Synthetic W-Graphs and Cells
Atlas of Lie Groups AIM Workshop V
16–21 July 2007

John R. Stembridge ⟨jrs@umich.edu⟩

Department of Mathematics
University of Michigan
Ann Arbor, Michigan 48109–1043

Contents

1. Definitions and Background

Let (W, S) be a Coxeter system, and $\mathcal{H} = \mathcal{H}(W, S)$ the corresponding Iwahori-Hecke algebra over $\mathbb{Z}[q^{\pm 1/2}]$. For convenience, we let $s_1, \dots, s_n$ denote the simple reflections (i.e., $S = \{s_1, \dots, s_n\}$), and $T_1, \dots, T_n$ the corresponding generators of $\mathcal{H}$. Recall that the defining relations of $\mathcal{H}$ are the quadratic relations $(T_i - q)(T_i + 1) = 0$ and the braid relations.

We define an *S-labeled graph* to be a triple $\Gamma = (V, m, \tau)$, where

(i) V is a vertex set,
(ii) m is a map $V \times V \to \{\text{scalars}\}$ (i.e., a matrix),
(iii) τ is a map $V \to \{\text{subsets of } S\}$.

The value of τ at a vertex v is referred to as the τ *invariant* of v.

For reasons that will soon be clear, we will always assume that the scalars in (ii) are nonnegative integers, but we have been deliberately vague here in recognition of the fact that there do exist contexts where more flexibility is necessary.

The matrix m implicitly defines a directed graph on the vertex set V; to emphasize this interpretation, we will use the notation $m(u \to v)$ to refer to the (u, v)-entry of m. One should regard this as the number of edges directed from vertex u to vertex v.

This work has been supported by NSF Grant DMS–0554278.

Let M_Γ denote the $\mathbb{Z}[q^{\pm 1/2}]$-module freely generated by V.

A *W-graph* is an S-labeled graph Γ such that the following defines an $\mathcal{H}$-module structure on M_Γ:

$$T_i(v) = \begin{cases} qv & \text{if } i \notin \tau(v), \\ -v + q^{1/2} \sum_{u: i \notin \tau(u)} m(v \to u)u & \text{if } i \in \tau(v). \end{cases} \tag{1}$$

It is easy to check that the quadratic relations $(T_i - q)(T_i + 1) = 0$ hold automatically in this context (presuming of course that the sums are finite); the content of the definition is that the braid relations hold as well.

Remark 1.1. (a) The matrices defined in (1) are transposed from the ones used in [**KL**] as well as in the `atlas` software. However, it is easy to see that the transpose of any matrix representation of $\mathcal{H}$ is again a matrix representation of $\mathcal{H}$, so the difference is harmless. The reason we bother to be different is that it seems more natural to have an $\mathcal{H}$-module in which T_i acts by propagating in the *forward* direction (i.e., following edges), rather than the reverse. Another way to fix this problem would be to reverse the conventions on edge orientation in the `atlas` documentation.

(b) If we set $q = 1$, the $\mathcal{H}$-action on M_Γ specializes to a W-action, so we can think of M_Γ as an $\mathcal{H}$-module as well as a W-module.

(c) There exist S-labeled graphs such that the operators defined in (1) do not satisfy the braid relations for generic q, but do satisfy them at $q = 1$ (thus yielding a W-module). For further information, see Remark 3.15(b).

It is important to note that if $\tau(v) \subseteq \tau(u)$, then the value of $m(v \to u)$ has no effect on the operators defined in (1). For this reason, we adopt the convention that

$$m(v \to u) = 0 \quad \text{whenever} \quad \tau(v) \subseteq \tau(u) \tag{2}$$

for all W-graphs. The `atlas` also follows this convention.

Remark 1.2. (a) Many of the W-graphs that occur in nature are symmetric; i.e., $m(u \to v) = m(v \to u)$. In light of our convention, we should modify this by saying that Γ has *symmetric edge weights* if

$$m(u \to v) = m(v \to u) \quad \text{whenever} \quad \tau(u) \perp \tau(v),$$

where $I \perp J$ means that I and J are incomparable (i.e., $I \not\subseteq J$ and $J \not\subseteq I$).

(b) Given $J \subset S$ and a W-graph Γ, there is an obvious way to define a W_J-graph Γ_J by restricting the τ invariants to J; i.e., replacing $\tau(v)$ with $J \cap \tau(v)$. Given our convention, this may also entail deleting edges (replacing certain matrix entries with 0's). We refer to Γ_J as a *parabolic restriction* of Γ.

A. Cells and subquotients.

Let $\Gamma = (V, m, \tau)$ be a W-graph. When does $U \subset V$ span an $\mathcal{H}$-submodule of M_Γ?

It is not hard to show that it is necessary and sufficient that U is an "arrow-closed" subset of V; i.e., for all $u \in U$ and $v \in V$, $m(u \to v) \neq 0$ implies $v \in U$. (This relies on (2) in an essential way, and justifies our use of this convention.)

Given that U spans an $\mathcal{H}$-submodule, it follows that the S-labeled subgraphs $\Gamma(U)$ and $\Gamma(V - U)$ induced by U and $V - U$ (with multiplicities and τ invariants inherited from Γ) are themselves W-graphs, and

$$M_{\Gamma(V-U)} \cong M_\Gamma / M_{\Gamma(U)}.$$

More generally, given a nested sequence $U_1 \subset U_2 \subset V$ of arrow-closed subsets, we call the W-graph $\Gamma(U_2 - U_1)$ a *subquotient* of Γ, the point being that the corresponding $\mathcal{H}$-module is a subquotient of M_Γ.

If Γ has no proper subquotients, then it is called a *cell.*

We remark that a subset U of V is the vertex set of a subquotient of Γ if and only if it is "convex" in the sense that for all $u, u' \in U$, every $v \in V$ that occurs along a directed path from u to u' also belongs to U.

Now define an equivalence relation on V by declaring $u \sim v$ if there are directed paths in Γ from u to v and v to u. In the graph theory literature, the equivalence classes are known as the *strongly connected components* of the graph Γ. It is clear that these equivalence classes are convex in the above sense, and moreover, these are the (unique) smallest subquotients of Γ: if $u \sim v$, then u and v must appear together or not at all in every subquotient of Γ.

Thus Γ is a cell if and only if Γ is strongly connected.

Note that for every W-graph Γ, the $\mathcal{H}$-module M_Γ has a "composition series" of $\mathcal{H}$-submodules in which the intermediate quotients are (isomorphic to) the modules $M_{\Gamma(U)}$, where U ranges over the strongly connected components of Γ.

Remark 1.3. The $\mathcal{H}$-modules corresponding to cells need not be irreducible. This persists for the W-modules one obtains at $q = 1$ (even for finite Weyl groups, where complete reducibility of W-modules over $\mathbb{Q}$ is available).

B. Duality.

The *dual* of an S-labeled graph $\Gamma = (V, m, \tau)$ is the S-labeled graph $\Gamma^* = (V, m^*, \tau^*)$, where $m^*(u \to v) := m(v \to u)$ and $\tau^*(v) := S - \tau(v)$ (i.e., reverse all edges and complement all τ invariants).

We have not checked in detail whether it is true that the dual of every W-graph is also a W-graph. A cautionary indication that this does require checking is that if we only require the braid relations to hold for Γ at $q = 1$, then these relations need not hold for Γ^*. The example discussed in Remark 3.15(b) has this property.

We could fix this defect by negating the edge weights of Γ^*. Indeed, if we modify the definition of Γ^* by setting $m^*(u \to v) := -m(v \to u)$, then it is not hard to show that the action of T_i on M_{Γ^*} can be obtained by twisting the action of T_i of M_Γ by the automorphism of $\mathcal{H}$ that sends $T_i \to -q^{-1}T_i$ and $q \to q^{-1}$. At $q = 1$, this corresponds to twisting by the sign representation of W. However, this takes us outside of the class of W-graphs with nonnegative edge weights.

On the other hand, many of the W-graphs that occur in nature are bipartite; i.e., the vertices may be partitioned into two sets, say V_0 and V_1, such that $m(u \to v) \neq 0$ only if $(u, v) \in (V_0 \times V_1) \cup (V_1 \times V_0)$. Given such a bipartition, the diagonal change of basis $v \mapsto \pm v$ (with sign depending only on membership of v in V_1) has the effect of negating the edge weights of Γ. Thus $\bar{\Gamma} = (V, -m, \tau)$ is also W-graph, and $M_{\bar{\Gamma}} \cong M_\Gamma$ (as $\mathcal{H}$-modules).

It follows that if Γ is a bipartite W-graph, then so is Γ^*.

C. The Kazhdan-Lusztig-Vogan story.

The (one-sided) W-graphs constructed by Kazhdan and Lusztig in [**KL**] are defined by taking $V = W$, $\tau(v) = \{i : \ell(s_i v) < \ell(v)\}$, and

$$m(u \to v) := \mu(u, v) + \mu(v, u),$$

where $\mu(u, v)$ is defined to be the coefficient of $q^{(\ell(v)-\ell(u)-1)/2}$ in the Kazhdan-Lusztig polynomial $P_{u,v}(q)$. In particular, either $\mu(u, v) = 0$ or $\mu(v, u) = 0$ (often both).

Note that the above definition should be modified to fit (2); Kazhdan and Lusztig do not follow our convention (and their W-graphs are not directed). In any case, after modification, these W-graphs will be edge-symmetric in the sense of Remark 1.2.

Assuming (W, S) is finite or crystallographic, one knows that $P_{u,v}(q)$ has nonnegative integer coefficients, so these W-graphs have nonnegative integer edge weights.

Since Kazhdan-Lusztig polynomials are polynomials in q (rather than $q^{1/2}$), it follows that $\mu(u, v)$ and $m(u \to v)$ can be nonzero only if the lengths of u and v have opposite parity. Therefore, these W-graphs are bipartite.

For each real Lie group G whose complex form has Weyl group W, Vogan has constructed more general (but similar) W-graphs on each block of irreducible representations of G; one block occurs for each choice of a real form for the dual group. The edge weights for the W-graph of a block are obtained from leading coefficients of certain Kazhdan-Lusztig-Vogan polynomials. Again, one knows that these polynomials have nonnegative integer coefficients, and the corresponding W-graphs are bipartite.

See also [**LV**] and the notes for Barbasch's lecture at Atlas IV [**B**].

The `atlas` command for describing the contents of a block is `block`; the command `wgraph` describes its W-graph, and `wcells` describes its partition into cells.

Assuming that W is a finite Weyl group, we define a W-graph to be *natural* if it occurs as a subquotient of one of the above block W-graphs, or a parabolic restriction of such a graph. This includes the Kazhdan-Lusztig W-graph as a special case by taking G to be a complex semisimple Lie group, viewed as a real Lie group.

2. Goals/Problems

Our main goal is to understand the structure of natural W-graphs.

WORKING HYPOTHESIS 2.1. *The class of natural W-graphs should be fairly rigid.*

It may be unrealistic to ask for an explicit description of all natural W-graphs, or even natural cells, but a direct combinatorial algorithm for generating cells (not involving the computation of Kazhdan-Lusztig polynomials) may be within reach.

As evidence in support of the rigidity hypothesis, we have run the `atlas` on most (all?) real forms of E_6, and obtained only 20 distinct natural cells up to isomorphism. For E_7, the analogous number we obtained is 42. It will be interesting to add complex E_6 to the picture, as well as the real forms of E_8.

QUESTION 2.2 (Algebraic versus Combinatorial Isomorphism). *Is it true that natural W-cells are isomorphic if and only if the corresponding W-modules are isomorphic?*

A positive answer to this question would strongly support our Working Hypothesis. It is also worth noting that the modules generated by the 20 isomorphism classes of $W(E_6)$-cells and 42 $W(E_7)$-cells we have identified are all distinct.

GOAL 2.3: CELL SYNTHESIS. *Given some fragment of a cell or W-graph (possibly void), construct all ways to complete it within an appropriate class of W-graphs.*

It is interesting to speculate that all natural cells occur in the blocks of split real groups (but a recent email from Vogan indicates that this fails for E_8). If true, this could be a big win computationally, since it is far easier to compute the blocks and cells of split groups by "honest" methods (i.e., `atlas`) than it is do the same for (say) the complex groups. In any case, it would be very interesting if we could determine the cells or full W-graph for complex E_8 by synthetic methods, since the direct calculation of all Kazhdan-Lusztig polynomials for complex E_8 is not feasible without major advances in the theory.

QUESTION 2.4 (Compressibility). *Given the W-graph Γ of a block or cell, understand how to specify the reconstruction of Γ from a parabolic restriction Γ_J for some $J \subset S$.*

For cells, the W-modules M_Γ tend to have a very small number of irreducible constituents, and one can usually choose J so that branching from W to W_J is multiplicity-free, so it is reasonable to speculate that natural cells have very short recursive descriptions.

GOAL 2.5. *Can we explain the counterexample of McLarnan and Warrington without using Kazhdan-Lusztig polynomials?*

This refers to the fact that there exist natural $W(A_n)$-cells with edge multiplicities > 1; the lowest rank in which such cells occur is $n = 9$ (see [**MW**]).

3. Admissible W-Graphs

We define an S-labeled graph $\Gamma = (V, m, \tau)$ to be *admissible* if it

(A1) has nonnegative integer edge weights,

(A2) is edge-symmetric (i.e., $m(u \to v) = m(v \to u)$ if $\tau(u) \perp \tau(v)$),

(A3) has a bipartition.

The point of this definition is that we know that every natural W-graph is admissible; if an admissible W-graph is not natural, then we say that it is *synthetic*.

Of course the above definition makes sense for any Coxeter system (or without any Coxeter system whatsoever), but we are only interested in admissible W-graphs for finite Weyl groups, so that natural W-graphs exist as we have defined them. We should probably also add the hypothesis that Γ is finite, although this would not be reasonable if we were considering (say) affine Weyl groups.

Remark 3.1. By work of Gyoja [**G**], it is known that for every finite Weyl group W, every irreducible W-module occurs as M_Γ for some W-graph Γ. (See also the explicit constructions for the exceptional groups obtained by Howlett and Yin [**HY**].) However, it should be noted that in general, these W-graphs require negative (integer) edge weights and are not edge-symmetric. For example, one cannot realize the reflection representation of $W(B_2)$ in an edge-symmetric way. So these W-graphs are not admissible.

Working Hypothesis 3.2. *The class of admissible W-cells should be fairly rigid.*

It is important to keep in mind that this hypothesis is not reasonable without some sort of assumptions about W. In the extreme case, if (W, S) is a free Coxeter system (so that $s_i s_j$ has infinite order for all $i \neq j$), then there are no braid relations, and every S-labeled graph is a W-graph.

However, assuming that W is a finite Weyl group, we expect that there should be relatively few synthetic W-cells. (Examples later.)

Question 3.3. *Does each finite Weyl group have only finitely many admissible cells?*

Now suppose that (W, S) is simply-laced, not necessarily finite.

In the following, we describe five simple rules that characterize when an admissible S-labeled graph is a W-graph (i.e., the operators in (1) satisfy the braid relations).

A. The Compatibility Rule and the Simplicity Rule.

An S-labeled graph $\Gamma = (V, m, \tau)$ is a W-graph if and only if all of its rank two parabolic restrictions $\Gamma_{\{i,j\}}$ are $\langle s_i, s_j \rangle$-graphs (i.e., T_i and T_j satisfy the appropriate braid relation). For $I \subseteq \{i, j\}$, let V_I denote the set of vertices of Γ whose τ invariant is I when restricted to $\{i, j\}$ (i.e., $\tau(v) \cap \{i, j\} = I$). Assuming that $\Gamma_{\{i,j\}}$ is an $\langle s_i, s_j \rangle$-graph, convention (2) implies that the only possible edge orientations in $\Gamma_{\{i,j\}}$ are as follows:

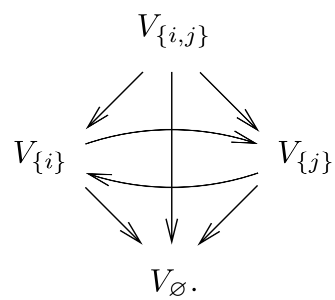

In particular, note that $V_{\{i\}} \cup V_{\{j\}}$ necessarily spans a subquotient of $\Gamma_{\{i,j\}}$. Letting A and B denote the respective (nonnegative integer) matrices encoding the multiplicities of edges from $V_{\{i\}}$ to $V_{\{j\}}$ and $V_{\{j\}}$ to $V_{\{i\}}$, we claim that having the braid relation involving T_i and T_j hold on the span of $V_{\{i\}} \cup V_{\{j\}}$ translates into a simple condition on A and B:

(R2) If $T_iT_j = T_jT_i$, then $A = 0$ and $B = 0$.

(R3) If $T_iT_jT_i = T_jT_iT_j$, then AB and BA are identity matrices.

(R4) If $(T_iT_j)^2 = (T_jT_i)^2$, then $ABA = 2A$ and $BAB = 2B$.

(R6) If $(T_iT_j)^3 = (T_jT_i)^3$, then $\begin{array}{c} ABABA - 4ABA + 3A = 0, \\ BABAB - 4BAB + 3B = 0. \end{array}$

In particular, (R2) implies that there cannot be edges between $V_{\{i\}}$ and $V_{\{j\}}$ in either direction if T_i and T_j commute. Translating this into a more direct statement about the structural features of Γ yields

FACT 3.4 (The Compatibility Rule). *If (V, m, τ) is a W-graph, then for all $u, v \in V$ such that $m(u \to v) \neq 0$, every $i \in \tau(u) - \tau(v)$ must be bonded to every $j \in \tau(v) - \tau(u)$.*

We use the term "bonded" here to mean that nodes i and j are adjacent in the diagram of (W, S); i.e., $T_iT_j \neq T_jT_i$.

REMARK 3.5. (a) The Compatibility Rule is valid for all W-graphs; we have used nothing (so far) about admissibility or being simply-laced.

(b) When $\tau(u)$ and $\tau(v)$ are comparable, the Compatibility Rule is vacuous.

For braid relations of length 3, note that (R3) is equivalent to having A and B be square matrices that are inverses of each other. This immediately shows that the problem of classifying W-cells, even in the A_2 case, is not reasonable without further assumptions.

Indeed, every choice of $A \in SL_m(\mathbb{Z})$ leads to a bipartite $W(A_2)$-graph with $2m$ vertices. Generically, this graph will be a cell with both positive and negative edges weights.

However, if Γ is admissible, then A and A^{-1} must be nonnegative integer matrices, and it is not hard to show that this is possible only if A is a permutation matrix. Thus,

FACT 3.6. *Assume (V, m, τ) is a W-graph satisfying (A1). If $T_iT_jT_i = T_jT_iT_j$, then there must be a bijection $\phi : V_{\{i\}} \to V_{\{j\}}$ such that for all $u \in V_{\{i\}}$ and $v \in V_{\{j\}}$,*

$$m(u \to v) = m(v \to u) = \begin{cases} 1 & \text{if } v = \phi(u), \\ 0 & \text{otherwise.} \end{cases}$$

It follows that the only admissible $W(A_2)$-cells are singletons with τ invariant $\{1, 2\}$ or $\varnothing$ and pairs $\{u, v\}$ with $\tau(u) = \{1\}$, $\tau(v) = \{2\}$ and $m(u \to v) = m(v \to u) = 1$.

Now assume that (W, S) is simply-laced, and consider the possible edge multiplicities involving a pair of vertices $u, v \in V$ in some W-graph Γ with nonnegative integer edge weights. If $\tau(u) \subset \tau(v)$, then we know that there can only be edges in the direction $v \to u$. We say that such edges are *arcs*.

Otherwise, if $\tau(u)$ and $\tau(v)$ are incomparable, then there is at least one $i \in \tau(u) - \tau(v)$ and one $j \in \tau(v) - \tau(u)$. By the Compatibility Rule, i and j must be bonded, so Fact 3.6 implies that there are no edges between u and v, or there is exactly one in each direction. In the latter case, we say that there is a *simple edge* between u and v. Summarizing,

FACT 3.7 (The Simplicity Rule). *If (W, S) is simply-laced then every W-graph (V, m, τ) satisfying (A1) consists of arcs and simple edges; i.e., for $u, v \in V$ such that $m(u \to v) \neq 0$, either $\tau(v) \subset \tau(u)$ and $m(v \to u) = 0$, or $\tau(u) \perp \tau(v)$ and $m(u \to v) = m(v \to u) = 1$.*

In particular, all such graphs are edge-symmetric.

Note that two vertices connected by a simple edge must belong to the same cell.

REMARK 3.8. (a) The matrix equations in (R4) have nonnegative integer solutions that yield $W(B_2)$-cells that are not edge symmetric, such as

Thus outside of the simply-laced case, nonnegativity alone does not imply edge-symmetry. If we add edge-symmetry to (R4) (i.e., $B = A^T$), one can show that the only admissible $W(B_2)$-cells are singletons with any of the four possible τ invariants, and the following:

Here as elsewhere, an undirected edge represents two edges, one for each direction.

(b) By classifying the nonnegative integer, edge-symmetric solutions of the matrix equations in (R6) (and assuming we made no mistakes in the somewhat tricky reasoning), the only admissible $W(G_2)$-cells other than singletons turn out to be

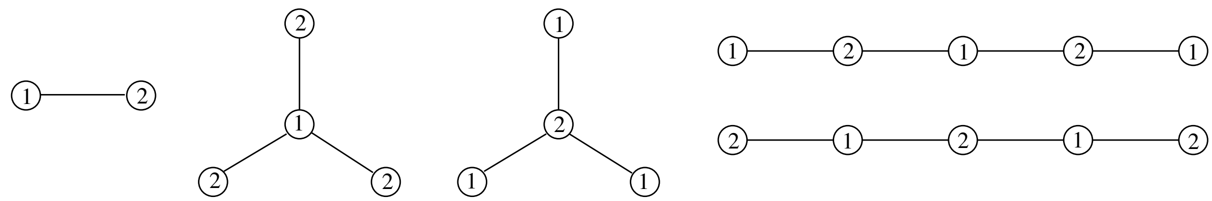

It appears that only the 5-vertex cells are natural; the other three are synthetic.

(c) Assuming we made no mistakes in (b), the above classifications show that if all braid relations in (W, S) have length 2, 3, 4, or 6 (e.g., if W is a finite Weyl group), then the Simplicity Rule holds for all admissible W-graphs.

Define the *compatibility graph* $\mathrm{Comp}(W, S)$ to be the simple graph with vertex set 2^S (i.e., subsets of S) and an edge between I and J if every $i \in I - J$ is bonded to every $j \in J - I$ in the diagram of (W, S). The compatibility graphs for A_3, A_4 and D_4 are displayed in Figure 1 with the always isolated vertices S and $\varnothing$ omitted.

Letting Γ_{sim} denote the graph on the vertex set V formed by the simple edges of Γ, note that the Compatibility Rule may be formulated as follows:

$$\tau \text{ is a graph homomorphism } \Gamma_{\mathrm{sim}} \to \mathrm{Comp}(W, S). \tag{3}$$

That is, if $\{u, v\}$ is a simple edge of Γ, then $\tau(u)$ is adjacent to $\tau(v)$ in $\mathrm{Comp}(W, S)$.

B. The Frontier Rule.

Continuing the hypothesis that (W, S) is simply-laced, the compatibility graph provides an interesting way to reformulate Fact 3.6. We define the *frontier* of a vertex v to be the set of bonds $\{i, j\}$ in the diagram of (W, S) such that $i \in \tau(v)$ and $j \notin \tau(v)$. For example in the A_8-diagram

$$1 \underset{a}{\text{———}} 0 \underset{b}{\text{———}} 1 \underset{c}{\text{———}} 1 \underset{d}{\text{———}} 1 \underset{e}{\text{———}} 0 \underset{f}{\text{———}} 0 \underset{g}{\text{———}} 1,$$

we have labeled the bonds a, b, c, d, e, f, g, and have replaced the nodes with 0's and 1's to indicate the τ invariant of a vertex v in some $W(A_8)$-graph. In this example, the frontier of v is the set $\{a, b, e, g\}$. If the endpoints of bond a are i and j, with (say) $i \in \tau(v)$ and $j \notin \tau(v)$, then Fact 3.6 says that there must be a *unique* vertex u connected to v by a simple edge such that $i \notin \tau(u)$ and $j \in \tau(u)$. Thus,

FACT 3.9 (The Frontier Rule). *If (W, S) is simply-laced and v is a vertex in some W-graph satisfying (A1), then for each bond $\{i, j\}$ in the frontier of v, there is a unique neighbor u of v (connected by a simple edge) such that $\{i, j\}$ is in the frontier of u.*

Conversely, the Compatibility Rule shows that every Γ_{sim}-neighbor u of v shares a nonempty set of bonds in their respective frontiers; namely, the bonds $\{i, j\}$ such that

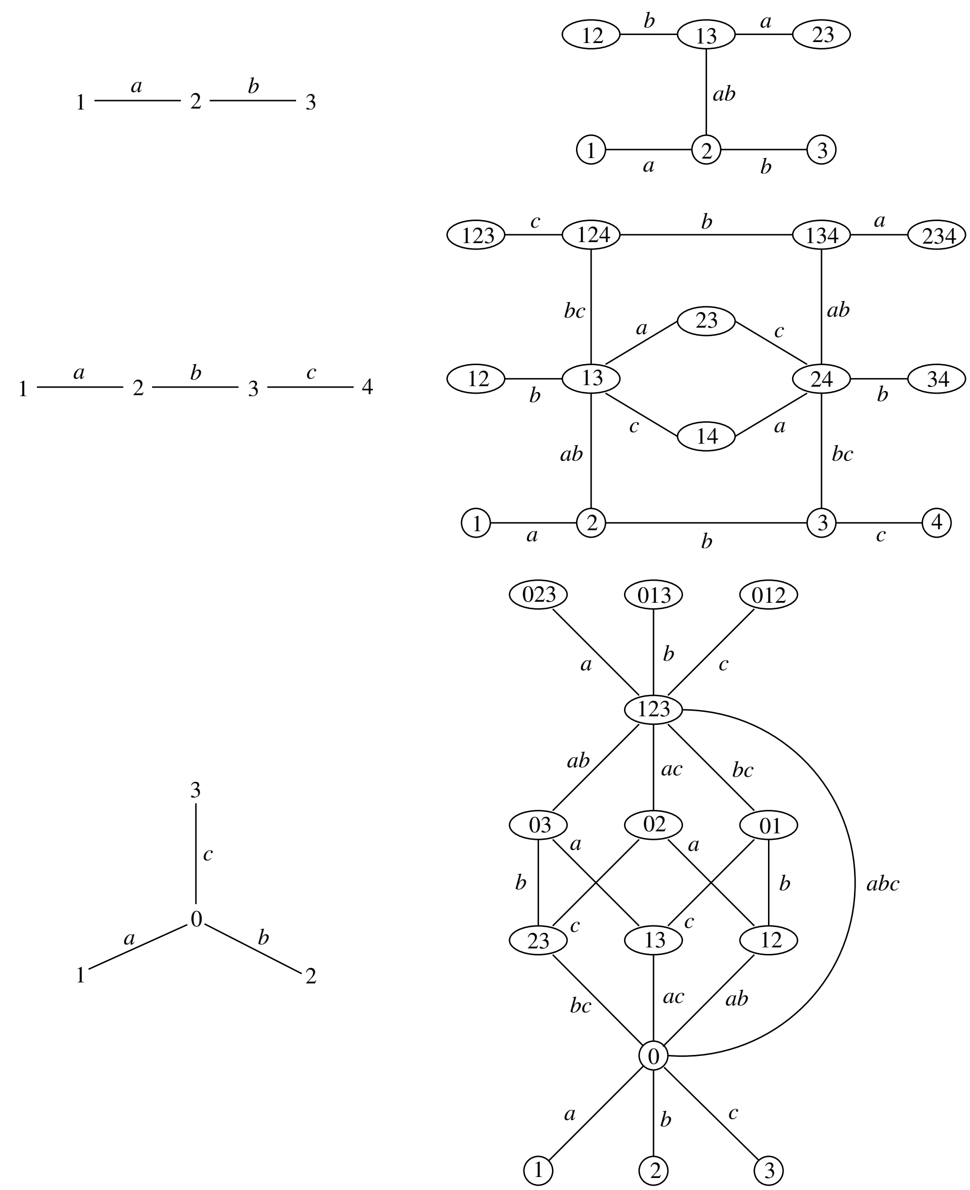

FIGURE 1: Diagrams and compatibility graphs for A_3, A_4, and D_4.

$i \in \tau(v) - \tau(u)$ and $j \in \tau(u) - \tau(v)$. Thus we can edge-color Γ_{sim} by assigning to (simple) edge $\{u, v\}$ the set of bonds shared by the frontiers of u and v. The Frontier Rule says that the sets of bonds distributed to the various edges incident to a vertex v by this edge coloring must form a partition of the frontier of v.

Noting that frontiers depend only on τ invariants, this edge coloring of Γ_{sim} descends to the compatibility graph; i.e., if $\{u, v\}$ is a simple edge, then the set of bonds shared by the frontiers of u and v depends only on $\tau(u)$ and $\tau(v)$. We have illustrated this in Figure 1, where each edge in $\text{Comp}(W, S)$ has been labeled by the appropriate subset of bonds in the diagram of (W, S).

For example, consider the vertex $\{1, 3\}$ in $\text{Comp}(A_4)$ (see Figure 1). A vertex v in Γ

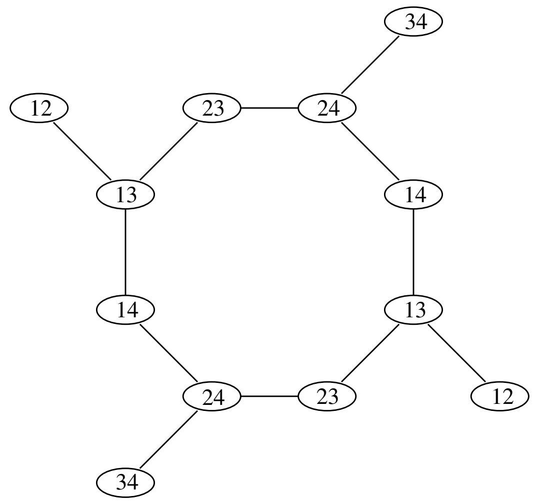

Figure 2: Not a $W(A_4)$-graph.

with τ invariant $\{1,3\}$ has frontier $\{a,b,c\}$, and the Frontier Rule says that the simple neighborhood of v must consist of vertices that "share" these three bonds, without overlap. Examining the neighborhood of $\{1,3\}$ in $\mathrm{Comp}(A_4)$, we conclude that the simple neighborhood of every vertex of Γ with τ invariant $\{1,3\}$ has one of three forms:

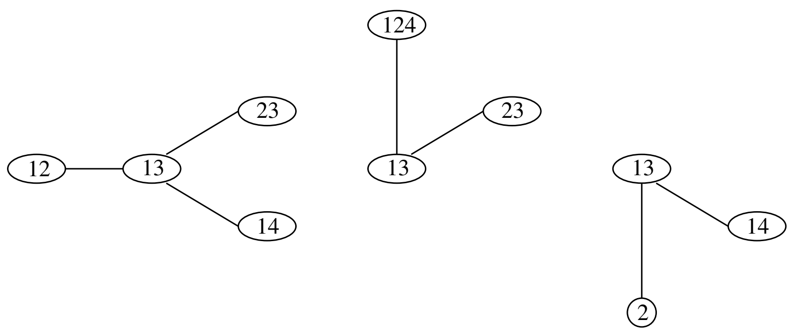

Remark 3.10. The three rules we have accumulated so far (Compatibility, Simplicity, and Frontier) are already quite powerful. For example, by staring at $\mathrm{Comp}(A_3)$ in Figure 1, it is not hard to deduce that if Γ is an admissible $W(A_3)$-cell, then Γ_{sim} must be a disjoint union of (perhaps multiple copies of) the following graphs:

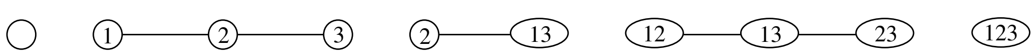

Of course, a $W(A_3)$-graph may also have arcs $u \leftarrow v$ such that $\tau(u) \subset \tau(v)$. However, note that all such arcs must be directed right-to-left among the above graphs, and there are no inclusions of τ invariants within these graphs, so we conclude that no admissible $W(A_3)$-cell can have arcs, and the above graphs are the only possible such cells. (Moreover, it is not hard to check that they are indeed $W(A_3)$-graphs.)

C. The Diamond Rule.

Before getting too excited about Remark 3.10, one should keep in mind that (R2)–(R6) are merely necessary conditions; they do not capture the full power of the braid relations.

Indeed, consider the 4-cycle in the center of $\mathrm{Comp}(A_4)$ (see Figure 1). By winding around this cycle one or more times, it is easy to construct arbitrarily large S-labeled graphs that are admissible and satisfy the Compatibility, Simplicity, and Frontier Rules, such as the graph in Figure 2. However, we shall see that this graph cannot be a $W(A_4)$-graph; an admissible $W(A_4)$-graph cannot have an 8-cycle (or 12-cycle, etc.) whose τ-image is the 4-cycle in $\mathrm{Comp}(A_4)$.

Fact 3.11 (The Diamond Rule). *If (V, m, τ) is a W-graph satisfying (A1) and (W, S) is simply-laced, then for all $i \neq j$ and all $v_\pm \in V$ such that $i, j \in \tau(v_+)$ and $i, j \notin \tau(v_-)$,*

$$\sum_{u:\, \substack{i \in \tau(u) \\ j \notin \tau(u)}} m(v_+ \to u) m(u \to v_-) = \sum_{u:\, \substack{i \notin \tau(u) \\ j \in \tau(u)}} m(v_+ \to u) m(u \to v_-). \tag{4}$$

That is, the number of 2-step paths from v_+ to v_- that pass through vertices with $i \in \tau$ and $j \notin \tau$ equals the number of such paths through vertices with $i \notin \tau$ and $j \in \tau$.

For example, consider one of the vertices in Figure 2 with $\tau = \{2, 4\}$. There is exactly one 2-step path from this vertex to each vertex with $\tau = \{1, 3\}$, and for one of these vertices, the intermediate vertex has $\tau = \{2, 3\}$. This violates the Diamond Rule (at $i = 2$ and $j = 4$), so the graph is not a W-graph.

Remark 3.12. The Diamond Rule can be used to determine the multiplicities of certain arcs. For example, consider the following fragment in an admissible $W(A_5)$-graph.

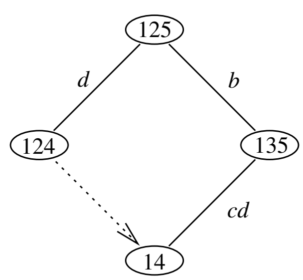

The edge labels indicate which bonds are shared by the frontiers of their endpoints, using the abbreviations $a = \{1, 2\}$, $b = \{2, 3\}$, $c = \{3, 4\}$, $d = \{4, 5\}$. This graph has a 2-step path $\{1, 2, 5\} \to \{1, 3, 5\} \to \{1, 4\}$. Any other 2-step path from $\{1, 2, 5\}$ to $\{1, 4\}$ must pass through the vertex with $\tau = \{1, 2, 4\}$ or else through other vertices in the graph that are not depicted. Since $\{1, 4\} \not\subset \{1, 2, 5\}$, these other vertices (if they exist) must be connected to $\{1, 2, 5\}$ or $\{1, 4\}$ by a simple edge (possibly both). However, the frontiers of $\{1, 2, 5\}$ and $\{1, 4\}$ are bd and acd, so by the Frontier Rule, $\{1, 2, 5\}$ has no other simple neighbors, and $\{1, 4\}$ must have exactly one additional simple neighbor that shares the

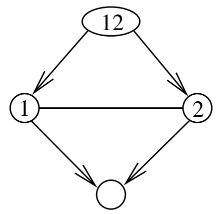

FIGURE 3: A non-bipartite $W(A_2)$-graph.

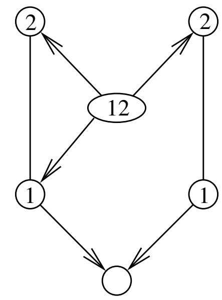

FIGURE 4: A $W(A_2)$-graph at $q = 1$ only.

bond a. Inspecting $\mathrm{Comp}(A_5)$, this extra neighbor necessarily has $\tau = \{2, 4\}$, and thus cannot be adjacent to $\{1, 2, 5\}$ by an arc or simple edge. We conclude that there are no other 2-step paths from $\{1, 2, 5\}$ to $\{1, 4\}$. Applying the Diamond Rule (with $i = 2$, $j = 5$), it follows that there must be exactly one arc $\{1, 2, 4\} \to \{1, 4\}$.

D. The Hexagon Rule.

There is one more graph-theoretic implication for braid relations of length 3.

FACT 3.13 (The Hexagon Rule). *Assume (V, m, τ) is a W-graph satisfying (A1). If $T_iT_jT_i = T_jT_iT_j$, then for all $v_\pm \in V$ such that $i, j \in \tau(v_+)$ and $i, j \notin \tau(v_-)$, we have*

$$\sum_{u \in V_{i/j}} m(v_+ \to u)m(\phi(u) \to v_-) = \sum_{u \in V_{i/j}} m(v_+ \to \phi(u))m(u \to v_-), \tag{5}$$

where $V_{i/j} := \{u \in V : i \in \tau(u), j \notin \tau(u)\}$, and $\phi(u)$ denotes the unique vertex in $V_{j/i}$ connected to $u \in V_{i/j}$ by a simple edge (as guaranteed by the Frontier Rule).

In other words, the number of 3-step paths from v_+ to v_- whose middle step follows a simple edge from $V_{i/j}$ to $V_{j/i}$ equals the number of such paths whose middle step follows a simple edge from $V_{j/i}$ to $V_{i/j}$.

THEOREM 3.14. *If (W, S) is simply-laced, then an S-labeled graph with nonnegative integer edge weights is a W-graph if and only if it satisfies*

(a) *the Compatibility Rule (Fact 3.4),*
(b) *the Simplicity Rule (Fact 3.7),*
(c) *the Frontier Rule (Fact 3.9),*
(d) *the Diamond Rule (Fact 3.11), and*
(e) *the Hexagon Rule (Fact 3.13).*

Moreover, all such W-graphs are edge-symmetric.

REMARK 3.15. (a) Not all W-graphs as above are necessarily admissible (i.e., such graphs may fail to be bipartite). A simple example of this is provided in Figure 3. To give a cellular example, take $W = \mathrm{Aff}(A_2)$ and view the diagram of (W, S) as an S-labeled graph in which node i is assigned τ invariant $\{i\}$. Perhaps there are no cellular examples in case W is a finite Weyl group, but this seems hard to prove.

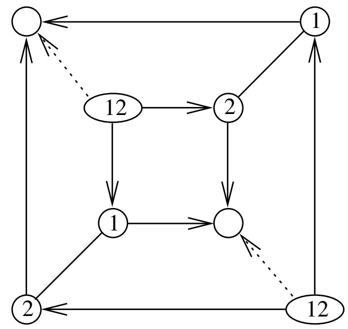

FIGURE 5: The generic admissible $W(A_2)$-graph.

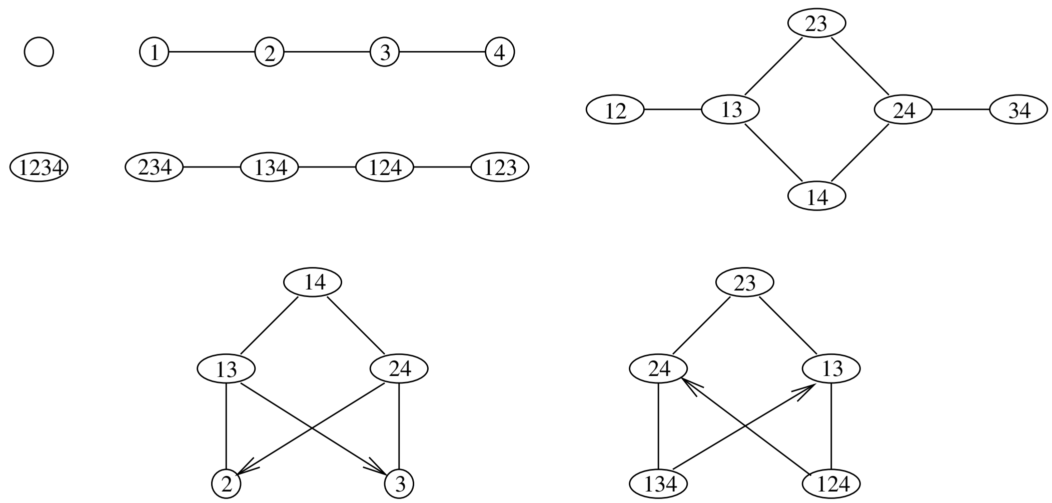

FIGURE 6: Admissible $W(A_4)$-cells.

(b) The (non-bipartite) S-labeled graph in Figure 4 produces operators that satisfy the braid relations at $q = 1$, but not for generic q. Note that it violates both the Diamond and Hexagon Rules.

(c) Using the above characterization, one can show that every admissible $W(A_2)$-graph is a commutative diagram as in Figure 5. More precisely, each oriented edge (both dotted and solid) should be viewed as a nonnegative integer matrix defining a map in the indicated direction, each unoriented edge should be viewed as an identity map, and the condition on the maps is that the subdiagram formed by the solid edges must be commutative.

(d) Using the above characterization, we have classified all admissible $W(A_4)$-cells and $W(D_4)$-cells (see Figures 6 and 7). It seems that work of Garfinkle and Vogan [**GV**] may be used to show that the two 8-vertex $W(D_4)$-cells are synthetic; the remainder are natural. Note that the two synthetic cells also have arcs with multiplicity > 1.

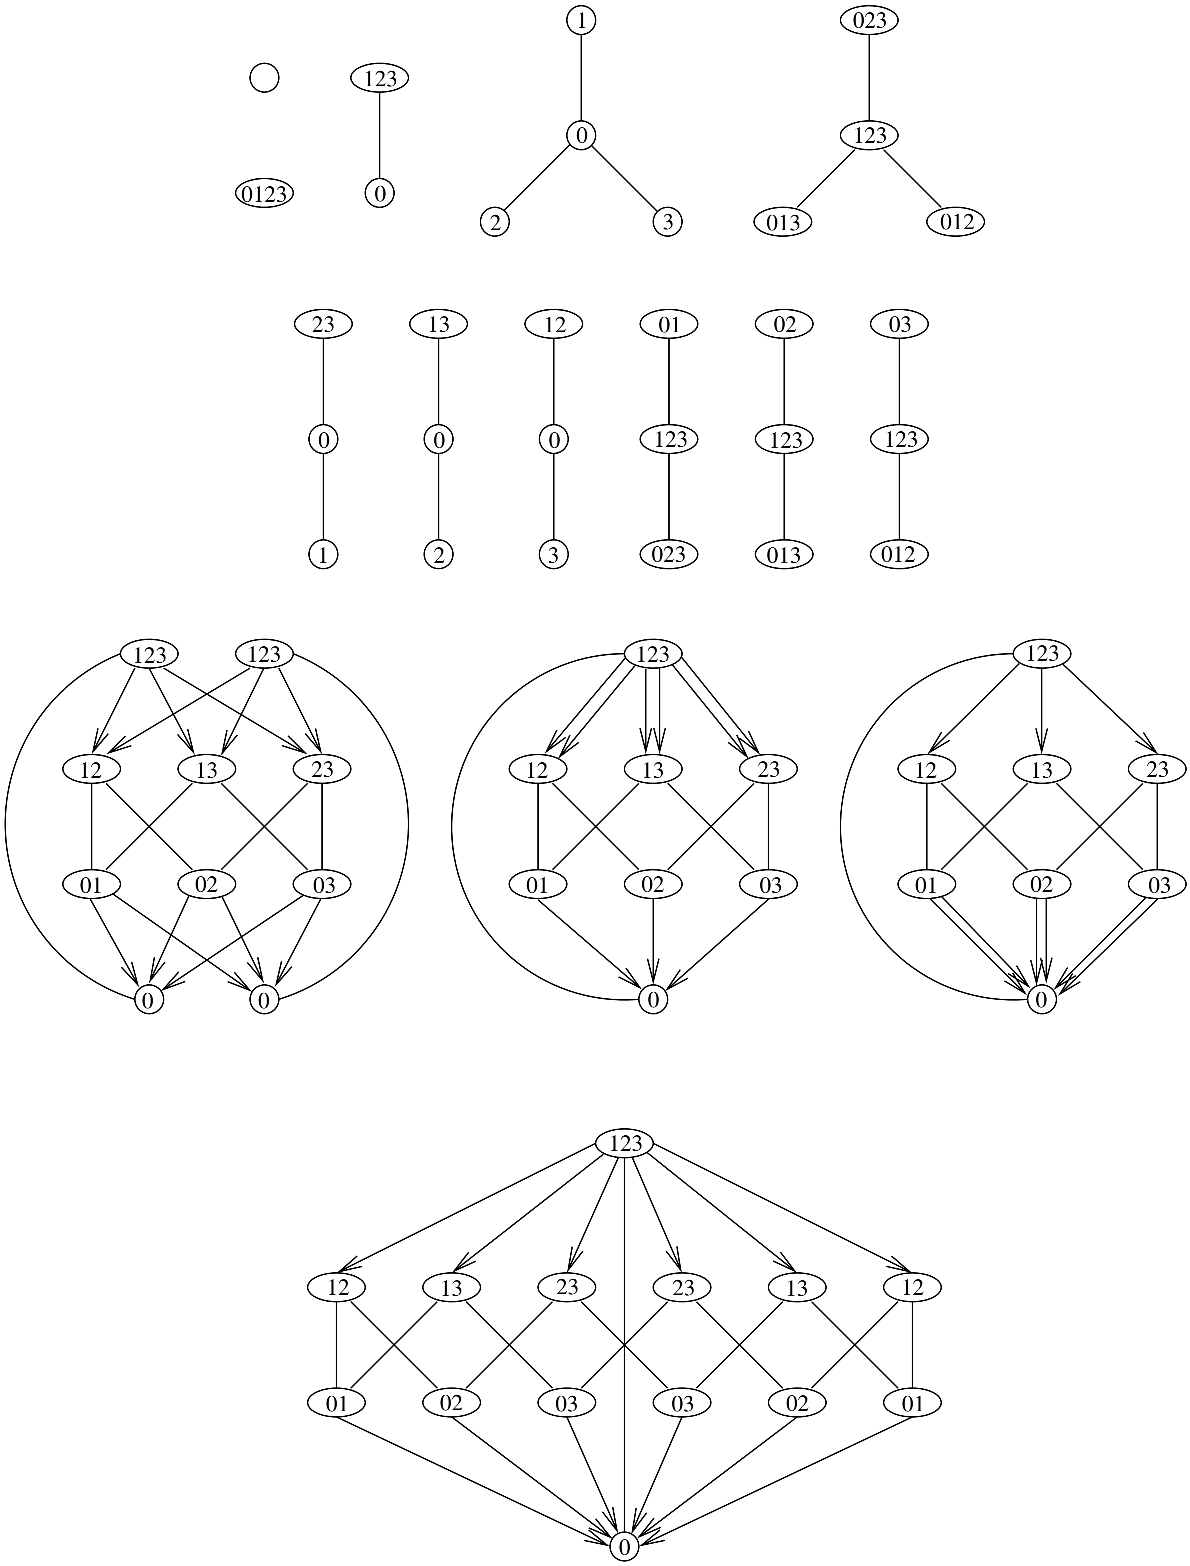

FIGURE 7: Admissible $W(D_4)$-cells.

4. Molecules (or Proteins?)

Toward the goal of constructing an algorithm to generate all admissible W-cells, we have the idea of a two-step process.

First, recognizing that the Compatibility, Simplicity and Frontier Rules provide strong constraints on the simple edges of a cell Γ, we would like to analyze the connected components of Γ_{sim} (recall that this denotes the subgraph of Γ formed by its simple edges), and the subgraphs of Γ that these components induce. We call these induced subgraphs *molecules* (or perhaps we should call them *proteins*), but see below for the precise definition. Up to this point, we have studied only the simply-laced versions of this problem, but it is a sensible idea for any finite Weyl group in light of the more general validity of the Simplicity Rule (see Remark 3.8(c)).

Once we have the ability to generate W-molecules, then we need to analyze the general problem of determining all possible ways that one or more molecules may be bound together with arcs to form a cell.

Definition 4.1. Assume (W, S) is simply-laced.
An admissible S-labeled graph $\Gamma = (V, m, \tau)$ is a *molecule* if it satisfies

(a) the Compatibility Rule (Fact 3.4),
(b) the Simplicity Rule (Fact 3.7),
(c) the Frontier Rule (Fact 3.9),
(d) the Localized Diamond Rule; i.e., (4) holds when $\tau(v_-) \not\subset \tau(v_+)$.
(e) the Localized Hexagon Rule; i.e., (5) holds when there exists $k, l \in \tau(v_-) - \tau(v_+)$ such that i and k are not bonded, and j and l are not bonded,

and Γ_{sim} is connected.

The point of the above definition is the following.

Fact 4.2. *If (W, S) is simply-laced and Γ is an admissible W-cell, then the subgraph of Γ induced by each connected component of Γ_{sim} is a molecule.*

Given that Γ is an admissible W-cell, it is easy to see that (a), (b), and (c) must hold automatically for any connected component of Γ_{sim}. For (d), note that in every 2-step path from v_+ to v_-, the intermediate vertex could belong to another component only if both steps of the path are arcs, and of course this requires $\tau(v_-) \subset \tau(v_+)$. The necessity of (e) follows by similar reasoning.

Remark 4.3. (a) It should be emphasized that a molecule may have arcs. Moreover, if one is given only the simple edges of a molecule Γ, one may regard the arc multiplicities within Γ (i.e., $m(u \to v)$ for all u, v such that $\tau(u) \supset \tau(v)$) as indeterminates, and view the Localized Diamond and Hexagon Rules as providing a system of equations in these variables. It is interesting to note that these equations turn out to be *linear*.

(b) One of the difficulties of this approach to studying admissible cells is that we have introduced yet another level of artifice: not all molecules can occur inside cells. For example, it is not hard to show that for each $m \geqslant 1$ there is a $6m$-cycle that forms a $W(D_4)$-molecule, but in examining Figure 7, one sees that the only cycles that occur as molecules in admissible $W(D_4)$-cells are 6-cycles and 12-cycles.

5. Computation (Cell Synthesis)

We have been developing software to generate molecules, but what we have at this point is code that imposes only a subset of the defining features of molecules (only parts of the Local Diamond Rule are implemented, and the Local Hexagon Rule is completely ignored). Nevertheless, the results so far do seem promising, even if the code is not ready for public consumption.

One general comment about computation we would like to make is that even though the existence of a bipartition (axiom (A3)) has not played an essential role in these notes, it is of central importance when the topic turns to algorithms.

To explain, one needs to begin with the observation that any bipartite graph $\Gamma = (V, E)$ may be viewed as a ranked poset. One fixes a base vertex $v_0 \in V$ and introduces a function $r : V \to \mathbb{Z}$ by defining $r(v)$ to be the distance from v to v_0. Having a bipartition (no odd cycles) is precisely the condition one needs to guarantee that adjacent vertices have r-values that differ by ± 1. The edges $\{u, v\}$ of Γ, ordered so that $u < v$ if $r(v) - r(u) = 1$, form the covering edges of a poset with rank function r and minimum element v_0.

With this in mind, the natural strategy for synthesizing a molecule is to start with a vertex v_0 and chosen τ invariant J (say), and then grow the remainder of the molecule rank by rank, by judicious use of the Compatibility and Frontier Rules. At various stages there will be ambiguities when two or more vertices at a given level might need to be identified. If the Local Diamond and Hexagon Rules cannot resolve the ambiguity, then the process can be forked, and additional molecules can be grown in parallel.

References

[**B**] D. Barbasch, Cells and Weyl groups in primitive ideals, Atlas IV lecture notes, ⟨`www.liegroups.org/papers/summer06/cells.pdf`⟩.

[**G**] A. Gyoja, On the existence of a W-graph for an irreducible representation of a Coxeter group, *J. Algebra* **86** (1984), 422–438.

[**HY**] ⟨`www.maths.usyd.edu.au/u/bobh/yy/`⟩.

[**KL**] D. Kazhdan and G. Lusztig, Representations of Coxeter groups and Hecke algebras, *Invent. Math.* **53** (1979), 165–184.

[**LV**] G. Lusztig and D. A. Vogan, Singularities of closures of K-orbits on flag manifolds, *Invent. Math.* **71** (1983), 365–379.

[**GV**] D. Garfinkle and D. A. Vogan, On the structure of Kazhdan-Lusztig cells for branched Dynkin diagrams, *J. Algebra* **153** (1992), 91–120.

[**MW**] T. J. McLarnan and G. S. Warrington, Counterexamples to the 0-1 conjecture, *Represent. Theory* **7** (2003), 181–195.