboring vertices. A commutes with any g acting on the left: rotate and then add neighbors, or add neighbors and then rotate — same difference.

The definition of Af does not obviously extend to an arbitrary $f \in \mathbf{C}[G]$, where a whisker-triplet could have three different values painted on it. Nonetheless, Sternberg gives a general argument on page 10 that it is always possible to extend A to a B which commutes with the left action of G and the right action of H. We need not linger over this proof, for on page 12 he tells us what B to use. The six 90° rotations carry each vertex to its three neighbors, two visits per neighbor. Say $\{r_i \mid i = 1\ldots6\}$ are these six rotations. Let

$$(Bf)(g) = \tfrac{1}{2}\sum_{i=1}^{6} f(r_i g)$$

It is easy and worthwhile to check that this B fills the bill.

OK, we have nice geometrical pictures for $\mathbf{C}[G]$, its subspace $\mathsf{F}(V)$, and the operator B whose restriction to $\mathsf{F}(V)$ is A, the adjacency operator. Now we shift gears: we use the decomposition of $\mathbf{C}[G]$ into its irreducible components.

Say $W_1,\ldots,W_k$ are the irreducible representations of G. Then a basic result of representation theory asserts:

$$\mathbf{C}[G] = W_1 \otimes W_1^* \oplus \ldots \oplus W_k \otimes W_k^*$$

We will develop a description of $\mathsf{F}(V)$ and B in terms of this decomposition. Once we possess that, all desired results will fall right out.

I'll start with some generalities before zooming in on our specific problem. No proofs, though you should be able to crank through the details if you care to.

Let W be the space of $n \times 1$ column vectors. Hom(W, W) is canonically isomorphic to the algebra of $n \times n$ matrices. Hom(W, W) can be written as a direct sum of "vertical spaces", symbolically:

$$\begin{bmatrix} * & * \\ * & * \end{bmatrix} = \begin{bmatrix} * & 0 \\ * & 0 \end{bmatrix} \oplus \begin{bmatrix} 0 & * \\ 0 & * \end{bmatrix}$$

and each vertical space is invariant under multiplication on the left. Moreover, each of these vertical spaces is irreducible under this action. If we multiply on the right, then we get an invariant direct sum decomposition into "horizontal spaces":

$$\begin{bmatrix} * & * \\ * & * \end{bmatrix} = \begin{bmatrix} * & * \\ 0 & 0 \end{bmatrix} \oplus \begin{bmatrix} 0 & 0 \\ * & * \end{bmatrix}$$

Hom(W, W) is also canonically isomorphic to $W \otimes W^*$, where W^* is the dual space of $1 \times n$ row vectors. If $\{\mathbf{e}_1, \ldots, \mathbf{e}_n\}$ is the standard basis for W, and $\{\mathbf{e}_1^*, \ldots, \mathbf{e}_n^*\}$ is the dual basis for W^*, then $W \otimes W^*$ has a basis of n^2 elements $\{\mathbf{e}_i \otimes \mathbf{e}_j^* \mid i,j = 1,\ldots,n\}$. As a matrix, $\mathbf{e}_i \otimes \mathbf{e}_j^*$ has a 1 in position (i,j) and zero elsewhere; as a transformation, $\mathbf{e}_i \otimes \mathbf{e}_j^*$ sends $\mathbf{e}_j$ to $\mathbf{e}_i$ and sends all other basis vectors to 0. So $W \otimes \mathbf{e}_j^*$ is the j-th vertical space, and $\mathbf{e}_i \otimes W^*$ is the i-th horizontal space.

Say t is an element of Hom(W,W). Let's multiply elements of Hom(W,W) on the left by t. Under the canonical isomorphism Hom(W,W) $\cong W \otimes W^*$, this amounts to multiplication of the left factor W: $t(\mathbf{v} \otimes \mathbf{w}^*) = (t\mathbf{v}) \otimes \mathbf{w}^*$. Likewise, multiplication on the right by t amounts to multiplication of the right factor W^*: $(\mathbf{v} \otimes \mathbf{w}^*)t = \mathbf{v} \otimes (t^* \mathbf{w}^*)$, where $t^*: W^* \to W^*$ is the adjoint of $t: W \to W$. (This makes it obvious, by the way, that the vertical spaces $W \otimes \mathbf{e}_j^*$ are invariant under left multiplication, and the horizontal spaces $\mathbf{e}_i \otimes W^*$ are invariant under right multiplication.)

Now let's turn to the decomposition $\mathbf{C}[G] = W_1 \otimes W_1^* \oplus \ldots \oplus W_k \otimes W_k^*$. Each element of $\mathbf{C}[G]$ can be regarded as a block diagonal matrix, with one block for each summand. $\mathbf{C}[G]$ has a basis consisting of the elements of G, by definition; we'll call this the G-basis. But $\mathbf{C}[G]$ has another basis consisting of all the $\mathbf{e}_i \otimes \mathbf{e}_j^*$'s from each of the W's; we'll call this the block diagonal basis. The block diagonal matrix for g in G simply expresses g as a linear combination of these $\mathbf{e}_i \otimes \mathbf{e}_j^*$'s.

$G \times G$ acts on $\mathbf{C}[G]$, with the first factor multiplying on the left and the second factor on the right; more precisely, the action is defined by $(g_1, g_2) \cdot x = g_1 x g_2^{-1}$. The block summands are irreducible for this action. However, when we restrict our attention to the subgroup $G \times 1$, i.e., multiply only on the left, these summands split up. Say $W \otimes W^*$ is one of the summands, with $\dim W = n$. As we saw above, $W \otimes W^*$ is the direct sum of n vertical spaces, each invariant for left multiplication by G (or by $\mathbf{C}[G]$). Ditto for right multiplication, *mutatis mutandis*.

Time to zero in on our cube. $\#G = 24$. G has five irreducible representations, which Sternberg lists on page 6: two 1-dimensional, two 3-dimensional, and one 2-dimensional. The arithmetic works out: $24 = 1^2 + 1^2 + 2^2 + 3^2 + 3^2$. So $\mathbf{C}[G]$) in block matrix form looks like this:

*									
	*								
		*	*						
		*	*						
				*	*	*			
				*	*	*			
				*	*	*			
							*	*	*
							*	*	*
							*	*	*

So $\mathbf{C}[G]$ is 24-dimensional, one dimension for each *; under the left action of G it splits into 10 irreps (counting multiplicities), one for each column; ditto for the right action of G, one irrep for each row; and there are 5 distinct irreps, one for each block.

We listed four of the five distinct irreps of G on page 2: the trivial, parity, coordinate, and crosswise reps. The remaining irrep is the 2-dimensional rep. You can get your hands on it by noting that G permutes the three axes of the cube, hence we have a homomorphism from G to S_3; we then have the obvious 2-dimension rep of S_3 where S_3 acts on the vertices of an equilateral triangle.

Our goal now is to see exactly how $\mathsf{F}(V)$ sits inside $\mathbf{C}[G]$ when expressed in block-matrix form. For example, consider the z function from the coordinate irrep, as illustrated on page 2. (On page 2 we didn't know yet about whiskers, so you will have to imagine them, with the colors copied from the vertices.) We know what z looks like in the G-basis: it is the sum of the whiskers on the top of the cube minus the sum of the whiskers on the bottom of the cube. The three whiskers at the apex are 1, h, and h^2, where h is the 120° rotation about the apex. Let r_z be a 90° rotation about the z axis, and let r_x be a 90° rotation about the x axis. A little thought shows that z equals the product $(1-r_x^2)(1+r_z+r_z^2+r_z^3)(1+h+h^2)$. So if we had the matrix representations of h, r_z and r_x^2, we could readily compute the block matrix form for z.

These matrix reps are easy to come by. Let's start with the coordinate irrep, carried by a vector space we'll call W. Let $\{\mathbf{e}_x, \mathbf{e}_y, \mathbf{e}_z\}$ be a basis for W. (Back on page 2, these were functions on the set V of cube vertices. But it doesn't really matter *what* the $\mathbf{e}_i$ are, just how the elements of G act on them.) Clearly h permutes the $\mathbf{e}_i$ cyclically. Also, $r_z\mathbf{e}_x = \mathbf{e}_y$, $r_z\mathbf{e}_y = -\mathbf{e}_x$, and $r_z\mathbf{e}_z = \mathbf{e}_z$. As for r_x^2, we have $r_x^2\mathbf{e}_x = \mathbf{e}_x$, $r_x^2\mathbf{e}_y = -\mathbf{e}_y$, $r_x^2\mathbf{e}_z = -\mathbf{e}_z$. Filling in the coordinate irrep block for these elements is now routine.

The story for the crosswise irrep is similar. Let $\{\mathbf{c}_x, \mathbf{c}_y, \mathbf{c}_z\}$ be the obvious basis; the basis element pictured on page 2 is $\mathbf{c}_y$. Again h permutes the basis elements cyclically, while r_z acts like this: $r_z\mathbf{c}_x = \mathbf{c}_y$, $r_z\mathbf{c}_y = -\mathbf{c}_x$, and $r_z\mathbf{c}_z = -\mathbf{c}_z$. I'll leave r_x^2 as an exercise.

The actions for the trivial and parity irreps are obvious. For the 2-dim rep, h gives a 120° rotation in the plane, while r_z interchanges two vertices; we can make this a reflection in the x-axis.

Putting it all together, and ordering the irreps thusly:
Trivial ⊕ Parity ⊕ 2Dim ⊕ Coordinate ⊕ Crosswise, we get for h and r_z (again leaving r_x^2 as an exercise):

$$h = \left[\begin{array}{c|c|cc|ccc|ccc}
1 & & & & & & & & & \\ \hline
 & 1 & & & & & & & & \\ \hline
 & & -\frac{1}{2} & -\frac{\sqrt{3}}{2} & & & & & & \\
 & & \frac{\sqrt{3}}{2} & -\frac{1}{2} & & & & & & \\ \hline
 & & & & & & 1 & & & \\
 & & & & 1 & & & & & \\
 & & & & & 1 & & & & \\ \hline
 & & & & & & & & & 1 \\
 & & & & & & & 1 & & \\
 & & & & & & & & 1 &
\end{array}\right]
\qquad
r_z = \left[\begin{array}{c|c|cc|ccc|ccc}
1 & & & & & & & & & \\ \hline
 & -1 & & & & & & & & \\ \hline
 & & 1 & & & & & & & \\
 & & & -1 & & & & & & \\ \hline
 & & & & & -1 & & & & \\
 & & & & 1 & & & & & \\
 & & & & & & 1 & & & \\ \hline
 & & & & & & & & -1 & \\
 & & & & & & & 1 & & \\
 & & & & & & & & & -1
\end{array}\right]$$

Or we can be more clever about the bases for the irreps. For the coordinate irrep, let ω be a primitive cube root of unity, and set:

$$\begin{aligned}
\mathbf{e}_1' &= \mathbf{e}_x + \mathbf{e}_y + \mathbf{e}_z \\
\mathbf{e}_2' &= \mathbf{e}_x + \omega^2\mathbf{e}_y + \omega\mathbf{e}_z \\
\mathbf{e}_3' &= \mathbf{e}_x + \omega\mathbf{e}_y + \omega^2\mathbf{e}_z
\end{aligned}$$

With respect to this basis, we have $h\mathbf{e}_1' = \mathbf{e}_1'; h\mathbf{e}_2' = \omega\mathbf{e}_2'; h\mathbf{e}_3' = \omega^2\mathbf{e}_3'$. We can pull the same trick for the crosswise irrep, and we can replace the **x**, **y** basis for the 2-dim irrep with **x**+*i***y**, **x**–*i***y**, so *h* again becomes multiplication by ω or ω². In other words, we can diagonalize *h*, and indeed *H*, illustrating the general fact that the irreps of an abelian group are all one-dimensional. So we get:

$$h$$

$$\left[\begin{array}{c|c|cc|ccc|ccc}
1 & & & & & & & & & \\ \hline
 & 1 & & & & & & & & \\ \hline
 & & \omega & & & & & & & \\
 & & & \omega^2 & & & & & & \\ \hline
 & & & & 1 & & & & & \\
 & & & & & \omega & & & & \\
 & & & & & & \omega^2 & & & \\ \hline
 & & & & & & & 1 & & \\
 & & & & & & & & \omega & \\
 & & & & & & & & & \omega^2
\end{array}\right]$$

As for $z = (1-r_x^2)(1+r_z+r_z^2+r_z^3)(1+h+h^2)$, in the "non-clever" basis we get:

$$(1-r_x^2)(1+r_z+r_z^2+r_z^3) \qquad\qquad z$$

$$\left[\begin{array}{c|c|cc|ccc|ccc}
 & & & & & & & & & \\ \hline
 & & & & & & & & & \\ \hline
 & & & & & & & & & \\
 & & & & & & & & & \\ \hline
 & & & & & & & & & \\
 & & & & & & & & & \\
 & & & & & & 8 & & & \\ \hline
 & & & & & & & & & \\
 & & & & & & & & & \\
 & & & & & & & & &
\end{array}\right]
\left[\begin{array}{c|c|cc|ccc|ccc}
 & & & & & & & & & \\ \hline
 & & & & & & & & & \\ \hline
 & & & & & & & & & \\
 & & & & & & & & & \\ \hline
 & & & & & & & & & \\
 & & & & & & & & & \\
 & & & & 8 & 8 & 8 & & & \\ \hline
 & & & & & & & & & \\
 & & & & & & & & & \\
 & & & & & & & & &
\end{array}\right]$$

I've given the partial product $(1-r_x^2)(1+r_z+r_z^2+r_z^3)$ as well, to show how it's just $8\mathbf{e}_z\otimes\mathbf{e}_z^*$.

OK, what does this tell us about $\mathsf{F}(V)$ in block matrix form?

We saw earlier that h acts trivially on $\mathsf{F}(V)$ when acting on the right. A little thought shows that this can serve as a definition of $\mathsf{F}(V)$: $\mathsf{F}(V) = \{x \in \mathbf{C}[G] \mid xh = x\}$. Acting on the left is a different story. For example, for the coordinate irrep, three basis elements are z, $x = hz$, and $y = h^2z$.

In the clever basis, h is diagonal. In general, if D is a diagonal matrix and A is any matrix, then DA multiplies the i-th row of A by d_{ii}, while AD multiplies the i-th column of A by d_{ii}. So we see immediately how $\mathsf{F}(V)$ fits into $\mathbf{C}[G]$ in block diagonal form, when the clever basis is used. Writing the multipliers atop each column, and marking elements of $\mathsf{F}(V)$ with •, we get this:

1	1	ω	ω^2	1	ω	ω^2	1	ω	ω^2
•									
	•								
		*	*						
		*	*						
				•	*	*			
				•	*	*			
				•	*	*			
							•	*	*
							•	*	*
							•	*	*

So $\mathsf{F}(V)$ is 8-dimensional, containing 4 irreps of G, all distinct.

It's also interesting to compare the left and right actions of H on $\mathsf{F}(V)$ in the non-clever basis. In the coordinate irrep, h is a permutation matrix. In general, if P is a permutation matrix and A is any matrix, then PA permutes the rows of A, while AP permutes the columns. We've already computed the matrix for z, and as noted above, $x = hz$ and $y = h^2z$. So:

$$x = \begin{bmatrix} 8 & 8 & 8 \\ & & \\ & & \end{bmatrix} \quad y = \begin{bmatrix} & & \\ 8 & 8 & 8 \\ & & \end{bmatrix} \quad z = \begin{bmatrix} & & \\ & & \\ 8 & 8 & 8 \end{bmatrix}$$

So we see why H, acting on the right, acts trivially on $\mathsf{F}(V)$.

Notice also that x, y, and z are zero outside the coordinate block. This illustrates the general fact that distinct irreps are orthogonal (Schur's lemma).

It's also worth looking at $1 + h + h^2$ in the clever basis:

$$\left[\begin{array}{c|c|cc|ccc|ccc}
3 & & & & & & & & & \\ \hline
 & 3 & & & & & & & & \\ \hline
 & & 0 & & & & & & & \\
 & & & 0 & & & & & & \\ \hline
 & & & & 3 & & & & & \\
 & & & & & 0 & & & & \\
 & & & & & & 0 & & & \\ \hline
 & & & & & & & 3 & & \\
 & & & & & & & & 0 & \\
 & & & & & & & & & 0
\end{array}\right]$$

In other words, the map $x \mapsto x(1+h+h^2)$ is just three times a projection map from $\mathbf{C}[G]$ to $\mathsf{F}(V)$.

As we mentioned a moment ago, $\mathsf{F}(V)$ contains 4 irreps of G, each once. Sternberg computes this fact twice. First on page 6, he does the usual "inner product of characters" calculation; for example, $(\varphi,\chi_{\text{coord}}) = 1$, so the coordinate irrep appears once. This makes no use of the embedding of $\mathsf{F}(V)$ in $\mathbf{C}[G]$.

On page 14, however, he uses a different argument: he says, "We want to know how many times the trivial representation of H occurs in an irreducible representation of S_4." As we've seen, each block corresponds to a direct summand $W\otimes W^*$, and the right action of G amounts to the action of G on W^*. In matrix terms, we are acting with G on a horizontal subspace. This subspace is irreducible for the action of G, but it splits into irreducible components under the action of H. For example, a row of the coordinate block splits into three one-dimensional spaces, with reps $h\mathbf{v} = 1\mathbf{v}$, $h\mathbf{v} = \omega\mathbf{v}$, $h\mathbf{v} = \omega^2\mathbf{v}$. We know this because we've just spent several pages working out the details of all the irreps. But to count how often the trivial irrep of H occurs, it's enough to take the inner product of characters: the character of the trivial H irrep with the G character of coordinate irrep restricted to H.

Time to bring our Laplacian back into the picture. Recall that we expressed L on $\mathsf{F}(V)$ as $I - \frac{1}{3}A$, and then we extended A (defined only for $\mathsf{F}(V)$) to B, defined on all of $\mathbf{C}[G]$:

$$Bg = \tfrac{1}{2}\sum_{i=1}^{6} r_i g$$

where $\{r_i \mid i = 1\ldots6\}$ are the six 90° rotations about the coordinate axes. Before we wrote $Bf(g)$, but that's when we were thinking of $\mathbf{C}[G]$ as $\mathsf{F}(G)$, functions on G. We see now that B is nothing more than $\frac{1}{2}\sum_{i=1}^{6} r_i$, i.e., one-half the sum of the elements in a complete conjugacy class.

B commutes with the left action of G on $\mathbf{C}[G]$, so by Schur's lemma, B restricted to the j-th irrep is $\lambda_j I$. Taking traces, if n_j is the dimension of the j-th irrep:

$$\mathrm{tr}(B|W_j) = \tfrac{1}{2}\sum_{i=1}^{6} \mathrm{tr}(r_i|W_j) = n_j\lambda_j$$

But the traces of all the r_i's are the same, since they are all conjugate, and this trace is just the χ_j character of r_i:

$$n_j\lambda_j = 3\chi_j(r_i)$$

So we can read off the eigenvalues from the character table for S_4, without ever working out the actual irreps. As we discovered earlier, by a more laborious but highly instructive process, the answers are –3, –1, +1, +3.