Some Notes on Kauffman's Graph Invariant

Notes by : L. R. Taylor

Fall 2002

The goal here is to flesh out the material in Flapan's description of one of Kauffman's graph invariants. Let G be an abstract graph, a collection of vertices and edges. The goal is to study embeddings of G in three space and in particular to produce invariants to distinguish different embeddings. We are particularly interested in distinguishing an embedding from its *mirror image*. The mirror image of an embedding $h: G \to \mathbf{R}^3$ is just the composition $r \circ h$ where $r(x, y, z) = (x, y, -z)$ from $\mathbf{R}^3$ to itself.

Recall that we say that two embeddings h_1, $h_2: G \to \mathbf{R}^3$ are equivalent provided that there is an automorphism of the graph, $\theta: G \to G$ such that the two embeddings $h_1 \circ \theta$ and h_2 are ambient isotopic: i.e. there is an isotopy of three space from the identity to H, such that $h_1 \circ \theta$ is equal to $H \circ h_2$.

We typically represent h by a picture or *regular projection* which is a collection of points and arcs in the plane. The points are the images of the vertices and are all distinct. The arcs intersect the vertices only at their endpoints. Each arc is immersed and both the self–crossings and the crossings with the other arcs are transverse. At each crossing we indicate an over or an under crossing just as we did with knots and links. Here is a simple example in which the arcs are actually embedded.

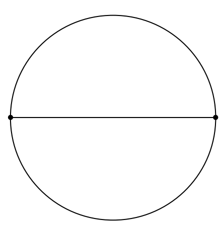

This a graph with two vertices and three edges, the so–called θ–curve. Sometimes one wants to only consider graphs with at most one edge between any two vertices. This graph clearly lacks this property, but this lack is usually easily rectified by adding a few additional vertices to divide the edges. An algorithm to do this is to add one new vertex for each edge at the middle of the edge. One can often get by with adding fewer vertices, but an algorithm is sometimes useful. Here are four pictures, the algorithm and three ways to achieve the property by adding two vertices.

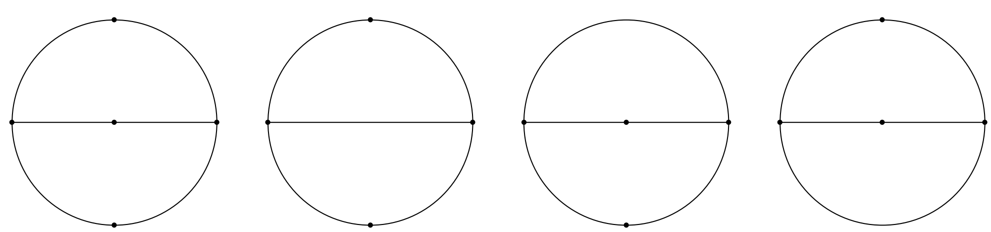

The algorithm is especially useful when issues of the automorphism group of the graph are important since the automorphism group of the original graph is always a subgroup of the automorphism group of the algorithmic graph, but not necessarily of the others.

Next we show a non–planar embedding of G together with its mirror image.

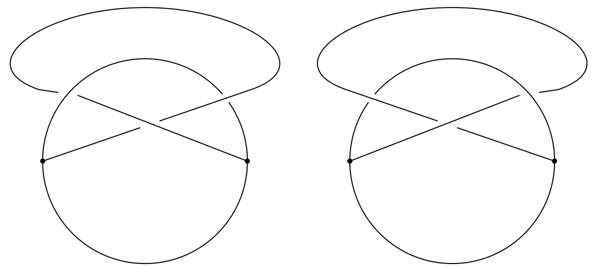

G_+ $\quad$ G_-

We suspect that these two embeddings are different and also different from the first embedding of G we showed above, but we need techniques to prove this. People have defined polynomial invariants analogous to the HOMFLY polynomial we talked about earlier, but Kauffman has an approach that allows us to use our link material directly.

We begin with the abstract graph G, and construct a set of subgraphs, $m_T(G)$. Before beginning the construction proper, we pause to discuss an operation on graphs that we will need shortly. Given a graph G, $\overline{G}$ is the subgraph obtained as follows. Delete all edges incidence to a vertex of valence 1. This may produce new edges incident to a vertex of valence 1 so repeat if necessary. Since the number of edges decreases as long as there are edges incident to a vertex of valence 1, this process must stop when there are no longer any vertices of valence 1. Then delete all the vertices of valence 0. It may very well be that $\overline{G} = \emptyset$, but if it is not, note that all the vertices of $\overline{G}$ are of valence ≥ 2.

The graph $\overline{G}$ has another description: it is *the* maximal subgraph of G all of whose vertices have valence ≥ 2. To see this, note first that the union of two subgraphs of G, all of whose vertices have valence ≥ 2, also has this property since valence can not decrease in taking the union. This shows that there is a unique maximal subgraph of G with this property and it contains every subgraph with the property. Temporarily, let H be the maximal subgraph of G all of whose vertices have valence ≥ 2, or equivalently the union of all the subgraphs with this property. Note $\overline{G} \subset H \subset G$. Now we get from G to $\overline{G}$ be a sequence of subgraphs, $G_0 = G$, G_1, $\ldots$, $G_r = \overline{G}$, where we pass from G_i to G_{i+1} either by removing an edge with a vertex of valence 1 or by removing an isolated vertex. If $K \subset G$ is a subgraph with all vertices of valence ≥ 2, note that $K \subset G_i$ for all i since neither operation removes a vertex or an edge of K. Hence $H \subset \overline{G}$ and we are done. We call this procedure *taking the maximal high valence subgraph*. We also call any subgraph all of whose vertices have valence ≥ 2 a *high valence subgraph*. In the special case of a subgraph all of whose vertices have valence precisely 2, we say the subgraph is a *sublink*.

Begin by numbering the vertices and let vertex i have valence v_i. It will increase the efficiency of our procedure if we number so that $v_i \geq v_{i+1}$: i.e. smaller integers mean larger valence. In any case, begin at the smallest vertex s whose valence is greater than 2. If there is no such vertex, our collection consists of the single element $\overline{G}$, which is a sublink.

Otherwise let e_1 , $\ldots$, e_{n_s} denote the edges incident to vertex s in some order. There are $\binom{n_s}{2}$ ways to pick two of the e_i's. Define a collection of subgraphs of G indexed by e_i, e_j for $i < j$ by first removing the open arcs which are the edges e_k for $k \neq i$ and $k \neq j$ to get $G'_{e_{i,1}, e_{j,1}}$ and then defining $G_{e_i,e_j} = \overline{G'}_{e_i,e_j}$. Notice that each G_{e_i,e_j} has at least one fewer vertex of valence > 2 than G has.

Replace each G_{e_i,e_j} in the collection obtained by applying the above procedure to each of the G_{e_i,e_j} and repeat until all the elements in the collection are sublinks. Denote this final collection by $m'(G)$. Let $m_T(G)$ be the collection of the maximal subgraphs in $m'(G)$. More explicitly, if $K \in m'(G)$, then there is at least one element $L \in m_T(G)$ with $K \subset L$, and if K and $H \in m'(G)$ with $K \subset H$ but $H \neq K$, then K is not in $m_T(G)$.

We claim $m_T(G)$ is the collection of maximal sublinks. To see this, first note that any element of $m'(G)$ is a subgraph of at least one of the elements of $m_T(G)$. Next let $H \subset G$ be a sublink. We want to show H is a sublink of at least one element of $m_T(G)$. Suppose by induction that H is a sublink of K, one of the elements obtained in an intermediate stage of the construction. (It is certainly in G, the initial stage of the construction.) To get the elements in the next stage coming from K, we pick a vertex of high valence; remove edges incident to that vertex in all possible ways so as to have valence 2; and then taking the maximal high valence subgraph of each subgraph. If that vertex is not in H, then none of the deleted edges lie in H and since H is a high valence subgraph, it lies in the maximal high valence subgraph: i.e. H is a sublink of all the graphs obtained. In the other case, the vertex under consideration is in H. In this case there is a unique way to remove edges so that the two edges of H incident to this vertex remain in the subgraph. It follows that H is in this subgraph and hence in the maximal high valence subgraph associated to this graph: i.e. H is in exactly one of the subgraphs obtained. Hence H is in at least one of the subgraphs obtained at the end of the construction and therefore H is a sublink of at least one element of $m_T(G)$.

Now let $K \subset G$ be a maximal sublink. Then there exists $H \in m_T(G)$ with $K \subset H$. But every element of $m_T(G)$ is a sublink so $K = H$. Conversely, if $K \in m_T(G)$, K must be maximal, since if H is a sublink of G with $K \subset H$, then there is an element $L \in m_T(G)$ with $H \subset L$, so $K \subset L$ are both in $m_T(G)$ and therefore $K = L$.

In general, $m_T(G)$ is smaller than the collection considered by Kauffman (and hence by Flapan). Consider the "dumbbell" graph:

Kauffman's collection has three subgraphs: the two circles and the union of the two circles. The collection you get by applying our procedure before taking the maximal ones, $m'(G)$, has two elements: the union of the two circles and the right hand circle. The collection $m_T(G)$ has one element: the disjoint union of the two circles.

You can probably get Kauffman's collection by calculating $m'(G)$ for all the different ways to number the vertices, although we shall not bother to investigate this. Actually of course, we only need to number the vertices of valence > 2 to produce our collection so this way of getting Kauffman's collection is probably at least as efficient as Kauffman's original procedure. However, we would argue that $m_T(G)$ is as good at detecting different embeddings as the full collection considered by Kauffman.

To get an invariant of an embedding $h: G \to \mathbf{R}^3$, apply h to each of the elements in $m_T(G)$. This gives a collection of links in $\mathbf{R}^3$ labelled by the elements of $m_T(G)$. If $K \in m_T(G)$, let $\mathrm{Link}_{K,h} \subset \mathbf{R}^3$ denote the link obtained by applying h to the graph $K \subset G$. Given a regular projection for h, it is easy to see a regular projection for $\mathrm{Link}_{K,h}$: just erase all the edges and vertices not in K. Let $m_T(G,h)$ denote the collection of labelled links obtained and let $m_T^o(G,h)$ denote the collection of links with oriented components.

If two embeddings, h_1 and h_2, are ambient isotopic, Link_{K,h_1} and Link_{K,h_2} are ambient isotopic links for each $K \in m_T(G)$. The need to consider automorphisms of the graph causes a wrinkle. The set of automorphisms of the graph forms a group, $\mathrm{Aut}(G)$. Since an automorphism permutes the vertices and the edges and only the identity fixes all the edges and all the vertices, there is an injective homomorphism

$$\mathrm{Aut}(G) \hookrightarrow \Sigma_v \times \Sigma_e$$

where Σ_v denotes the permutation group on the vertices and Σ_e denotes the permutation group on the edges. If there is at most one edge between any two vetices, then the composition $\mathrm{Aut}(G) \hookrightarrow \Sigma_v \times \Sigma_e \to \Sigma_v$ is still an injection. For the complete graph on n vertices, it is an isomorphism. For the θ–curve above, the map $\mathrm{Aut}(G) \hookrightarrow \Sigma_v \times \Sigma_e$ is an isomorphism.

If there are no isolated vertices, we might hope that the composition $\mathrm{Aut}(G) \hookrightarrow \Sigma_v \times \Sigma_e \to \Sigma_e$ is an injection, but the example of the θ–curve shows this is false. The problem is that there are two essentially different ways to map an edge to an edge. For an isolated edge, this problem is insuperable without knowledge of which vertex goes where, and even if we glue several edges together at a pair of vertices (like the θ–curve) we will need vertex information. *But*, if the graph has no isolated vertices or isolated edges and if there is at most one edge between any two vetices, then the composition $\mathrm{Aut}(G) \hookrightarrow \Sigma_v \times \Sigma_e \to \Sigma_e$ is injective.

An automorphism θ of the graph induces a permutation of elements of $m_T(G)$, say K goes to $\theta[K]$ and an explicit isomorphism of graphs, $\theta_K: K \to \theta[K]$. You need to check three things. The first is that an automorphism of a graph takes any subgraph to a (possibly different) subgraph. The second is that an automorphism preserves valence: if v has valence n, then $\theta(v)$ must also have valence n. It follows that an automorphism takes subgraphs of high valence to subgraphs of high valence and sublinks to sublinks. The third thing to check is that if $K_1 \subset K_2$, then $\theta(K_1) \subset \theta(K_2)$, so that an automorphism takes maximal sublinks to maximal sublinks.

Hence if h_1 and $h_2 \circ \theta$ are ambient isotopic, the links $\mathrm{Link}_{K,,h_1}$ and $\mathrm{Link}_{\theta[K],h_2}$ are equivalent and these are the links which are easy to see from the pictures for h_1 and h_2. The function $K \mapsto \theta_K$ gives a homomorphism

$$\Psi: \mathrm{Aut}(G) \to \Sigma_{m_T(G)}$$

where $\mathrm{Aut}(G)$ denote the automorphism group of G and $\Sigma_{m_T(G)}$ denotes the permutation group on the elements of $m_T(G)$.

One would now like to apply our previous techniques for links to these particular links. Unfortunately, most of our earlier work applies to oriented links. As an example, the left and right Hopf links are different as oriented links but the same as unoriented ones. We say this by observing that the two HOMFLY polynomials are different. Since the mirror image of the right Hopf link is the left Hopf link, this will be an annoyance in studying chirality questions.

As Flapan remarks, sometimes the chemistry orients the sublinks for you. A generalization of this remark is that sometimes the chemistry of the compound restricts the automorphisms. If some of the vertices are carbon and others are silicon, then no chemically realizable automorphism can exchange a silicon vertex with a carbon vertex. Hence in our study of equivalent embedded graphs, we can insist that only chemically realizable automorphisms are permitted.

Anyway, even abstractly we can proceed as follows. Each element in $m_T(G)$ is a link. Label the components and remember that each component has two possible orientations. Let $m_T^o(G)$ denote the collection consisting of the elements of $m_T(G)$ labeled and oriented in all possible ways. Given an embedding $h\colon G \to \mathbf{R}^3$, let $m_T^o(G, h)$ denote the labelled collection of oriented links obtained by applying h to each element in $m_T^o(G)$. Given two ambient isotopic embeddings of G, the associated links are ambient isotopic as oriented links, so $m_T^o(G, h)$ really only depends on the ambient isotopy class of h.

Passing from $m_T(G)$ to $m_T^o(G)$ greatly increases the number of elements in the collection. Even if all the elements of $m_T(G)$ are knots, $m_T^o(G)$ has twice as many elements. Each k component link turns into $2^k \cdot k!$ objects. Its virtue is that it remains canonical. Given an automorphism of the graph, we get an automorphism of $m_T(G)$ and since we have an explicit isomorphism $K \to \theta[K]$, θ induces an automorphism of $m_T^o(G)$ as well. We can describe the map induced on $m_T^o(G)$ as follows. Send K with an orientation and label to $\theta[K]$ with the unique orientations and labels so that the map $K \to \theta[K]$ preserves labels and the orientation on each component. Once you know the effect of θ on one K with labels and orientations, calculating θ on the other labels and orientations of K is easy. You have already calculated which component of K goes to which component of $\theta[K]$ and when you switch an orientation on a component of K, switch the orientation on the corresponding component of $\theta[K]$. There is a function $m_T^o(G) \to m_T(G)$ which comes from forgetting the orientations. Let $\mathrm{Aut}\big(m_T^o(G) \to m_T(G)\big)$ denote the automorphisms of the set $m_T^o(G)$ which induce automorphisms of $m_T(G)$. Check that there is a homomorphism

$$\Psi^o\colon \mathrm{Aut}(G) \to \mathrm{Aut}\big(m_T^o(G) \to m_T(G)\big)$$

To give one of our subinks, it suffices to give its edges. To orient a sublink it suffices to orient its components. Suppose the vertices and the edges are ordered. There is usually no natural way to do this so just do it. Each component of a sublink is a circle, so the order of the edges in it is fixed once a starting edge and an orientation are chosen. The elements of $m_T^o(G)$ can be given as ordered lists of edges, where each component begins with an edge incident to the "smallest" vertex in the component and is oriented so that we travel from this vertex along the initial edge. Each component has a "preferred" orientation given by going in the direction of the smallest edge out of the smallest vertex. The "preferred" is in quotes since it depends on the chosen order: it is not canonical. We can also label the components by ordering them, so we may speak of the first, second and so on. Since no two components of a sublink can share an edge, we can label the components of a sublink by just giving the smallest edge in the component.

So elements of $m_T(G)$ are just sets of edges while elements of $m^o_T(G)$ are ordered tuples of elements with an ordering on the subtuples that make up the components of the sublink. There is a "preferred" representative of each element in $m_T(G)$ up in $m^o_T(G)$ given by taking the "preferred" orientation on each component and the lexicographical ordering on the components. If there is at most one edge between two vertices, then just listing the edges in the order they occur around the circle beginning with the smallest will orient the circle by just starting at the vertex of the first edge which is not incident to the second. There is a "preferred" orientation in this scheme as well.

If the elements of $m^o_T(G)$, or $m_T(G)$, are written like this, the effect of θ can be worked out from the permutation induced on the edges and vertices. A more compact way to describe $m^o_T(G)$ is as triples. The first element in each triple is a "preferred" element. If this sublink has k components, the second element in the triple is a tuple of length k, $(\pm, \cdots, \pm)$ (the *orientation tuple*), and the third element in the triple is an element of the symmetric group on k elements, τ. This triple represents an element in $m^o_T(G)$ by first orienting the components using the $\pm$–tuple: if the i^{th} entry is a $+$, use the orientation on the i^{th} component coming from the "preferred" orientation: otherwise, use the opposite orientation. Then permute the components using the permutation from their "preferred" order. Hence to describe how θ acts on $m^o_T(G)$ it is enough to apply θ to the "preferred" elements in $m^o_T(G)$ and just give the new orientation tuple and permutation, so in practice on ends up working with lists the size of $m_T(G)$.

Let us see how this technique works in practice on the θ–curve, where the numbers order the nearby vertices and the lower case letters order the edges. Here we are NOT in the case of at most one edge between two vertices.

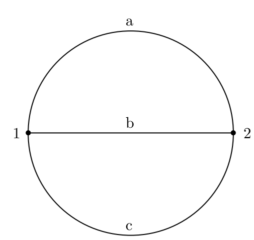

Look at vertex 1 and remove one edge. If we remove a, we get the graph $\{b, c\}$; if we remove b, we get the graph $\{a, c\}$; if we remove c, we get the graph $\{a, b\}$. Each of these is already a sublink so we are done. The set $m^o_T(G)$ is $\big\{(b, c), (c, b), (a, c), (c, a), (a, b), (b, a)\big\}$. The ordered set (b, c) means the circle and orientation starting at vertex 1, going out along b to vertex 2 and then coming back along c. Note that (c, b), starting at 1, going out along c and coming back along b, gives the same circle but with the opposite orientation. Since each sublink has a single component it is just as efficient to list all the elements in $m^o_T(G)$ as to use the more compact notation, but for completeness, it would be written

$$\Big\{ \big((b,c),(+),\text{id}\big), \big((b,c),(-),\text{id}\big), \big((a,c),(+),\text{id}\big), \big((a,c),(-),\text{id}\big), \big((a,b),(+),\text{id}\big), \big((a,b),(-),\text{id}\big) \Big\} .$$

The "preferred" representatives are (b, c), (a, c) and (a, b).

The group of automorphisms of θ was worked out above. There is an automorphism which takes each edge to itself but permutes the two vertices. The symmetric group on three letters permutes the edges and leaves the vertices fixed. The permutation that switches the vertices commutes with all the permutations that swap the edges. In other words, the natural map $\mathrm{Aut}(G) \to \Sigma_v \times \Sigma_e$ is an isomorphism.

Permuting the two vertices takes each element of $m_T(G)$ to itself, although it reverses the orientation on each knot: i.e. it switches (b, c) and (c, b) and has a similar effect on the other two pairs. A permutation (automorphism) of the edges acts on $m_T^o(G)$ by apply the permutation to the edges in the ordered pairs. As an example, the permutation τ which takes a to b, b to c and c to a takes (a, b) to (b, c), (b, c) to (c, a), etc. The homomorphism Ψ^0 above is an isomorphism for this graph.

Now recall our other embeddings

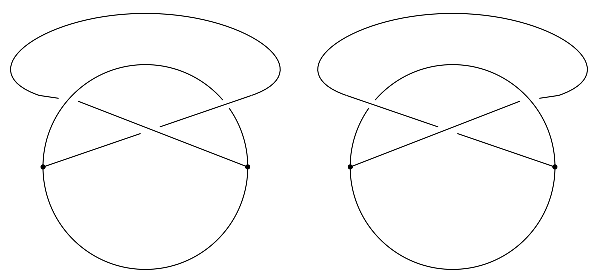

G$_+$ G$_-$

The set of labeled links (in this case knots) for G_+ is

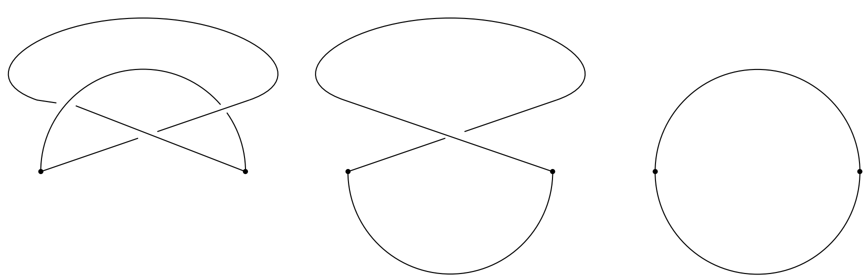

{a,b} {b,c} {a,c}

Note $\{a, c\}$ and $\{b, c\}$ are unknotted and $\{b, c\}$ is a right handed trefoil. It follows that G_+ is not equivalent to our first embedding of the θ–curve. If we consider G_- instead, we get a left handed trefoil for $\{b, c\}$ so the embedding and unknots for the other two. Hence G_- is yet a third embedding.

We can also use $m_T^o(G)$ to restrict which automorphisms of the graph can be realized by ambient isotopies of $\mathbf{R}^3$. In other words, $m_T^o(G)$ gives information on the question for which θ is h ambient isotopic to $h \circ \theta$?

This subset is actually a subgroup of $\mathrm{Aut}(G)$ which we will denote $\mathrm{Aut}_h(G)$ even though it only depends on the ambient isotopy class of h. To see this, recall that if h is ambient isotopic to $h\circ\theta$ then there is an isotopy of three space H_θ such that $H_\theta(z,0)=z$ and $H_\theta\big(h(x),1\big)=h\circ\theta(x)$. If θ_1 and θ_2 are in $\mathrm{Aut}_h(G)$, there are isotopies H_{θ_i} as above. Consider $H_{\theta_1}\big(h(x),1\big)=h\circ\theta_1(x)$ and substitute $\theta_2(x)$ for x to get $H_{\theta_1}\big(h(\theta_2(x)),1\big)=h\circ\theta_1(\theta_2(x))$, which says $H_{\theta_1}\big(h(\theta_2(x)),1\big)=h\circ(\theta_1\circ\theta_2)(x)\big)$. But $H_{\theta_2}\big(h(x),1\big)=h\circ\theta_2(x)$, so $H_{\theta_1}\big(h(\theta_2(x)),1\big)=H_{\theta_1}\big(H_{\theta_2}\big(h(x),1\big),1\big)=h\circ(\theta_1\circ\theta_2)(x)\big)$. Define $H_{\theta_1\circ\theta_2}(x,t)=H_{\theta_1}(H_{\theta_2}(x,t),t)$, so $H_{\theta_1\circ\theta_2}\big(h(x),1\big)=h\circ(\theta_1\circ\theta_2)(x)$. Check that $H_{\theta_1\circ\theta_2}(z,0)=z$ and notice that for a fixed t $H_{\theta_1\circ\theta_2}(z,t)$ is the composition of the diffeomorphisms $H_{\theta_1}(z,t)$ and $H_{\theta_2}(z,t)$. This shows $\theta_1\circ\theta_2\in\mathrm{Aut}_h(G)$. Check that the identity isotopy gives the required equation for the identity automorphism of the graph. Finally check that the isotopy H_θ^{-1} gives the required equation for the automorphism θ^{-1}. We have checked $\mathrm{Aut}_h(G)$ is a subgroup.

The same ideas show that for any $\theta\in\mathrm{Aut}(G)$

$$\mathrm{Aut}_{h\circ\theta}(G)=\theta^{-1}\mathrm{Aut}_h(G)\theta$$

as subgroups of $\mathrm{Aut}(G)$. This follows from the observation that if h and $h\circ\theta_1$ are ambient isotopic, so are $h\circ\theta$ and $h\circ(\theta_1\circ\theta)$ for all $\theta\in\mathrm{Aut}(G)$. Hence $h\circ\theta$ and $(h\circ\theta)\circ(\theta^{-1}\circ\theta_1\circ\theta)$ are ambient isotopic, so $\mathrm{Aut}_{h\circ\theta}(G)\supset\theta^{-1}\mathrm{Aut}_h(G)\theta$. But $\mathrm{Aut}_h(H)=\theta\circ\big(\theta^{-1}\mathrm{Aut}_h(G)\theta\big)\circ\theta^{-1}\subset\theta\circ\big(\mathrm{Aut}_{h\circ\theta}(G)\big)\circ\theta^{-1}\subset\mathrm{Aut}_{h\circ\theta\circ\theta^{-1}}(G)=\mathrm{Aut}_h(G)$ so the inclusion is an equality.

The mirror image of h is the embedding $r\circ h\colon G\to\mathbf{R}^3\ \to\mathbf{R}^3$ where $r\colon\mathbf{R}^3\to\mathbf{R}^3$ is $r(x,y,z)=(x,y,-z)$. We show

$$\mathrm{Aut}_h(G)=\mathrm{Aut}_{r\circ h}(G)\ .$$

Suppose $\theta\in\mathrm{Aut}_h(G)$, or $H_\theta\big(h(x),1\big)=h\circ\theta(x)$. Then $r\circ H_\theta\big(h(x),1\big)=(r\circ h)\circ\theta(x)$. Define $\hat{H}_\theta(z,t)=r\circ H_\theta(r(z),t)$ and note $\hat{H}_\theta(z,0)=z$. Check $\hat{H}_\theta\big((r\circ h)(x),1\big)=(r\circ h)\circ\theta(x)$, so $\theta\in\mathrm{Aut}_{r\circ h}(G)$, showing $\mathrm{Aut}_h(G)\subset\mathrm{Aut}_{r\circ h}(G)$. Since $r\circ r$ is the identity, this proves the desired equality.

Next suppose h is achiral, and let θ_- be any automorphism such that $(r\circ h)$ is ambient isotopic to $h\circ\theta_-$. In other words, there is an ambient isotopy of three space H so that $H\big(r\circ h(x),t\big)=(h\circ\theta_-)(x)$. Hence for any $\theta\in\mathrm{Aut}(G)$, $H\big(r\circ h(\theta(x)),t\big)=(h\circ\theta_-)(\theta(x))$, or $H\big(r\circ(h\circ\theta)(x)),t\big)=(h\circ(\theta_-\circ\theta))(x)$.

First note that this implies that if h is achiral, so is $h\circ\theta$ for any $\theta\in\mathrm{Aut}(G)$ and an automorphism which displays the achirality is $\theta^{-1}\circ\theta_-\circ\theta$.

Let θ'_- be any automorphism such that $h\circ\theta'_-$ is ambient isotopic to $r\circ h$: $\theta'_-=\theta_-$ is one possibility, but there may be others. Let H' be the promised ambient isotopy. With $\theta=\theta'_-$, we get $H\big(r\circ(h\circ\theta'_-)(x)),t\big)=h\circ(\theta_-\circ\theta'_-)(x)$, or $H\big(r\circ H'\big(r\circ h(x),t\big),t\big)=h\circ(\theta_-\circ\theta'_-)(x)$. Define $K(z,t)=H\big(r\circ H'\big(r(z),t\big),t\big)$ so $K\big(h(x),t\big)=h\circ(\theta_-\circ\theta'_-)(x)$. For $t=0$, K is the identity and for a fixed t, K is the composition of $H(r(\_\ ),t)$ with $H'(r(\_\ ),t)$. Since $H(r(\_\ ),t)$ and $H'(r(\_\ ),t)$ are diffeomorphisms for each fixed t, so is K. Hence $\theta_-\circ\theta'_-\in\mathrm{Aut}_h(G)$. It follows that the set of θ such that $h\circ\theta$ is ambient isotopic to h or

$r \circ h$ is also a subgroup, $\mathrm{Aut}^{\pm}_h(G)$, and $\mathrm{Aut}_h(G) \subset \mathrm{Aut}^{\pm}_h(G)$ is a subgroup of index 1 or 2. It is entirely possible that θ_- itself is in $\mathrm{Aut}_h(G)$, as for example when h is planar when we may take θ_- to be the identity. In other words, it is possible that $\mathrm{Aut}_h(G) = \mathrm{Aut}^{\pm}_h(G)$.

Theorem 1: Let h be achiral. The groups $\mathrm{Aut}_h(G)$ and $\mathrm{Aut}^{\pm}_h(G)$ are equal if and only if h and its mirror image are ambient isotopic. If $m^o_T(G,h)$ contains a chiral element then $\mathrm{Aut}^{\pm}_h(G)/\mathrm{Aut}_h(G) = \mathbf{Z}/2\mathbf{Z}$.

Remark: If G is chiral, it is immediate that $\mathrm{Aut}_h(G) = \mathrm{Aut}^{\pm}_h(G)$.

Proof: If h and its mirror image are ambient isotopic, then θ_- can be taken to be the identity and the two subgroups are equal. If the two subgroups are equal and if h is achiral, there is an automorphism, θ_- displaying the achirality: i.e. $r \circ h$ is ambient isotopic to $h \circ \theta_-$. But since $\theta_- \in \mathrm{Aut}_h(G)$, $h \circ \theta_-$ is ambient isotopic to h. This proves the first part. If $m^o_T(G,h)$ contains a chiral element, then $r \circ h$ and h are not ambient isotopic. Since we are assuming h is achiral the two subgroups can not be equal. ∎

As Flapan points out later in the book, we have a way to construct new embeddings of G as soon as we have one embedding h for which $\mathrm{Aut}_h(G) \neq \mathrm{Aut}(G)$. If θ is a representative for a left coset $\mathrm{Aut}_h(G)\backslash\mathrm{Aut}(G)$ which is not the coset of the identity, then $h \circ \theta$ is not ambient isotopic to h. More generally,

Theorem 2: Given two automorphisms θ_1 and θ_2, $h \circ \theta_1$ is ambient isotopic to $h \circ \theta_2$ if and only if θ_1 and θ_2 lie in the same coset of $\mathrm{Aut}_h(G)\backslash\mathrm{Aut}(G)$.

Proof: If $h \circ \theta_1$ is ambient isotopic to $h \circ \theta_2$ then $h = h \circ \theta_1 \circ \theta_1^{-1}$ is ambient isotopic to $h \circ \theta_2 \circ \theta_1^{-1}$, so $\theta_2 \circ \theta_1^{-1} \in \mathrm{Aut}_h(G)$, so θ_1 and θ_2 are in the same coset of $\mathrm{Aut}_h(G)\backslash\mathrm{Aut}(G)$. Conversely, if $\theta_2 = \theta \circ \theta_1$ for some $\theta \in \mathrm{Aut}_h(G)$, then $h \circ \theta_2 = h \circ \theta \circ \theta_1$ is ambient isotopic to $h \circ \theta_1$. ∎

Remark: This coset space is what Flapan calls the set of *topological sterioisomers* (page 155). It is worth recalling that any diffeomorphism (or homeomorphism) that preserves orientation is ambient isotopic to the identity.

For our first embedding of the θ–curve, all automorphisms can be realized by ambient isotopies. For the embeddings $G_{\pm}$, our techniques say nothing about the automorphism which just switches the vertices, and indeed, these automorphisms can be realized. If the vertices are fixed, then the only permutation of the edges which *might* be realizable is the permutation which switches a with b and leaves c fixed. The easiest proof I know that this permutation *is* realizable is to tie the embedded graph out of string and put it on the table so as to represent both pictures. Note this really is a proof, not just highly suggestive.

By tying lots of knots in the edges of G we can produce embeddings h for which $\mathrm{Aut}_h(G)$ is very small. If there is at most one edge between any two vertices, then by tying a different knot in each edge, we can make $\mathrm{Aut}_h(G)$ the trivial group. Hence there are no non–trivial "universal" automorphisms of a such graph. There are as we shall see later (Flapan p. 158) some elements $\theta \in \mathrm{Aut}(G)$ which are in no $\mathrm{Aut}_h(G)$. Let $\mathcal{N}(G)$ be this set. Note $\mathcal{N}(G)$ is closed under conjugation: if $\theta \in \mathcal{N}(G)$ and $\nu \in \mathrm{Aut}(G)$, then

$\nu \circ \theta \circ \nu^{-1} \in \mathcal{N}(G)$. Additionally, if $x \in \mathrm{Aut}(G)$ satisfies $x^n \in \mathcal{N}(G)$, then $x \in \mathcal{N}(G)$. A corollary of this is that if $x \in \mathcal{N}(G)$, $x^r \in \mathcal{N}(G)$ for all r relatively prime to the order of x. Thus if $\mathcal{N}(G) \neq \emptyset$, it tends to contain quite a bit.

Next we work out an example for which there is a multi–component link. It is the example from class. See Figure 1 below.

The top line is the graph and the next four lines are the elements of $m_T(G)$. The $\pm$ inside each square explains the orientation. All the sublinks are planar as drawn and any circle in the plane can be oriented by travelling so that the inside is on your left. The sign is $+$ if this orientation and the "preferred" one agree, the sign is $-$ if they do not. The elements of $m_T(G)$ will be referred to by the capital letter on their line. Hence the unique 2 component sublink in $m_T(G)$ will be called sublink D). For now ignore the dotted axes in the picture of the graph.

There are four vertices of valence 3 and four of valence 2. The four of valence 4 span a square. One maximal sublink has 4 vertices, two have 6 and one has 8.

Automorphisms (of practically anything) can be confusing at first. We know abstractly that for a graph with at most one edge between two vertices (such as this one) an automorphism is determined by its action on the vertices. Or if we prefer, on a graph with no isolated vertices, isolated edges and with at most one edge between two vertices (such as this one), an automorphism is determined by its action on the edges. The issue is that not all permutations of the edges actually give automorphisms of the graph. To check if a particular permutation does give an automorphism, try to relabel the graph using the new edges and see that no contradictions of the vertex relations occur. (*Vertex relations* mean things like a and d share a vertex but a and c do not and that both vertices of a have valence 2.) Hence vertex relations will only permit $c \mapsto c$ or $c \mapsto g$ and similarly for g. For valence reasons, c can only be mapped to itself or g or k or l. Suppose it were possible to map c to k. Then k would have to go to c or to g. Suppose k goes c. Then d would have to map to itself and then a would have to map to itself and finally b would have to map to itself. But then k and b would have to share a common vertex and they don't. Similarly, k can not map to g. A similar argument shows c can not map to l. If $c \mapsto c$, then vertex relations also force $k \mapsto k$ or $k \mapsto l$ and similarly for l, since k can not map to b or d since both b and d are incident to a vertex of valence 2 and k is not. If $c \mapsto c$ and $k \mapsto k$ then the automorphism is the identity.

We can also use the induced action on the elements of $m_T(G)$ or $m_T^o(G)$ to help determine the automorphism group. In our case, C) is the only element of $m_T(G)$ with 4 vertices so it must be invariant and we get a homomorphism $\mathrm{Aut}(G) \to D_8$, where D_8 is the dihedral group of order 8 or, more relevantly, the symmetry group of the square. The valence relation example in the last paragraph implies that $\mathrm{Aut}(G)$ is a subgroup of D_8. The valence relation example also shows that the 90° rotation of the square in D_8 is not in the image of $\mathrm{Aut}(G)$. Both reflections and the 180° rotation of the square do extend to automorphisms of the entire graph, so $\mathrm{Aut}(G) = \mathbf{Z}/2\mathbf{Z} \oplus \mathbf{Z}/2\mathbf{Z} \subset D_8$. Using the dotted lines as a set of axes, we can describe the four elements in $\mathrm{Aut}(G)$ as e, the identity; r_x, reflection in the x–axis, r_y, reflection in the y–axis; and R_{180}, rotation about the origin through 180° degrees. Each of the non–identity elements has order 2; all the elements commute with each other and $r_x \circ r_y = R_{180}$.

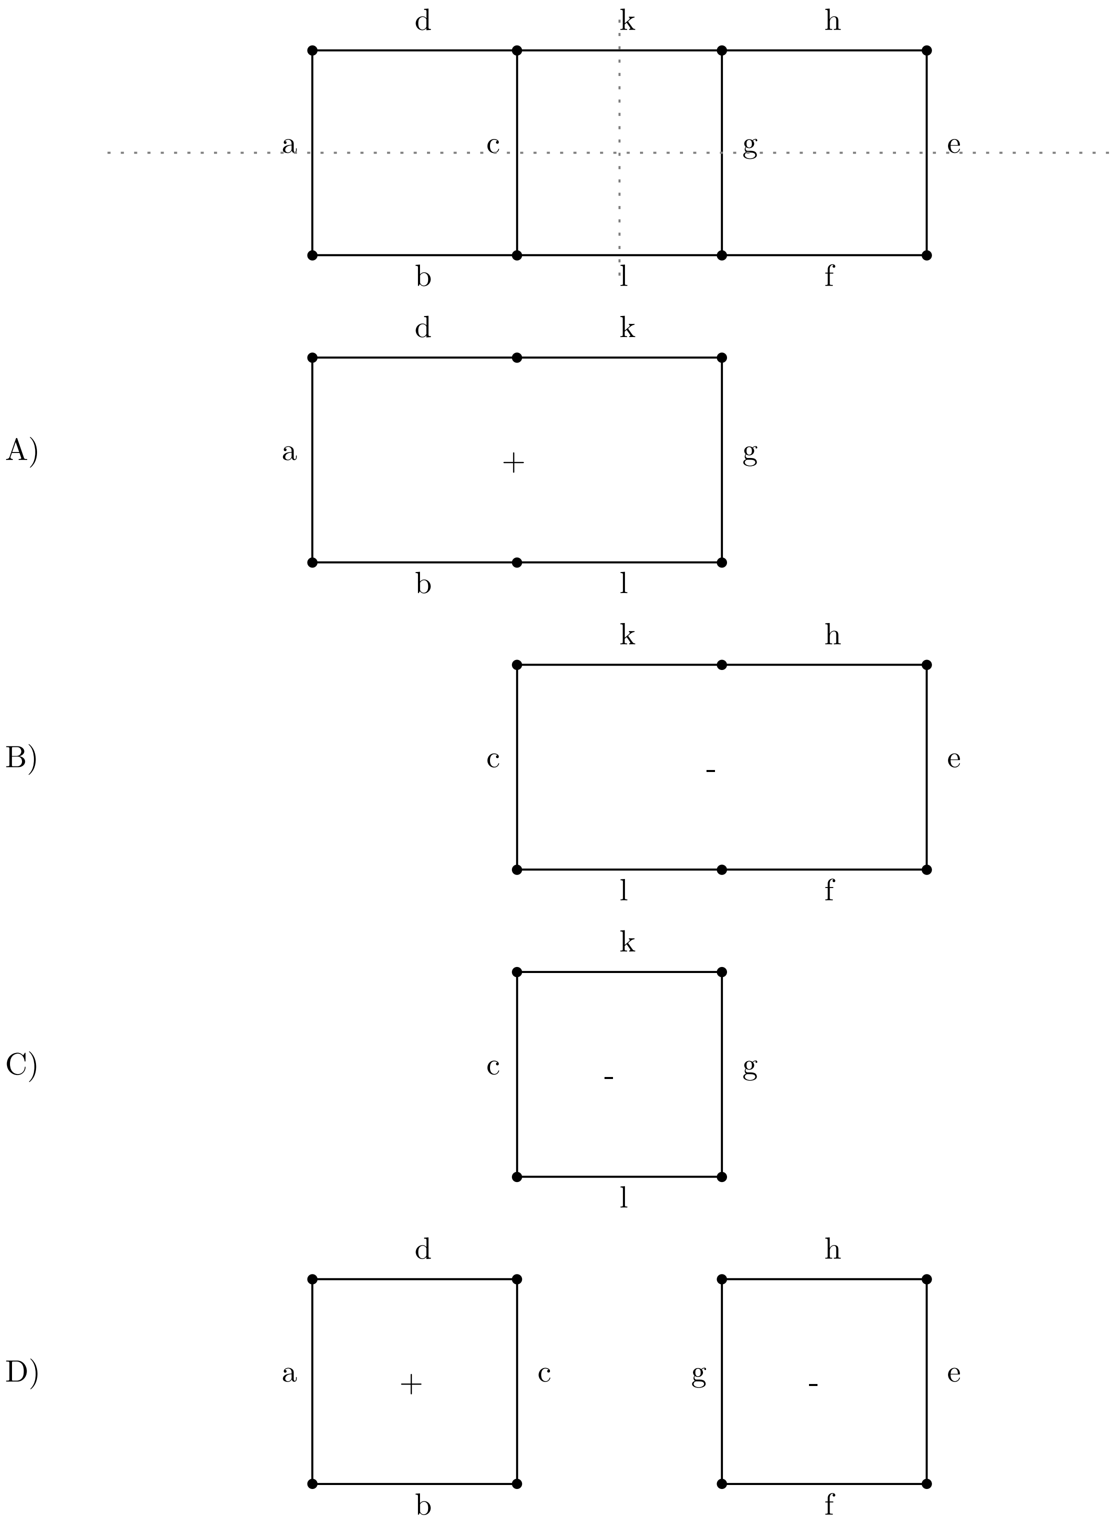

Figure 1

Let us work out how these automorphism act on the 2 component link D) in $m_T^o(G)$.

Rotation by 180° interchanges the two components as does r_y; r_x takes each component to itself. As for orientations, we compute

r_x takes each component to itself and reverses orientation in each of them;
r_y interchanges the components but preserves both orientations;
R_{180} interchanges the components and reverses both orientations.

Now look at a different embedding of G side by side with the mirror image of that embedding.

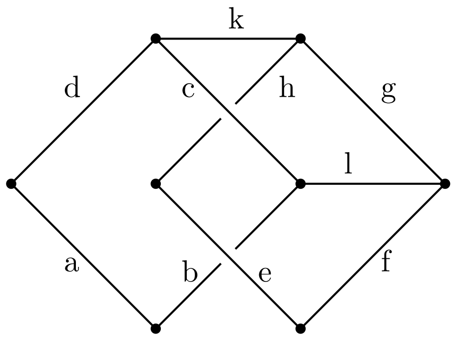

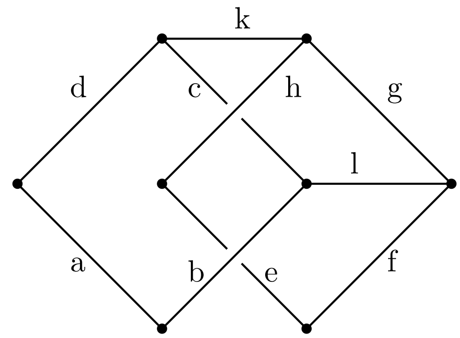

Are these two embeddings equivalent? Or equivalently, are each of these embeddings chiral? If we let h denote the embedding on the left in the picture, then we are inquiring whether h and $r \circ h$ are equivalent, or whether there is an automorphism θ such that $h \circ \theta$ and $r \circ h$ are ambient isotopic.

Since every sublink will have a regular projection with at most two crossing, all the subknots in $m_T^o(G)$ are unknots. However, sublink D) is a labelled, oriented Hopf link. On the left, the linking number of the two components is $+2$, on the right -2. Recall that the linking number is unchanged if we switch the order of the components or if we reverse both orientations. Hence the linking number of D) after applying $h \circ \theta$ for any $\theta \in \mathrm{Aut}(G)$ remains $+1$. Hence these embeddings are chiral.

Here is a different projection of the embedding h.

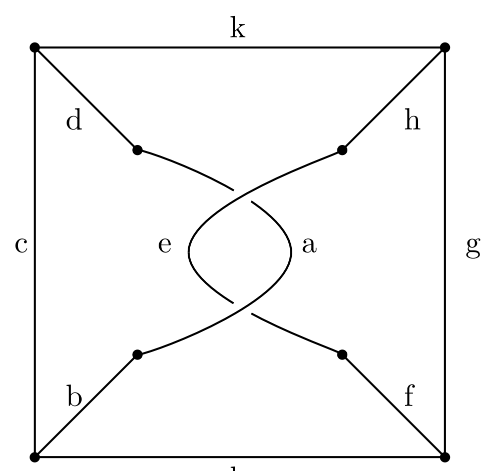

Figure 2

From this projection it is easy to see lots of symmetry. Consider the following two step construction.

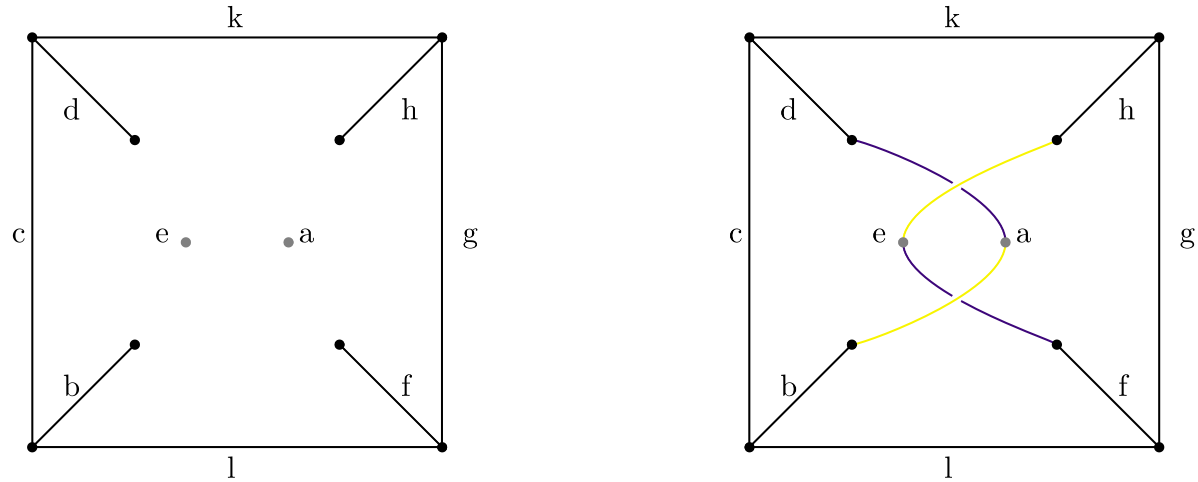

First draw the left hand picture in the xy–plane. The x–axis is horizontal and the y–axis is vertical: the positive z–axis points straight up from the page. The two gray dots labelled a and e are temporary to help in step 2: in particular, they are not vertices. Draw the picture so that rotation by 180° degrees about any of the three axes is a symmetry. As an aside, note that the full symmetry group of the square almost acts on this picture. If it were not for those pesky gray dots rotation by 90° around the z–axis would be a symmetry as well.

Anyway, now add the yellow and purple arcs between the indicated vertices and the gray dots. Make the yellow arcs lie in the upper z half space except at their end points. Add the purple arcs similarly except make them go down. If you do this carefully, rotation about the three axes remain symmetries of the embedding. Now just erase the gray dots and change the colored arcs to black to get the projection of h in Figure 2.

With these preliminaries, rotation about the three axes induce automorphisms of the graph. Specifically, rotation about the z axis is R_{180}; rotation about the y–axis is r_x; and rotation about the x–axis is r_y. In particular, $\mathrm{Aut}_h(G) = \mathrm{Aut}(G)$.

This gives another proof of the chirality of h by using Theorem 1. The oriented Hopf link is chiral and occurs in $m^o_T(G, h)$. By Theorem 1, if h is chiral, $\mathrm{Aut}^{\pm}_h(G)/\mathrm{Aut}_h(G) = \mathbf{Z}/2\mathbf{Z}$. But since $\mathrm{Aut}^{\pm}_h(G) \subset \mathrm{Aut}(G)$, $\mathrm{Aut}^{\pm}_h(G) = \mathrm{Aut}_h(G)$.

Homework:

1) Apply the procedure describe above to the graph G below to calculate $m_T(G)$. Use the ordering of the vertices and edges given. Describe the sublinks you get by just listing the edges in each sublink.

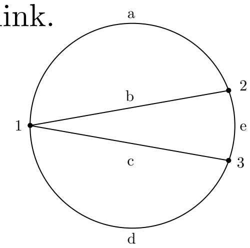

2) Consider the embedding of G given here.

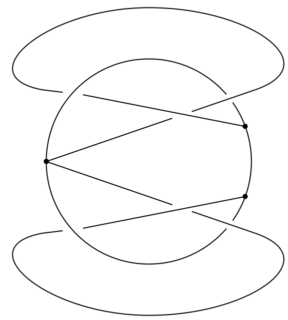

   a) Draw the mirror image of this embedding, and show this embedding is achiral.
   b) Draw the set of labelled links, $\mathrm{Link}_{K,h}$ for $K \in m_T(G)$, where h is the embedding shown. Show that several of these are chiral. How can this be reconciled with a)?

3) What is the maximal high valence subgraph of the graph below? Argue that the embedding of G given below is chiral.