Admissible W-Graphs

John R. Stembridge ⟨jrs@umich.edu⟩

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

June 2, 2008
minor revisions June 9, 2008; August 25, 2008

Contents

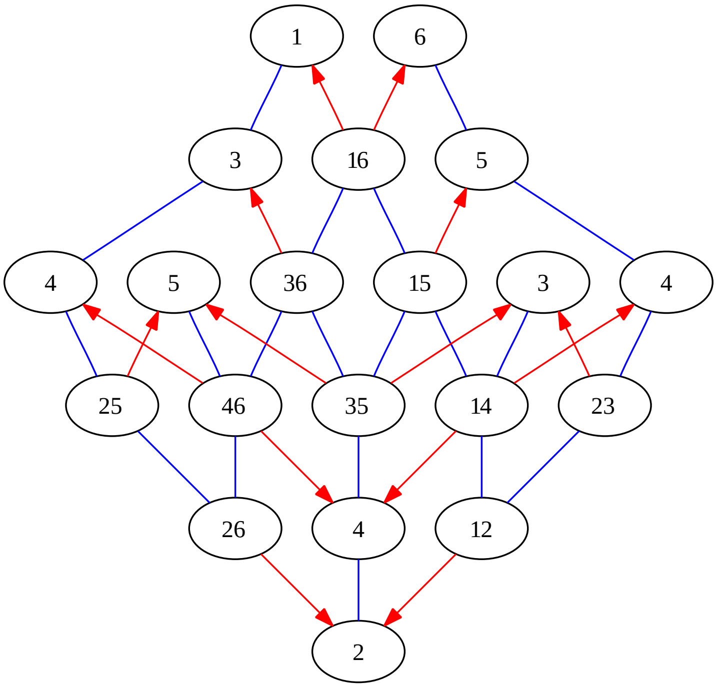

This work has been supported by NSF Grant DMS–0554278.

0. Introduction

Given a Coxeter group W, a W-graph is a combinatorial structure that encodes a representation of W, or more generally, a representation of the associated Iwahori-Hecke algebra. Of special interest are the W-graphs that encode the Kazhdan-Lusztig cell representations of Hecke algebras [**KL**], and more generally, the "Harish-Chandra cells" associated to blocks of irreducible representations of reductive real Lie groups (e.g., see [**M**]).

In this paper, we introduce a class of "admissible" W-graphs that has a minimal set of defining features: nonnegative integer edge weights, edge symmetry, and 2-colorability. These features are designed to include all of the W-graphs and cells that occur in Kazhdan-Lusztig theory as well as those in the Harish-Chandra category. Remarkably, the admissible W-graphs can be characterized by a simple set of combinatorial rules (see Theorem 4.9), and yet there seem to be few admissible W-cells other than Kazhdan-Lusztig cells. (A W-graph is a cell if and only if it is strongly connected.) In particular, (1) most Harish-Chandra cells for irreducible W appear to be Kazhdan-Lusztig cells (for contrary examples, see [**M**]), and (2) the evidence we have accumulated supports the hypothesis that there are only finitely many admissible W-cells for each finite Coxeter group W. This evidence also makes plausible the hope that Kazhdan-Lusztig cells might be computable (as explicit graphs, not merely as representations up to isomorphism) without the need to first compute Kazhdan-Lusztig polynomials. Similarly, it may be possible to determine Harish-Chandra cells without computing Kazhdan-Lusztig-Vogan polynomials.

The paper is structured as follows.

After two preliminary sections on W-graphs and their properties, we show in Section 3 that the nontrivial admissible cells in rank 2 Coxeter groups are A-D-E Dynkin diagrams (Theorem 3.4). In Section 4, we provide the combinatorial characterization of admissible W-graphs mentioned previously; it is interesting to note that similar properties were used by Lusztig in [**L**] to define a class of infinite W-graphs. In the final section, we use this characterization to classify the admissible cells for several low-rank groups, including B_3, H_3, A_4 and D_4. In a sequel to this paper, we will analyze the structure of admissible W-graphs in much greater detail, and use this analysis to classify admissible cells for various higher rank groups.

1. General W-Graphs

Fix a finite index set I. An *I-labeled graph* is a triple $\Gamma = (V, m, \tau)$, where

(i) V is a finite vertex set,
(ii) m is a map $V \times V \to \mathbb{Z}[q^{\pm 1/2}]$ (i.e., a matrix),
(iii) τ is a map $V \to \{\text{subsets of } I\}$.

The value of τ at a vertex v is referred to as the τ *invariant* of v. (This terminology is intended to suggest a connection with the Borho-Jantzen-Duflo τ invariant associated with annihilators of irreducible representations of semisimple Lie algebras.)

To emphasize the interpretation of these objects as directed graphs, we use the notation

$m(u \to v)$ to refer to the (u, v)-entry of the matrix m. This entry should be viewed as the weight of an edge directed from vertex u to vertex v. (If $m(u \to v) = 0$, then there is no such edge.) Many of the I-labeled graphs under consideration here will have nonnegative integer edge weights; for these graphs one may alternatively view the edges as unweighted, and instead interpret $m(u \to v)$ as the number of edges directed from u to v.

In illustrations of I-labeled graphs, such as Figure 1.1, we adopt the convention that unlabeled edges have weight 1, and an edge without an orientation represents a pair of edges of equal weight in both directions.

Now let W be a Coxeter group relative to a generating set of the form $S = \{s_i : i \in I\}$, and let $\{T_i : i \in I\}$ denote the corresponding set of generators of the associated Iwahori-Hecke algebra $\mathcal{H} = \mathcal{H}(W, S)$ over the ground ring $\mathbb{Z}[q^{\pm 1/2}]$. Letting p_{ij} denote the order of $s_i s_j$ in W, one should recall that the defining relations of $\mathcal{H}$ are the quadratic relations $(T_i - q)(T_i + 1) = 0$ and the braid relations

$$\begin{aligned} (T_i T_j)^{p/2} &= (T_j T_i)^{p/2} && \text{if } p = p_{ij} \text{ is even,} \\ T_i (T_j T_i)^{(p-1)/2} &= T_j (T_i T_j)^{(p-1)/2} && \text{if } p = p_{ij} \text{ is odd,} \end{aligned}$$

for all distinct $i, j \in I$ such that $p_{ij} < \infty$.

An I-labeled graph Γ is a W*-graph* if the $\mathbb{Z}[q^{\pm 1/2}]$-module M_Γ freely generated by the vertex set V may be given an $\mathcal{H}$-module structure such that for all $u \in V$,

$$T_i(u) = \begin{cases} qu & \text{if } i \notin \tau(u), \\ -u + q^{1/2} \sum_{v : i \notin \tau(v)} m(u \to v) v & \text{if } i \in \tau(u). \end{cases} \tag{1.1}$$

It is easy to check that the quadratic relations $(T_i - q)(T_i + 1) = 0$ hold on M_Γ for all graphs Γ; the content of this definition is that the braid relations should hold as well.

Remark 1.1. (a) The matrices implicit in (1.1) are transposed from the ones used in the original Kazhdan-Lusztig definition of W-graphs [**KL**]. However, the class of finite dimensional matrix representations of $\mathcal{H}$ is transpose-closed, so this difference is harmless.

(b) The Coxeter systems we are primarily interested in are finite Weyl groups, so our requirement that all I-labeled graphs are finite is reasonable. However, for more general Coxeter systems, this assumption is likely to eliminate too much that is of interest. One possible relaxation is to assume that each vertex has finite out-degree, so that the sums in (1.1) are finite. For an even weaker hypothesis, see the appendix to [**L**].

(c) If we set $q = 1$, the $\mathcal{H}$-action on M_Γ specializes to a W-action, so if Γ is a W-graph, one may think of M_Γ as an $\mathcal{H}$-module and a $\mathbb{Z}W$-module simultaneously.

(d) There exist I-labeled graphs such that the operators defined in (1.1) do not satisfy the braid relations for generic q, but do satisfy them at $q = 1$ (thus yielding a $\mathbb{Z}W$-module). An example of this for the Weyl group A_2 is provided in Figure 1.1.

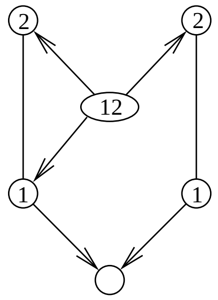

FIGURE 1.1: A fake A_2-graph at $q = 1$.

It is important to note that if $\tau(u) \subseteq \tau(v)$, then the value of $m(u \to v)$ has no effect on the operators defined in (1.1). Thus there is no loss of generality in restricting our attention to W-graphs that are *reduced* in the sense that

$$m(u \to v) = 0 \quad \text{whenever} \quad \tau(u) \subseteq \tau(v). \tag{1.2}$$

Henceforth, all W-graphs in this paper will be implicitly assumed to be reduced.

For many W-graphs of interest, the matrix of edge weights is symmetric, at least to the extent permitted by (1.2). Thus if

$$m(u \to v) = m(v \to u) \quad \text{whenever} \quad \tau(u) \not\subseteq \tau(v) \text{ and } \tau(v) \not\subseteq \tau(u), \tag{1.3}$$

then we say that Γ is *edge-symmetric.*

A. Parabolic restriction.

Given $J \subset I$ and an I-labeled graph $\Gamma = (V, m, \tau)$, let $\Gamma{\downarrow_J}$ denote the J-labeled graph obtained from Γ by: (1) restricting the τ invariants to J; (i.e., replacing $\tau(v)$ with $J \cap \tau(v)$ for all $v \in V$), and (2) deleting those edges of Γ (replacing certain entries of the matrix m with 0's) necessary to bring the resulting graph into compliance with (1.2). It is clear that if Γ is a W-graph, then the restricted graph $\Gamma{\downarrow_J}$ is a W_J-graph, where W_J denotes the parabolic subgroup of W generated by $\{s_i : i \in J\}$.

B. Outer products.

If the Coxeter system (W, S) is reducible, then there is some $J \subset I$ such that W is the direct product of the parabolic subgroups W_J and W_{I-J}, and $\mathcal{H}$ is isomorphic to the tensor product of the corresponding subalgebras $\mathcal{H}_J$ and $\mathcal{H}_{I-J}$. Under these circumstances, if Γ is a W_J-graph on the vertex set V and Γ' is a W_{I-J}-graph on the vertex set V', then we may define a W-graph $\Gamma \times \Gamma'$ on the Cartesian product VV' by setting

$$\tau(vv') = \tau(v) \cup \tau(v'),$$
$$m(uu' \to vv') = m(u \to v)\delta_{u',v'} + m(u' \to v')\delta_{u,v}$$

for all $u, v \in V$ and $u', v' \in V'$. We call $\Gamma \times \Gamma'$ the *outer product* of Γ and Γ'. Note that this construction is designed so that $M_{\Gamma \times \Gamma'}$ is isomorphic to the tensor product $M_\Gamma \otimes M_{\Gamma'}$ (as a module for $\mathcal{H} = \mathcal{H}_J \otimes \mathcal{H}_{I-J}$).

C. 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 for this to happen it is necessary and sufficient that U is a "forward-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 (1.2) in an essential way, and justifies the use of this convention.)

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

$$M_{\Gamma(V-U)} \cong M_\Gamma / M_{\Gamma(U)} \quad \text{(as } \mathcal{H}\text{-modules)}.$$

More generally, given a nested sequence $U_1 \subset U_2 \subset V$ of forward-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 quotient of a submodule of M_Γ.

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

A subset U of V is the vertex set of a subquotient of Γ if and only if it is path-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. If we 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, then the equivalence classes are path-convex, 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. In the language of graph theory, the equivalence classes are the strongly connected components of Γ.

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

Note that M_Γ has a filtration $0 = M^0 \subset M^1 \subset \cdots \subset M^l = M_\Gamma$ of $\mathcal{H}$-submodules in which the subquotients M^i / M^{i-1} are isomorphic to the modules induced by the strongly connected components (i.e., cells) of Γ. However, this will generally not be a composition series, since the $\mathcal{H}$-modules corresponding to cells need not be irreducible.

D. Duality.

Let $a \mapsto a^*$ denote the unique ring automorphism of $\mathbb{Z}[q^{\pm 1/2}]$ such that $q^{1/2} \mapsto -q^{-1/2}$.

We define the *dual* of an I-labeled graph $\Gamma = (V, m, \tau)$ to be the I-labeled graph $\bar{\Gamma} = (V, (m^*)^T, \bar{\tau})$, where $(m^*)^T(u \to v) = m(v \to u)^*$ and $\bar{\tau}(v) = I - \tau(v)$ (i.e., reverse all edges, apply * to all edge weights, and complement all τ invariants).

Note that if Γ has (nonnegative) integer edge weights, then so does its dual.

Proposition 1.2. *If Γ is a W-graph, then $\bar{\Gamma}$ is also a W-graph.*

Proof. If we order the vertex set of $\Gamma = (V, m, \tau)$ so that those $v \in V$ with $i \notin \tau(v)$ precede those with $i \in \tau(v)$, then the matrix of T_i on M_Γ has the form

$$E_i = \begin{bmatrix} q & q^{1/2} m_i \\ 0 & -1 \end{bmatrix},$$

where m_i is a submatrix of m. If we maintain this vertex ordering, but recognize that $\bar{\tau}$ reverses membership, we obtain that the matrix of T_i on $M_{\bar{\Gamma}}$ is

$$-q(E_i^*)^T = \begin{bmatrix} -1 & 0 \\ q^{1/2}(m_i^*)^T & q \end{bmatrix}.$$

Given that the matrices $\{E_i : i \in I\}$ satisfy the braid relations, it follows easily that the same must be true for $\{-q(E_i^*)^T : i \in I\}$, and thus $\bar{\Gamma}$ is a W-graph. $\square$

REMARK 1.3. If Γ has integer edge weights (so that $m^* = m$), one can see from the above proof that the W-representations afforded by M_Γ and $M_{\bar{\Gamma}}$ at $q = 1$ are sign-twisted duals of each other.

2. Admissible W-Graphs

In this section, we assume W is finite for simplicity.

Following [**KL**], one knows that the Iwahori-Hecke algebra $\mathcal{H}$ has a distinguished basis $\{C_w : w \in W\}$ (the Kazhdan-Lusztig basis), and the set of matrices representing the left action of the generators T_i on this basis (or rather the transposed matrices, given our conventions) may be encoded by a W-graph Γ_W as follows. The vertex set is the Coxeter group W, the τ invariant of a vertex $v \in W$ is the left descent set; i.e.,

$$\tau(v) := \tau_L(v) = \{i \in I : \ell(s_i v) < \ell(v)\},$$

where $\ell(\cdot)$ denotes the length function on (W, S), and the edge weights are

$$m(u \to v) := \begin{cases} \mu(u,v) + \mu(v,u) & \text{if } \tau_L(u) \not\subseteq \tau_L(v), \\ 0 & \text{if } \tau_L(u) \subseteq \tau_L(v), \end{cases} \tag{2.1}$$

where $\mu(u, v)$ is by definition the coefficient of $q^{(\ell(v)-\ell(u)-1)/2}$ in the Kazhdan-Lusztig polynomial $P_{u,v}(q)$. Note that since $P_{v,v}(q) = 1$, and $P_{u,v}(q) \neq 0$ only if $u \leqslant v$ in the Bruhat order, we have either $\mu(u, v) = 0$ or $\mu(v, u) = 0$ for all $u, v \in W$ (often both).

It should be noted that the exceptional (zero) values for $m(u \to v)$ that occur when $\tau_L(u) \subseteq \tau_L(v)$ are necessary only to fit our convention that all W-graphs must be reduced as in (1.2). Kazhdan and Lusztig do not follow this convention and use symmetrized edge multiplicities of the form $\mu(u, v) + \mu(v, u)$ without exception. Regardless of convention, the Kazhdan-Lusztig W-graphs are edge-symmetric in the sense of (1.3).

It is well known that for finite W, the polynomials $P_{u,v}(q)$ have 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.

Abstracting these common features, we are led to the following.

Definition 2.1. An I-labeled graph is said to be *admissible* if it

(a) has nonnegative integer edge weights,
(b) is edge-symmetric as in (1.3), and
(c) has a bipartition.

Of course, the Kazhdan-Lusztig graphs Γ_W are admissible. Note also that restrictions, outer products, subquotients, and duals of admissible graphs are again admissible.

Remark 2.2. Gyoja has proved that for every finite Weyl group W, every irreducible W-representation may be realized by means of a W-graph with integer edge weights [**G**]. Furthermore, thanks to recent work of Y. Yin and R. Howlett [**Y**], we now have explicit graphs available for download[1] that realize the irreducible representations of E_7 and E_8. However, it should be noted that these W-graphs cannot be admissible in general. For example, one cannot realize the reflection representation of B_2 in an edge-symmetric way (over $\mathbb{Z}$), and the 6-dimensional irreducible representation of D_4 cannot be realized with nonnegative integer edge weights.

Remark 2.3. For each real Lie group G whose complex form has Weyl group W, Lusztig and Vogan have constructed W-graphs that are similar to the Kazhdan-Lusztig graphs Γ_W [**LV**]. The vertex sets of these W-graphs are blocks of irreducible representations of G. One block occurs for each choice of a real form for the dual group, and the edge weights are obtained from the leading coefficients of certain Kazhdan-Lusztig-Vogan polynomials. Furthermore, these W-graphs are admissible.

Remark 2.4. The W-graph Γ_W is constructed by Kazhdan and Lusztig as a parabolic restriction of a $W \times W$-graph $\Gamma_W^{LR} = (W, m, \tau_{LR})$ arising from the left and right actions of the generators T_i on $\mathcal{H}$. Here, $\tau = \tau_{LR}$ tracks descents on both the left and right; i.e.,

$$\tau_{LR}(v) = \{i_L : \ell(s_i v) < \ell(v)\} \cup \{i_R : \ell(v s_i) < \ell(v)\},$$

where $\{i_L : i \in I\}$ and $\{i_R : i \in I\}$ denote disjoint index sets for the left and right copies of W in the direct product $W \times W$. Also, the matrix of edge weights m is defined as in (2.1), but with τ_{LR} replacing τ_L. If G is a complex semisimple Lie group with Weyl group W, the "two-sided" graph Γ_W^{LR} is a Lusztig-Vogan graph as in the previous remark, if we view G as real Lie group; the complexification of this real group has Weyl group $W \times W$.

In our project to understand and classify admissible W-graphs, one of the fundamental questions raised by the data we have gathered so far is the following.

Question 2.5. *If W is a finite Coxeter group, are there only finitely many admissible W-cells up to isomorphism?*

We have affirmative answers to this question for W of rank $\leqslant 3$, as well as for the Weyl groups A_n $(n \leqslant 9)$, D_n $(n \leqslant 6)$, and E_6.

[1] See ⟨`www.maths.usyd.edu.au/u/bobh/yy/`⟩.

REMARK 2.6. If W is reducible, it is not true in general that all admissible W-cells are outer products of cells for the irreducible factors of W.[2] For example, the cells of the $W \times W$-graph Γ_W^{LR} discussed in Remark 2.4 need not be outer products of pairs of W-cells. Indeed, any cell of Γ_W that generates a reducible representation of W cannot be contained in a $W \times W$-cell that is an outer product, since the $W \times W$-action on $\mathbb{C}W$ decomposes into the sum of outer tensor squares of all irreducible $\mathbb{C}W$-modules. Thus in order to resolve questions about the classification of admissible W-graphs, it is *not* sufficient to consider only irreducible groups.

The data we have accumulated so far indicates that there is a distinct tendency for the W-representations generated by admissible cells to be *combinatorially rigid*. By this we mean that admissible cells that generate isomorphic representations of W (over $\mathbb{C}$, say) tend to be isomorphic as I-labeled graphs. Although there are instances where this tendency fails (e.g., see Remark 5.2), the data we have supports an affirmative answer to the following.

QUESTION 2.7. *Assume that Γ and Γ' are both cells of the Kazhdan-Lusztig graph Γ_W such that $M_\Gamma \cong M_{\Gamma'}$ (as $\mathbb{C}W$-modules). Is it true that $\Gamma \cong \Gamma'$?*

It is interesting to note that this fails if we allow one of the cells to merely be admissible, not necessarily a Kazhdan-Lusztig cell. The one instance of this failure we have found is a 150-dimensional representation of E_6 that is realized by nine combinatorially distinct admissible cells, only one of which occurs as a cell of the Kazhdan-Lusztig graph.

The empirical results we have for the symmetric groups $W = A_n$ are quite striking. In this case, it is a well-known result of Kazhdan and Lusztig that the representations generated by the cells of Γ_W are precisely the irreducible representations of W, and this particular class of cells is combinatorially rigid (see Theorem 1.4 of [**KL**]). On the other hand, we have found that there are no other admissible A_n-cells for $n \leqslant 9$.

QUESTION 2.8. *Is every admissible A_n-cell a Kazhdan-Lusztig cell?*

Although it is tempting to conjecture an affirmative answer to this question, the work of McLarnan and Warrington shows that Kazhdan-Lusztig cells in type A have interesting features that emerge only in high rank [**MW**]. In particular, they have shown that the Kazhdan-Lusztig graph Γ_W for $W = A_n$ has edge weights > 1 for $n \geqslant 9$, and there are individual cells of Γ_W that have edge weights > 1 for $n \geqslant 15$.

[2]We thank D. Vogan for pointing this out.

$$
\begin{tikzcd}
 & V_{\{1,2\}} \arrow[dl] \arrow[dr] \arrow[dd] & \\
V_{\{1\}} \arrow[rr, bend left] \arrow[dr] & & V_{\{2\}} \arrow[ll, bend left] \arrow[dl] \\
 & V_{\varnothing} &
\end{tikzcd}
$$

FIGURE 3.1.

3. Cells in Rank Two

Let $W = I_2(p)$ be the rank 2 Coxeter group defined by a braid relation of length $p \geqslant 2$. In order to classify the admissible W-cells, let us first consider necessary and sufficient conditions for an arbitrary (reduced, but not necessarily admissible) $\{1,2\}$-labeled graph Γ to be a W-cell. (Note that the infinite rank 2 Coxeter group $I_2(\infty)$ has no braid relation, so every $\{1,2\}$-labeled graph is an $I_2(\infty)$-graph.)

For $J \subseteq \{1,2\}$, let V_J denote the set of vertices of Γ with τ invariant J. Given that Γ is reduced, the only possible edge orientations in Γ are as indicated in Figure 3.1. If Γ is a W-graph, each vertex in $V_\varnothing$ has out-degree 0 and thus necessarily spans a singleton cell affording the trivial representation of W, and each vertex in $V_{\{1,2\}}$ has in-degree 0 and spans a singleton cell affording the sign representation. More generally in any Coxeter system, a cell with a vertex whose τ invariant is full or empty is necessarily a singleton.

Leaving aside these trivial cells, all other W-cells must be strongly connected graphs whose τ invariants (either $\{1\}$ or $\{2\}$) provide a proper 2-coloring (i.e., the endpoints of all edges have distinct labels). In particular, all such cells are necessarily bipartite.

PROPOSITION 3.1. *Let Γ be a reduced $\{1,2\}$-labeled graph such that every vertex has τ invariant $\{1\}$ or $\{2\}$. If m is the matrix of edge weights of Γ, then Γ is an $I_2(p)$-graph if and only if $\phi_p(m) = 0$, where $\phi_r(t)$ is defined by the recurrence*

$$\phi_{r+1}(t) = t\phi_r(t) - \phi_{r-1}(t) \tag{3.1}$$

and the initial conditions $\phi_0(t) = 0$, $\phi_1(t) = 1$.

Note that $\phi_r(2t)$ is a Chebyshev polynomial of the second kind.

To prove the above result, let us first consider an arbitrary pair of operators T_1 and T_2 that satisfy the quadratic relations $(T_i - q)(T_i + 1) = 0$. For each $r \geqslant 0$, let Δ_r denote the length r "braid commutator" of T_1 and T_2; i.e.,

$$\Delta_r = \begin{cases} (T_1T_2)^{r/2} - (T_2T_1)^{r/2} & \text{if } r \text{ is even,} \\ T_1(T_2T_1)^{(r-1)/2} - T_2(T_1T_2)^{(r-1)/2} & \text{if } r \text{ is odd.} \end{cases}$$

In particular, $\Delta_0 = 0$ and $\Delta_1 = T_1 - T_2$.

Lemma 3.2. *If T_1, T_2 and Δ_r are as above, then for all $r \geqslant 0$, we have*

$$\Delta_r = (-1)^{r-1} q^{(r-1)/2} \phi_r(A)(T_1 - T_2),$$

where $A = q^{-1/2}(T_1 + T_2 - (q-1))$.

Proof. The result is clear when $r = 0$ or 1. For $r \geqslant 1$, consider that

$$\begin{aligned}(T_1 + T_2)\Delta_r &= (T_1 + T_2)((T_1 T_2 \cdots) - (T_2 T_1 \cdots)) \\ &= -\Delta_{r+1} + (T_1^2 T_2 \cdots) - (T_2^2 T_1 \cdots) \\ &= -\Delta_{r+1} + (((q-1)T_1 + q)T_2 T_1 \cdots) - (((q-1)T_2 + q)T_1 T_2 \cdots) \\ &= -\Delta_{r+1} + (q-1)\Delta_r - q\Delta_{r-1}.\end{aligned}$$

Thus by induction with respect to r, we obtain

$$\Delta_{r+1} = -q^{1/2} A\, \Delta_r - q\Delta_{r-1} = (-1)^r q^{r/2} (A\phi_r(A) - \phi_{r-1}(A))(T_1 - T_2).$$

Since $\phi_{r+1}(A) = A\phi_r(A) - \phi_{r-1}(A)$ by definition, the result follows. $\square$

Proof of Proposition 3.1. If we order the vertex set of Γ so that the vertices with τ invariant $\{1\}$ precede those with τ invariant $\{2\}$, then the edge weight matrix m, as well as the representing matrices for T_1 and T_2 on M_Γ, have a 2×2 block structure; namely,

$$m = \begin{bmatrix} 0 & B \\ A & 0 \end{bmatrix}, \quad T_1 = \begin{bmatrix} -1 & 0 \\ q^{1/2}A & q \end{bmatrix}, \quad T_2 = \begin{bmatrix} q & q^{1/2}B \\ 0 & -1 \end{bmatrix}, \tag{3.2}$$

for certain matrices A and B. Hence $m = q^{-1/2}(T_1 + T_2 - (q-1))$, and Lemma 3.2 implies that T_1 and T_2 satisfy the braid relation of length p if and only if

$$\phi_p(m)(T_1 - T_2) = 0.$$

To see that this is equivalent to the vanishing of $\phi_p(m)$, note that $\phi_r(t)$ is either an even or odd polynomial function of t depending on the parity of r, therefore

$$\phi_r(m) = \begin{bmatrix} 0 & B_r \\ A_r & 0 \end{bmatrix} \text{ or } \begin{bmatrix} A_r & 0 \\ 0 & B_r \end{bmatrix}$$

for certain matrices A_r and B_r. On the other hand, $T_1 - T_2$ has the scalars $\pm(q+1)$ on its two diagonal blocks, so included among the block entries of $\phi_r(m)(T_1 - T_2)$ are scalar multiples of A_r and B_r, for all $r \geqslant 1$. $\square$

FIGURE 3.2: Asymmetric B_2-cells.

REMARK 3.3. (a) Since $\phi_2(t) = t$, the above proposition implies that all $A_1 \times A_1$-cells have no edges (and thus are singletons), and conversely, a singleton with any of the four possible τ invariants is an $A_1 \times A_1$-cell. Note that all such cells are admissible.

(b) For $A_2 = I_2(3)$, we have $\phi_3(t) = t^2 - 1$. It follows that if Γ has an edge weight matrix m as in (3.2), then Γ is an A_2-cell if and only if $AB = BA = 1$ (i.e., $m^2 = 1$). This condition is fulfilled precisely when A and B are an inverse pair of (necessarily square) matrices over $\mathbb{Z}[q^{\pm 1/2}]$. Generically, such pairs have no nonzero entries, and thus the corresponding graphs are generically cells. This suggests that it would be futile to attempt a classification of all Weyl group cells without regard to admissibility.

However, if m has nonnegative integer entries, then A and B must be permutation matrices such that $A = B^T$, and thus Γ is edge symmetric. As a graph, this amounts to a disjoint union of bi-directed edges. Breaking such a graph into strongly connected components, we conclude that there is a unique nontrivial admissible A_2-cell: it consists of two vertices u and v with $\tau(u) = \{1\}$, $\tau(v) = \{2\}$, and $m(u \to v) = m(v \to u) = 1$.

(c) For $B_2 = I_2(4)$, we have $\phi_4(t) = t^3 - 2t$. Unlike the previous case, nonnegative integer edge weights in a B_2-cell are not sufficient to force edge symmetry. For example, the edge weight matrices of the two asymmetric graphs in Figure 3.2 are both roots of ϕ_4, and it is easy to check that they both generate the reflection representation of B_2.

Now suppose that Γ is an admissible $I_2(p)$-cell. As noted above, if we assume that Γ is not one of the trivial singletons with full or empty τ invariant, then Γ is properly 2-colored: every vertex has τ invariant $\{1\}$ or $\{2\}$, and the τ invariants at the endpoints of every edge are distinct. In particular, since there are no nontrivial inclusions of τ invariants, it follows that edge-symmetry in this case amounts to the condition that $m = m^T$.

Encoding each pair of edges $u \to v$ and $v \to u$ of weight $k \in \mathbb{Z}^{\geqslant 0}$ with k *undirected* edges between u and v, one may thus regard Γ as a properly 2-colored multigraph.

THEOREM 3.4. *A properly 2-colored multigraph is an (admissible) $I_2(p)$-cell if and only if it is a connected simply-laced Dynkin diagram of finite type whose Coxeter number divides p. In particular, no such graph has edge multiplicities > 1.*

More explicitly, the Dynkin diagrams in question are those for A_n ($n \geqslant 1$), D_n ($n \geqslant 4$), E_6, E_7, and E_8; their respective Coxeter numbers are $n + 1$, $2n - 2$, 12, 18, and 30.

Considering $B_2 = I_2(4)$ and $G_2 = I_2(6)$ for example, the Dynkin diagrams with Coxeter number dividing 4 are those for A_1 and A_3; for 6, they are A_1, A_2, D_4, and A_5. The corresponding cells are illustrated in Figures 3.3 and 3.4.

The nontrivial Kazhdan-Lusztig cells for $I_2(p)$ are 2-colorings of A_{p-1}-diagrams.

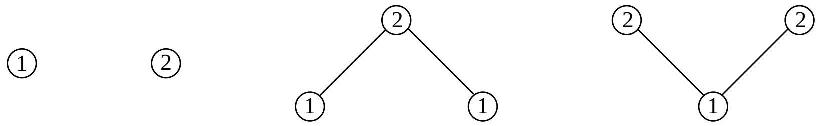

FIGURE 3.3: The nontrivial admissible B_2-cells.

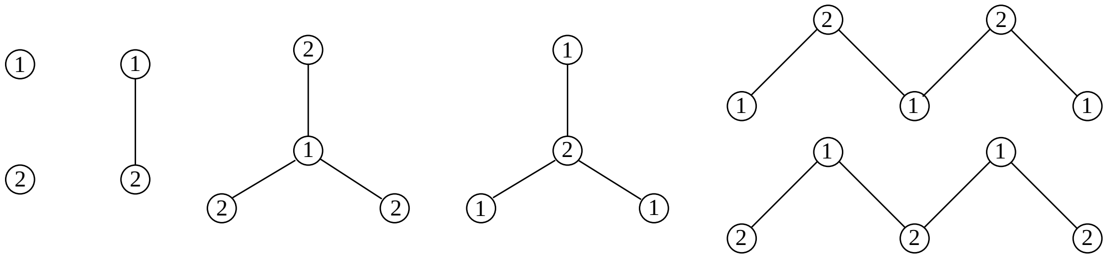

FIGURE 3.4: The nontrivial admissible G_2-cells.

Proof. Let Γ be a properly 2-colored multigraph, viewed as an admissible $\{1,2\}$-labeled graph with a symmetric edge-weight matrix m. By Proposition 3.1, we have that Γ is an $I_2(p)$-cell if and only if Γ is connected and $\phi_p(m) = 0$. Since

$$\sin(r+1)\theta = 2\cos\theta\sin r\theta - \sin(r-1)\theta,$$

it follows that $\phi_p(2\cos\theta) = \sin p\theta/\sin\theta$. Bearing in mind that $\phi_p(t)$ is a polynomial of degree $p-1$, we conclude that its roots are $2\cos(k\pi/p)$ for $1 \leqslant k < p$.

Noting that the roots of $\phi_p(t)$ are strictly less than 2, it follows that if $\phi_p(m) = 0$, then the eigenvalues of $2-m$ are strictly positive, and thus $2-m$ is a symmetric Cartan matrix of finite type. In other words, m is the adjacency matrix of an A-D-E Dynkin diagram.

Conversely, if Γ is a connected, properly 2-colored simply-laced Dynkin diagram of finite type, it remains to determine for which p (if any) the corresponding graph operators T_1 and T_2 satisfy the braid relation of length p. Given Proposition 3.1 and our knowledge of the roots of $\phi_p(t)$, this is equivalent to determining when the eigenvalues of the Cartan matrix $2-m$ of Γ are of the form $2-2\cos(k\pi/p)$ for suitable integers k.

The easiest way to resolve this is to use Exercise V.6.4 in [**B**]: the eigenvalues of an irreducible Cartan matrix C of finite type (simply-laced or not) are $2-2\cos(e_j\pi/h)$, where $e_1, e_2, \ldots$ are the exponents of C and h is the Coxeter number. Since 1 is always an exponent, we conclude that the length p braid relation holds if and only if h divides p.

Alternatively, let W_Γ be a Weyl group with Dynkin diagram Γ, and let c_1 and c_2 denote the products of the generators of W_Γ corresponding to vertices that are colored with the τ invariants $\{1\}$ and $\{2\}$, respectively. It follows that c_1c_2 is a Coxeter element of W_Γ, and thus (by definition) has order equal to the Coxeter number h. In particular, the actions of c_1 and c_2 in any faithful representation of W_Γ satisfy the length p braid relations if and only if h divides p.

Now consider the (faithful) reflection representation of W_Γ. Since c_1 and c_2 both act as products of commuting reflections, it is an easy calculation to show that the matrices for the action of c_1 and c_2 on a set of simple roots for W_Γ have the form

$$c_1 = \begin{bmatrix} -1 & A^T \\ 0 & 1 \end{bmatrix}, \quad c_2 = \begin{bmatrix} 1 & 0 \\ A & -1 \end{bmatrix}, \quad \text{where} \quad m = \begin{bmatrix} 0 & A^T \\ A & 0 \end{bmatrix}.$$

Comparing with (3.2), we see that the matrices for c_1 and c_2 are transposed from the $q = 1$ specializations of T_1 and T_2. However, given that m does not depend on q, one can see from the proof of Proposition 3.1 that T_1 and T_2 satisfy a length r braid relation for some special value of $q \notin \{0, -1\}$ only if they satisfy the same braid relation for generic q. Thus, the pairs (c_1, c_2) and (T_1, T_2) satisfy the same braid relations. □

4. Combinatorial Characterization

Now assume that (W, S) is a general Coxeter system, not necessarily finite.

In the following, we will describe four rules that characterize in simple combinatorial terms when an admissible (and reduced) I-labeled graph $\Gamma = (V, m, \tau)$ is a W-graph.

As a starting point, note that Γ is a W-graph if and only if for all distinct $i, j \in I$, the restriction $\Gamma{\downarrow}_{\{i,j\}}$ is a $W_{\{i,j\}}$-graph; i.e., T_i and T_j satisfy the appropriate braid relation. Thus Theorem 3.4 provides strong constraints on the structure of Γ, but it is not a complete characterization, since it concerns only the cells of such restrictions.

To simplify the exposition, we will assume that (W, S) is *braid-finite*; i.e., every pair of generators satisfies a braid relation of finite length. If there do exist pairs $i, j \in I$ such that s_i and s_j do not satisfy a braid relation, then there are no constraints to impose on the $\{i, j\}$-restrictions of I-labeled graphs, and it is an easy exercise for the reader to modify the rules below to account for this possibility.

A. The Compatibility and Simplicity Rules.

Suppose temporarily that Γ is a general W-graph, not necessarily admissible.

If there is an edge $u \to v$ (i.e., $m(u \to v) \neq 0$), then for every $i \in \tau(u) - \tau(v)$ and $j \in \tau(v) - \tau(u)$ (which is possible only if $\tau(u)$ and $\tau(v)$ are incomparable), the vertices of $\Gamma{\downarrow}_{\{i,j\}}$ with τ invariant $\{i\}$ or $\{j\}$ span a subquotient that includes u and v. Applying Proposition 3.1 with $p = 2$, we see that this can happen only if s_i and s_j do not commute. That is, nodes i and j must be bonded[3] in the diagram of (W, S). This condition also holds vacuously when $\tau(u)$ and $\tau(v)$ are comparable, since either i or j cannot exist in such cases. In summary, this yields

Proposition 4.1 (The Compatibility Rule). *If $u \to v$ is an edge in any W-graph, then every $i \in \tau(u) - \tau(v)$ must be bonded to every $j \in \tau(v) - \tau(u)$.*

The above rule may be reformulated in terms of a graph homomorphism as follows. Define the *compatibility graph* $\mathrm{Comp}(W, S)$ to be the directed graph whose vertices are the

[3]In order to avoid confusion with edges in I-labeled graphs, we refer to adjacencies in the Coxeter diagram of (W, S) as bonds.

subsets of I and whose edges are of the form $J \to K$, where J and K range over all subsets of I such that $J \not\subseteq K$ and every $j \in J - K$ is bonded to every $k \in K - J$ in the diagram of (W, S). Given that Γ is reduced, the Compatibility Rule is equivalent to the statement that τ is a graph homomorphism $\Gamma \to \mathrm{Comp}(W, S)$. That is, if $u \to v$ is an edge of Γ with nonzero weight, then $\tau(u) \to \tau(v)$ is an edge of $\mathrm{Comp}(W, S)$.

The compatibility graphs for A_3, A_4 and D_4 are displayed in Figure 4.1. In these figures, we have omitted all of the edges $J \to K$ such that $J \supset K$, since this Boolean algebra of set inclusions is present in the compatibility graph of every Coxeter system. All other edges are between incomparable pairs and are symmetric (i.e., $J \to K$ implies $K \to J$); these pairs of edges are represented in the figures by single unoriented edges. We have also omitted the vertices I and $\varnothing$, since they are not the endpoints of any symmetric edges.

Now assume that Γ is admissible (and reduced).

Continuing the hypothesis that u and v are vertices such that $m(u \to v) \neq 0$ in Γ, note that $\tau(u) \supset \tau(v)$ forces $m(v \to u) = 0$. In such cases, we say that $u \to v$ is an *arc* of Γ.

Since Γ is admissible and hence edge-symmetric, the remaining possibility is that $\tau(u)$ and $\tau(v)$ are incomparable, and $m(v \to u) = m(u \to v) \neq 0$. It follows that u and v belong to the same cell of $\Gamma{\downarrow}_{\{i,j\}}$ for all choices of indices $i \in \tau(u) - \tau(v)$ and $j \in \tau(v) - \tau(u)$ (and such choices exist). Given that (W, S) is braid-finite, we see that Theorem 3.4 forces the edge multiplicity in both directions to be exactly one. In this case, we say that there is a *simple* edge between u and v. This yields

PROPOSITION 4.2 (The Simplicity Rule). *If (W, S) is braid-finite, then every admissible W-graph $\Gamma = (V, m, \tau)$ is comprised of arcs and simple edges; i.e., for all $u, v \in V$ such that $m(u \to v) \neq 0$, either*

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

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

REMARK 4.3. If (W, S) is simply-laced, then every W-graph with nonnegative integer edge weights satisfies the Simplicity Rule, even if it fails to be admissible. Indeed, this follows by the above reasoning, together with the observation from Remark 3.3 that all $A_1 \times A_1$-cells, and all A_2-cells with nonnegative integer edge weights, are admissible.

B. The Bonding Rule.

We now introduce an edge labeling of $\mathrm{Comp}(W, S)$ as follows. For each bonded pair of nodes in the diagram of (W, S), say j and k, we use the bond $b = \{j, k\}$ as a label for all edges $J \to K$ of $\mathrm{Comp}(W, S)$ such that $j \in J - K$ and $k \in K - J$ or vice-versa. As a result, only the symmetric edges of $\mathrm{Comp}(W, S)$ are assigned labels, and some edges may receive multiple labels as illustrated in Figure 4.1.

We may transport this edge labeling to every W-graph Γ by declaring that each edge $u \to v$ inherits the labels of the corresponding edge $\tau(u) \to \tau(v)$ of $\mathrm{Comp}(W, S)$.

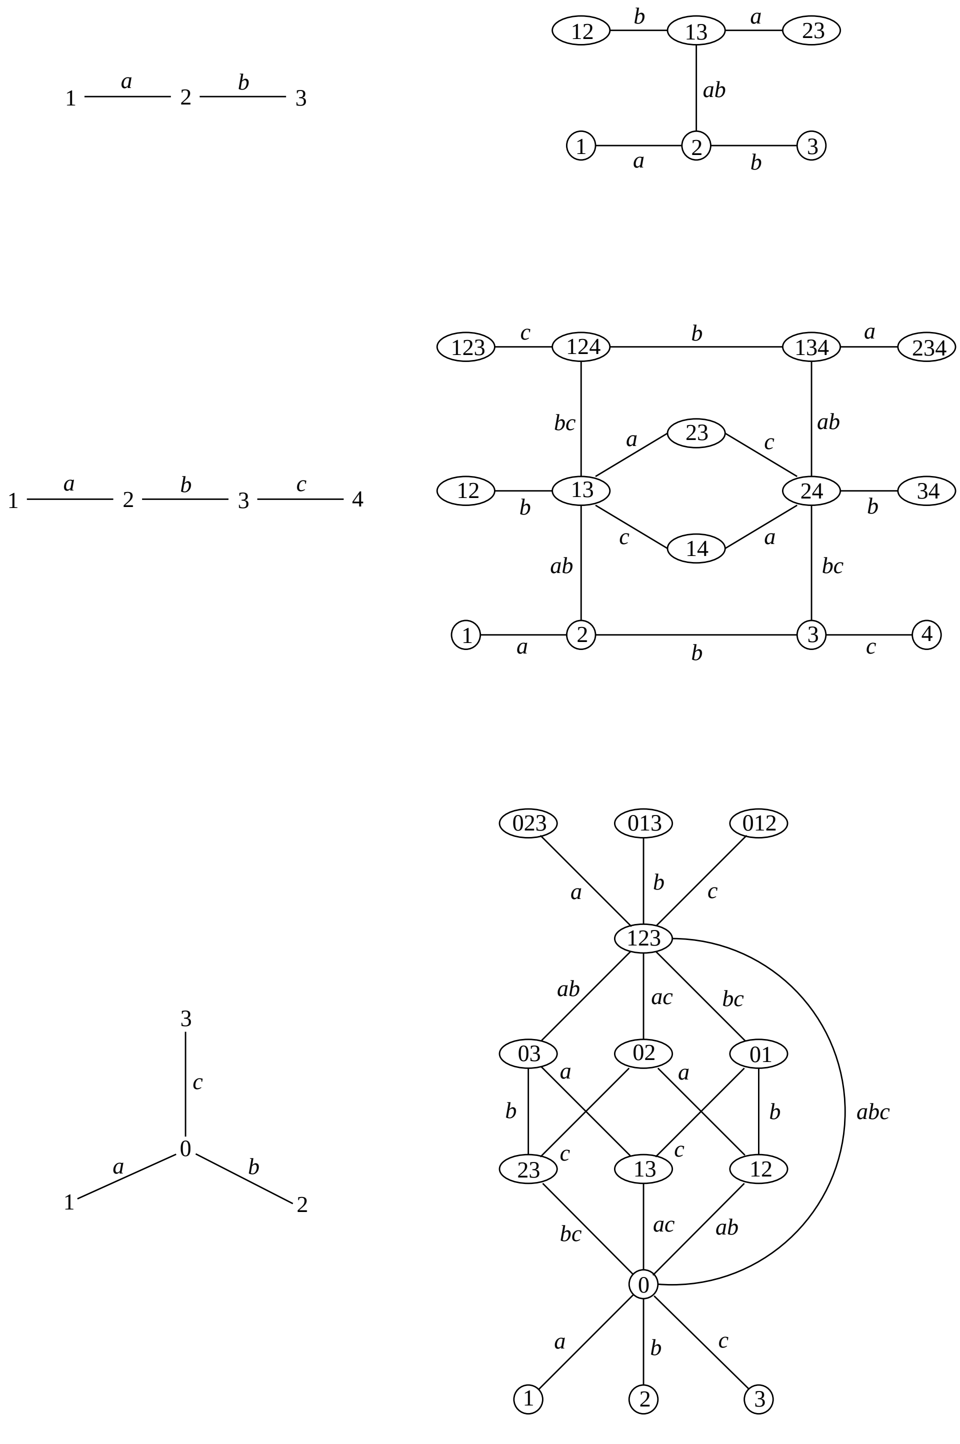

FIGURE 4.1: Diagrams and compatibility graphs for A_3, A_4, and D_4.

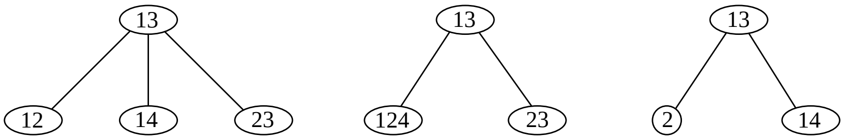

FIGURE 4.2: Simple neighborhoods in an admissible A_4-graph.

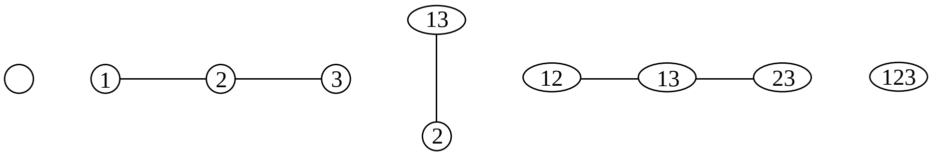

FIGURE 4.3: The admissible A_3-cells.

Define the *frontier* of a vertex v in Γ, denoted $\mathrm{Fr}(v)$, to be the set of bonded pairs $\{i, j\}$ in the diagram of (W, S) such that $i \in \tau(v)$ and $j \notin \tau(v)$ or vice-versa. Equivalently, this is the set of labels that occur on the edges of $\mathrm{Comp}(W, S)$ incident to $\tau(v)$. Each vertex that has $\{i, j\}$ in its frontier restricts to a vertex with τ invariant $\{i\}$ or $\{j\}$ in $\Gamma{\downarrow}_{\{i,j\}}$, so one may reformulate the remaining consequences of Theorem 3.4 as follows.

PROPOSITION 4.4 (The Bonding Rule). *If Γ is an admissible W-graph and $\{i, j\}$ is a bond in the diagram of (W, S) such that $s_i s_j$ has order $p < \infty$ in W, then the vertices of Γ with $\{i, j\}$ in their frontier, together with the edges of Γ that include the label $\{i, j\}$, form a disjoint union of A-D-E Dynkin diagrams whose Coxeter numbers divide p.*

In particular, if $p = 3$, then for every vertex u such that $i \in \tau(u)$ and $j \notin \tau(u)$, there is a unique adjacent vertex v such that $j \in \tau(v)$ and $i \notin \tau(v)$.

For example, referring to Figure 4.1, suppose that v is a vertex of an admissible A_4-graph with $\tau(v) = \{1, 3\}$. The frontier of v is $\{a, b, c\}$, so by examining the neighborhood of $\{1, 3\}$ in $\mathrm{Comp}(A_4)$, we see that the possible τ invariants of the neighbors of v connected along simple edges are $\{2\}$, $\{1, 2\}$, $\{1, 4\}$, $\{2, 3\}$ and $\{1, 2, 4\}$, and the respective labels along these edges are ab, b, c, a and bc. By the Bonding Rule, each of the labels a, b, and c in $\mathrm{Fr}(v)$ must each occur exactly once among the edges incident to v, so there are three possibilities for the simple neighborhood of v as indicated in Figure 4.2.

REMARK 4.5. The rules we have accumulated so far (Compatibility, Simplicity, and Bonding) are already quite strong. For example, by examining $\mathrm{Comp}(A_3)$ in Figure 4.1, one sees that these rules force the simple edges of any admissible A_3-graph to be disjoint unions of copies of the graphs in Figure 4.3. Of course, an A_3-graph may also have arcs $u \to v$ such that $\tau(u) \supsetneq \tau(v)$. However, all such arcs must be directed right-to-left among the graphs in Figure 4.3, and there are no inclusions of τ invariants internal to these graphs, so we conclude that no admissible A_3-cell can have arcs, and the graphs in Figure 4.3 are the only possibilities for admissible cells. Conversely, each of these graphs *is* an A_3-graph (e.g., by Theorem 4.9 below), so this classifies the admissible A_3-cells.

C. The Polygon Rule.

If one orders the vertices of a $\{1,2\}$-labeled graph so that the τ invariants occur in the order $\varnothing$, $\{1\}$, $\{2\}$ and $\{1,2\}$, then the matrices representing T_1 and T_2 have a 4×4 block structure; namely,

$$T_1 = \begin{bmatrix} q & q^{1/2}B_1 & 0 & q^{1/2}D \\ 0 & -1 & 0 & 0 \\ 0 & q^{1/2}A_1 & q & q^{1/2}C_1 \\ 0 & 0 & 0 & -1 \end{bmatrix}, \quad T_2 = \begin{bmatrix} q & 0 & q^{1/2}B_2 & q^{1/2}D \\ 0 & q & q^{1/2}A_2 & q^{1/2}C_2 \\ 0 & 0 & -1 & 0 \\ 0 & 0 & 0 & -1 \end{bmatrix}, \tag{4.1}$$

for certain submatrices $A_1, A_2, B_1, B_2, C_1, C_2$ and D of the edge weight matrix m.

Lemma 4.6. *If T_1 and T_2 are as above, then their length r braid commutator is*

$$\Delta_r = (-1)^{r-1}q^{r/2}\begin{bmatrix} 0 & B\phi_r(A)J & -\psi_r \\ 0 & \phi_r(A)(A-t)J & -\phi_r(A)JC \\ 0 & 0 & 0 \end{bmatrix},$$

where $t = q^{1/2}+q^{-1/2}$, ϕ_r is defined as in (3.1), $\psi_r = \sum_{i=1}^{r-1} B\phi_i(t)\phi_{r-i}(A)JC$,

$$A = \begin{bmatrix} 0 & A_2 \\ A_1 & 0 \end{bmatrix}, \quad J = \begin{bmatrix} 1 & 0 \\ 0 & -1 \end{bmatrix}, \quad B = [\,B_1 \quad B_2\,], \quad C = \begin{bmatrix} C_2 \\ C_1 \end{bmatrix},$$

and the block sizes of J match those of A.

Proof. If we define matrices E_1 and E_2 by setting

$$E_1 = \frac{T_1+T_2-(q-1)}{q^{1/2}} = \begin{bmatrix} t & B & 2D \\ 0 & A & C \\ 0 & 0 & -t \end{bmatrix}, \quad E_2 = \frac{T_1-T_2}{q^{1/2}} = \begin{bmatrix} 0 & BJ & 0 \\ 0 & (A-t)J & -JC \\ 0 & 0 & 0 \end{bmatrix},$$

then Lemma 3.2 implies $\Delta_r = (-1)^{r-1}q^{r/2}\phi_r(E_1)E_2$. Thus it suffices to evaluate $\phi_r(E_1)E_2$.

The evaluation is trivial when $r=0$ or 1, and an induction for $r \geqslant 1$ yields

$$\begin{aligned}
\phi_{r+1}(E_1)E_2 &= E_1\phi_r(E_1)E_2 - \phi_{r-1}(E_1)E_2 \\
&= \begin{bmatrix} t & B & 2D \\ 0 & A & C \\ 0 & 0 & -t \end{bmatrix}\begin{bmatrix} 0 & B\phi_r(A)J & -\psi_r \\ 0 & \phi_r(A)(A-t)J & -\phi_r(A)JC \\ 0 & 0 & 0 \end{bmatrix} \\
&\qquad - \begin{bmatrix} 0 & B\phi_{r-1}(A)J & -\psi_{r-1} \\ 0 & \phi_{r-1}(A)(A-t)J & -\phi_{r-1}(A)JC \\ 0 & 0 & 0 \end{bmatrix} \\
&= \begin{bmatrix} 0 & B(t\phi_r(A)+\phi_r(A)(A-t)-\phi_{r-1}(A))J & -t\psi_r - B\phi_r(A)JC + \psi_{r-1} \\ 0 & A\phi_r(A)(A-t)J - \phi_{r-1}(A)(A-t)J & -(A\phi_r(A)-\phi_{r-1}(A))JC \\ 0 & 0 & 0 \end{bmatrix}.
\end{aligned}$$

The blocks in positions $(1,2)$, $(2,2)$ and $(2,3)$ above simplify respectively to $B\phi_{r+1}(A)J$, $\phi_{r+1}(A)(A-t)J$ and $-\phi_{r+1}(A)JC$. For the $(1,3)$-block, we have

$$\begin{aligned} t\psi_r + B\phi_r(A)JC - \psi_{r-1} &= B\Big(\phi_r(A) + \sum_{i=1}^{r-1} t\phi_i(t)\phi_{r-i}(A) - \sum_{i=1}^{r-1} \phi_{i-1}(t)\phi_{r-i}(A)\Big)JC \\ &= B\Big(\phi_r(A) + \sum_{i=1}^{r-1} \phi_{i+1}(t)\phi_{r-i}(A)\Big)JC = \psi_{r+1}, \end{aligned}$$

which completes the induction. $\square$

Given an I-labeled graph $\Gamma = (V, m, \tau)$ and disjoint subsets $J, K \subseteq I$, let

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

In the following, we may also use abbreviated notations such as (for example) $V_{i/j}$ in place of $V_{\{i\}/\{j\}}$. For distinct $i, j \in I$, we define a directed path $u \to v_1 \to \cdots \to v_{r-1} \to v$ in Γ of length $r \geqslant 2$ to be *alternating of type* (i,j) if

$$u \in V_{ij/\varnothing},\ \ v_k \in V_{i/j} \text{ for } k \text{ odd},\ \ v_k \in V_{j/i} \text{ for } k \text{ even},\ \text{ and } v \in V_{\varnothing/ij}. \tag{4.2}$$

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

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

where the vertices v_k are restricted as in (4.2). Note that if Γ is an admissible W-graph (and (W,S) is braid-finite) then the Simplicity Rule implies that 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.

Proposition 4.7 (The Polygon Rule). *If Γ is a W-graph with integer edge weights and s_is_j has finite order $p \geqslant 2$ in W, then $N^r_{ij}(\Gamma; u, v) = N^r_{ji}(\Gamma; u, v)$ for $2 \leqslant r \leqslant p$ and all vertices u, v with $i, j \in \tau(u)$ and $i, j \notin \tau(v)$; 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 $r \leqslant p$.*

Proof. The matrices representing the action of T_i and T_j have the form indicated in (4.1), if we identify T_i with T_1 and T_j with T_2 (say). By hypothesis, T_i and T_j satisfy a braid relation of length p, so Lemma 4.6 implies

$$\psi_p = \sum_{r=1}^{p-1} B\phi_r(t)\phi_{p-r}(A)JC = 0,$$

using the notation introduced previously.

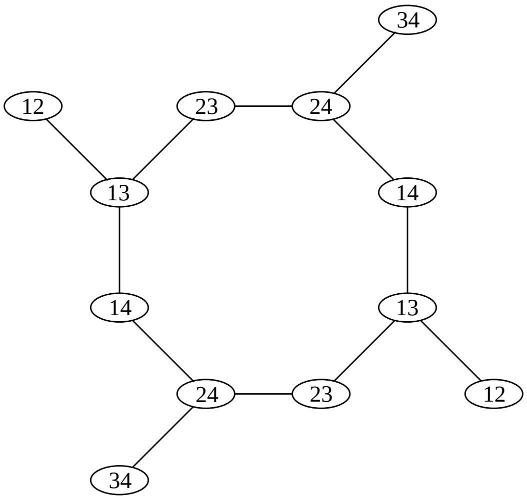

FIGURE 4.4: Not an A_4-graph.

Since Γ is assumed to have integer edge weights, the entries of A, B, C (and J) are scalars independent of q or t. Noting that $\phi_r(t)$ is a polynomial in t of degree $r-1$, it follows that ψ_p vanishes only if $B\phi_r(A)JC = 0$ for $1 \leqslant r < p$. Using the same feature of ϕ_r a second time, we conclude that ψ_p vanishes only if $BA^rJC = 0$ for $0 \leqslant r \leqslant p-2$.

However, a simple calculation reveals that

$$BA^rJC = \begin{cases} B_1(A_2A_1)^{r/2}C_2 - B_2(A_1A_2)^{r/2}C_1 & \text{if } r \text{ is even,} \\ B_2A_1(A_2A_1)^{(r-1)/2}C_2 - B_1A_2(A_1A_2)^{(r-1)/2}C_1 & \text{if } r \text{ is odd,} \end{cases}$$

and the matrix entries of these products are edge-weighted counts of certain restricted paths in Γ. For example, A_1 is the submatrix of m obtained by selecting the columns from $V_{i/j}$ and rows from $V_{j/i}$, and similarly C_2 has the columns of $V_{ij/\varnothing}$ and the rows of $V_{i/j}$. From this it is easy to deduce that the (v,u)-entry of BA^rJC is $N_{ij}^{r+2}(\Gamma;u,v) - N_{ji}^{r+2}(\Gamma;u,v)$ for all $u \in V_{ij/\varnothing}$ and all $v \in V_{\varnothing/ij}$. $\square$

Note that in the special case $r = 2$, the Polygon Rule applies to all distinct pairs $i, j \in I$ (assuming that (W,S) is braid-finite), and amounts to a Diamond Rule that relates counts of alternating 2-step paths forming a diamond shape.

EXAMPLE 4.8. Consider the 4-cycle in the center of Comp(A_4) (see Figure 4.1). By winding around this cycle one or more times, it is easy to construct arbitrarily large admissible I-labeled graphs that satisfy the Compatibility, Simplicity, and Bonding Rules, such as the graph in Figure 4.4. However, if u is one of the vertices with $\tau = \{2,4\}$, then there is exactly one alternating path of length 2 and type $(2,4)$ starting at u, and one of type $(4,2)$, but these two paths terminate at two distinct vertices with τ invariant $\{1,3\}$. This violates the 2-step Polygon Rule, so this graph is not an A_4-graph.

We claim that the rules collected so far yield the following characterization of admissible W-graphs. Note that the ingredients of this characterization (especially the Polygon Rule) are closely related to the W-graph definition used by Lusztig in the appendix to [**L**].

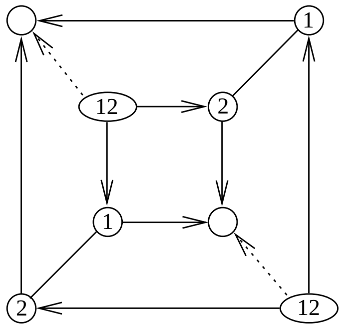

FIGURE 4.5: The generic admissible A_2-graph.

THEOREM 4.9. *If (W, S) is braid-finite, then an admissible I-labeled graph is a W-graph if and only if it satisfies*

(a) *the Compatibility Rule (Proposition 4.1),*
(b) *the Simplicity Rule (Proposition 4.2),*
(c) *the Bonding Rule (Proposition 4.4), and*
(d) *the Polygon Rule (Proposition 4.7).*

Proof. We know that the four stated conditions are necessary. For sufficiency, we need to show that if $\Gamma = (V, m, \tau)$ is admissible and satisfies (a)–(d), then for all $i, j \in I$, the operators T_i and T_j satisfy the braid relations on M_Γ. From the proof of Proposition 4.7, it is clear that if the Polygon Rule is satisfied, then the quantity ψ_p defined in Lemma 4.6 vanishes, where p is the order of $s_i s_j$ in W. Furthermore, if the Bonding and Simplicity Rules are satisfied, then Theorem 3.4 implies that the subgraph Γ' of $\Gamma{\downarrow_{\{i,j\}}}$ spanned by the vertices in $V_{i/j} \cup V_{j/i}$ is a $W_{\{i,j\}}$-graph. (In the case $p = 2$, this requires use of the Compatibility Rule.) Letting A denote the edge weight matrix of Γ', it follows that $\phi_p(A) = 0$ (Proposition 3.1), and hence Lemma 4.6 implies that T_i and T_j satisfy the length p braid relation. $\square$

REMARK 4.10. Recalling that admissible W-graphs are bipartite, it follows that in the generic case, an admissible A_2-graph Γ has eight types of vertices, corresponding to the four possible τ invariants and a choice of parity. The possible orientations of edges among these eight types are illustrated in Figure 4.5. In these terms, Γ may be interpreted as a special type of commutative diagram. 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 in both directions (as forced by the Simplicity and Bonding Rules), and the Polygon Rule amounts to the condition that the subdiagram formed by the solid edges is commutative.

5. Admissible Cells in Low Rank

As an application of Theorem 4.9, we claim the following classification results.

PROPOSITION 5.1. *A complete set of admissible W-cells for $W = A_3$, B_3, H_3, A_4 and D_4 appear in Figures 4.3 and 5.1–4, respectively.*

Before beginning the proof, we first outline a general classification strategy. In a sequel to this paper, we will describe a much more refined approach to the cell classification problem that is suitable for larger Coxeter systems.

Step 1. Compile a list $\mathcal{L} = \mathcal{L}(W, S)$ of all isomorphism classes of I-labeled graphs that are (a) admissible, (b) comprised only of simple edges (no arcs), (c) connected, and (d) satisfy the Compatibility and Bonding Rules. This list will usually be infinite, even though there may only be finitely many admissible W-cells.

Every admissible W-cell may be obtained by selecting a multiset $\mathcal{M} = \{\Gamma_1, \ldots, \Gamma_l\}$ from $\mathcal{L}$, properly 2-coloring the vertices of each component Γ_i, inserting arcs between various pairs of vertices of opposite color (the coloring is used to guarantee that the result is bipartite), and adjusting the weights on the arcs so that the Polygon Rule is satisfied.

In some cases, by judicious application of the Polygon Rule, we can eliminate in advance the use of certain members of $\mathcal{L}$ in the multisets $\mathcal{M}$.

Step 2. To narrow the range of possibilities for $\mathcal{M}$, impose a directed graph structure on the list $\mathcal{L}$ by declaring $\Gamma \prec \Gamma'$ if there is a strict inclusion $\tau(u) \subset \tau(u')$ between a vertex u in Γ and u' in Γ'. Every arc in an admissible cell supported on the multiset $\mathcal{M}$ must be directed from a vertex in Γ_i to a vertex in Γ_j such that $\Gamma_j \prec \Gamma_i$, so a necessary condition on $\mathcal{M}$ is that it spans a strongly connected subgraph of $\mathcal{L}$.

The Polygon Rule may be used in some cases to delete edges from $\mathcal{L}$ in advance.

For Coxeter systems of low rank, most of the strongly connected subgraphs of $\mathcal{L}$ are singletons $\{\Gamma\}$. In such cases, one cannot use more than one copy of the graph Γ to construct a W-cell unless there is a loop at Γ (i.e., $\Gamma \prec \Gamma$).

Step 3. Use the Polygon Rule to identify all remaining multisets $\mathcal{M}$ that can be bound together by arcs into W-cells.

Proof of Proposition 5.1 (sketch). We have previously described how to classify the admissible A_3-cells (see Remark 4.5). In this case, the list $\mathcal{L}(A_3)$ is finite and its graph structure has no loops or cycles. Thus the members of $\mathcal{L}$ are the only possible admissible A_3-cells, and they *are* A_3-cells because the Polygon Rule holds vacuously: there are no alternating paths of length $\geqslant 2$ of any type in any of these graphs. (See Figure 4.3.)

B_3 and H_3. For these groups, we can number the generators so that they share the same compatibility graph with A_3 (see Figure 4.1) and $(s_1s_2)^4 = 1$ (for B_3) or $(s_1s_2)^5 = 1$ (for H_3). Here, according to the Bonding Rule, the nontrivial cells that may occur in $\{1, 2\}$-restrictions are Dynkin diagrams of types A_1 or A_3 (for B_3) or A_4 (for H_3).

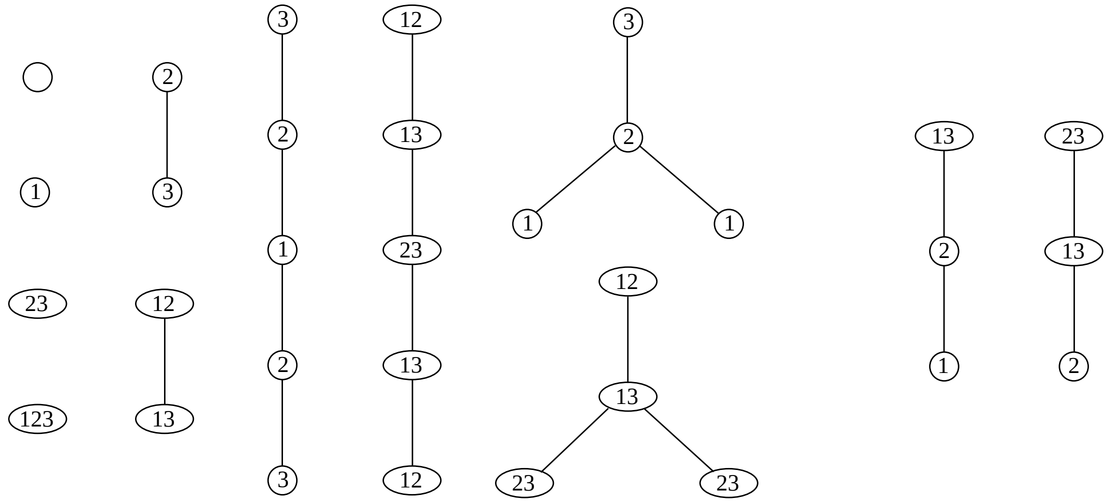

(a) uniform (b) non-uniform

FIGURE 5.1: The admissible B_3-cells.

In generating $\mathcal{L}$, it is helpful to separately identify the graphs whose τ invariants have uniform cardinality from those that are non-uniform (and therefore have at least one simple edge of the form $\{2\}-\{1,3\}$). Given the limited possibilities for inclusions of τ invariants in rank 3, one sees *a priori* that the only possible loops or cycles in $\mathcal{L}$ must occur among the non-uniform graphs. Also, the uniform graphs have no alternating paths (a special feature in rank 3), so each one is an admissible cell; these are listed in Figures 5.1(a) and 5.2(a).

The non-uniform graphs in $\mathcal{L}(B_3)$ are

$$\Gamma_1 = \{1\}-\{2\}-\{1,3\},$$
$$\Gamma_2 = \{2\}-\{1,3\}-\{2,3\},$$

and in $\mathcal{L}(H_3)$ the three non-uniform graphs are

$$\Gamma_3 = \{3\}-\{2\}-\{1\}-\{2\}-\{1,3\},$$
$$\Gamma_4 = \{1\}-\{2\}-\{1,3\}-\{2,3\},$$
$$\Gamma_5 = \{2\}-\{1,3\}-\{2,3\}-\{1,3\}-\{1,2\}.$$

We have $\Gamma_2 \not\prec \Gamma_1$, so the only cycles in $\mathcal{L}(B_3)$ are the two loops involving Γ_1 and Γ_2. The one inclusion of τ invariants in Γ_1 (namely $\{1\} \subset \{1,3\}$) involves two vertices at an even distance, so any arcs must be directed between *distinct* copies of Γ_1. However, such an arc would create an alternating path of length 2 and type $(1,3)$, with no possibility of an alternating path of type $(3,1)$. Thus by the Polygon Rule, no such arc can occur. Similar reasoning applies to Γ_2. Since these two graphs have no alternating paths, we conclude that both are cells, and that they are the only remaining possibilities for B_3.

In the case of H_3, we claim that the edges $\Gamma_4 \prec \Gamma_3$ and $\Gamma_5 \prec \Gamma_4$ may be deleted from $\mathcal{L}$. In the former case, the one inclusion of τ invariants has the form $\{1\} \subset \{1,3\}$, and an arc

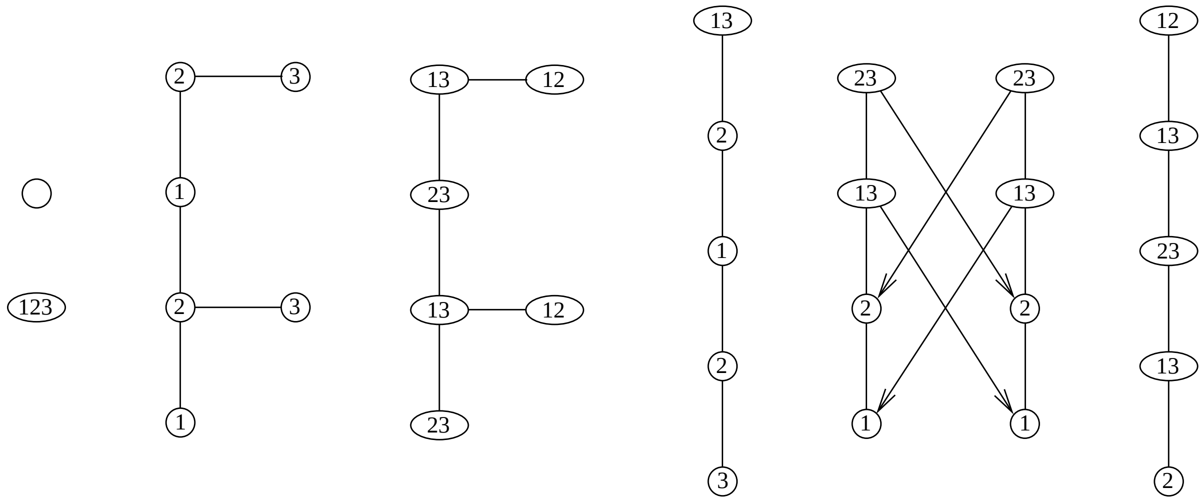

(a) uniform (b) non-uniform

FIGURE 5.2: The admissible H_3-cells.

directed between the corresponding pair of vertices creates an alternating path of length 2 and type $(1,3)$, with no possibility of an alternating path of type $(3,1)$. Similar reasoning applies in the second case.

The remaining possibilities for the support of an admissible H_3-cell involve multiple copies of a single non-uniform graph. In the case of Γ_3 or Γ_5, one cannot add arcs between two copies of either graph without violating the Polygon Rule, and without arcs, there are no alternating paths, so both graphs are H_3-cells. On the other hand, the graph Γ_4 has an alternating 3-step path of type $(3,2)$. The only way to create an alternating 3-step path of the opposite type is to add arcs linking this graph with a second copy of Γ_4; further applications of the Polygon Rule force a configuration of arcs as indicated in Figure 5.2(b). Also, one can show that there cannot be additional copies of Γ_4 linked to this graph without violating the Polygon Rule. For example, an arc of the form $\{1,3\} \to \{1\}$ from the first copy of Γ_4 to a third copy would create an alternating 3-step path of type $(2,3)$ from the second copy to the third without the possibility of any alternating paths of type $(3,2)$. After checking that the 8-vertex graph in Figure 5.2(b) satisfies the Polygon Rule, one concludes that it is the only remaining H_3-cell.

A_4. In this case, we claim that every graph Γ in $\mathcal{L}(A_4)$ that has more than one vertex with $\tau = \{2,4\}$ (or by symmetry, $\tau = \{1,3\}$) can be eliminated. Indeed, the only closed paths through $\{2,4\}$ in $\mathrm{Comp}(A_4)$ that do not violate the Bonding Rule necessarily contain the 4-cycle $\{2,4\}$—$\{1,4\}$—$\{1,3\}$—$\{2,3\}$—$\{2,4\}$ (see Figure 4.1). If the endpoints of this cycle lift to distinct vertices in Γ, then there are alternating 2-step paths of type $(1,3)$ and $(3,1)$ that originate at the vertex with $\tau = \{1,3\}$ and terminate at distinct vertices, and one can show that there is no possibility of alternating paths of the opposite types with the same endpoints.

It is easy to check that the only members of $\mathcal{L}(A_4)$ that are not eliminated by the above

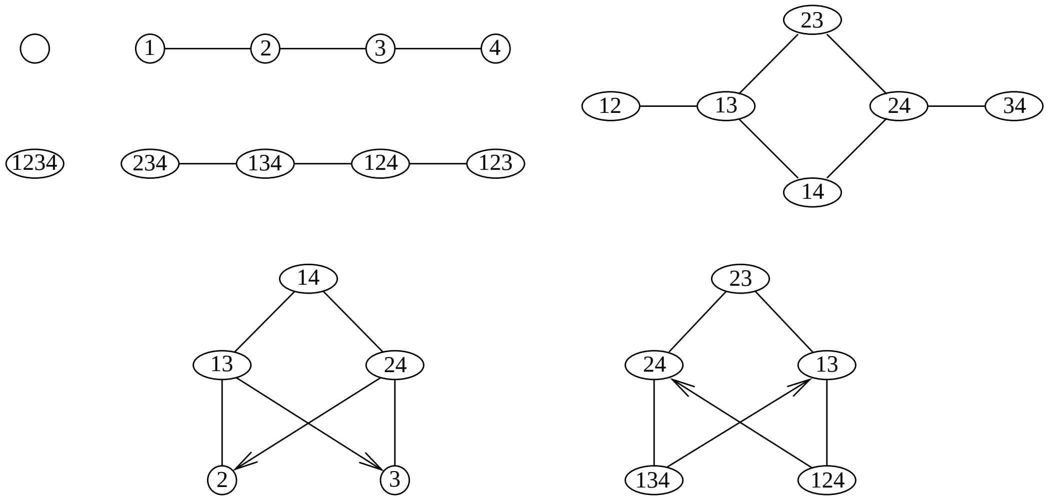

FIGURE 5.3: The admissible A_4-cells.

restriction are the graphs listed in Figure 5.3 (omitting the arcs from the two 5-vertex graphs). In the $\prec$-subgraph they span in $\mathcal{L}(A_4)$, the only cycles are loops involving the two 5-vertex graphs. By reasoning similar to the previous cases, one can use the Polygon Rule to eliminate the possibility of binding two or more copies of either graph into a cell and confirm that the graphs in Figure 5.3 are the only admissible A_4-cells.

D_4. Using the compatibility graph in Figure 4.1, it is not hard to check that $\mathcal{L}(D_4)$ consists of the 11 graphs with $\leqslant 3$ vertices in Figure 5.4, together with a cycle of length $6k$ for each $k \geqslant 1$ in which the τ invariants follow the period 6 pattern

$$\cdots - \{0,1\} - \{1,2\} - \{0,2\} - \{2,3\} - \{0,3\} - \{1,3\} - \{0,1\} - \cdots.$$

The graph structure of $\mathcal{L}(D_4)$ has no loops, and there are three strongly connected components that are not singletons: two components each with three 3-vertex graphs, and a third involving the 2-vertex graph $\{0\} - \{1,2,3\}$ and the $6k$-cycles. One may use the Polygon Rule to eliminate the possibility of arcs binding together any two or more of the 3-vertex graphs in the first two components. Furthermore, since the 3-vertex graphs have no alternating paths, they each form cells.

In the third component, note that the 2-vertex graph is itself a cell (it has no alternating paths), whereas the $6k$-cycles each have alternating 3-step paths of type $(1,0)$ and none of type $(0,1)$ (for example). Thus any such cycle must be bound to at least one copy of the 2-vertex graph by arcs of the form $\{0,1\} \to \{0\}$ and $\{1,2,3\} \to \{2,3\}$, thereby creating alternating 3-step paths of type $(0,1)$. Application of the 2-step Polygon Rule shows that all of the arcs of the form $\{0,i\} \to \{0\}$ $(1 \leqslant i \leqslant 3)$ binding this pair of graphs together must have the same weight, and the same is true for the arcs of the form $\{1,2,3\} \to \{i,j\}$ $(1 \leqslant i < j \leqslant 3)$. This prevents the 2-vertex graph from being bound by arcs to any other cycle; otherwise, there would be alternating 3-step paths of type $(0,1)$ whose endpoints are in different cycles, with no possibility of a 3-step path of type $(1,0)$ connecting them. This

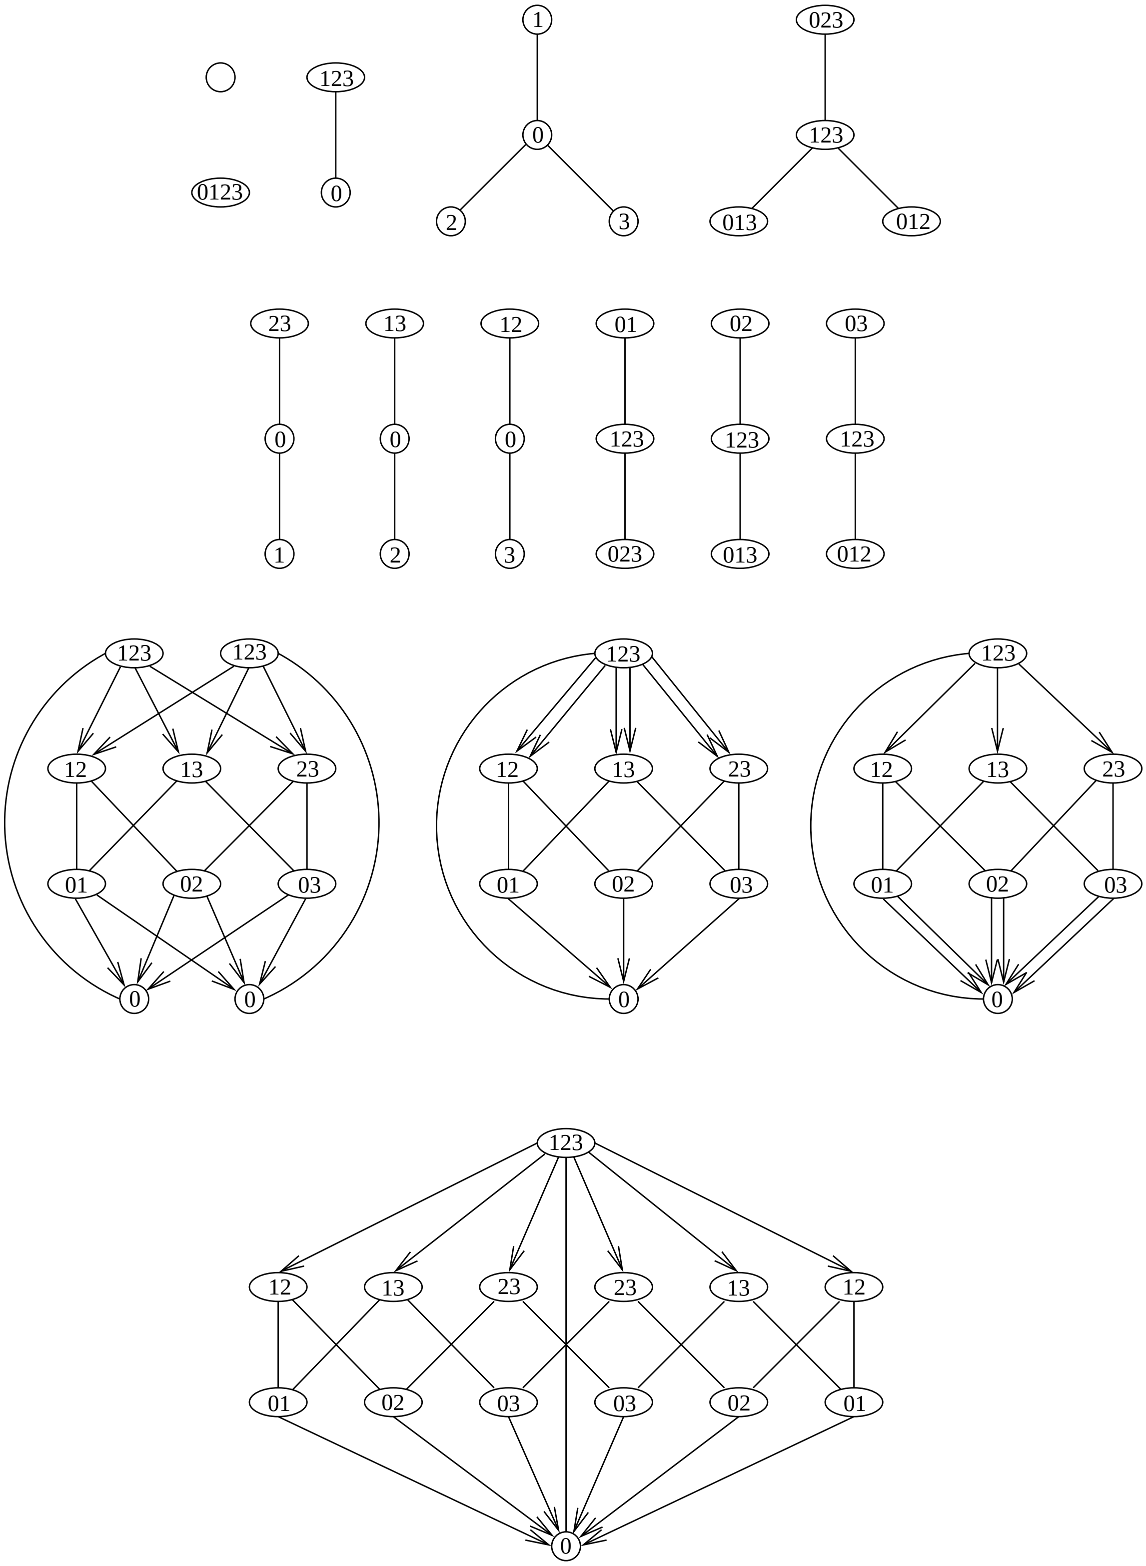

FIGURE 5.4: The admissible D_4-cells.

also forces the cycle to have length 6 or 12; otherwise, there would be vertices sufficiently far apart that cannot be connected by 3-step paths of type $(1,0)$, even though they are endpoints of 3-step paths of type $(0,1)$.

If the cycle has length 6, then there are two 3-step paths of type $(1,0)$ between the vertices with $\tau = \{0,1\}$ and $\tau = \{2,3\}$. Thus there must either be two 2-vertex graphs bound to it by arcs, all with weight 1, or there is only one such graph, with inbound arcs having weight 2 and outbound arcs having weight 1, or vice-versa. Otherwise, if the cycle has length 12, then there is a unique 3-step path of type $(1,0)$ connecting each pair of vertices with $\tau = \{0,1\}$ and $\tau = \{2,3\}$, so the cycle must be bound by arcs of weight 1 to one copy of the 2-vertex graph.

This accounts for all of the remaining graphs in Figure 5.4. $\square$

Remark 5.2. Among the admissible cells classified here, the only ones that are not Kazhdan-Lusztig cells are the two singleton B_3-cells with $\tau = \{1\}$ and $\tau = \{2,3\}$, the two doubleton B_3-cells, the doubleton D_4-cell, and the two 8-vertex D_4-cells. The latter two cells are noteworthy not only for the fact that they have edge weights > 1, but also for the fact that the representations of D_4 that they generate are isomorphic, thus providing an example where combinatorial rigidity fails (cf. Question 2.7). It is also interesting that the work of Garfinkle and Vogan [**GV**] on coherent continuation representations shows that these two cells cannot occur in any D_4-restrictions of Kazhdan-Lusztig cells or Harish-Chandra cells.

References

[**B**] N. Bourbaki, "Groupes et Algèbres de Lie, Chp. IV–VI," Masson, Paris, 1981.

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

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

[**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.

[**L**] G. Lusztig, Periodic W-graphs, *Represent. Theory* **1** (1997), 207–279.

[**M**] W. M. McGovern, Cells of Harish-Chandra modules for real classical groups, *Amer. J. Math.* **120** (1998), 211–228.

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

[**Y**] Yunchuan Yin, "W-graph representations for Coxeter groups and Hecke algebras," Ph. D. Thesis, University of Sydney, 2004.