More W-Graphs and Cells: Molecular Components and Cell Synthesis Atlas of Lie Groups AIM Workshop VI 7–11 July 2008

John R. Stembridge ⟨jrs@umich.edu⟩

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

Contents

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

1. Graphs, Cells and Admissibility

We begin with a review of the basic concepts associated with admissible W-graphs, following the lecture notes of last year [**S1**] and the more detailed recent paper [**S2**].

We assume throughout that W is a finite Weyl group, with distinguished generators $\{s_i : i \in I\}$, where I is some finite index set (usually $\{1, 2, \ldots, n\}$). We let p_{ij} denote the order of $s_i s_j$ in W, or equivalently, the length of the braid relation satisfied by s_i and s_j.

Recall that an *I-labeled graph* is a triple $\Gamma = (V, m, \tau)$, where

(1) V is a finite vertex set,
(2) τ is a map $V \to \{\text{subsets of } I\}$, and
(3) $m : V \times V \to \mathbb{Z}[q^{\pm 1/2}]$ is a matrix of edge weights.

The (u, v)-entry of the matrix m is denoted $m(u \to v)$ and interpreted as the weight of the edge $u \to v$; having $m(u \to v) = 0$ means that there is no such edge.

All I-labeled graphs we consider will also be required to be *reduced*; i.e.,

$$\tau(u) \subseteq \tau(v) \;\Rightarrow\; m(u \to v) = 0 \text{ for all } u, v \in V. \tag{1.1}$$

The (reduced) I-labeled graph Γ is *admissible* if it is bipartite, all edge weights are nonnegative integers, and it is *edge-symmetric*; i.e.,

$$m(u \to v) = m(v \to u) \;\text{ whenever }\; \tau(u) \not\subseteq \tau(v) \text{ and } \tau(v) \not\subseteq \tau(u).$$

Let M_Γ denote the free $\mathbb{Z}[q^{\pm 1/2}]$-module generated by the vertices of Γ.

The I-labeled graph Γ is a *W-graph*[1] if the following operators on M_Γ satisfy the braid relations for (W, S):

$$T_i(u) = \begin{cases} qu & \text{if } i \notin \tau(u), \\ -u + q^{1/2} \sum_{v : i \notin \tau(v)} m(u \to v) v & \text{if } i \in \tau(u). \end{cases} \tag{1.2}$$

A W-graph is a *cell* if it is strongly connected.

Any operator having the form in (1.2) satisfies the quadratic relation $(T_i - q)(T_i + 1) = 0$, so M_Γ carries an action of the Iwahori-Hecke algebra of (W, S) (if Γ is a W-graph), and also provides (at $q = 1$) an integral representation of W.

Our initial motivation for studying admissible W-graphs is based on the simple observation that the W-graphs in Kazhdan-Lusztig theory, as well as in the representation theory of real reductive groups, are admissible. The results obtained so far support our...

Main Contention 1.1. *The study of admissible W-graphs and cells can provide considerable insight into the structure of Kazhdan-Lusztig cells, Harish-Chandra cells, and the larger W-graphs into which these cells are embedded.*

[1] We have transposed the Kazhdan-Lusztig definition, but this modification has no effect on what constitutes a W-graph.

2. The Four Combinatorial Rules

There are four rules that characterize when an admissible I-labeled graph is a W-graph.

DEFINITION 2.1. An I-labeled graph $\Gamma = (V, m, \tau)$ satisfies the *W-Compatibility Rule* if for every edge $u \to v$, every $i \in \tau(u) - \tau(v)$ is bonded to every $j \in \tau(v) - \tau(u)$ in the Dynkin diagram of W.

REMARK 2.2. (a) If $\tau(u)$ contains $\tau(v)$, then the above condition is vacuous.

(b) Every W-graph, admissible or not, obeys the W-Compatibility Rule.

(c) Recall that the W-Compatibility Rule can be reformulated in terms of the *compatibility graph* $\mathrm{Comp}(W, S)$: the vertices of this graph are the subsets of I, and there is an edge $J \to K$ for all $J, K \subseteq I$ such that $J \nsubseteq K$ and every $j \in J - K$ is bonded to every $k \in K - J$ in the Dynkin diagram. To impose the W-Compatibility Rule is to require that $\tau : \Gamma \to \mathrm{Comp}(W, S)$ is a graph homomorphism.

Note that if $J \supsetneq K$ then there is always an edge $J \to K$. Only the edges between incomparable pairs J, K depend on the Dynkin diagram, and these latter edges are symmetric (i.e., $J \to K$ if and only if $K \to J$). The symmetric parts of various compatibility graphs are illustrated in Figure 1.

DEFINITION 2.3. An I-labeled graph $\Gamma = (V, m, \tau)$ satisfies the *Simplicity Rule* if for all edges $u \to v$ (i.e., $m(u \to v) \neq 0$), either

(a) $\tau(u) \supsetneq \tau(v)$ and $m(v \to u) = 0$, or
(b) $\tau(u)$ and $\tau(v)$ are incomparable and $m(u \to v) = m(v \to u) = 1$.

We say that the edge $u \to v$ is an *arc* if (a) holds; it is a *simple edge* if (b) holds.

To describe the third rule, first consider what happens when an admissible I-labeled graph Γ is restricted to a parabolic subgroup of rank 2, say $\langle s_i, s_j \rangle$, then reduced as in (1.1) and broken into strongly connected components. (The components are cells if Γ happens to be a W-graph.) In the restricted graph, there are four possible τ invariants: $\{i, j\}$, $\{i\}$, $\{j\}$, and $\varnothing$. Each node with $\tau = \{i, j\}$ or $\tau = \varnothing$ forms a singleton component (we regard these as trivial by convention), so all other components must be composed of zero or more symmetric edges each of which connects a vertex with $\tau = \{i\}$ to a vertex with $\tau = \{j\}$.

DEFINITION 2.4. An admissible I-labeled graph Γ satisfies the *W-Bonding Rule* if for all $i, j \in I$ that are bonded in the diagram of (W, S) (i.e., $p_{ij} \geqslant 3$), the nontrivial strongly connected components in the $\{i, j\}$-restriction of Γ are A-D-E Dynkin diagrams whose Coxeter numbers evenly divide p_{ij}.

REMARK 2.5. (a) We could allow $p_{ij} = 2$ in the Bonding Rule, but this would be redundant if the W-Compatibility Rule is also being imposed. (Only A_1 has a Coxeter number that divides 2.)

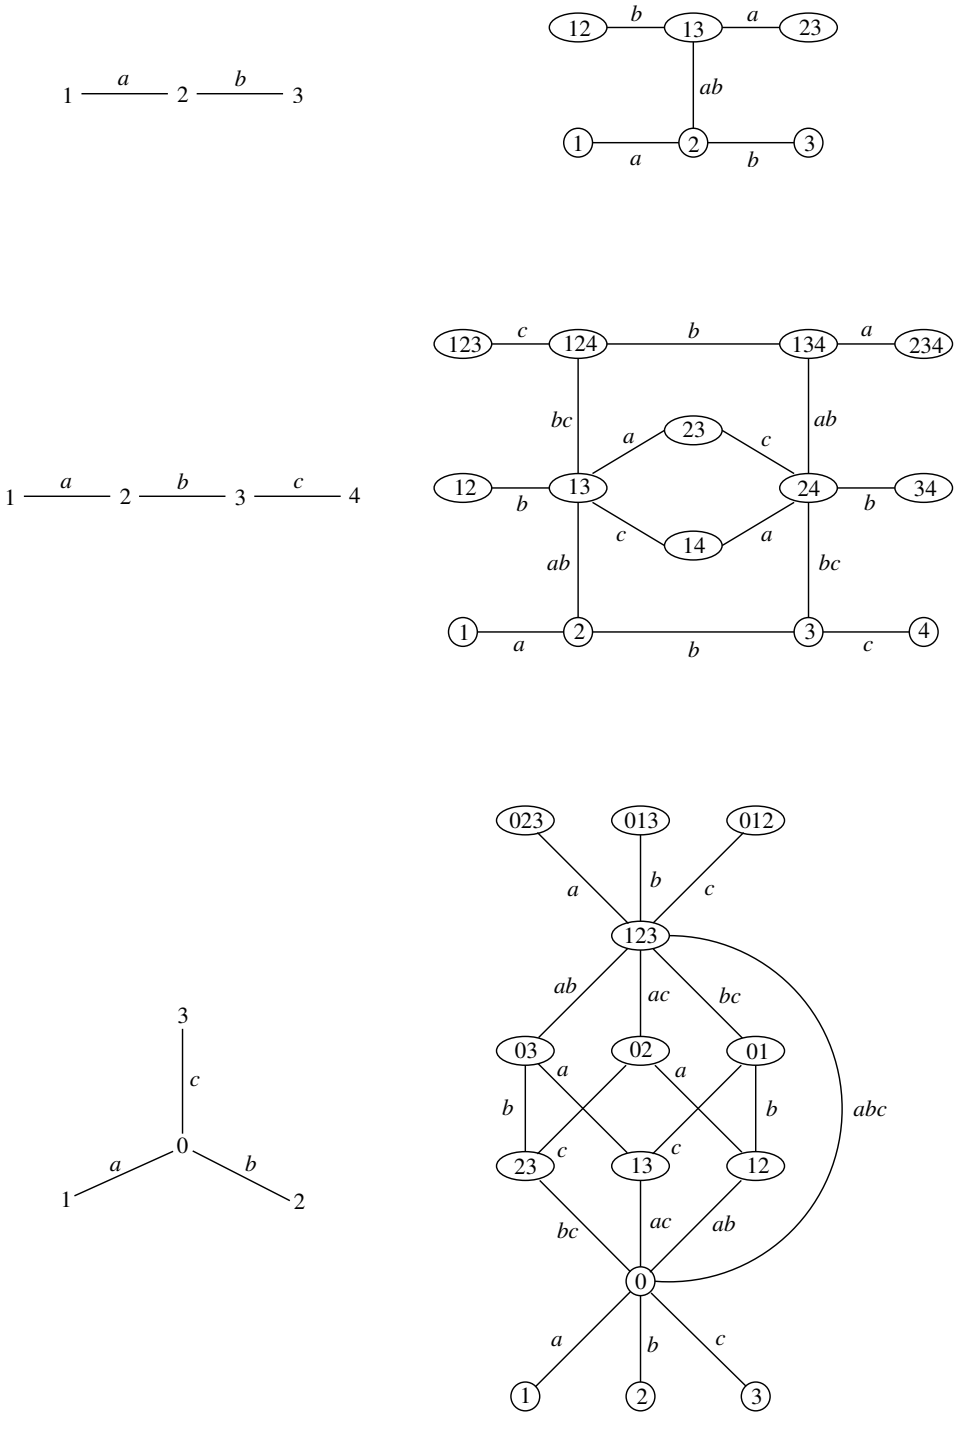

FIGURE 1: Diagrams and compatibility graphs for A_3, A_4, and D_4.

(b) Similarly, only A_2 has a Coxeter number that divides 3. Thus if the Bonding Rule is imposed for a pair i, j with $p_{ij} = 3$, then for every vertex u with $i \in \tau(u)$ and $j \notin \tau(u)$, there must be a *unique* vertex v adjacent to u with $i \notin \tau(v)$ and $j \in \tau(v)$.

(c) The "Frontier Rule" in [**S1**] is the W-Bonding Rule for simply-laced W.

It seems highly likely that the following is true, but a proof seems to be elusive.[2]

QUESTION 2.6. *Given a rank 2 parabolic subgroup $W_J \cong I_2(p)$, is every cell in the W_J-restriction of a Harish-Chandra or Kashzdan-Lusztig W-cell also a Kazhdan-Lusztig cell for W_J (i.e., either trivial or a Dynkin diagram of type A_{p-1})?*

If so, this suggests that we should add this requirement to the defining properties for admissible W-graphs. Note that this would have an effect only for the multiply-laced cases such as B_n or F_4, but we have not gathered enough data on these cases to know how significant it might be.

REMARK 2.7. Although it is only a wild guess at this point, the connection between admissible W-graphs and *symmetric* (i.e., *A-D-E*) Cartan matrices suggests that it might be worthwhile to explore a variation of admissibility in which edge-symmetry is replaced with *edge-symmetrizability*. Taking a cue from symmetrizable Cartan matrices, one could hypothesize the existence of a diagonal change of basis that symmetrizes the edge weight matrix m. More precisely, there should exist nonzero scalars λ_v $(v \in V)$ such that

$$\lambda_u m(u \to v) = \lambda_v m(v \to u) \quad \text{whenever} \quad \tau(u) \not\subseteq \tau(v) \text{ and } \tau(v) \not\subseteq \tau(u).$$

In any case, it is easy to show that this brings multiply-laced Dynkin diagrams into play as cells for rank two groups, and in particular, allows the realization of the reflection representation of B_2, something that is impossible in the world of admissible cells.

Following up on a remark made by Peter Trapa at ATLAS V, this suggests

SPECULATION 2.8. *Are the "geometric cells" that arise in Springer's construction edge-symmetrizable?*

Given an I-labeled graph $\Gamma = (V, m, \tau)$ and disjoint subsets $J, K \subseteq I$, let $V_{J/K}$ denote the set of vertices that include J and exclude K from their τ invariant; i.e.,

$$V_{J/K} = \{v \in V : J \subseteq \tau(v),\ \tau(v) \cap K = \varnothing\}.$$

We will also use abbreviated forms of this notation, such as $V_{i/j}$ in place of $V_{\{i\}/\{j\}}$.

For distinct $i, j \in I$, a directed path $u \to v_1 \to \cdots \to v_{r-1} \to v$ in Γ of length $r \geqslant 2$ is defined to be *alternating of type* (i, j) if

$$u \in V_{ij/\varnothing},\ \ v_k \in V_{i/j} \text{ for } k \text{ odd},\ \ v_k \in V_{j/i} \text{ for } k \text{ even},\ \text{ and } v \in V_{\varnothing/ij}. \tag{2.1}$$

[2]David Vogan informs me that this is true for $p = 4$, and that it follows from IC1.

We let $N_{ij}^r(\Gamma; u, v)$ denote the edge-weighted count of all such paths from u to v; i.e.,

$$N_{ij}^r(\Gamma; u, v) := \sum_{v_1,\ldots,v_{r-1}} m(u \to v_1)m(v_1 \to v_2)\cdots m(v_{r-2} \to v_{r-1})m(v_{r-1} \to v),$$

where the vertices v_k are restricted as in (2.1). Note that if Γ obeys the Simplicity Rule, then all of the internal edges in an alternating path are simple and thus have unit weight; only the initial and terminal edges $u \to v_1$ and $v_{r-1} \to v$ may be arcs of weight > 1.

DEFINITION 2.9. An admissible I-labeled graph Γ satisfies the *W-Polygon Rule* if for all distinct pairs $i, j \in I$ and all vertices u, v with $i, j \in \tau(u)$ and $i, j \notin \tau(v)$, we have

$$N_{ij}^r(\Gamma; u, v) = N_{ji}^r(\Gamma; u, v) \qquad \text{for } r = 2, 3, \ldots, p_{ij};$$

i.e., the weighted counts of alternating paths of length r of types (i, j) and (j, i) from u to v are the same for all $r \leqslant p_{ij}$.

The "Diamond Rule" and "Hexagon Rule" used in [**S1**] are together equivalent to the W-Polygon Rule when W is simply-laced.

THEOREM 2.10 [**S2**]. *An admissible I-labeled graph is a W-graph if and only if it satisfies the W-Compatibility, Simplicity, W-Bonding, and W-Polygon Rules.*

Although we are assuming that W is a finite Weyl group, it should be noted that the above result is more generally valid as long as W is braid-finite (i.e., all pairs of generators satisfy a braid relation of finite length). If W is not braid finite, there is a similar characterization but it is messier to state.

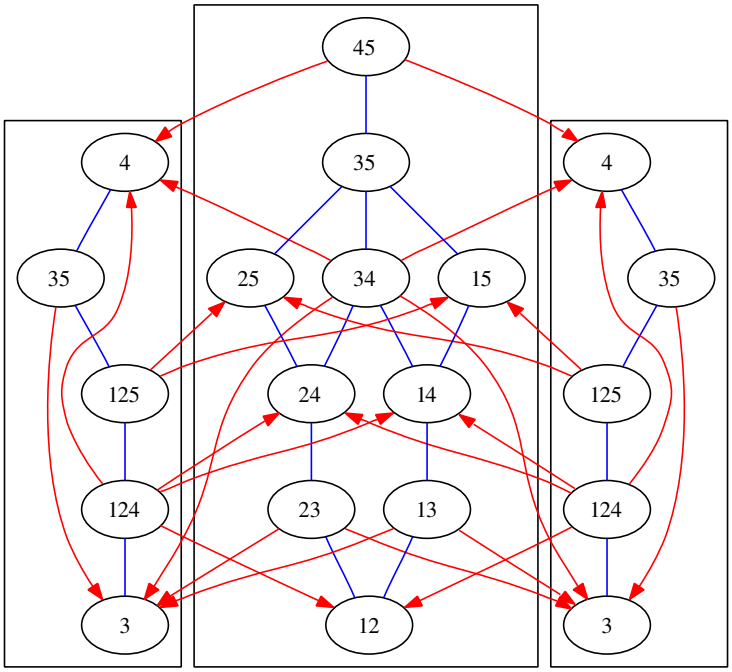

FIGURE 2: A D_5-cell with three molecular components.

3. Molecules

Consider the problem of generating admissible W-graphs or W-cells from scratch. The Compatibility, Simplicity and Bonding Rules provide fairly tight constraints on the formation of simple edges, and we can write software to generate all connected simple graphs that meet those constraints. However, except for groups of low rank, there will generally be infinitely many such graphs, most of which cannot occur in any W-graph.

The main point of this section is to describe how the W-Polygon Rule may be used to impose further constraints on these graphs (and the weights of arcs connecting them). Note that since these graphs are not full W-graphs, the application of this rule is subtle.

A. The Local Polygon Rule.

Let $\Gamma = (V, m, \tau)$ be an admissible W-graph, and define an equivalence relation on V by declaring that $u \sim v$ if there is a path from u to v via simple edges. Since simple edges are bi-directional, these equivalence classes are compatible with (and in general finer than) the decomposition of Γ into cells. For example, there is a 20-vertex admissible D_5-cell with three such equivalence classes—see Figure 2.

Let $U \subset V$ be one of the equivalence classes of vertices, and let $\Gamma(U)$ be the corresponding induced subgraph of Γ. ("Induced" means that all edges of Γ whose endpoints both belong to U are included in $\Gamma(U)$.) We call $\Gamma(U)$ a *simple component* of Γ; later, it will also make sense to call it a *molecular component*. As the example in Figure 2 illustrates, some arcs of Γ may connect distinct simple components; other arcs may be internal to a given component.

Note that the simple components of Γ are admissible, and obey the W-Compatibility, Simplicity, and W-Bonding Rules.

FACT 3.1 (The Local Polygon Rule). *If $\Gamma(U)$ is a simple component of an admissible W-graph Γ, then for all $i, j \in I$ and all r such that $2 \leqslant r \leqslant p_{ij}$, and all $u, v \in U$ such that $i, j \in \tau(u)$ and $i, j \notin \tau(v)$, we have*

$$N_{ij}^r(\Gamma(U); u, v) = N_{ji}^r(\Gamma(U); u, v)$$

under any of the following conditions:

(a) *$r = 2$, and there is an index $k \in \tau(v) - \tau(u)$;*
(b) *$r = 3$, and there exist indices $k, l \in \tau(v) - \tau(u)$ such that k is not bonded to i and l is not bonded to j in the Dynkin diagram of W (possibly $k = l$); or*
(c) *$r \geqslant 4$, and there is an index $k \in \tau(v) - \tau(u)$ such that k is not bonded to i or j in the Dynkin diagram of W (i.e., (b) except that $k = l$ is required).*

That is, under these limitations, the Polygon Rule holds for each simple component of Γ.

We omit the proof, except to say that conditions (a)–(c) are designed to force either the first or last step in every alternating r-step path of type (i, j) or (j, i) to be a simple edge. Since all internal edges in such paths are simple (recall the discussion in §2), it follows that such paths have at most one arc; all other edges are simple and have weight 1. Moreover, every internal vertex in such a path is necessarily in the same simple component as either the initial or terminal vertex. This yields

FACT 3.2. *Every subgraph of an admissible W-graph that is induced by a union of simple components obeys the Local Polygon Rule.*

DEFINITION 3.3. An admissible I-labeled graph Γ that satisfies the W-Compatibility, Simplicity, W-Bonding and Local W-Polygon Rules is said to be a *W-molecular graph*.

If Γ has only one simple component, then it is said to be a *W-molecule*.[3]

The content of Fact 3.1 is that the simple components of admissible W-graphs are W-molecules (thus we may also call them molecular components), and the content of Fact 3.2 is that any union of simple components in an admissible W-graph induces a W-molecular subgraph. The converse generally fails: for a given a W-molecular graph Γ, there need not exist an admissible W-graph into which Γ embeds. In such cases, we say that Γ is *unstable*; otherwise it is *stable*.

QUESTION 3.4. *Are there only finitely many stable W-molecules?*

EXAMPLE 3.5. There are 13 stable D_4-molecules (see Figure 3); none of them have internal arcs, and 11 of them are also D_4-cells. Only the 6-cycle and 12-cycle are not cells by themselves, but they do occur in cells with more than one molecular component. The remaining (unstable) D_4-molecules consist of a family of $6k$-cycles, one for each $k \geqslant 3$.

[3]We are open to suggestions for better terminology. Proteins?

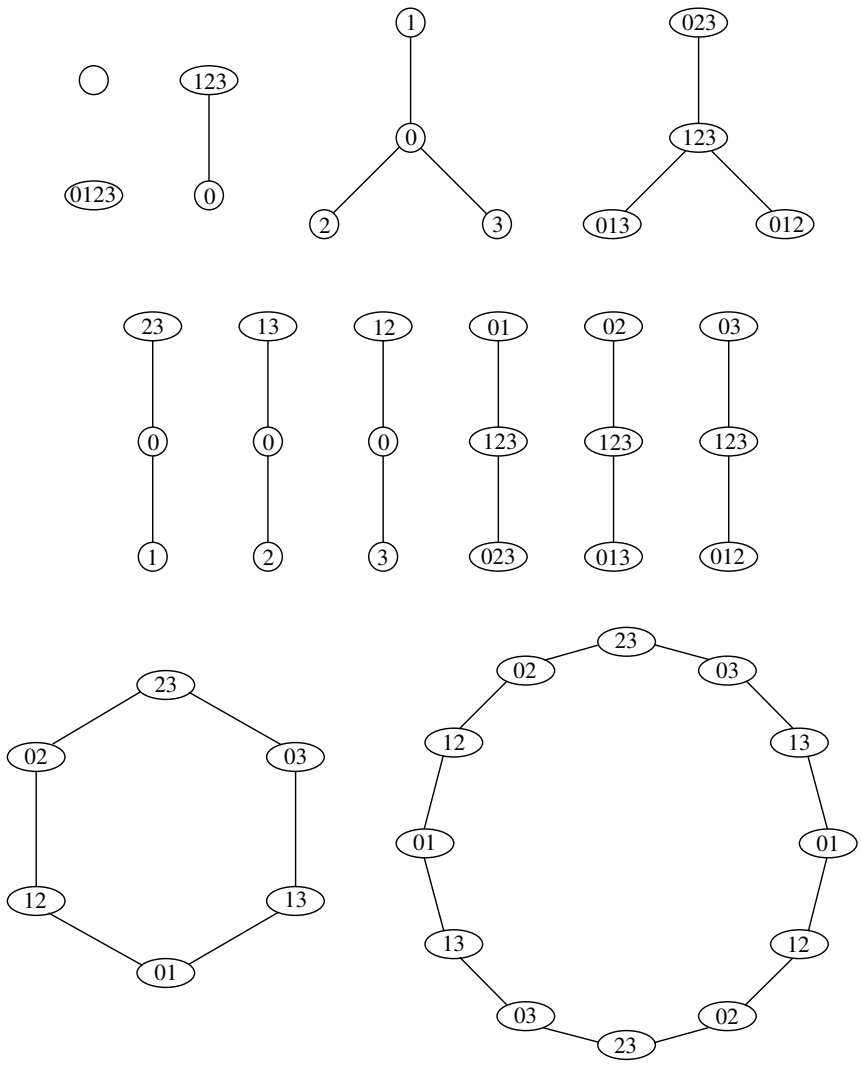

FIGURE 3: The stable D_4-molecules.

Every admissible W-graph may be constructed by binding together suitable collections of (stable) molecules with arcs, so one strategy for constructing all admissible W-cells from scratch is to first construct all W-molecules. However, as the D_4 example indicates, we still have the prospect of having infinitely many W-molecules to analyze, and the problem of identifying which ones are stable.

For $W = A_n$, we have generated all A_n-molecules for $n \leqslant 9$; remarkably, we found that not only are all of them stable, but they are precisely the Kazhdan-Lusztig cells! It is natural to ask whether this trend continues for all larger n, but Remark 3.8 below explains why it cannot.

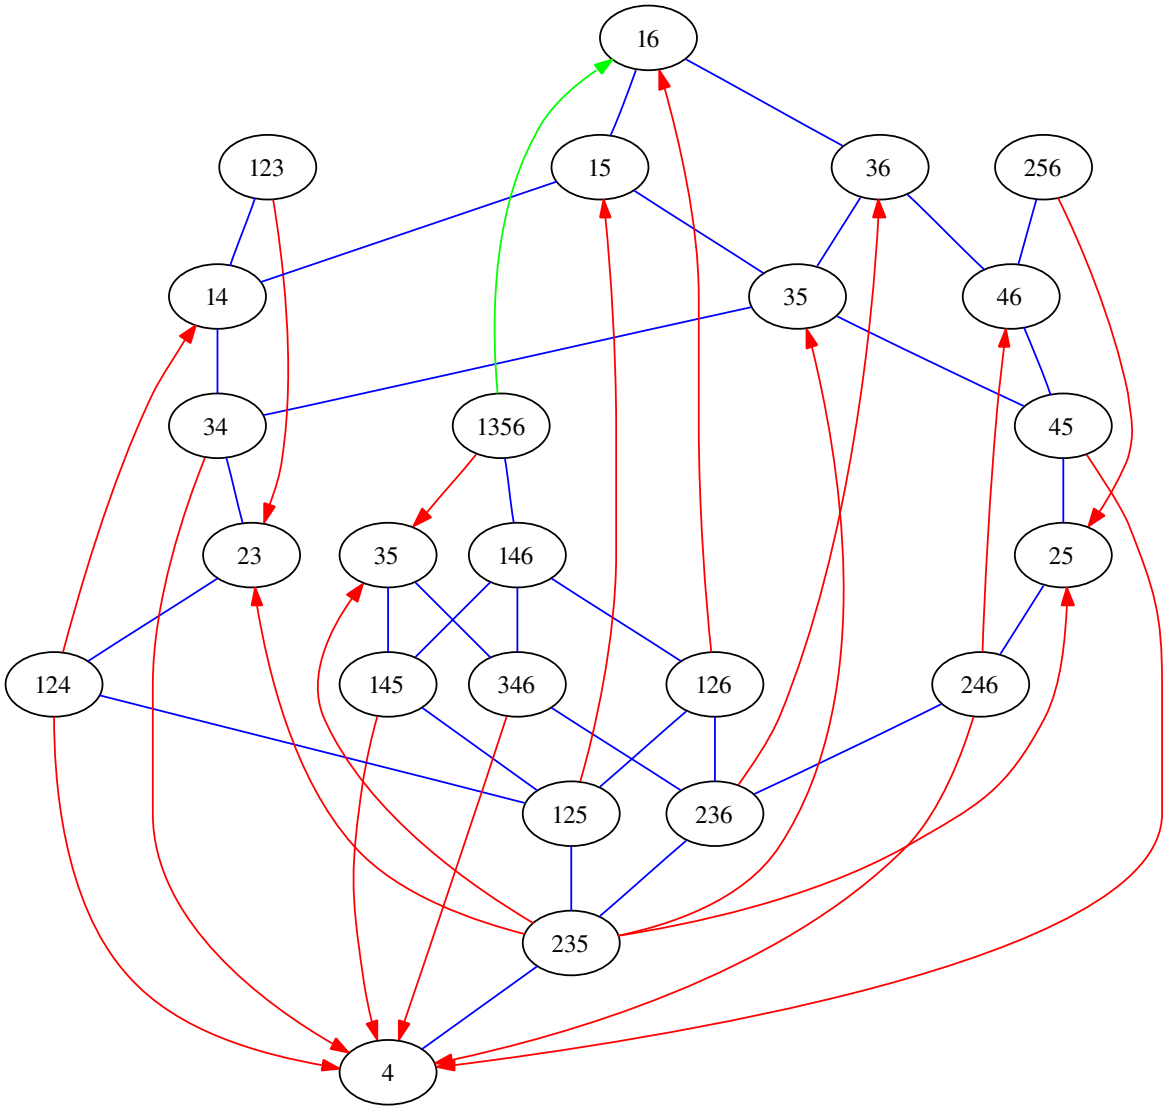

FIGURE 4: A generic E_6-molecule.

B. Generic molecules.

Suppose we are given an admissible I-labeled graph that satisfies the W-Compatibility, Simplicity and W-Bonding Rules, and consists only of simple edges. Assuming it is also connected, can we add suitable arcs so that the result is a W-molecule (i.e., satisfies the Local W-Polygon Rule)?

In order to add an arc $u \to v$, it must be the case that $\tau(u) \supsetneq \tau(v)$. In addition, given that the original graph is connected, it must also be the case that u and v are an even distance apart, or the result would fail to be bipartite. If we treat the weights $m(u \to v)$ as indeterminates, then every alternating path contributing to an instance of the Local W-Polygon Rule has at most one such arc, so this rule amounts to a nonhomogeneous linear system in these variables. One may regard the general solution of this system as a "generic" molecule (assuming of course that the system is consistent).

More precisely, we define Γ to be a *generic* W-molecule if it satisfies all of the defining conditions for a W-molecule, except that

(1) it has arc weights that are affine-linear functions in some set of variables, and
(2) these variables parameterize the most general solution of the Local Polygon Rule with the given set of vertices and simple edges.

We say that Γ is *stable* if some specialization of its variables is a stable W-molecule; otherwise, it is *unstable*.

The *dimension* of a generic W-molecule is the dimension of the parameter space for its arc weights; i.e., the number of independent variables.

In low ranks, all generic W-molecules are 0-dimensional and all arc weights are 0 or 1. The smallest generic molecule we have found that has positive dimension is the 24-vertex E_6-molecule displayed in Figure 4 (and its dual). In this illustration, the red edges are arcs of weight 1; the unique green edge is an arc of indeterminate weight—any nonnegative integer assigned to it will produce an E_6-molecule. If the assigned weight is 0, then the resulting graph is an E_6-cell (in fact, a Kazhdan-Lusztig cell), and no other specialization yields a stable molecule.

It should be emphasized that the arcs of a generic molecule are determined from its simple edges by straightforward linear algebra. Furthermore, these linear equations have an especially simple form.

Probable Fact 3.6. *The constraints imposed by the Local Polygon Rule may be reduced to a linear system in which each equation has one of the forms* $m(u \to v) = 0$, $m(u \to v) = 1$, *or* $m(u \to v) = m(u' \to v')$ *for various arcs* $u \to v$ *and* $u' \to v'$.

We call this a "probable fact" because we have (so far) proved it only in the simply-laced case. Even if it turns out to be false for (say) B_8, the ramifications would be minimal; nothing that follows will be significantly affected beyond the fact that it may be more cumbersome to impose the condition that all arc weights in a molecular graph must be nonnegative integers.

Example 3.7. Consider the A_3-graph in Figure 5. If we regard the arc weights in this graph as indeterminates and impose the Local Polygon Rule, an analysis of the alternating paths of type $(1, 3)$ and $(3, 1)$ from $u = 13$ to $v = 2$ yields the relation

$$m(13 \to 1) + m(12 \to 2) = m(13 \to 3) + m(23 \to 2), \tag{3.1}$$

which seems to contradict Probable Fact 3.6. However, analyzing the alternating paths of type $(1, 2)$ and $(2, 1)$ from $u = 12$ to $v = 3$ yields the relation $m(12 \to 2) = m(13 \to 3)$, and a similar analysis of paths from $u = 23$ to $v = 1$ yields $m(13 \to 1) = m(23 \to 2)$, so (3.1) is redundant. In fact, any solution of the latter pair of equations satisfies the *full* Polygon Rule and thus yields an A_3-graph.

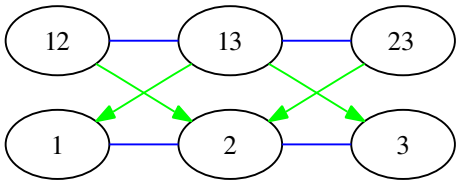

FIGURE 5: An A_3-graph with variable arc weights.

REMARK 3.8. Recall that McLarnan and Warrington [**MW**] have shown that there are Kazhdan-Lusztig cells for A_{15} that have edge-weights > 1. Furthermore, it is known that all type A Kazhdan-Lusztig cells consist of single molecules.[4] Comparing this with Probable Fact 3.6, we see that there must be at least one (stable) generic A_{15}-molecule with at least one free parameter. It is conceivable that all (nonnegative integer) specializations of these parameter(s) yield stable molecules, but Kazhdan-Lusztig cells for A_n are known to be combinatorially rigid, so even if by some miracle all A_{15}-molecules turn out to be stable, they cannot all occur in Kazhdan-Lusztig cells.

REMARK 3.9. Not every stable molecule occurs in some Kazhdan-Lusztig cell. For example, there is a 40-vertex generic E_6-molecule that is stable and 0-dimensional, but occurs only in cells that are not Kazhdan-Lusztig cells.

[4]Thanks to David Vogan for pointing this out to me.

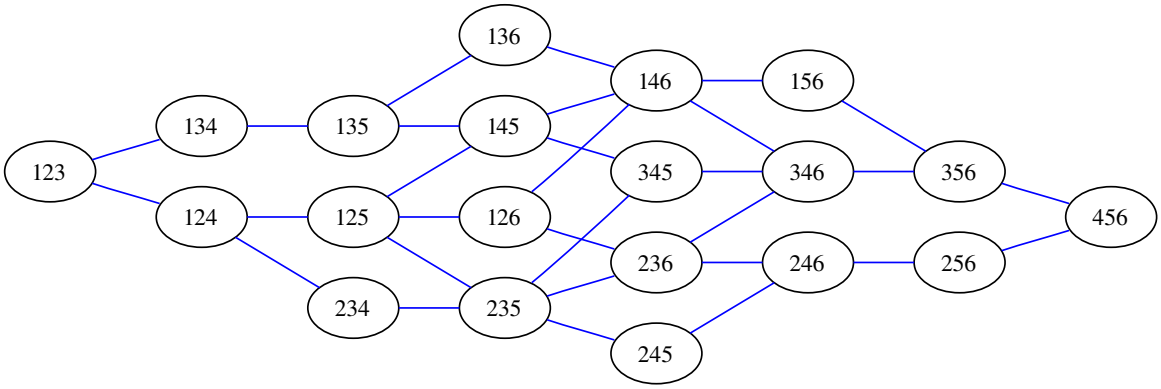

FIGURE 6: The 3-uniform compatibility graph for E_6.

4. Uniform Molecules

In this section, we assume that W is simply-laced and irreducible.

We say that a W-molecule $\Gamma = (V, m, \tau)$ is *k-uniform* if the τ invariant of every vertex has cardinality k. Since there are no proper inclusions of k-subsets, it follows immediately that uniform molecules cannot have arcs. In particular, they are trivially generic.

If Γ is k-uniform, one can think of the τ invariant of a vertex v as a configuration of k marbles occupying distinct nodes of the Dynkin diagram of W. According to the W-Compatibility Rule, following a simple edge from v corresponds to moving one marble to an adjacent unoccupied node.[5] Conversely, by the W-Bonding Rule, for each legal way to move a marble in the configuration of $\tau(v)$, there must be a unique vertex u adjacent to v whose τ invariant corresponds to that move.

It follows that if there are r vertices in Γ with τ invariant J, then there must also be r vertices with τ invariant K for each k-subset K adjacent to J in the compatibility graph of W, and the (simple) edges between these two r-tuples of vertices must form a perfect matching. Since W is assumed to be irreducible, we can reach any k-subset from any other k-subset by moving marbles around, so there must be exactly r vertices in Γ with τ invariant K for every k-subset $K \subset I$.

Thus, one may view every k-uniform W-molecule as an r-fold cover of $\mathrm{Comp}_k(W, S)$, the subgraph of $\mathrm{Comp}(W, S)$ induced by the k-subsets of I. For example, every 3-uniform E_6-molecule is an r-fold cover of the graph in Figure 6, for some $r \geqslant 1$.

On the other hand, since we have not yet imposed the Local Polygon Rule, it should be expected that not every r-fold cover of $\mathrm{Comp}_k(W, S)$ is a W-molecule. Indeed, let us consider the possibilities for a 2-step alternating path of type (i, j) for some covering graph Γ. This can be achieved only by moving two marbles, one occupying the j-th node (and moved first), and the second occupying the i-th node (and moved second, to some position other than i or j). One can move the marbles in either order to achieve the same result, and it is not hard to see that the Local Polygon Rule forces the endpoints of the

[5] This requires the Dynkin diagram of W to have no 4-cycles.

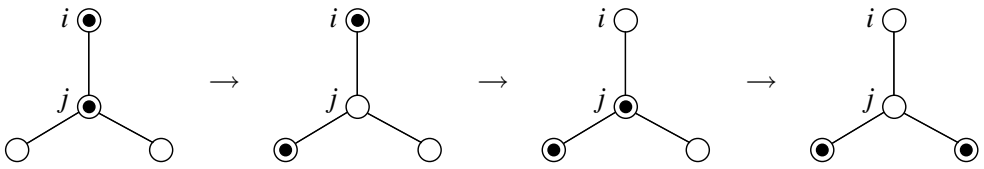

FIGURE 7: An alternating 3-step path.

two alternating paths to be the same vertex in the covering graph.

In other words, each 4-cycle in $\mathrm{Comp}_k(W,S)$ that is obtained by moving two marbles to two new nodes and back[6] is *flat*; i.e., it lifts to r disjoint 4-cycles in the covering graph Γ. In particular (taking $r=1$), the graph $\mathrm{Comp}_k(W,S)$ itself always satisfies the 2-step Local Polygon Rule.

Since we are assuming W is simply-laced, the only remaining consequences of the Local Polygon Rule involve 3-step alternating paths. It is not hard to show that such a path of type (i,j) can exist only if the Dynkin diagram has a fork at node j, and two unoccupied nodes are bonded to j in addition to (occupied) node i (see Figure 7). Given that W is a finite Weyl group, there cannot be a second fork at node i, so Γ cannot have any alternating 3-step paths of type (j,i) (and hence, cannot be a W-graph). However, this does not contradict the Local Polygon Rule, since the terminal vertex of the 3-step path illustrated in Figure 7 has a τ invariant that contains exactly two nodes not in the initial τ invariant, and both of these nodes are bonded to j, so these alternating paths are ignored by the Local Polygon Rule (see part (b) of Fact 3.1). Summarizing,

FACT 4.1. *If W is simply-laced and irreducible, then the k-uniform W-molecules are the connected r-fold covers of $\mathrm{Comp}_k(W,S)$ in which every 4-cycle is flat. Furthermore,*

(a) *the k-subset graph $\mathrm{Comp}_k(W,S)$ is a W-molecule,*
(b) *if the diagram of W is linear (i.e., $W=A_n$), then the only uniform W-molecules are the k-subset graphs (i.e., $r=1$), and these molecules are W-cells, and*
(c) *if the diagram of W has a fork and $2 \leqslant k \leqslant |I|-2$, then there are infinitely many k-uniform W-molecules, and none of them are cells.*

The remaining issues for uniform molecules are: (1) how to identify in a systematic way the "flat" r-fold covers of $\mathrm{Comp}_k(W,S)$, and (2) how to decide which ones are stable.

Flat covers.

To generate r-fold covers with flat 4-cycles, one approach is to number the vertices with a given τ invariant from 1 to r, and think of these numbers as defining a partition of the covering graph into r layers. The extra information needed to specify the cover is a collection of permutations of $\{1,\ldots,r\}$, one for each edge in $\mathrm{Comp}_k(W,S)$.

On grounds of sanity alone, we would prefer that as many of these permutations as possible are trivial, and only connect vertices in the same layer. So we can start by

[6]One can show that *every* 4-cycle in $\mathrm{Comp}_k(W,S)$ has this form.

choosing a spanning tree of $\mathrm{Comp}_k(W,S)$, and insist that the permutation associated to each edge in the spanning tree is trivial. We can then start adding additional edges of $\mathrm{Comp}_k(W,S)$ to this spanning tree as long as the new edge creates a 4-cycle, and repeat until this is no longer possible. The point is that the flatness of the 4-cycle forces the permutation associated to each new edge to also be trivial.

If we are lucky, the process will end when we run out of edges, thereby forcing every edge-permutation to be trivial. In this case, there could only be one layer; otherwise, the covering graph would be disconnected. This is exactly what happens in type A, and explains part (b) of Fact 4.1. (It also happens when $k = 1$ or $|I| - 1$.)

On the other hand, if W has a forked diagram and $2 \leqslant k \leqslant |I|-2$, then there will be one or more 6-cycles that need not be flat.[7] Indeed, one can see from Figure 7 that alternating 3-step paths in $\mathrm{Comp}_k(W,S)$ occur in pairs: one has a choice of moving the marble at the fork into either of the two adjacent unoccupied nodes. These two paths have the same endpoints in $\mathrm{Comp}_k(W,S)$, thereby creating a 6-cycle. However, these endpoints need not be the same in the covering graph, and it is not hard to construct explicit examples (with flat 4-cycles) where this happens.

Example 4.2. (a) In the case of D_4, the 2-uniform compatibility graph is a 6-cycle, and a spanning tree consumes all but one of its edges. Thus, an r-fold cover of $\mathrm{Comp}_2(D_4)$ is parameterized by a single permutation. However, since a molecule must be connected, this single permutation must be an r-cycle, and all such permutations yield the same D_4-molecule up to isomorphism.

(b) Consider the 3-uniform compatibility graph for E_6 (recall Figure 6). It is not hard to discover that one can choose a spanning tree of this graph and add 4-cycle-creating edges until (say) only the edges $123—124$ and $356—456$ remain. Thus the 3-uniform E_6-molecules may be parameterized by pairs of permutations of r objects for various $r \geqslant 1$. Of course, not all such pairs will parameterize connected graphs, and it is tricky to decide when two such graphs will be isomorphic. However, these are moot points, since it turns out that only four of the 3-uniform E_6-molecules are stable (see Fact 4.4).

Stability.

Suppose Γ is a k-uniform W-molecule of index r. How do we decide if Γ is stable?

Of course we should assume that $2 \leqslant k \leqslant |I| - 2$ and that W has a forked diagram; otherwise, Γ is a cell and the question is easy.

To answer this in general, we need to know about all of the other W-molecules that could occur together with Γ in an admissible W-cell. While the details are nontrivial (and omitted), there are a few easy comments in this direction that we should make, since they narrow the scope of this problem significantly.

The main point to make is that, as mentioned previously in Example 3.5, the only 2-uniform D_4-molecules that are stable are the 1-fold cover and the 2-fold cover; i.e., a

[7] The exact number of such 6-cycles is $\binom{n-2}{k-2}$, where $n = |I|$.

6-cycle and a 12-cycle. (The fact that the r-fold covers for $r \geqslant 3$ are not stable is explained in Section 5 of [**S2**].) Thus by restriction, this carries over to any of the "type D_4" 6-cycles in $\mathrm{Comp}_k(W, S)$ afforded by moving marbles around the fork as in Figure 7. Thus,

FACT 4.3. *In a stable uniform W-molecule Γ, every 6-cycle of type D_4 in the k-uniform compatibility graph of W lifts to a disjoint union of 6-cycles and 12-cycles in Γ.*

For example, recall from Example 4.2(b) that a 3-uniform E_6-molecule Γ of index r is parameterized by pairs of permutations of r objects, say (π_1, π_2). The first permutation specifies the twisting of edges between layers that connect the r vertices with τ invariants 123 and 124, and the second permutation specifies the analogous edge-twisting between τ invariants 356 and 456.

According to Fact 4.3, a necessary condition for Γ to be stable is that π_1 and π_2 must be involutions. This severely cuts down the number of distinct possibilities. If we contract all other edges in the graph (their structure is completely determined), what remains consists of r points, together with a collection of black edges (from π_1) and white edges (from π_2). The fact that the permutations are involutions means that each point has one black edge or loop incident to it, and one white edge or loop. Given that the graph must be connected, this means that there are only two possibilities: an alternating black-white r-cycle for even r (we call these *cyclic* r-fold covers), or an alternating black-white path of length $r - 1$, with loops of the appropriate colors at both endpoints (we call these *linear* r-fold covers).

For example, the cyclic 6-fold cover is obtained by choosing

$$\pi_1 = (1, 2)(3, 4)(5, 6), \quad \pi_2 = (2, 3)(4, 5)(6, 1),$$

and the linear 7-fold cover is obtained by choosing

$$\pi_1 = (1, 2)(3, 4)(5, 6)(7), \quad \pi_2 = (1)(2, 3)(4, 5)(6, 7).$$

Using the binding voodoo described in Section 5, we are able to show

FACT 4.4. *The stable 3-uniform E_6-molecules are the linear 1-fold and 3-fold covers and the cyclic 2-fold and 6-fold covers of the 3-uniform compatibility graph.*

The classification of stable 4-uniform E_8-molecules is likely to be very interesting.

We should add that a fortuitous accident that has made the classification of admissible cells for E_6 and D_6 easier is that there are only finitely many non-uniform generic molecules, and all of them turn out to be stable! We are skeptical that this holds in E_7 and E_8; nevertheless, we should at least raise

PESSIMISTIC QUESTION 4.5. *Are there finitely many non-uniform generic molecules? Is every unstable generic molecule uniform?*

Bear in mind that we haven't touched on the multiply-laced cases.

PROBLEM 4.6. *What is the analogue of a uniform molecule in the multiply-laced cases?*

5. Bindings

We now turn to the problem of using arcs to bind a collection of W-molecules into a W-graph. In keeping with the philosophy that we should exhaust linear algebraic methods first before confronting nonlinear ones (or worse, Diophantine issues), we first address the easier problem of binding a collection of W-molecules into a molecular graph; i.e., we impose only the Local Polygon Rule on the output, rather than the full Polygon Rule.

A. Binding spaces and p-molecules.

Suppose Γ_1 and Γ_2 are W-molecules. If we attempt to create a molecular graph by adding arcs between Γ_1 and Γ_2, we cannot do so arbitrarily, or there is a risk that we will create a graph that is not bipartite, and hence inadmissible.

Thus to be careful, we should be working in a more refined category of *molecules-with-parity* (or *p-molecules*, for short): each vertex must have a designated parity (odd or even), and all edges, both simple edges and arcs, must have endpoints of opposite parity. Of course, molecules are internally connected by simple edges, so designating the parity of any one vertex determines the parity of all vertices within the molecule. More generally, the vertex parities of a p-molecular graph are determined by a choice of parity for one vertex from each weak connected component.

Now suppose that $\Gamma_1 = (V_1, m, \tau)$ and $\Gamma_2 = (V_2, m, \tau)$ are two W-molecules with parity. For each opposite-parity vertices $v_1 \in V_1$ and $v_2 \in V_2$ such that $\tau(v_1) \supsetneq \tau(v_2)$, we may hypothesize the existence of an arc $v_1 \to v_2$ with some unknown weight $m(v_1 \to v_2)$, possibly 0. We define the *W-binding space* $B(\Gamma_1 \to \Gamma_2)$ to be the set of assignments for these arc weights that satisfy the Local W-Polygon Rule.

Recall that the Local Polygon Rule imposes linear equations on all unknown arc weights. However, unlike the case of arcs internal to a single molecule, all alternating paths from Γ_1 to Γ_2 relevant to the Local Polygon Rule will necessarily use exactly one arc of unknown weight. Thus the resulting linear equations are homogeneous, and the trivial solution is always available. Furthermore, no alternating path from Γ_1 to Γ_2 may use any arc internal to Γ_1 or Γ_2, so the binding space is unchanged if we delete all such arcs from Γ_1 or Γ_2, or replace them with their generic counterparts.

If we want the p-molecular graph corresponding to some point in the binding space to be admissible, we should also insist that the assigned arc weights are nonnegative integers. However, we prefer to think of the W-binding space as a $\mathbb{Z}$-module. In any case, Probable Fact 3.6 shows that every point in a binding space is obtained by setting certain arc weights equal to 0, and certain arc weights equal to each other, so in practice there is no extra difficulty involved in requiring the use of nonnegative integers.

Example 5.1. Consider the D_5-cell in Figure 2. It is a binding of three p-molecules: two copies of a 5-vertex molecule Γ_0, and one copy of the 2-uniform D_5-molecule Γ_1 of index 1. The fact that there are arcs in this cell between Γ_1 and both copies of Γ_0 (and in both directions) provides proof that the binding spaces $B(\Gamma_0 \to \Gamma_1)$ and $B(\Gamma_1 \to \Gamma_0)$ are

both nonzero. Moreover, it is not hard to show that both binding spaces are 1-dimensional; all of the arcs from a copy of Γ_0 to a copy of Γ_1 (say) that have unit weight in the D_5-cell must have equal weight in the binding space.

If Γ is any p-molecule, we let $-\Gamma$ denote the p-molecule obtained by reversing the parity of all vertices in Γ. Note that the binding spaces $B(\Gamma_1 \to \Gamma_2)$ and $B(-\Gamma_1 \to -\Gamma_2)$ are naturally isomorphic. Of special interest are *self-bindings*; i.e., nontrivial points in the binding spaces $B(\Gamma \to \Gamma)$ (*even* self-bindings) and $B(\Gamma \to -\Gamma)$ (*odd* self-bindings).

To avoid confusion, we should clarify that the even self-binding spaces $B(\Gamma \to \Gamma)$ involve two distinct copies of Γ, not a single copy. We could also define internal binding spaces involving a single copy of Γ, but the resulting spaces would amount to nothing but an alternative way to define generic molecules.

Recall that a generic molecule is defined to be stable if some specialization of it occurs as a component of some admissible W-graph. Analogously, we define a binding space $B(\Gamma_1 \to \Gamma_2)$ to be *stable* if there is a nonzero element of this space such that the resulting two-molecule graph is stable; otherwise, it is *unstable*.

B. The bindability graph.

We define the *bindability graph* $\mathrm{BG}(W, S)$ to be the directed graph whose vertices are the isomorphism classes of generic W-molecules and whose edges are of the form

$$\Gamma \to \Gamma' \quad \text{whenever} \quad B\big(\pm\Gamma \to \pm\,\Gamma'\big) \neq 0;$$

i.e., there is an edge whenever there exists a choice of parities for both Γ and Γ' so that the corresponding binding space is nonzero. Similarly, we define the *stable bindability graph* $\mathrm{BG}_{\mathrm{st}}(W, S)$ to be the subgraph of $\mathrm{BG}(W, S)$ obtained by keeping only the vertices corresponding to stable molecules and the edges corresponding to stable bindings.

Easy Fact 5.2. *Every admissible W-cell may be constructed by binding together one or more generic W-molecules from some strongly connected component of $\mathrm{BG}(W, S)$ and specializing the variables. Similarly, the same is true for $\mathrm{BG}_{\mathrm{st}}(W, S)$, and more generally for any graph that interpolates between $\mathrm{BG}_{\mathrm{st}}(W, S)$ and $\mathrm{BG}(W, S)$.*

The stable bindability graph is the most useful of these graphs, but it is difficult to pin it down exactly until *after* one has a classification of admissible W-cells. Starting from scratch, one would initially build something close to the full bindability graph, and refine it while accumulating information about which molecules and bindings are stable.

Regardless, the stable bindability graph provides a natural way to partition admissible W-cells into families: for each strongly connected component $\mathcal{C}$ of $\mathrm{BG}_{\mathrm{st}}(W, S)$, family $\mathcal{C}$ consists of all admissible W-cells whose constituent molecules all belong to $\mathcal{C}$.

Recall that when the strongly connected components of any graph are contracted to points, the remaining edges form an acyclic directed graph whose transitive closure is a poset. In Figure 8, we have displayed a few of these posets for the subgraph of $\mathrm{BG}(W, S)$

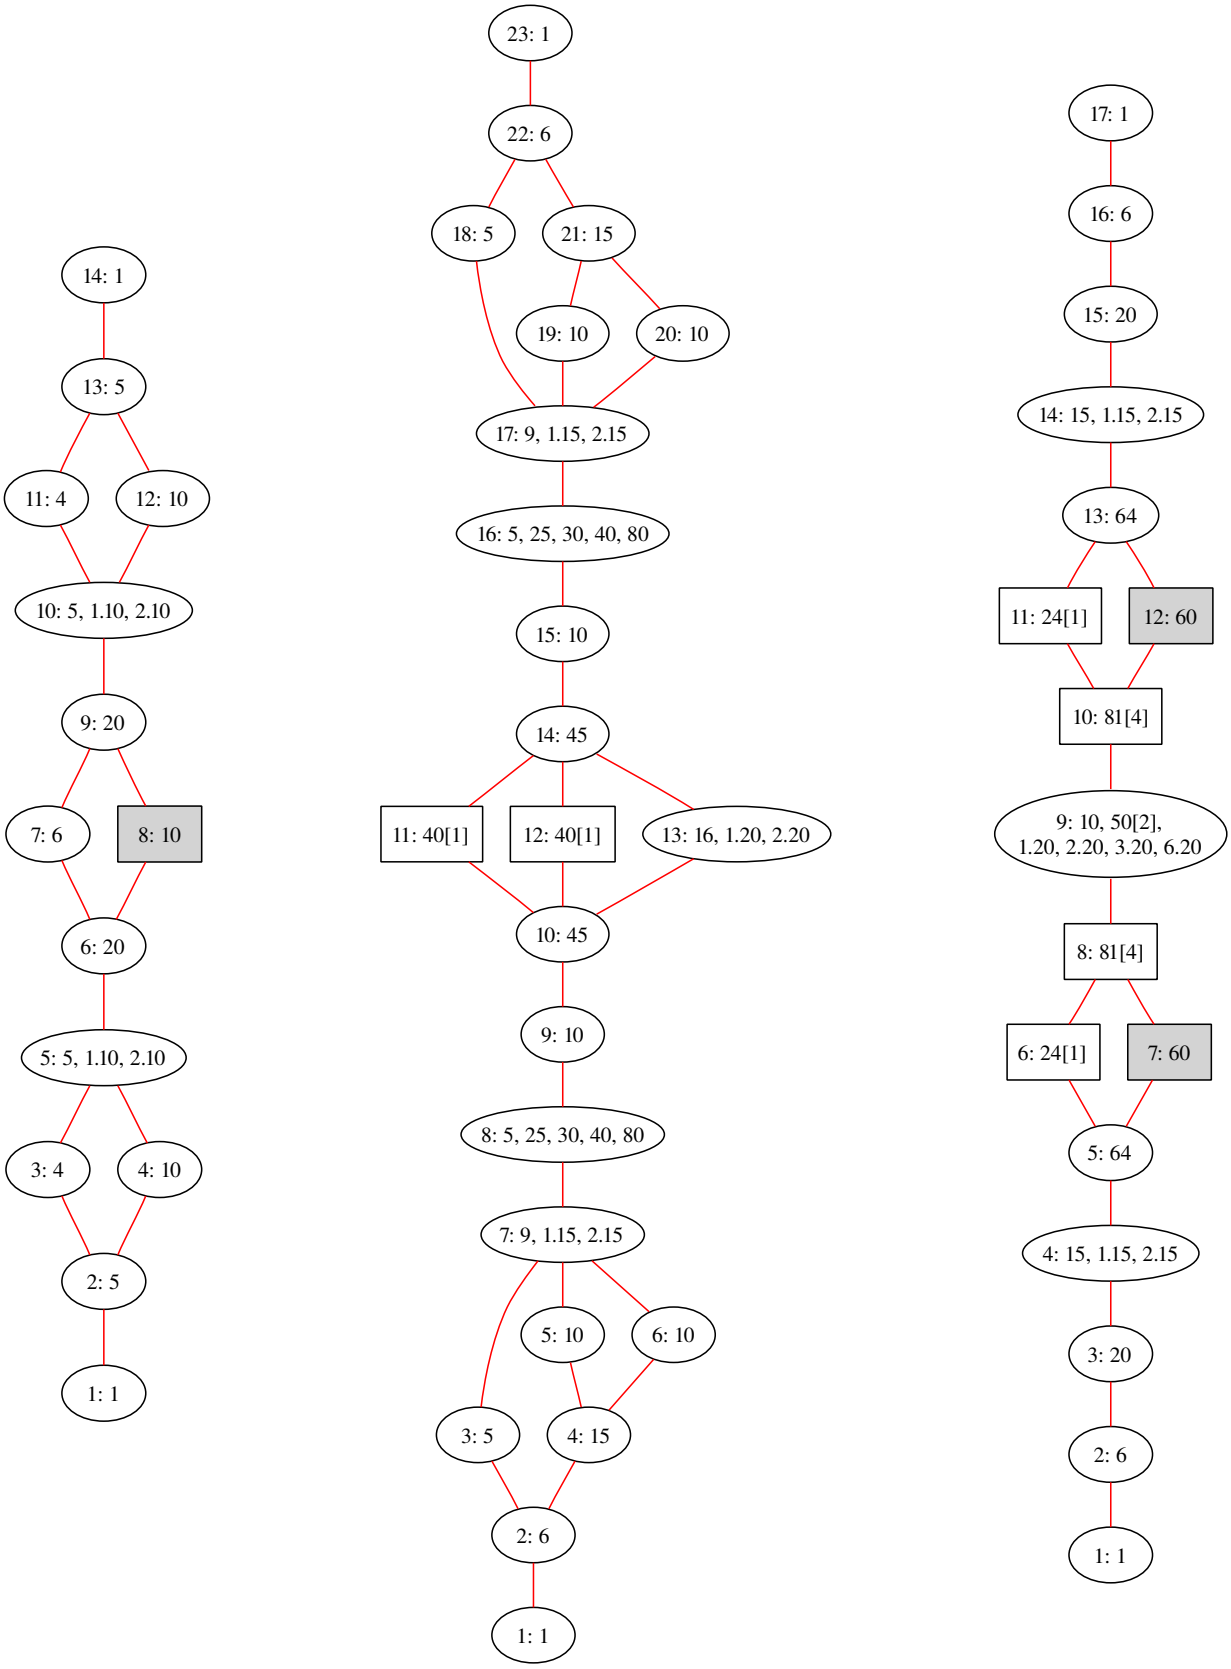

FIGURE 8: The partial ordering of molecular families for D_5, D_6, and E_6.

induced by the stable generic W-molecules. In general, this graph may have more edges and fewer strongly connected components than $\mathrm{BG}_{\mathrm{st}}(W,S)$ (see Remark 5.3(b) below) but it is easier to compute.

Regarding the annotations in Figure 8, each strongly connected component is labeled by a number, followed by a list of vertex cardinalities for the generic molecules in that component. The vertex cardinalities are further refined by suffixes indicating the dimension of the molecule (no suffix means 0-dimensional), and (when the molecule is uniform) prefixes indicating the index of the covering map. For example, the label $7:9$, 1.15, 2.15 means that component 7 contains three 0-dimensional molecules: one with 9 vertices, and a 1-cover and 2-cover of a uniform molecule with 15 vertices. The label $12:40[1]$ means that component 12 consists of a single 1-dimensional generic molecule with 40 vertices. Also, some of the components are drawn as boxes rather than ovals. These boxed components are singletons with non-trivial self-bindings: odd if shaded, even if unshaded.

REMARK 5.3. (a) If a component of $BG(W,S)$ is a singleton $\mathcal{C} = \{\Gamma\}$ and Γ has no self-bindings, then Γ can only appear by itself in an admissible W-cell. In that case, one simply needs to check that Γ specializes to a W-graph (thereby confirming it to be stable). One may then conclude that specializations of Γ are the only W-cells in family $\mathcal{C}$. This leaves only components 5, 8, 10 in D_5, 7, 8, 11–13, 16, 17 in D_6, and 4, 6–12, 14 in E_6.

(b) For D_5 and E_6 the extra (unstable) edges we used in constructing Figure 8 have no effect on the partition into strongly connected components, but in the case of D_6, it turns out that components 8 and 16 split into two components each in the stable bindability graph: the 30-vertex molecules are each in singleton components and have odd self-binding.

These W-cell families are quite remarkable; we mention here a few of the questions raised by the data we have gathered so far.

QUESTION 5.4. *Let $\mathcal{C}$ be a strongly connected component of* $\mathrm{BG}_{\mathrm{st}}(W,S)$.

(a) *Is there a unique generic molecule $\Gamma = \Gamma_{\mathcal{C}}$ in $\mathcal{C}$ that occurs as a molecular component of every admissible W-cell in family $\mathcal{C}$?*
(b) *Is there a specialization of $\Gamma_{\mathcal{C}}$ that yields an admissible W-cell?*
(c) *Are the W-representations generated by admissible W-cells from distinct families orthogonal (i.e., have no irreducible constituents in common)?*

REMARK 5.5. (a) We know that Question 5.4(b) fails for $W = H_3$. There is a component of $\mathrm{BG}_{\mathrm{st}}(H_3)$ that consists of a single 4-vertex molecule Γ with odd self-binding, and the only admissible H_3-cell in this family is a binding of two copies of Γ. So we are not proposing to extend the domain of this question beyond the crystallographic groups.

(b) Regarding (c), we could construct an undirected graph on the set of admissible W-cells by declaring two such cells adjacent if their W-representations have at least one irreducible constituent in common. It is reasonable to speculate that the W-cell families defined above are the connected components of this graph.

Another remarkable property of the W-representations generated by admissible W-cells is that they seem to depend (up to isomorphism) only on the multiset of generic molecules used to construct them, and not on the weights of the arcs that bind them. If true, this would explain (in my opinion) the phenomenon of "combinatorial rigidity" for cell representations discussed briefly in [**S1**] and [**S2**]. An even stronger version of this is

QUESTION 5.6. *Is it possible to attach a virtual W-character to each (stable) generic W-molecule so that the character of every admissible W-graph is the sum of the virtual characters of its molecular components?*

There is a well-known partition of Kazhdan-Lusztig cells in which two left cells are placed in the same block if and only if they occur in the same 2-sided cell, or equivalently, if the W-representations they generate include the same special representation of W among their constituents. In particular, a positive answer to Question 5.4(c) would imply a positive answer to the following.

QUESTION 5.7. *Are the molecular components of the left Kazhdan-Lusztig cells in a two-sided cell all contained in a single strongly connected component of* $\mathrm{BG}_{\mathrm{st}}(W,S)$*?*

It would be nice if there were a bijection between two-sided cells and the components of $\mathrm{BG}_{\mathrm{st}}(W,S)$, but there are 27 two-sided D_6-cells whereas $\mathrm{BG}_{\mathrm{st}}(D_6)$ has 25 components. In particular, the two sets of D_6-molecules in components 8 and 16 in Figure 8 are each unions of two components of $\mathrm{BG}_{\mathrm{st}}(D_6)$ (recall Remark 5.3(b)). On the other hand, they contain the (left) molecular components of 3 two-sided Kazhdan-Lusztig cells.

6. Cell Synthesis

Once we have collected enough information about W-molecules and their bindings, we are ready to classify admissible W-cells.

A. Encoding the parameter space.

Suppose $\Gamma_1, \Gamma_2, \ldots$ is a list of generic molecules with parity that are selected from a single component of the bindability graph for W. Although we may not know yet whether any molecules in our list are stable, we want to determine all admissible W-graphs that can be built using n_1 copies of Γ_1, n_2 copies of Γ_2, and so on. We would expect that there will be infinitely many such graphs, but perhaps only finitely many that are cells.

By Probable Fact 3.6, we can select bases for each binding space $B(\Gamma_k \to \Gamma_l)$ and the internal arc spaces for each Γ_k so that the points in these spaces with nonnegative integer arc weights are precisely the nonnegative integer combinations of the basis elements. In this way, the free parameters needed to specify a W-molecular graph Γ with the desired components may be encoded by a collection of $n_k \times n_l$ matrices of nonnegative integers.

More precisely, if the binding space $B(\Gamma_k \to \Gamma_l)$ has dimension d, then the arcs from each copy of Γ_k to each copy of Γ_l are encoded by a d-tuple of $n_k \times n_l$ matrices, say $A^1_{kl}, \ldots, A^d_{kl}$. If $d = 1$, we may omit the superscript and simply write A_{kl}.

It is important to note that the case $k = l$ requires special considerations. Here, the matrices $A^1_{kk}, A^2_{kk}, \ldots$ are square, but the diagonals must be regarded as identically 0 by convention. The point is that a diagonal entry would encode the weights of certain arcs from one copy of Γ_k to the *same* copy. While there is certainly a natural set of arcs internal to each Γ_k; namely, the ones that participate in its internal binding space, the dimension of this space need not be the same[8] as the dimension of the even self-binding space $B(\Gamma_k \to \Gamma_k)$. Thus, we should encode the choice of internal arc weights in the n_k copies of Γ_k by an e-tuple of $n_k \times n_k$ diagonal matrices $B^1_k, \ldots, B^e_k$, where $e = \dim \Gamma_k$.

If nature is kind to us, there will be a natural identification between the internal arc space of the molecule Γ_k and its even self-binding space, in which case we can drop the convention of zero diagonals and replace (say) A^1_{kk} with $B^1_k + A^1_{kk}$.

While the structure of these parameter sets may seem complicated, one finds that in practice, if molecules are chosen from within a strongly connected component of the bindability graph, then most binding spaces have dimension 0 or 1, and if $B(\Gamma_k \to \Gamma_l)$ is nonzero, then the opposite space $B(\Gamma_k \to -\Gamma_l)$ is usually 0; i.e., the relative parity of all molecules is often forced.

Example 6.1. In Figure 9, we provide a graphic representation of the binding space dimensions for a few of the more complicated components of the bindability graphs for D_5, D_6 and E_6. Each molecule is represented by two nodes, corresponding to the two

[8]Update: This is not clear. Certainly one needs to keep in mind that the internal binding space is affine, whereas the even self-binding space is central. It may still be that there *is* a natural identification between the two, in which case this entire section may be more complicated than it needs to be.

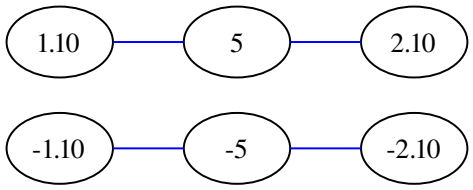

(a) Component 5 of D_5.

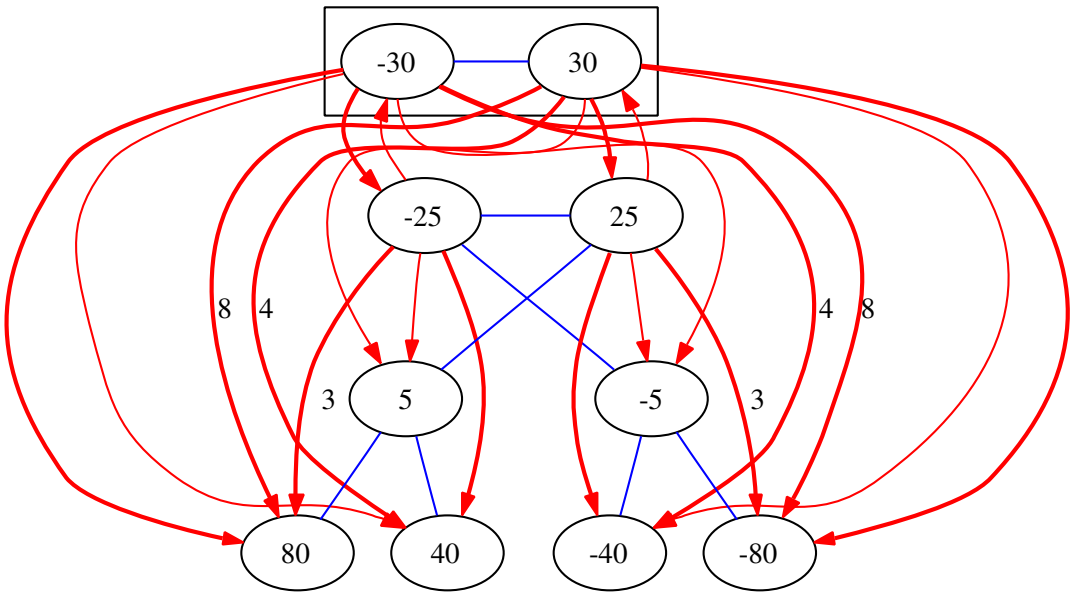

(b) Component 8 of D_6.

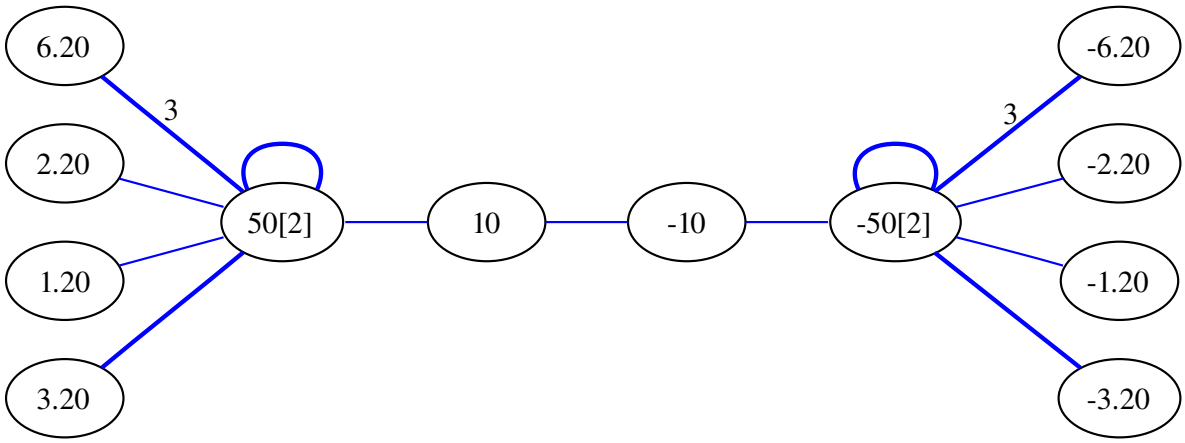

(c) Component 9 of E_6.

FIGURE 9: Binding space dimensions for various components of BG(W, S).

ways of assigning parity. Blue edges are abbreviations for spaces of equal dimension in both directions. Thin edges are used for 1-dimensional spaces; thick edges for spaces with dimension $\geqslant 2$. If the dimension is $\geqslant 3$, it is labeled. Note that an edge (necessarily blue) between a node and its opposite indicates a nonzero odd self-binding space; a loop indicates an even self-binding space. Nontrivial internal binding spaces are indicated in the node labels; e.g., the suffix "[2]" indicates an internal space of dimension 2.

In Figure 9(b), the two 1-dimensional binding spaces from the 25-vertex p-molecules to the 30-vertex p-molecules turn out to be unstable. This isolates the two 30-vertex p-molecules in a separate component of the *stable* bindability graph for D_6, as noted previously in Remark 5.3(b). In Figure 9(c), the odd-self binding space for the 10-vertex p-molecule also turns out to be unstable; deleting the corresponding edge disconnects the graph and cuts in half the number of distinct p-molecules one needs to consider.

B. Imposing the braid relations.

Any nonnegative integer specialization of the above parameter space will produce an admissible W-molecular graph Γ. Determining which specializations are W-graphs amounts to imposing the full Polygon Rule. To analyze this, fix a distinct pair of indices $i, j \in I$ and consider the evaluation of weighted counts of alternating r-step paths of type (i, j) between pairs of vertices u, v; i.e., the quantities $N_{ij}^r(\Gamma; u, v)$.

By identifying the vertices in the various copies of Γ_k, we can think of u varying over the vertices in a single abstract copy of Γ_k and similarly v varying over Γ_l, with the result that for fixed k and l, one should view $N_{ij}^r(\Gamma; u, v)$ as an $n_k \times n_l$-dimensional matrix.

Recall that only the first and last steps in an alternating path may be arcs—all other steps must be along simple edges. Thus if $k \neq l$, an alternating path from an instance of u in some copy of Γ_k to an instance of v in some copy of Γ_l must either

(1) jump from u along an arc immediately into some third molecule, say a copy of Γ_t, follow a simple path in that molecule, and then jump from that molecule along an arc that terminates at v, or
(2) jump from u along an arc immediately into the molecule containing v, and then follow a simple path to v (the last step may optionally be an internal arc), or
(3) follow a simple path starting at u (optionally the first step may be an internal arc), and then jump from that molecule along an arc that terminates at v.

Since the net contributions of these paths is multilinear in the parameterizing arc weights, it follows that the matrix $N_{ij}^r(\Gamma; u, v)$ may be expressed as an explicit nonnegative integer combination of the matrices

$$A_{kt}^x A_{tl}^y, \quad A_{kl}^x B_l^y, \quad B_k^x A_{kl}^y, \quad A_{kl}^x \tag{6.1}$$

for all sensible values of x, y, t. Furthermore, the coefficients in this linear combination may be computed from a single generic binding of one copy each of Γ_k, Γ_l, and Γ_t; in particular, they depend only on x, y, k, l, t (and of course u, v, i, j and r).

When $k = l$, the same considerations apply when we are computing the off-diagonal entries $N^r_{ij}(\Gamma; u, v)$ (i.e., when u and v are selected from distinct copies of Γ_k). However, in the diagonal case, cases (2) and (3) above do not occur, and this is reflected in the fact that the diagonal entries of the last three groups of matrices in (6.1) are necessarily zero. Instead, an alternating path from u to v may use edges internal to the molecule, and optionally start or finish with an arc internal to Γ_k. Thus, the diagonal matrices

$$B^x_k B^y_k, \quad B^x_k, \quad 1$$

must be used to supplement the list in (6.1) when $k = l$.

In these terms, the Polygon Rule amounts to the condition that the matrices $N^r_{ij}(\Gamma; u, v)$ depend symmetrically on i and j.

The key point of this analysis is that the unbounded multiplicities n_k affect only the sizes of the matrices, not the number or form of the matrix equations that must be imposed. While the number of such equations may seem potentially large, it should be noted that in any instance of the Polygon Rule, we must have $i, j \in \tau(u)$ and $i, j \notin \tau(v)$. Secondly, since the Local Polygon Rule has already been imposed throughout our parameter space, we may further assume that none of the conditions (a)–(c) listed in Fact 3.1 apply. Third, many of the equations are redundant; in practice, the number of independent equations tends to be less than the square of the number of distinct molecules.

C. Case study: a family of D_5-cells.

Consider the problem of classifying all of the admissible D_5-cells that can be constructed out of molecules from component 5 of D_5 (see Figure 8). There are three such molecules: a non-uniform 5-vertex molecule and a single and double cover of the 10-vertex 2-uniform compatibility graph for D_5. We let Γ_0, Γ_1 and Γ_2 denote these graphs.

Note that the D_5-cell in Figure 2 is composed of two copies of Γ_0 and one copy of Γ_1.

The dimensions of the binding spaces for these generic molecules are illustrated in Figure 9(a). Note that they are all 0-dimensional, and their relative parities are forced whenever they appear together in a molecular graph, so we may assume that Γ_0, Γ_1 and Γ_2 have been assigned fixed parities that are mutually compatible.

Following the parameterization described in §6A, one may use four nonnegative integer matrices A_{10}, A_{01}, A_{20} and A_{02} to describe every admissible molecular graph that may be constructed out of these components. When the Polygon Rule is converted to a system of matrix equations as in §6B, one obtains

$$A_{10}A_{01} = 2, \quad A_{20}A_{02} = 1, \quad A_{10}A_{02} = 0, \quad A_{20}A_{01} = 0$$

as necessary and sufficient conditions for the result to be a D_5-graph.

A consequence of the last two equations is that there cannot be a directed path from any copy of Γ_1 to any copy of Γ_2. However, given that the graphs we are aiming to construct must be strongly connected, it follows that one cannot use both graphs. Furthermore,

since $A_{10}A_{01}$ and $A_{20}A_{02}$ are diagonal matrices, there cannot be directed paths between any two copies of Γ_1 or any two copies of Γ_2. Again by connectivity considerations, it follows that the graph has at most one copy of Γ_1 or Γ_2, and cannot have both.

If we have one copy of Γ_2 (and none of Γ_1), then the matrices A_{10} and A_{01} do not exist, and our system reduces to a single matrix equation: $A_{20}A_{02} = [1]$. Up to row and column permutations, all solutions have the form $A_{20} = [1, a_2, a_3, \dots]$, $A_{02}^T = [1, b_2, b_3, \dots]$, where $a_k b_k = 0$ for all $k \geqslant 2$. However, any 0's in these matrices would correspond to copies of Γ_0 that cannot be reached from Γ_2 or vice-versa, so $A_{20} = A_{02} = [1]$ is the unique strongly connected solution.

If we have one copy of Γ_1 (and none of Γ_2), then our system reduces to the matrix equation $A_{10}A_{01} = [2]$. Recognizing there cannot be 0's in A_{10} or A_{01} in order to produce a strongly connected solution, it is not hard to see that there are three cellular solutions: $A_{10} = A_{01}^T = [1, 1]$, $A_{10} = 2A_{01} = [2]$, and $2A_{10} = A_{01} = [2]$.

Finally, if there are no copies of Γ_1 or Γ_2, then there are no equations. This means that Γ_0 is itself a cell, and since the self-binding space of Γ_0 is zero, it follows that the only D_5-graphs whose molecular components are copies of Γ_0 are disjoint unions.

Summarizing, we see that there are five admissible D_5-cells in this family:

- one 25-vertex cell with molecular components Γ_2 and Γ_0,
- one 20-vertex cell with one copy Γ_1 and two copies of Γ_0 (see Figure 2),
- two 15-vertex cells with components Γ_1 and Γ_0, and
- one 5-vertex uni-molecular cell, Γ_0.

Only the 25-vertex and 20-vertex cells are Kazhdan-Lusztig cells.

7. Miscellany

We should say a bit more about how we have determined complete lists of generic molecules. The short answer is software. We have code (so far, only for the simply-laced case) that will take as input a set of allowed τ invariants and a seed graph Γ_0, and generate all generic molecules that contain Γ_0 and use only the allowed τ invariants.

If there are infinitely many such molecules, it will run forever.

However, by the fortuitous accident related to Pessimistic Question 4.5, the only infinite families of generic molecules we have encountered are uniform, and as outlined in §4, we have a pretty good handle on their structure. So we only need software for the non-uniform molecules. In these cases, we can examine the compatibility graph for symmetric edges between subsets of I of unequal cardinality, and uses these as seeds for the software. As long as each of these seeds terminate with a finite list of molecules, we are in business.

We are not planning for E_7 and E_8 to be as cooperative.

8. What's Next?

A. Molecular induction? We think that a more promising approach to identifying the stable generic molecules for E_7 and E_8 is to exploit restriction. Every W_J-restriction of a stable W-molecule is a stable W_J-molecular graph. Starting with a list of stable generic W_J-molecules, one can analyze all of the ways to bind these molecules into W-molecules, and from there hope to prune out the unstable ones. Note that the Howlett-Yin papers [**HY1**] and [**HY2**] on inducing W-graphs might be relevant.

B. Lacking a complete list of stable molecules for E_7 and E_8, it may be worthwhile to solve a lazier version of the classification problem: extract the generic molecules that occur in the real forms of E_7 and E_8 and determine all admissible cells that can be built from them. This would be an especially worthwhile exercise if it leads to the discovery that there are infinitely many admissible E_8-cells, for example.

C. Does the "irreducible" case live up to its name? Less cryptically, can we prove that if there are finitely many admissible W_1-cells and W_2-cells, then there are finitely many admissible $W_1 \times W_2$-cells?

The classification of admissible $I_2(p) \times I_2(q)$-cells looks very interesting. It is roughly equivalent to classifying all commuting pairs of symmetric Cartan matrices of finite type. We do not yet know the complete answer, but in any case, we can show that there are only finitely many for a given choice of p and q.

D. Extend the classifications of admissible cells deeper into the list of multiply-laced Weyl groups: F_4, B_4, $B_5, \ldots$, and perhaps also H_4.

E. Possibly transform the (Maple) software for manipulation, generation and display of W-graphs into something more presentable in public. It is hard to draw the line between parts of the code that should be considered private one-time hacks versus code that could have wider utility. The amount of work involved depends on whether the potential user base is 1 or 2 others, or 10, or 50.

References

[**HY1**] R. B. Howlett and Y. Yin, Inducing W-graphs, *Math Z.* **244** (2003), 415–431.

[**HY1**] R. B. Howlett and Y. Yin, Inducing W-graphs II, *Manuscripta Math.* **115** (2004), 495–511.

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

[**S1**] J. R. Stembridge, Synthetic W-graphs and cells, lecture notes from Atlas V ⟨`www.liegroups.org/papers/summer07/Synthetic.pdf`⟩.

[**S2**] J. R. Stembridge, Admissible W-graphs, preprint