REPRESENTATION THEORY
An Electronic Journal of the American Mathematical Society
Volume 8, Pages 267–289 (July 8, 2004)
S 1088-4165(04)00236-5

EXPLICIT MATRICES FOR IRREDUCIBLE REPRESENTATIONS OF WEYL GROUPS

JOHN R. STEMBRIDGE

Abstract. We present algorithms for constructing explicit matrices for every irreducible representation of a Weyl group, with particular emphasis on the exceptional groups. The algorithms we present will produce representing matrices in either of two forms: real orthogonal, with matrix entries that are square roots of rationals, or rational and seminormal. In both cases, the matrices are "hereditary" in the sense that they behave well with respect to restriction along a chosen chain of parabolic subgroups.

Introduction

The goal of this paper is to describe algorithms for generating explicit matrices representing the simple reflections in each of the irreducible representations of a Weyl group W. For the symmetric groups, such matrices are well known, the most prominent being the seminormal and orthogonal matrix models constructed by Alfred Young [Y] (see also [G], [JK], [OV], and [Ru]), and it is possible to extend these models to cover the remaining classical Weyl groups (e.g., see [F1] and [R]). Here, we are primarily concerned with the five exceptional groups.

An alternative approach to the representations of a Weyl group involves the W-graph construction of Kazhdan and Lusztig [KL1]. In this approach, the representing matrices are encoded (mainly) by a single edge-weighted graph whose vertices correspond to basis elements of the representation. The original W-graphs in [KL1] provide $\mathbf{Z}W$-modules for each Kazhdan-Lusztig cell, although not all irreducible representations are afforded by such cells. Later work of Gyoja [Gy] (see also the discussion in Chapter 11 of [GP]) demonstrates that there is a W-graph affording every irreducible representation of every Weyl group, but knowing the existence of a W-graph is not the same as having explicit matrices.

In order to clarify what we mean by "explicit," we should explain that our goal is not simply a description or algorithm, but a construction that is completely detailed down to the level of having computer files of representing matrices available for computation. Considering that the largest irreducible representation of the Weyl group of type $\mathcal{E}_8$ has dimension 7168, this places a premium on solutions in which the representing matrices are sparse and the entries are small in terms of the number of bits used to represent them. As far as we are aware, the solutions we have obtained are the first ones available that provide this level of explicitness for $W(\mathcal{E}_8)$.

Received by the editors March 12, 2004.
2000 *Mathematics Subject Classification.* Primary 20F55, 20C40; Secondary 05E15, 20-04.
This work was supported by NSF grants DMS–0070685 and DMS–0245385.

Our motivation for this work originates with the Atlas of Lie Groups Project[1], one of whose goals is to understand the structure and classification of the unitary representations of real and p-adic semisimple Lie groups. For example, in the split p-adic case, it is known from the work of Barbasch and Moy [BM] that the unitarity of a spherical representation may be detected by testing an element of the group algebra $\mathbf{R}W$ for positive semi-definiteness in the regular representation. By passing to the simple components of the group algebra, one may reduce this to a positivity test involving each irreducible W-representation. In the case of real groups, it is known there are necessary conditions for unitarity involving the positivity of an operator on some subset of the irreducible W-representations, and in the split real cases, there is hope that these necessary conditions are sufficient. (In the split classical cases, recent work of Barbasch confirms the sufficiency [B].) We plan to use the explicit matrix models reported on here to apply these tests for unitarity in the exceptional cases, with the ultimate goal being the classification of the spherical unitary duals of the exceptional real and p-adic groups.

The philosophy of our approach follows that of Young—we construct matrix models that are "hereditary" in the sense that they behave well when the action is restricted along a chosen chain of parabolic subgroups of W. It happens that for every Weyl group except $W(\mathcal{E}_8)$, one may choose the subgroup chain so that branching from each level to the next is multiplicity-free; this renders the models essentially unique up to a diagonal change of basis. Among the various possible diagonal rescalings, our algorithms single out two: one in which the representing matrices are real orthogonal and have matrix entries whose squares are rational (these are unique up to diagonal rescaling by factors of ± 1), and a second whose matrix entries are rational. The latter "rational seminormal" models require non-canonical choices to be made, but may be optimized for quality.

The Maple programs we developed to implement the algorithms described here, the resulting matrices that these programs produced, and many tables of statistics, such as measures of the sparseness and quality of the matrices, are available at

⟨`www.math.lsa.umich.edu/~jrs/archive.html`⟩, and
⟨`atlas.math.umd.edu/unitarity/weyl/hereditary`⟩.

In a project parallel to ours, Adams has developed an alternative approach that has yielded integral matrix models for most of the irreducible representations of the exceptional groups, but not yet all of $W(\mathcal{E}_8)$. See

⟨`atlas.math.umd.edu/unitarity/weyl/integral`⟩.

An outline of the paper follows.

We first discuss general features of hereditary models for representations of finite groups; for example, we show that under mild conditions, if a representation is realizable over some subfield F of $\mathbf{C}$, then it has a unitary hereditary model with matrix entries whose squares belong to F (Proposition 1.1).

In Section 2, we specialize to the case of Weyl groups. One of the peculiar features that develops in this case is that for certain "graceful" chains of parabolic subgroups, it is possible to compute traces of products of distinct simple reflections by pointwise multiplication of the diagonals of the corresponding matrices (Corollary 2.6). This trick, first exploited by Rutherford in the symmetric group case [Ru] (see also Greene [G]), plays a key role in our approach.

[1] See ⟨`atlas.math.umd.edu`⟩.

In the final two sections, we describe the algorithms. We do not provide low-level, line-by-line details of the implementation; the intent is to provide the reader with sufficient information to write his or her own implementation.

Among the numerous computational issues we address, the core problem is one of efficiently generating and solving a large system of quadratic equations that define a 0-dimensional variety whose points are orthogonal matrix models for a chosen irreducible representation. For example, the system we use to identify matrices for the largest representation of $W(\mathcal{E}_8)$ has (roughly) 15,000 equations and 600 variables over a finite extension of $\mathbf{Q}$. In our experience, general-purpose Gröbner basis packages are not adequate for a computation of this scale, so we devised a special algorithm that uses Gröbner-like reductions to find a solution. While we are unable to prove *a priori* that this algorithm will necessarily find a solution, the fact remains that the algorithm did succeed in finding a solution for every irreducible representation of every Weyl group of rank $\leqslant 8$ (and, in particular, all of the exceptional groups), and the full calculation took only a few hours of CPU time and 50MB of memory on a 2.8GHz Pentium IV running Maple 9.

Further problems. It would be interesting to convert the rational matrix models produced by our algorithm to integral form by identifying a basis for the lattice generated by the W-orbit of the natural coordinates. What is not clear is whether a sparse model of this type exists.

Another interesting problem would be to explicitly determine the hereditary orthogonal models for $W(\mathcal{D}_n)$-representations relative to the parabolic chain

$$W(\mathcal{D}_2) \leqslant W(\mathcal{D}_3) \leqslant \cdots \leqslant W(\mathcal{D}_n).$$

Combined with the exceptional group results reported here and the analogous (known) results for $W(\mathcal{A}_n)$ and $W(\mathcal{B}_n)$, this would yield a complete set of hereditary models for all Weyl group representations. Note that Ram [R] provides explicit (but non-hereditary) matrices for the irreducible representations of $W(\mathcal{D}_n)$ via branching from $W(\mathcal{B}_n)$.

Acknowledgment. I would like to thank Jeffrey Adams for helpful discussions, and for providing access to additional computer hardware during the development of the software.

1. Hereditary models

Let $s_1, \ldots, s_n$ be an ordered list of generators for some finite group W, and let W_k denote the subgroup of W generated by $s_1, \ldots, s_k$. Eventually we will take W to be a Weyl group and $s_1, \ldots, s_n$ an ordered list of simple reflections, but we can afford to begin in this more general context.

Let V be a finite-dimensional W-module over the complex field $\mathbf{C}$. We say that a basis $\mathcal{B}$ for V is *hereditary* (relative to $s_1, \ldots, s_n$) if for all $k \leqslant n$, the basis may be partitioned into disjoint blocks $\mathcal{B}_1, \ldots, \mathcal{B}_l$ so that

(i) each block spans an irreducible W_k-submodule, and
(ii) if $\mathcal{B}_i$ and $\mathcal{B}_j$ span isomorphic W_k-modules, then the representing matrices for W_k relative to $\mathcal{B}_i$ and $\mathcal{B}_j$ are identical (equivalently, there is a bijection $\mathcal{B}_i \to \mathcal{B}_j$ that extends to a W_k-module isomorphism).

Analogously, we say that a matrix representation of W is hereditary if the basis formed by the natural coordinates has this property.

It is an easy consequence of complete reducibility over $\mathbf{C}$ that every W-module V has a hereditary basis. Indeed, when $n = 0$ every basis is hereditary. Otherwise for $n \geqslant 1$, we may proceed by induction and assume the existence of a hereditary basis $\mathcal{B}$ for V as a W_{n-1}-module. We may then decompose V into irreducible W-modules V_i, and each V_i into irreducible W_{n-1}-modules V_{ij}. For each summand V_{ij}, there is at least one block $\mathcal{B}_{ij}$ of $\mathcal{B}$ that spans a W_{n-1}-module isomorphic to V_{ij}, and we may obtain a W-hereditary basis for V by selecting an isomorphic image of $\mathcal{B}_{ij}$ from within each V_{ij}.

We say that a basis $\mathcal{B}$ for V (hereditary or not) is *unitary* if the representing matrices for W relative to $\mathcal{B}$ are unitary. A basis is *seminormal* if it may be converted to a unitary basis via some diagonal transformation. Equivalently, this means that $\mathcal{B}$ is orthogonal with respect to some positive definite W-invariant inner product on V.

Given a subfield F of $\mathbf{C}$, we say that V is *realizable over* F if there is a basis $\mathcal{B}$ of V such that the representing matrices for W with respect to $\mathcal{B}$ have entries in F; in that case, we say that $\mathcal{B}$ is an *F-basis.* Also, let $\bar{F}$ denote the smallest extension of F closed under complex conjugation; i.e., the field generated by the real and imaginary parts of F. In all cases of interest, such as $F = \mathbf{Q}$, $\mathbf{R}$, or an algebraic number field, we have $\bar{F} = F$.

Proposition 1.1. *If every irreducible W_k-submodule of V is realizable over F for all $k \leqslant n$, then V has a hereditary $\bar{F}$-basis that is seminormal. Moreover, V has a hereditary unitary basis in which the diagonals of all representing matrices, as well as the squares of all matrix entries, are in $\bar{F}$.*

Proof. Without loss of generality, we may assume that $\bar{F} = F$.

Let $\mathcal{B}$ be an F-basis for V and $\langle\,,\rangle$ a positive definite Hermitian inner product that is W-invariant and F-valued on the F-span of $\mathcal{B}$. A standard way to construct the latter is to start with the inner product $B(\,,)$ relative to which $\mathcal{B}$ is orthonormal and then average over W, setting

$$\langle u, v\rangle := \frac{1}{|W|}\sum_{w\in W} B(wu, wv) \quad (u, v \in V).$$

This is F-valued on the F-span of $\mathcal{B}$ since $\bar{F} = F$.

Given any irreducible W-submodule U of V, there must be a W-module embedding $U \to V$ over the ground field F (or indeed, over any ground field of characteristic 0 where U and V may both be realized), so there is an image of an F-basis for U in the F-span of $\mathcal{B}$. Since $\langle\,,\rangle$ is F-valued, the same must be true for the orthogonal complement of U. Thus we may replace $\mathcal{B}$ with an F-basis $\mathcal{B}'$ that may be partitioned into blocks $\mathcal{B}'_j$ that span irreducible, orthogonal W-submodules V_j. Furthermore, we may arrange it so that isomorphic submodules have isomorphic bases.

Similarly, we may decompose each V_j into orthogonal, irreducible W_{n-1}-modules, and we may select an F-basis from one member U of each isomorphism class that occurs, and make an F-linear change of basis within each block $\mathcal{B}'_j$ to obtain a new F-basis $\mathcal{B}''$ that may be partitioned into bases for orthogonal, irreducible W_{n-1}-submodules. Again, we may arrange it so that isomorphic submodules have isomorphic bases. Continuing this process down to the level of 1-dimensional W_0-modules, we obtain an orthogonal (i.e., seminormal) hereditary F-basis for V.

Finally, we may convert this seminormal F-basis to unitary form by rescaling the vectors to unit length relative to $\langle\,,\rangle$. Since there is a unique W_k-invariant form on each irreducible W_k-module (up to a scalar multiple), it follows that this renormalization preserves the basis isomorphisms between the blocks that span isomorphic W_k-submodules, and hence the basis remains hereditary. Note also that this change of basis creates matrix entries whose squares belong to F, but (as with any diagonal change of basis) has no effect on the diagonals of the representing matrices. □

We say that V is *totally free* (relative to $s_1, \dots, s_n$) if for all $k \leqslant n$, every irreducible W_k-submodule of V is multiplicity-free as a W_{k-1}-module.

Proposition 1.2. *If V is irreducible and totally free, then all hereditary bases for V are diagonal transformations of each other, and hence seminormal. In particular, the diagonals of all representing matrices are independent of the choice of hereditary basis.*

Proof. The hypotheses force V to be multiplicity-free as a W_{n-1}-module, so the (canonical) W_{n-1}-isotypic components of V provide the unique decomposition of V into irreducible W_{n-1}-modules. Each of these submodules is multiplicity-free as a W_{n-2}-module, so their decompositions into irreducible W_{n-2}-submodules are unique, and so on. Since irreducible W_0-modules are trivial, it follows that each element of a hereditary basis is uniquely determined up to a choice of scalar, and hence, all such bases are related by diagonal transformations. By Proposition 1.1, at least one such basis is seminormal, hence all hereditary bases are seminormal. □

It is a general principle that unique or canonical objects are easier to construct than those that require choices to be made. In this sense, the best matrix models for totally free W-modules are those arising from unitary $\mathbf{R}$-bases. In this case, the representing matrices are (real) orthogonal, and any two such (hereditary) bases are related by a diagonal orthogonal transformation; i.e., there is a unique such basis up to factors of ± 1.

Corollary 1.3. *If V is totally free, then the matrix entries relative to any unitary hereditary $\mathbf{R}$-basis are canonical up to sign.*

Given an irreducible, totally free V that has hereditary $\mathbf{R}$-bases, consider the problem of constructing real orthogonal matrices representing the action of W on V. By taking direct sums of matrix models for irreducible W_{n-1}-modules of the appropriate multiplicity, we may recursively assume that orthogonal matrices $A_1, \dots, A_{n-1}$ representing the action of $s_1, \dots, s_{n-1}$ on V have been previously constructed.

Proposition 1.4. *Let V be a W-module that is realizable over $\mathbf{R}$ and multiplicity-free as a W_{n-1}-module. Given real orthogonal matrices $A_1, \dots, A_{n-1}$ as above, the number of real orthogonal matrices A_n such that $s_k \mapsto A_k$ $(1 \leqslant k \leqslant n)$ extends to a representation isomorphic to the W-action on V is 2^{m-l}, where l and m denote the number of irreducible constituents in the actions of W and W_{n-1} on V, respectively.*

Proof. Any two orthogonal matrices that could represent the action of s_n on V, while at the same time being compatible with having s_k represented as A_k for $k < n$, must be related by an orthogonal change of basis that commutes with the action of W_{n-1}. However, V is multiplicity-free as a W_{n-1}-module, so this change of basis must act as the scalar ± 1 on each irreducible W_{n-1}-submodule of V (Schur's Lemma). Since V has m such constituents, there are 2^m such changes of basis, and the action of this group of base changes has a kernel of order 2^l. $\square$

We say that two W-modules are *clones* if they are isomorphic as W_{n-1}-modules.

Corollary 1.5. *Given V and $A_1, \dots, A_{n-1}$ as above, the number of real orthogonal matrices A_n such that $s_k \mapsto A_k$ $(1 \leqslant k \leqslant n)$ extends to a W-representation is $\sum 2^{m-l_i}$, where the sum ranges over isomorphism classes of clones of V, and l_i is the number of irreducible constituents of the ith clone, as a W-module.*

If we take the matrices $A_1, \dots, A_{n-1}$ as (recursively) granted, the advantage of imposing only the condition that A_n should be orthogonal and generate (with $A_1, \dots, A_{n-1}$) a representation of W is that it depends only on the group structure of W, rather than prior knowledge of the W-action on V. The disadvantage is the possibility of spurious solutions arising from clones, but these may be eliminated if the W-character of V is known.

2. Weyl groups

Henceforth, we assume that W is a Weyl group; i.e., a finite crystallographic group generated by reflections in a real Euclidean space, and that $s_1, \dots, s_n$ is an ordering of a set of simple reflections for W. We could possibly replace $s_1, \dots, s_n$ with any sequence of (not necessarily simple) reflections that generate W, but it will develop that this affords no particular advantage.

A. Realizability. As a starting point, it should be noted that every irreducible representation of a Weyl group is realizable over $\mathbf{Q}$, a result that was first obtained on a case-by-case basis. For the classical Weyl groups, it can be traced back to the work of Young [Y] (in particular, see QSA V for types $\mathcal{B}$ and $\mathcal{D}$); the exceptional groups were settled by Kondo [K] and Benard [Be]. Later, Springer's construction provided a more unified approach to the subject (see [Sp] and [KL2]). In view of Proposition 1.1, we may conclude the following.

Theorem 2.1. *Every representation of a Weyl group has a hereditary $\mathbf{Q}$-basis that is seminormal, as well as a hereditary $\mathbf{R}$-basis that is unitary and has matrix entries whose squares are rational.*

B. Naming conventions. It will be convenient to have a canonical name attached to each irreducible representation of each exceptional Weyl group. In the context of a given group W, we will use names such as

$$R_m, \quad R_m(t), \quad R_m^{\varepsilon}(t), \quad R_m^{\varepsilon}(t_1, t_2)$$

for an irreducible representation of dimension m in which a reflection has trace t (or there are two conjugacy classes of reflections, having traces t_1 and t_2), and the sign of the trace of the longest element is ε (one of $+$, 0, or $-$). For each exceptional group, this is nearly sufficient to uniquely identify every irreducible representation, the only exceptions being the two 6-dimensional representations of $W(\mathcal{F}_4)$.

We remark that the exceptional groups with two conjugacy classes of reflections (namely, $W(\mathcal{F}_4)$ and $W(\mathcal{G}_2)$) have outer automorphisms that interchange the two classes, so they need not be distinguished in any particular way, as long as the usage is consistent.

C. Standard chains. For each irreducible W, we fix a particular ordering of the simple reflections as follows.

In the classical cases, we require $(s_k s_{k+1})^3 = 1$ for all $k < n$, except $(s_1 s_2)^4 = 1$ in $W(\mathcal{B}_n)$, and $(s_1 s_2)^2 = (s_1 s_3)^3 = 1$ in $W(\mathcal{D}_n)$. All other pairs of simple reflections commute. In this way, W_{n-1} is a classical Weyl group in the same series as W (aside from a few small degeneracies). For $W(\mathcal{E}_8)$, we order the simple reflections according to the following numbering of its Coxeter graph:

$$\begin{array}{c} \quad\;\; 2 \\ \quad\;\; | \\ 1{-}3{-}4{-}5{-}6{-}7{-}8. \end{array}$$

In turn, this induces orderings for $W(\mathcal{E}_6)$ and $W(\mathcal{E}_7)$. In $W(\mathcal{F}_4)$, we follow a linear ordering of the Coxeter graph (so $(s_1 s_2)^3 = (s_2 s_3)^4 = (s_3 s_4)^3 = 1$); for $W(\mathcal{G}_2)$ there is only one ordering up to automorphisms.

We refer to these orderings as *standard*.

It will be convenient to adopt the practice that non-conventional Weyl group names, such as $W(\mathcal{E}_5)$, $W(\mathcal{D}_3)$, or $W(\mathcal{F}_3)$, refer to the subgroup of the appropriate rank in the standard chain suggested by its name. To be more pedantic, we are working in a category of *ordered* Coxeter systems; from this standpoint, $W(\mathcal{E}_5)$ and $W(\mathcal{D}_5)$ are objects in this category that are not equivalent, even though the underlying groups are isomorphic.

D. Branching. We will take as granted that the character table of each Weyl group W and the fusion maps between the conjugacy classes of W and its reflection subgroups (especially the subgroups W_k) are known, in the sense that they are readily available for computing branching multiplicities for restriction to reflection subgroups, as well as traces.

Of course, for the classical Weyl groups of types $\mathcal{A}$, $\mathcal{B}$, $\mathcal{D}$ (and $\mathcal{G}_2$), these are well known and easy to compute. Among the exceptional groups, the character table of $W(\mathcal{F}_4)$ was first obtained by Kondo [K], and $W(\mathcal{E}_n)$ ($n = 6, 7, 8$) by Frame [F2], [F3]. Modern computer algebra packages such as *GAP* and *Magma* can easily generate these character tables starting from a permutation representation of the group; the Maple package *coxeter* provides character tables and fusion maps for all of the finite Coxeter groups [S2].

It is well known that branching from a classical Weyl group to the previous Weyl group in the same series is multiplicity-free. Less well known is that multiplicity-free branching is also found among the exceptional groups (for example, Ram [R] notes that branching from $W(\mathcal{F}_4)$ to $W(\mathcal{B}_3)$, $W(\mathcal{E}_7)$ to $W(\mathcal{E}_6)$, and $W(\mathcal{E}_6)$ to $W(\mathcal{D}_5)$ is multiplicity-free).

Empirical Fact 2.2. Every irreducible representation of every irreducible Weyl group is totally free with respect to the standard chain, except for the following nine representations of $W(\mathcal{E}_8)$: $R_{3240}(\pm 594)$, $R_{4536}(\pm 378)$, $R_{5600}(\pm 280)$, $R_{6075}(\pm 405)$, and R_{7168}.

Remark 2.3. (a) For a given Weyl group, there may be many parabolic subgroups W' such that branching from W to W' is multiplicity-free, and hence, the irreducible representations of W may be totally free relative to many different orderings of the simple reflections. For example, an easy special case of the Littlewood-Richardson Rule shows that branching from $W(\mathcal{A}_n)$ to $W(\mathcal{A}_{n-2}) \times W(\mathcal{A}_1)$ is multiplicity-free, and every representation of $W(\mathcal{A}_4)$ is totally free with respect to every ordering of the simple reflections.

(b) None of the nine representations of $W(\mathcal{E}_8)$ listed above are multiplicity-free with respect to any reflection subgroup, so there is no generating set of reflections relative to which they are totally free. Also, the $W(\mathcal{E}_7)$-actions on these nine modules are nearly multiplicity-free: none of the irreducible constituents have a multiplicity that exceeds 2, and eight of the nine have just one constituent of multiplicity 2; R_{7168} has two.

E. Graceful chains and para-Coxeter classes. Recall that a Coxeter element of W is a product of the simple reflections, taken in any order. All such elements belong to a single conjugacy class. More generally, we define a *para-Coxeter element* of W to be a product of some subset of the simple reflections; i.e., a Coxeter element of some parabolic subgroup of W.

In an expository account of Young's work (see §23 of [Ru], as well as [G]), Rutherford observed that if $A_1, \ldots, A_n$ are the matrices representing $s_1, \ldots, s_n$ in Young's (hereditary) orthogonal or seminormal models for representations of the symmetric group $W(\mathcal{A}_n)$, then for distinct indices $i_1, \ldots, i_k$, we have

$$\delta(A_{i_1} \cdots A_{i_k}) = \delta(A_{i_1}) \cdots \delta(A_{i_k}),$$

where $\delta(A)$ denotes the diagonal of A (i.e., the matrix obtained by replacing all off-diagonal entries of A with 0's).

Given the representing matrices A_i, this provides a fast way to evaluate traces of para-Coxeter elements that avoids full matrix multiplication. In the symmetric group, every conjugacy class has a para-Coxeter element, so this provides a way to determine the entire character table of each of the symmetric groups.

Rutherford's observation may be generalized to hereditary representations of Weyl groups, but not without restrictions. For example, consider the symmetric group $W(\mathcal{A}_3)$, with the simple reflections in the non-standard order s_2, s_1, s_3. Relative to the (essentially unique) orthonormal basis that is hereditary for this order, the representing matrices for the reflection representation of $W(\mathcal{A}_3)$ are

$$A_2 = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & -1 \end{bmatrix}, \quad A_1 = \begin{bmatrix} 1 & 0 & 0 \\ 0 & -1/2 & 3\alpha \\ 0 & 3\alpha & 1/2 \end{bmatrix}, \quad A_3 = \begin{bmatrix} -1/3 & \beta & 6\alpha\beta \\ \beta & 5/6 & -\alpha \\ 6\alpha\beta & -\alpha & 1/2 \end{bmatrix},$$

where $\alpha^2 = 1/12$ and $\beta^2 = 2/9$. However, one may check that $\delta(A_1 A_3) \neq \delta(A_1)\delta(A_3)$.

Returning to the general case, we say that an ordering $s_1, \ldots, s_n$ of the simple reflections is *graceful* if for all $i < j < k$, we have i and k adjacent in the Coxeter graph (i.e., s_i and s_k do not commute) only if j and k are also adjacent. It is easy to check that the standard order we have chosen for each Weyl group is graceful, but the non-standard order in the above example for $W(\mathcal{A}_3)$ is not.

Lemma 2.4. *If $s_1, \ldots, s_n$ is graceful, then every irreducible, totally free W_{n-1}-module is multiplicity-free as a W_r-module, where r is the largest index $(0 \leqslant r < n)$ such that s_n centralizes W_r.*

Proof. For clarity, we will assume that $r = n - 3$; the general case follows by essentially the same argument. Given that the ordering is graceful, s_{n-1} and s_{n-2} must be the only simple reflections that do not commute with s_n. Since the Coxeter graph of every finite Weyl group is acyclic, it follows that W_{n-1} is a direct product of parabolic subgroups, say $W_I \times W_J$, with generating sets I and J that include s_{n-1} and s_{n-2}, respectively.

Each irreducible W_{n-1}-module V is therefore a tensor product of irreducible modules for W_I and W_J, say $U \otimes U'$. As a W_{n-2}-module, V must be a direct sum of the form $(U_1 \otimes U') \oplus \cdots \oplus (U_l \otimes U')$, where $U_1, \ldots, U_l$ are irreducible modules for the parabolic subgroup generated by $I - \{s_{n-1}\}$; if V is totally free, these modules must be distinct. Similarly, the freeness of V forces $U_1 \otimes U'$ to be multiplicity-free as a W_{n-3}-module, so it must be a direct sum of $U_1 \otimes U'_1, \ldots, U_1 \otimes U'_m$, where $U'_1, \ldots, U'_m$ are distinct irreducible modules for the parabolic subgroup generated by $J - \{s_{n-2}\}$. Hence, V is the direct sum of the W_{n-3}-modules $U_i \otimes U'_j$ ($1 \leqslant i \leqslant l$, $1 \leqslant j \leqslant m$), and is therefore multiplicity-free. □

Proposition 2.5. *Let $A_1, \ldots, A_n$ be the matrices of $s_1, \ldots, s_n$ relative to a hereditary basis for some W-module V. If $s_1, \ldots, s_n$ is a graceful ordering and V is totally free as a W_{n-1}-module, then $\delta(A_{i_1} \cdots A_{i_k}) = \delta(A_{i_1}) \cdots \delta(A_{i_k})$ for all distinct $i_1, \ldots, i_k$.*

Proof. More generally, we claim that the result remains true if we choose invertible diagonal matrices $D_1, \ldots, D_n$ and replace each A_i with $A'_i = D_i A_i$. We may further assume that n occurs among the indices $i_1, \ldots, i_k$; otherwise, replace n by $n-1$ and proceed by induction. It follows that $A'_{i_1} \cdots A'_{i_k} = BA'_nC$, where B, C, and BC are matrix products not involving A'_n (possibly B or C is an identity matrix). Since the coordinates are hereditary, each of B and C are block diagonal matrices, say $B = B_1 \oplus \cdots \oplus B_l$ and $C = C_1 \oplus \cdots \oplus C_l$, with the blocks corresponding to a decomposition of V into irreducible W_{n-1}-submodules V_i. It follows that

$$\delta(BA'_nC) = \delta\big(B_1(A'_n)_1C_1\big) \oplus \cdots \oplus \delta\big(B_l(A'_n)_lC_l\big),$$

where the notation $(\,\cdot\,)_i$ refers to the i-th block along the diagonal of a matrix.

As an operator on V_i, $(A_n)_i$ commutes with the action of W_r, where r denotes the largest index $< n$ such that s_n centralizes W_r. However, Lemma 2.4 implies that V_i is multiplicity-free as a W_r-module, so by Schur's Lemma, $(A_n)_i$, and hence also $(A'_n)_i$, must be a diagonal matrix. By adjusting the choice of diagonal matrices, we may therefore replace C_i with $C'_i := (A'_n)_iC_i$, yielding

$$\delta\big(B_i(A'_n)_iC_i\big) = \delta(B_iC'_i) = \delta(B_i)\delta(C'_i) = \delta(B_i)\delta((A'_n)_i)\delta(C_i),$$

using induction for the second equality. Hence $\delta(BA'_nC) = \delta(B)\delta(A'_n)\delta(C)$, and the result follows by induction. □

A consequence of Empirical Fact 2.2 is that every representation of a Weyl group W is totally free as a W_{n-1}-representation, relative to the standard chain. Hence,

Corollary 2.6. *In every matrix representation of a Weyl group that is hereditary relative to the standard chain, the matrices $A_1, \ldots, A_n$ for the simple reflections satisfy $\delta(A_{i_1} \cdots A_{i_k}) = \delta(A_{i_1}) \cdots \delta(A_{i_k})$ for all distinct $i_1, \ldots, i_k$.*

F. Clones. Recall that two W-modules are clones if they are isomorphic as W_{n-1}-modules. Among the symmetric groups, non-isomorphic clones are rare relative to the standard chain. Indeed, it is an easy consequence of Young's Rule for branching from $W(\mathcal{A}_n)$ to $W(\mathcal{A}_{n-1})$ that the only irreducible representations with clones are the reflection representations of $W(\mathcal{A}_n)$ for $n \leqslant 3$, and their sign twists.

Empirical Fact 2.7. If V is an irreducible W-module, then for every clone V' of V (relative to the standard chain) that is not isomorphic to V, there is a para-Coxeter element $w \in W$ such that the traces of w on V and V' differ.

Sketch of Proof. For the exceptional groups, this is an easy calculation involving the character tables. For the symmetric groups, the result is immediate, since every conjugacy class in a symmetric group has a para-Coxeter element.

For the case $W = W(\mathcal{B}_n)$, there is an irreducible W-module $V_{\mu,\nu}$ for each pair of partitions (μ, ν) of total size n. As a W_{n-1}-module, $V_{\mu,\nu}$ is the direct sum of the modules obtained by decreasing one part of μ or ν by 1 in all possible ways. The key consequence of this branching rule is that if $V_{\mu,\nu}$ has two or more constituents as a W_{n-1}-module, then any two of these constituents may be used to reconstruct (μ, ν). Thus if V' is a clone of an irreducible W-module V, then V' must be a direct sum of W-modules that are irreducible as W_{n-1}-modules. The latter are necessarily of the form $V_{\mu,\varnothing}$ or $V_{\varnothing,\nu}$, and hence either $n = 2$ (a case that may be checked separately), or V is also of this form. These modules restrict irreducibly to the symmetric group generated by $s_2, \dots, s_n$, so both V and V' would be clones relative to $W(\mathcal{A}_{n-1})$, whence $n \leqslant 4$ and the para-Coxeter classes in $W(\mathcal{A}_{n-1})$ suffice to distinguish V and V'.

For the case $W = W(\mathcal{D}_n)$, the $W(\mathcal{B}_n)$-modules $V_{\mu,\nu}$ and $V_{\nu,\mu}$ restrict irreducibly to the same W-module (if $\mu \neq \nu$), or to a sum of two distinct W-modules $V^{\pm}_{\mu,\mu}$ (if $\mu = \nu$ and n is even). The latter pairs are clones, and it is known (e.g., by Theorem A.1 of [S1]) that the traces of the para-Coxeter element $w = s_2 s_4 \cdots s_n$ on these two clones must differ. Via the branching rule for the B-series, one may deduce that the only other clones of an irreducible W-module V occur when $n \leqslant 4$, and these cases may be checked separately. □

Remark 2.8. For the exceptional groups, we have the following inventory of clones.

- $\mathcal{F}_4$: Seven irreducible representations have clones; in particular, R_{16} has 5 clones.
- $\mathcal{E}_5$: Eleven irreducible representations have clones, including $R_{20}(\pm 2)$ with 5 each.
- $\mathcal{E}_6$: No irreducible representations have clones.
- $\mathcal{E}_7$: R^+_{512} and R^-_{512} are clones of each other.
- $\mathcal{E}_8$: R_{5670} is a clone of $R_{1134} \oplus R_{4536}(0)$, and R_{7168} is a clone of $R_{2688} \oplus R_{4480}$.

3. Models for totally free representations

Now we consider the problem of constructing explicit matrices representing the action of the simple reflections in every irreducible representation of a Weyl group W. We will present algorithms for producing hereditary models that are real orthogonal, as well as for converting each real orthogonal model to an optimal rational seminormal form.

In this section, we consider the representations that are totally free relative to the standard chain; in the next section, we consider the non-free representations.

As the starting point for the algorithms, we will need only the information in the character tables and the fusion maps between the various subgroups W_k.

A. Sparsity issues. The largest irreducible representation of $W(\mathcal{E}_8)$ that is totally free is R_{5670}. A naive approach that allocated 32 bits for each matrix entry (assuming this is sufficient), and took no advantage of sparseness, would need about

1GB of memory to store the representing matrices for the 8 simple reflections, and multiplication of two such matrices (via the naive algorithm) would entail about 180 billion scalar multiplications.

In our Maple implementation of the algorithm, we use a sparse representation of matrices in which each row is stored as a linear form in a fixed but potentially infinite list of variables, say $e_1, e_2, \ldots$. For example, we would represent the 2×3 matrix whose rows are $(3, 4, 5)$ and $(6, 0, 7)$ as the list

```
[ 3*e1 + 4*e2 + 5*e3 , 6*e1 + 7*e3 ].
```

This has the advantage that the amount of storage space required is proportional to the number of nonzero entries, regardless of whether the matrix is sparse or dense.

Multiplication of two matrices A and B in this format amounts to composition of linear forms; in Maple, this may be achieved via the substitution

```
subs({ seq(var[i]=B[i], i=1..m)} , A),
```

where `m` denotes the number of rows of B, and `var=[e1, e2, ...]` is the list of variables. This has the advantage that only the nonzero matrix entries are multiplied, and when the coefficients of A are rational, the expansion and collection of terms in each row is done automatically by the (fast) Maple kernel.

In the hereditary models for R_{5670}, the total number of nonzero entries in the matrices representing the 8 simple reflections is 135496, an average of about 3 nonzero entries per row; Maple uses a total of about 1.5MB of memory to store the 8 matrices in the rational seminormal model found by our algorithm, and it takes roughly 3 seconds on a 2.8GHz Pentium IV to multiply the matrices representing s_7 and s_8.

B. Economies of space. Beyond the question of sparsity, there is extensive redundancy in the representing matrices for any hereditary model. Indeed, for a given $r \leqslant n$, one may partition a hereditary basis $\mathcal{B}$ into blocks $\mathcal{B}_1, \ldots, \mathcal{B}_m$ that span irreducible W_r-modules. If A_k is the matrix of s_k relative to $\mathcal{B}$ and s_k centralizes W_r, then Schur's Lemma implies that the submatrix of A_k formed by the rows indexed by $\mathcal{B}_i$ and the columns indexed by $\mathcal{B}_j$ must be a scalar multiple of the identity, say a_{ij}. Furthermore, we must have $a_{ij} = 0$ unless $\mathcal{B}_i$ and $\mathcal{B}_j$ span isomorphic W_r-modules. It follows that the $m \times m$ matrix $[a_{ij}]$ encodes all of the data needed to recover A_k, but in a compact form that has at most $\sum m_i^2$ nonzero entries, where $m_1, m_2, \ldots$ denote the multiplicities of the irreducible W_r-submodules in the given representation of W.

Definition 3.1. Given k and r as above, define $\phi_r(A_k)$ to be the matrix $[a_{ij}]$.

We remark that the map $s_k \mapsto \phi_r(A_k)$ extends to a representation of the parabolic subgroup of W generated by those simple reflections s_k that centralize W_r.

Given the recursive nature of $\mathcal{B}$, we may (recursively) assume the existence of previously constructed hereditary models for each irreducible W_{n-1}-submodule that occurs in the desired representation of W. The correct inventory of models may be identified via branching calculations derived from the character tables. By taking direct sums, one may thus obtain representing matrices A_k for s_k ($1 \leqslant k < n$), and the only new information required is the matrix A_n representing s_n. In turn, this is completely determined by the matrix $\phi_r(A_n)$, taking r to be the largest index such that s_n centralizes W_r.

In most cases, the standard chains have the property that s_n centralizes W_{n-2} and we may take $r = n-2$. The only exceptions occur in low ranks.

For example, we have $r = 6$ in $W(\mathcal{E}_8)$, and the representation R_{5670} decomposes into 88 irreducible summands relative to $W(\mathcal{E}_6)$; the multiplicities are 2 (four times), 3, 4, 5, 7, 9 (twice each), 10, and 14. Thus, aside from the storage required for the irreducible representations of smaller Weyl groups, the information contained in a hereditary matrix model for R_{5670} may be stored in a sparse 88×88 matrix with at most $672 = \sum m_i^2$ nonzero entries. In this way, we are able to save the data needed to reproduce orthogonal hereditary matrix models for every irreducible representation of $W(\mathcal{E}_n)$ for $n \leqslant 8$ (including the non-free cases) in a Maple table whose size is about 1MB.

C. The Coxeter relations. Continuing the above notation, we assume that r is the largest index such that s_n centralizes W_r, and that $A_1, \ldots, A_{n-1}$ are (recursively obtained) real orthogonal matrices representing $s_1, \ldots, s_{n-1}$ relative to a hereditary basis for an irreducible W-module V.

The real orthogonal matrices A_n that may be used to represent s_n must satisfy the Coxeter relations; i.e., $(A_i A_n)^{m(i,n)} = 1$, where $m(i,n)$ denotes the order of $s_i s_n$ in W. In particular, $m(n,n) = 1$, so the condition that A_n is orthogonal may be replaced with the condition that A_n is symmetric.

It is important to note that V may have clones, so the Coxeter relations alone are generally not sufficient to characterize A_n. However, given that V is totally free, there can only be finitely many choices for A_n that obey these relations (Corollary 1.5), and the ones that generate models for V must be among them.

In any case, we may view $\phi_r(A_n) = [a_{ij}]$ as a symmetric matrix of indeterminates, and take the Coxeter relations as a collection of polynomial conditions on the variables a_{ij}. In these terms, the matrix A_n induced by a given choice for $[a_{ij}]$ will necessarily commute with $A_1, \ldots, A_r$, so we need only to impose the relations $(A_i A_n)^{m(i,n)} = 1$ for $r < i \leqslant n$. To economize further, it suffices merely to require $\phi_r(A_n)^2 = 1$ and the braid relations

$$\phi_k(A_i)\phi_k(A_n)\phi_k(A_i)\cdots = \phi_k(A_n)\phi_k(A_i)\phi_k(A_n)\cdots \tag{3.1}$$

for $r < i < n$, where the number of factors on both sides is $m(i,n)$, and $k = k_i$ denotes the largest index such that s_i and s_n both centralize W_k. These conditions are either linear in the entries of A_n (if $m(i,n) = 2$), or quadratic (if $m(i,n) = 3$ or 4), or cubic (if $m(i,n) = 6$), so in almost all cases of interest, the equations will be quadratic.

Note that $\phi_k(A_i)$ and $\phi_k(A_n)$ may both be arranged into block diagonal form, the block sizes being the multiplicities of the irreducible W_k-submodules of V, say $n_1, n_2, \ldots$. Furthermore, since $\phi_k(A_i)$ and $\phi_k(A_n)$ are both symmetric, it follows that the difference between the two sides in (3.1) is either symmetric (if $m(i,n)$ is odd) or skew-symmetric (if $m(i,n)$ is even), so the total number of independent scalar conditions implicit in (3.1) is at most $\sum n_j(n_j+1)/2$ or $\sum n_j(n_j-1)/2$, depending on the parity of $m(i,n)$.

Similarly, the requirement that $\phi_r(A_n)^2 = 1$ is equivalent to $\sum m_j(m_j+1)/2$ scalar conditions, where $m_1, m_2, \ldots$ are the multiplicities of the irreducible W_r-modules in V.

As noted previously, it will usually be the case that $r = n-2$. Moreover, among these cases, it is most common that $m(n-1,n) = 3$, and that $k = n-3$ is the largest index such that s_{n-1} and s_n both centralize W_k. Thus, the Coxeter relations involving A_n will most typically amount to the conditions that $\phi_{n-2}(A_n)^2 = 1$ and

$$\phi_{n-3}(A_{n-1})\phi_{n-3}(A_n)\phi_{n-3}(A_{n-1}) = \phi_{n-3}(A_n)\phi_{n-3}(A_{n-1})\phi_{n-3}(A_n).$$

For example, when $W = W(\mathcal{E}_8)$, we have $r = 6$, $k = 5$, $m(7, 8) = 3$, and in the representation R_{5670}, the matrix $\phi_6(A_8)$ has $380 = \sum m_j(m_j+1)/2$ indeterminates. Also, the multiplicities $n_1, n_2, \dots$ are 9 (four times), 24 (three times), 12, 36, 48, 60 (twice each), and 18, so the Coxeter relations involving A_8 amount to $9131 = 380 + \sum n_j(n_j + 1)/2$ quadratic equations in 380 variables.

D. The orthogonal algorithm. Let χ be the character of an irreducible, totally free W-module V. To construct an orthogonal hereditary matrix model for V, we proceed as follows:

1. Decompose χ into irreducible characters relative to W_{n-1}, and recursively build orthogonal matrix models for each W_{n-1}-summand. Taking direct sums, we obtain representing matrices $A_1, \dots, A_{n-1}$ for the action of $s_1, \dots, s_{n-1}$ on V.

2. Following the techniques described in the previous subsection, use the Coxeter relations to generate a system of equations for the symmetric matrix $\phi_r(A_n) = [a_{ij}]$.

3. Test for clones; i.e., identify all W-characters χ' whose restriction to W_{n-1} agrees with that of χ. Once each irreducible W-character is decomposed into irreducible W_{n-1}-characters, this amounts to a simple partitioning problem that may be quickly solved by brute force. For each clone $\chi' \neq \chi$, one knows that there is a para-Coxeter element w such that $\chi'(w) \neq \chi(w)$ (Empirical Fact 2.7). Moreover, by Corollary 2.6, the trace of w on V is expressible in terms of the diagonals of $A_1, \dots, A_n$, and the condition that $\chi(w)$ is the trace of w on V amounts to a linear equation in the diagonal entries a_{ii} of $\phi_r(A_n)$.

4. Combine the equations in Step 2 with zero or more linear equations that eliminate the clones identified in Step 3, thereby obtaining a polynomial system whose solutions encode the possible (orthogonal) matrices representing the action of s_n on V; by Proposition 1.4, the number of solutions is exactly 2^{m-1}, where m denotes the number of irreducible constituents of V as a W_{n-1}-module. Find a solution of this system via the reduction algorithm described below.

For the totally free $W(\mathcal{E}_8)$-representation R_{5670}, we have seen that the system that determines A_8 consists of 9131 quadratic equations (and one linear equation to eliminate a clone—see Remark 2.8) in 380 variables. As an added complication, the entries of $A_1, \dots, A_7$ are square roots of rationals (Theorem 2.1), so the ground field for this system is an extension of $\mathbf{Q}$ by certain square roots of integers.

As we noted in the introduction, systems of this size may be too large to be handled by general-purpose Gröbner basis packages. Instead, we employ a sequence of Gröbner-like reductions that exploit the special features of these systems and mitigate against internal expression swell.

First, we order the variables so that a_{ij} $(i \leqslant j)$ precedes a_{kl} $(k \leqslant l)$ if either $i < k$, or $i = k$ and $j < l$. This implicitly assumes an ordering of the rows and columns of $\phi_r(A_n)$, or equivalently, an ordering of the blocks of coordinates that span irreducible W_r-modules. In all cases, we sort these by isomorphism class. The ordering of the classes (usually by increasing dimension), and the orderings within each class, are determined by the choices made during the recursive constructions in Step 1. In any case, we expect that the performance of the algorithm is relatively insensitive to these choices.

Once the variable ordering is established, we employ a degree-lex term ordering for monomials in these variables; i.e., monomial m_1 precedes monomial m_2 if m_1 has higher total degree, or if they have the same total degree and the first variable appearing in m_1/m_2 has positive degree.

Now we are ready to describe the reduction algorithm. Given the polynomials $q_1, \dots, q_l$ whose vanishing identifies the possible solutions for A_n, we proceed as follows:

1. If any of the polynomials q_i is linear, choose one with the fewest number of variables; set q_i aside, and use it to eliminate the first variable of q_i from the remaining system. (If any q_i is a nonzero constant, then a branch with no solutions has been encountered via Step 6 or 7; return a failure flag.)
2. Repeat Step 1 until no remaining equations are linear.
3. If all variables have been eliminated, then the set of saved linear equations forms a triangular system for a particular solution; solve it by back substitution and halt.
4. Otherwise, sort the remaining (nonlinear) polynomials by increasing number of dependent variables. For each $i = 1, \dots, l$, if there is a $j < i$ such that q_i and q_j have the same leading term,[2] replace q_i with $q_i - q_j$ (renormalized). Repeat this until q_i vanishes, or the leading term of q_i does not match those of $q_1, \dots, q_{i-1}$.
5. Repeat steps 1–4 until a solution is found, or no changes occur in the system.
6. If the system is unchanged, and any of the remaining polynomials factors over $\mathbf{Q}$, say $q_i = \ell_1 \ell_2$, where ℓ_1 and ℓ_2 are linear, replace q_i with ℓ_1 and return to Step 1. If no solution is found, then replace q_i with ℓ_2 and return to Step 1.
7. Otherwise, if any of the remaining polynomials has the form $a_{ij}^2 - c$, where c is a positive rational, follow Step 6 with $\ell_1 = a_{ij} - c^{1/2}$ and $\ell_2 = a_{ij} + c^{1/2}$.
8. If neither of the conditions in Step 6 or Step 7 apply, then the reduction algorithm halts and fails.

It is not at all obvious that this algorithm will succeed in all cases; nevertheless, we were able to use it to construct hereditary orthogonal models for every totally free irreducible representation of every exceptional Weyl group. We also tested it successfully on every irreducible representation of every classical Weyl group of rank $\leqslant 8$.

For example, using this algorithm to determine the matrix A_8 in the representation R_{5670} of $W(\mathcal{E}_8)$ takes about 16 minutes on a 2.8GHz Pentium IV running Maple 9.

Remark 3.2. Although the Coxeter relations involving A_n have only finitely many solutions (Corollary 1.5), it is important that the linear equations that eliminate clones are part of the initial system. Otherwise, there would in general be exponentially many solutions, of which a large fraction would have to be discarded. Instead, we have a system of equations for which any solution suffices, and hence the second branches in Step 6 or 7 will often be unnecessary.

E. The seminormal algorithm. Once we have orthogonal matrices $A_1, \dots, A_n$ representing the simple reflections relative to some hereditary basis $\mathcal{B}$, we may

[2]At all times we assume that each polynomial q_i is normalized to be monic.

consider the problem of finding a diagonal transformation D so that the matrices DA_iD^{-1} are rational, seminormal, and hereditary. Note that Propositions 1.1 and 1.2 show that such a change of basis necessarily exists; however, there is no obvious canonical choice for D, and poor choices will lead to matrix entries with large numerators and denominators.

If we partition $\mathcal{B}$ into blocks $\mathcal{B}_1, \ldots, \mathcal{B}_l$ that span irreducible W_{n-1}-modules, then we may (recursively) assume that diagonal transformations D_i for each block $\mathcal{B}_i$ have been identified that produce rational hereditary matrix models for the action of W_{n-1}. Moreover, we may assume that these models are optimal with respect to some measure to be specified later. The remaining problem is to identify scalars $x = (x_1, \ldots, x_l)$ so that the change of basis $D(x) = x_1D_1 \oplus \cdots \oplus x_lD_l$ converts the matrices $A_1, \ldots, A_n$ to an optimal rational form. Since the choices for x have no effect on A_i for $i < n$, this amounts to rationalizing and optimizing $D(x)A_nD(x)^{-1}$.

It should be noted that there will necessarily be rationalizing choices for x, regardless of the previous choices made for $D_1, \ldots, D_l$. Indeed, there must be some diagonal transformation $D' = D'_1 \oplus \cdots \oplus D'_l$ that converts $\mathcal{B}$ to a $\mathbf{Q}$-basis, so D'_i and D_i are both rescalings that yield $\mathbf{Q}$-bases for the irreducible W_{n-1}-module spanned by $\mathcal{B}_i$, and hence for each i there is a scalar x_i such that $x_iD_i(D'_i)^{-1}$ is rational. In other words, there is an x such that $D(x)$ is a rational diagonal multiple of D', and the space of rationalizing choices for x forms a single $(\mathbf{Q}^*)^l$-orbit.

To simplify the set of constraints, note that the (nonzero) entries of $\phi_r(A_n) = [a_{ij}]$ are the same as those of A_n, except that they occur with lower multiplicity. Moreover, the effect of the diagonal change of basis $D(x)$ on $\phi_r(A_n)$ is to replace a_{ij} with $a'_{ij}x_{b(i)}x_{b(j)}^{-1}$, where $b(i)$ denotes the index of the W_{n-1}-block that contains the i-th W_r-block, and $[a'_{ij}] = \phi_r(D(1)A_nD(1)^{-1})$.

To identify a rationalizing x, we treat $x_1, \ldots, x_l$ as variables and proceed as follows. Selecting any nonzero entry of $\phi_r(D(x)A_nD(x)^{-1})$ from rows and columns belonging to distinct W_{n-1}-blocks, say $a'_{ij}x_{b(i)}x_{b(j)}^{-1}$, one sees that it is necessary for $x_{b(j)}$ to be a rational multiple of $a'_{ij}x_{b(i)}$; conversely, since the solution space is $(\mathbf{Q}^*)^l$-stable, any rational multiple will suffice. We may therefore substitute $x_{b(j)} = |a'_{ij}|x_{b(i)}$ and eliminate all occurrences of the variable $x_{b(j)}$. The effect of this substitution will be to rationalize some entries of $\phi_r(D(x)A_nD(x)^{-1})$; these may be ignored henceforth. We continue by choosing another nonzero entry that depends on two remaining variables, and eliminate one of them in the same way, and so on. Since the support graph of nonzero entries in $\phi_r(A_n)$ is necessarily connected (given that V is irreducible), this process ends when only one variable remains. This last variable may be specialized arbitrarily; it has no effect on the matrix entries.

We remark that absolute values are used in the above substitutions so that we are able to produce a positive solution for x, and hence by induction, a positive rescaling of $\mathcal{B}$.

F. Optimization. Having identified a (positive) rationalizing choice for x, say x_0, we now describe a method for finding a point $y \in (\mathbf{Q}^+)^l$ that minimizes the least common denominator of the off-diagonal entries in $A_n(y) := D(yx_0)A_nD(yx_0)^{-1}$, or equivalently in $\phi_r(A_n(y))$.

In practice, the matrix entries of $\phi_r(A_n(1)) = [a''_{ij}]$ will have denominators involving relatively few primes. In most cases, the only primes are those that divide

$|W|$ (see also Remark 3.5(b) below). It therefore suffices to solve the following localized version of the denominator-minimization problem for each of these primes.

Problem 3.3. Given a prime p, find $v \in \mathbf{Z}^l$ so that $y = (p^{v_1}, \ldots, p^{v_l})$ maximizes the lowest exponent of p in $A_n(y)$. More precisely, find $v \in \mathbf{Z}^l$ so that the objective function

$$\min_{i,j}(e_{ij} + v_{b(i)} - v_{b(j)})$$

is maximized, where e_{ij} denotes the exponent for the power of p involved in a''_{ij}, and the minimum is taken only over those pairs i, j such that $i \neq j$ and $a''_{ij} \neq 0$.

Note that the above optimization problem is necessarily bounded; indeed, since A_n is symmetric, the nonzero entries of $\phi_r(A_n(1))$ are symmetrically placed. It follows that

$$\min(e_{ij} + v_{b(i)} - v_{b(j)},\, e_{ji} + v_{b(j)} - v_{b(i)}) \leqslant (e_{ij} + e_{ji})/2,$$

whence

$$\min_{i,j}(e_{ij} + v_{b(i)} - v_{b(j)}) \leqslant \min_{i,j}\lfloor (e_{ij} + e_{ji})/2 \rfloor. \tag{3.2}$$

It should also be noted that this upper bound is intrinsic to V; it does not depend on the initial rationalizing choice x_0. If some diagonal change of basis attains this upper bound, we say that it is *strongly p-optimal*.

Although Problem 3.3 appears to involve integer optimization with a nonlinear objective function, the following result allows us to reduce it to linear programming.

Proposition 3.4. *Given a loopless, symmetric (i.e., $(i,j) \in \Gamma \Rightarrow (j,i) \in \Gamma$), connected digraph $\Gamma \subset [l] \times [l]$ and an integer c_{ij} for each $(i,j) \in \Gamma$, let*

$$P(t) := \{v \in \mathbf{Q}^l : \ v_1 = 0 \text{ and } c_{ij} + v_i - v_j \geqslant t \text{ for all } (i,j) \in \Gamma\}.$$

If t_0 is the maximum value for t in the polyhedron $Q = \{(v,t) \in \mathbf{Q}^{l+1} : v \in P(t)\}$, then

$$\lfloor t_0 \rfloor = \max_{v \in \mathbf{Z}^l} \min_{(i,j) \in \Gamma} c_{ij} + v_i - v_j, \tag{3.3}$$

and every vertex of the (nonempty) polytope $P(\lfloor t_0 \rfloor)$ is a lattice point v that achieves the above maximum.

Proof. By following paths from 1 to i and i to 1 in Γ, it is easy to derive upper and lower bounds for v_i proving that $P(t)$ is bounded for all t.

The matrix whose rows are the linear forms defining the polytope $P(t)$ is a $0, \pm 1$-matrix with at most one 1 and one -1 per row, and hence each of its invertible submatrices is invertible over the integers. (Indeed, graphic matroids are totally unimodular.) It follows that for each integer t, every vertex of $P(t)$ is a lattice point.

Now since the linear forms $c_{ij} + v_i - v_j$ are invariant under translation by $v = (1, \ldots, 1)$, it follows that their value ranges are unaffected by dropping the constraint $v_1 = 0$. Hence

$$t_0 = \max_{v \in \mathbf{Q}^l} \min_{(i,j) \in \Gamma} c_{ij} + v_i - v_j,$$

and this maximum is finite, by the same reasoning used to establish (3.2). It follows that $\lfloor t_0 \rfloor$ is an upper bound for the maximum in (3.3); in particular, $P(\lfloor t_0 \rfloor)$ is nonempty, and hence any vertex of this polytope provides a lattice point where this upper bound is attained. □

To solve Problem 3.3, we proceed by defining c_{ij} to be the lowest power of p appearing among the nonzero matrix entries of $\phi_r(A(1))$ in the rows and columns belonging to W_r-blocks in $\mathcal{B}_i$ and $\mathcal{B}_j$ (respectively) for all $i \neq j$. If there are no nonzero entries, we leave c_{ij} undefined. Note that the underlying support graph Γ necessarily fits the hypothesis of Proposition 3.4; in particular, connectedness follows from the irreducibility of V.

We may thus proceed by using linear programming methods, such as the simplex algorithm, to determine the maximum value for t in the polyhedron Q. Once the maximum t_0 is obtained, we then use a second call to a linear program solver to find a vertex v of the polytope $P(\lfloor t_0 \rfloor)$. If t_0 happens to be an integer, then the second linear program may be omitted; in that case, any extreme point $(v, t_0) \in Q$ produced by the first linear program will provide an integer solution v for Problem 3.3.

Once we have found $v \in \mathbf{Z}^l$ that optimizes the lowest exponent of p in the off-diagonal entries of $A_n(y)$, a secondary constraint we may impose is that among all optimizing v, we should minimize the highest power of p appearing among the same matrix entries.

To describe this secondary optimization problem more explicitly, let b_{ij} denote the highest power of p appearing among the nonzero matrix entries of $\phi_r(A(1))$ in the rows and columns belonging to W_r-blocks in $\mathcal{B}_i$ and $\mathcal{B}_j$ (respectively) for all $(i, j) \in \Gamma$. We then seek to minimize the objective function

$$\max_{(i,j)\in\Gamma} b_{ij} + v_i - v_j$$

over all lattice points $v \in P(m)$, where $m = \lfloor t_0 \rfloor$ denotes the optimal power of p in (3.3).

This optimization problem may also be solved by linear programming methods. Indeed, having first determined m as described earlier, let t_1 denote the minimum value for t among all $v \in \mathbf{Q}^l$ and $t \in \mathbf{Q}$ such that

$$t - b_{ij} \geqslant v_i - v_j \geqslant m - c_{ij} \tag{3.4}$$

for all $(i, j) \in \Gamma$. Of course, we may easily determine t_1 via linear programming. It follows by reasoning similar to Proposition 3.4 that

$$\lceil t_1 \rceil = \min_{v\in\mathbf{Z}^l\cap P(m)} \max_{(i,j)\in\Gamma} b_{ij} + v_i - v_j,$$

and if we add the constraints $t = \lceil t_1 \rceil$ and $v_1 = 0$ to the polyhedron defined by (3.4), then any vertex v of the resulting polytope will necessarily be a lattice point, and hence a solution of our primary and secondary optimization problems.

For example, the representation R_{5670} of $W(\mathcal{E}_8)$ decomposes into a sum of $l = 16$ irreducible $W(\mathcal{E}_7)$-modules, and in its orthogonal hereditary model, the least common denominator of the squares of the off-diagonal matrix entries in A_8 is $(2^5 3^4 5^1 7^1)^2$. It follows that the strong p-bounds for the rational rescalings of A_8 (see (3.2)) are $-5, -4, -1, -1$ for the primes $p = 2, 3, 5, 7$ (respectively). The initial rationalizing choice for x found by our algorithm yielded a rational seminormal model for R_{5670} with an off-diagonal least common denominator of $2^6 3^8 5^3 7^2$. We then used Maple's *simplex* package to solve the linear programs needed to produce an optimal diagonal rescaling (in both the primary and secondary sense) with respect to the primes $2, 3, 5, 7$; this took about 12 seconds on a 2.8GHz Pentium IV, and yielded a model that matched the strong bounds for each prime.

Remark 3.5. (a) Most of the rational seminormal models for $W(\mathcal{E}_8)$-representations we have produced are strongly optimal with respect to each prime (i.e., the bounds in (3.2) are equalities). Even the cases that fail to be strongly optimal are typically off by a single prime factor. In fact, we suspect that strongly optimal models exist in all cases, but it may be difficult to confirm this—there are generally many optimal solutions to choose from, and the choices made when optimizing the models for the subgroups W_k for $k < n$ have an effect on the optimum values that can be achieved for $k = n$.

(b) For each of the exceptional groups, the only primes that occur in the denominators of the optimized rational seminormal models that we produced, even in the non-free cases discussed below, are those that divide $|W|$.

4. Models for non-free representations

We now turn to the problem of constructing explicit matrices representing the simple reflections in an irreducible W-module V that is not totally free. The main complication in this situation is that orthonormal hereditary bases for V are not canonical up to sign. Indeed, if we proceed recursively and fix hereditary orthogonal matrices $A_1, \dots, A_{n-1}$ representing the action of W_{n-1} on V, then the collection of all orthonormal W-hereditary bases for V that are compatible with this choice forms a single orbit under the group of isometric W_{n-1}-module automorphisms of V, which by Schur's Lemma is isomorphic to $O(m_1, \mathbf{R}) \times \cdots \times O(m_l, \mathbf{R})$, where $m_1, \dots, m_l$ denote the multiplicities of the irreducible W_{n-1}-modules in V. This is a discrete orbit only if V is multiplicity-free.

The complexity of this orbit space has several negative consequences. First, the variety of solutions for the matrix A_n representing s_n, as defined by the Coxeter relations and the clone equations (see Section 3C), is no longer 0-dimensional. In particular, this means that the Gröbner-like reduction algorithm of Section 3D cannot be used without modifications. Second, even if we manage to find a solution, there is no guarantee that it will be convertible to a rational seminormal solution by means of a diagonal transformation. Third, even if we find a solution that is convertible to rational form, the lack of a canonical solution means that it is likely to have poor quality (i.e., the matrix entries are likely to have large numerators and denominators).

On the other hand, our primary goal is not to construct hereditary bases with respect to *every* ordering of the simple reflections; rather, we are seeking (optimal) hereditary bases for the irreducible representations of W with respect to the standard order. Thus we are practically concerned only with the nine remaining representations of $W = W(\mathcal{E}_8)$ that are non-free (see Empirical Fact 2.2). This allows us to make several simplifying assumptions (see Remark 2.3(b)), the most important of which are

(i) V is totally free as a W_{n-1}-module, and
(ii) each irreducible W_{n-1}-module has multiplicity $\leqslant 2$ (and usually $\leqslant 1$) in V.

Under these circumstances, we shall see that it is possible to make small adjustments to the algorithms of Section 3 and still produce suitably optimal orthogonal and rational seminormal hereditary models for V.

A. Orthogonal models of rational type. We say that an orthonormal hereditary basis $\mathcal{B}$ for V (or equivalently, the corresponding matrix model) is of *rational type* if some diagonal transformation of $\mathcal{B}$ is a (necessarily seminormal) hereditary $\mathbf{Q}$-basis.

Following the approach of Section 3, we may recursively assume that orthogonal matrix models for the irreducible W_{n-1}-submodules of V have been previously constructed. Using branching data derived from the character tables of W and W_{n-1}, we may thus form direct sums of these models so as to obtain orthogonal matrices A_k representing the action of s_k on V for $1 \leqslant k < n$; the problem is to identify one or more possible orthogonal matrices A_n representing the action of s_n. As in the totally free case, A_n may be recovered from the matrix $\phi_r(A_n)$ that records the action of s_n on W_r-blocks, where r denotes the largest index such that s_n centralizes W_r.

Proposition 4.1. *Assume V is totally free as a W_{n-1}-module, and that there is one irreducible W_{n-1}-module of multiplicity 2 in V, at least one of multiplicity 1, and none of multiplicity > 2. Given orthogonal matrices $A_1, \ldots, A_{n-1}$ as above, there exist pairs i, j such that there is a unique solution for A_n (up to choices of sign) in which the i, j-entry of $\phi_r(A_n)$ vanishes. Furthermore, this solution is necessarily of rational type.*

Proof. By Proposition 1.1 and Theorem 2.1, we know that V has a hereditary $\mathbf{Q}$-basis $\mathcal{B}$ that is orthogonal with respect to a positive definite W-invariant inner product $\langle\,,\rangle$ that is $\mathbf{Q}$-valued on the $\mathbf{Q}$-span of $\mathcal{B}$. By replacing each $v \in \mathcal{B}$ with $\bar{v} := v/\langle v, v\rangle^{1/2}$, we thereby obtain an orthonormal hereditary basis $\bar{\mathcal{B}}$ of rational type. Furthermore, since V is totally free as a W_{n-1}-module, we may apply sign changes to $\bar{\mathcal{B}}$ (if necessary) so that A_k is the matrix of s_k with respect to this basis, for $k = 1, \ldots, n-1$ (Corollary 1.3).

Now let $\mathcal{B}_1 = \{u_i : i \in I\}$ and $\mathcal{B}_2 = \{v_i : i \in I\}$ denote the two blocks of $\mathcal{B}$ that span copies of the same irreducible W_{n-1}-module in V, indexed so that $u_i \mapsto v_i$ extends to an isomorphism. By Schur's Lemma, the group G of isometric W_{n-1}-module automorphisms of V consists of sign changes and a copy of $SO(2, \mathbf{R})$ that intertwines $\bar{\mathcal{B}}_1$ and $\bar{\mathcal{B}}_2$. Furthermore, any hereditary orthonormal basis for V that represents s_k by A_k for $k < n$ is in the G-orbit of $\bar{\mathcal{B}}$. Leaving aside sign changes, it follows that every possible matrix A_n representing s_n may be obtained from some (necessarily hereditary, orthonormal) basis for V generated from $\bar{\mathcal{B}}$ by selecting a point (a, b) on the circle $a^2 + b^2 = 1$ and replacing

$$\bar{\mathcal{B}}_1 \to \{a\bar{u}_i + b\bar{v}_i : i \in I\}, \quad \bar{\mathcal{B}}_2 \to \{-b\bar{u}_i + a\bar{v}_i : i \in I\}.$$

By hypothesis, $\mathcal{B}_1 \cup \mathcal{B}_2$ spans a proper subspace of the irreducible W-module V, so it cannot be the case that $\langle s_n\bar{u}_i, v\rangle = \langle s_n\bar{v}_i, v\rangle = 0$ for all $i \in I$ and $v \in \mathcal{B} - (\mathcal{B}_1 \cup \mathcal{B}_2)$. Hence, there must exist $i \in I$ and $v \in \mathcal{B} - (\mathcal{B}_1 \cup \mathcal{B}_2)$ such that $\langle as_n\bar{u}_i + bs_n\bar{v}_i, v\rangle$ is a nontrivial linear form in a and b. If we impose the condition that this linear form should vanish, then there will be a unique solution for (a, b) (up to a choice of sign), and hence a unique orthonormal hereditary basis $\mathcal{B}'$ (up to sign) such that the representing matrix for s_n relative to $\mathcal{B}'$ has a zero in the row and column corresponding to $\bar{v}$ and $a\bar{u}_i + b\bar{v}_i$.

To complete the proof, we show that $\mathcal{B}'$ is necessarily of rational type. Returning to the orthogonal $\mathbf{Q}$-basis $\mathcal{B}$, one knows that there is a unique invariant bilinear form (up to scalar multiples) for any irreducible W_{n-1}-module, so there must be a positive rational q such that $q = \langle v_i, v_i\rangle / \langle u_i, u_i\rangle$ for all $i \in I$. Furthermore, the vanishing condition for a and b may be rewritten in terms of the $\mathbf{Q}$-basis $\mathcal{B}$ as

$$a\sqrt{q}\langle s_n u_i, v\rangle + b\langle s_n v_i, v\rangle = 0,$$

whence $t := (a/b)\sqrt{q}$ is rational. (If $b = 0$, then $\bar{\mathcal{B}} = \mathcal{B}'$ and there is nothing further to prove.) Now consider the rational change of basis obtained by replacing

$$\mathcal{B}_1 \cup \mathcal{B}_2 \to \{tu_i + v_i : i \in I\} \cup \{-qu_i + tv_i : i \in I\}.$$

It is easy to see that this yields another hereditary $\mathbf{Q}$-basis for V. Moreover, it is not hard to check that it remains orthogonal, and that it normalizes to $\mathcal{B}'$ (up to sign), whence $\mathcal{B}'$ is of rational type. □

Remark 4.2. (a) In order to uniquely specify an orthogonal matrix model of rational type for V via the above argument, we need to identify a pair of isomorphic W_r-blocks, one in $\mathcal{B}_1 \cup \mathcal{B}_2$, and one not in $\mathcal{B}_1 \cup \mathcal{B}_2$, so that the corresponding entry of $\phi_r(A_n)$ does not vanish identically as we vary the point chosen from $SO(2, \mathbf{R})$. Although it requires *a posteriori* verification, it turns out that among the nine representations of $W(\mathcal{E}_8)$ that are not totally free with respect to the standard chain, it is usually the case that *all* of the matrix entries of this type do not vanish identically, and hence any such choice will suffice. The only exception is $R_{6075}(\pm 405)$; each member of this pair of representations has 102 eligible matrix entries in $\phi_6(A_8)$; of these, 96 are not identically zero.

(b) One may show more generally that if there are k irreducible W_{n-1}-modules that occur with multiplicity 2 in V, at least one of multiplicity 1, and none of multiplicity > 2, then it is possible to specify a unique orthonormal hereditary basis of rational type by forcing k entries of $\phi_r(A_n)$ to vanish. However, the sets of entries that suffice for this purpose are determined by the pattern of generically nonzero entries in $\phi_r(A_n)$, and hence difficult to predict *a priori.* In any case, the only instance of this problem with $k > 1$ that is of interest involves the $W(\mathcal{E}_8)$-representation R_{7168}, which has two $W(\mathcal{E}_7)$-constituents of multiplicity 2. And as noted in the previous remark, all of the relevant matrix entries of $\phi_6(A_8)$ turn out to be generically nonzero in this case.

For each of the nine non-free representations of $W(\mathcal{E}_8)$, we constructed an orthogonal hereditary matrix model as follows. First, we selected an arbitrary matrix entry of $\phi_r(A_n)$ that met the requirements described in Remark 4.2(a). (In the case of R_{7168}, we selected two such entries, one for each $W(\mathcal{E}_7)$-constituent that occurs with multiplicity 2.) We then combined the condition that these matrix entries should vanish with the Coxeter relations and clone equations of Section 3C, and passed the resulting system to the reduction algorithm of Section 3D. A more robust approach would include trapping for errors that would be generated if the chosen matrix entries vanished identically, but as noted above, this can happen only in two of the nine cases, and is unlikely even for these two.

It was not clear in advance that the reduction algorithm would necessarily succeed; nevertheless, in each case we did obtain a solution. For example, the equations defining the matrix for s_8 in R_{7168} include 14597 Coxeter relations in 593 variables, one clone equation (see Remark 2.8), and two vanishing matrix entries. It took the reduction algorithm about 1.25 hours to find a solution on a 2.8GHz Pentium IV running Maple 9.

B. Optimizing an orthogonal solution. Once we have produced an orthonormal hereditary basis for V of rational type, there is no reason to expect that the representing matrices corresponding to this particular basis will have good quality. Thus we consider the problem of making an optimal choice among all such bases of rational type.

For simplicity, we continue the hypotheses of Proposition 4.1; i.e., that V is totally free as a W_{n-1}-module, and multiplicity-free with the exception of one W_{n-1}-isotypic component of multiplicity two. Let us also assume that we have identified a particular orthonormal hereditary basis for V of rational type. The algorithm in Section 3E shows that it is easy to construct a diagonal change of basis D that converts an orthogonal hereditary matrix model to a rational seminormal hereditary model (given that one exists), so we may equivalently take an orthogonal hereditary $\mathbf{Q}$-basis $\mathcal{B}$ for V as given.

As in the previous subsection, let $\mathcal{B}_1 = \{u_i : i \in I\}$ and $\mathcal{B}_2 = \{v_i : i \in I\}$ denote the blocks of $\mathcal{B}$ that span copies of the one W_{n-1}-component of V that has multiplicity two, labeled so that $u_i \mapsto v_i$ extends to an isomorphism. As we noted in the proof of Proposition 4.1, the quantity $q = \langle v_i, v_i\rangle / \langle u_i, u_i\rangle \in \mathbf{Q}^+$ is necessarily independent of $i \in I$. Furthermore, we may easily compute q by recognizing that the matrix of the W-invariant form $\langle\, ,\rangle$ with respect to $\mathcal{B}$ is D^{-2}, where D denotes the transformation used to convert the original orthogonal matrix model to rational seminormal form.

Let $\bar{\mathcal{B}} = \{\bar{v} : v \in \mathcal{B}\}$ denote the orthonormal basis corresponding to $\mathcal{B}$.

Proposition 4.3. *Given V, $\mathcal{B}$, and q as above, every orthonormal hereditary basis for V of rational type may be obtained (up to a choice of sign) from the orthonormal basis $\bar{\mathcal{B}}$ by replacing $\bar{\mathcal{B}}_1 \cup \bar{\mathcal{B}}_2$ with*

$$\left\{ \frac{1}{\sqrt{1+qt^2}} \big(\bar{u}_i + t\sqrt{q}\,\bar{v}_i\big) : i \in I \right\} \bigcup \left\{ \frac{1}{\sqrt{1+qt^2}} \big(t\sqrt{q}\,\bar{u}_i - \bar{v}_i\big) : i \in I \right\} \tag{4.1}$$

for some rational $t \geqslant 0$.

Proof. From Schur's Lemma, the decompositions of V into irreducible $\mathbf{Q}W_{n-1}$-modules form a single orbit relative to a $GL(2, \mathbf{Q})$ action that intertwines the W_{n-1}-submodules spanned by $\mathcal{B}_1$ and $\mathcal{B}_2$. Furthermore, the hereditary $\mathbf{Q}$-bases for any irreducible, totally free W_{n-1}-module are diagonal transformations of each other (Proposition 1.2). Hence, a member of each diagonal equivalence class of hereditary $\mathbf{Q}$-bases for V may be generated from $\mathcal{B}$ by replacing $\mathcal{B}_1 \cup \mathcal{B}_2$ with

$$\{au_i + bv_i : i \in I\} \cup \{cu_i + dv_i : i \in I\}$$

for suitable $a, b, c, d \in \mathbf{Q}$ with $ad - bc \neq 0$. If we add the condition that the resulting basis should remain orthogonal; i.e., $ac + qbd = 0$, then we may rescale $cu_i + dv_i$ if necessary so that $(c, d) = (qb, -a)$.

If $a = 0$, then the resulting basis is in the same diagonal equivalence class as $\mathcal{B}$, so we may assume henceforth that $a \neq 0$. Replacing (a, b) with $(-a, -b)$ if necessary, we may further assume $a > 0$. Rescaling by the factor a and setting $t := b/a \in \mathbf{Q}$, we conclude that replacing $\mathcal{B}_1 \cup \mathcal{B}_2$ with

$$\{u_i + tv_i : i \in I\} \cup \{qtu_i - v_i : i \in I\}$$

allows one to generate at least one member from each equivalence class of orthogonal hereditary $\mathbf{Q}$-bases. It is also possible to restrict to the case $t \geqslant 0$ (if $t < 0$, replace $t \to -1/qt$ and rescale). By normalizing these bases to unit length, we thereby obtain all orthonormal hereditary bases of rational type (up to a choice of sign), and it is not hard to see that the normalizations of these bases coincide with (4.1). $\square$

Remark 4.4. The above result easily generalizes for W-modules that have several W_{n-1}-components of multiplicity two. In these cases, any two orthonormal hereditary bases of rational type are related by a sequence of base changes each in the form of (4.1), one for each doubleton component.

Once we have an initial orthogonal hereditary model for V of rational type, Proposition 4.3 allows us to search through the space of all such models by varying a nonnegative rational parameter t. (Or k such parameters, if k of the W_{n-1}-components of V have multiplicity two.) In order to identify an optimal value for these parameter(s), we focus on the diagonal entries of the matrix A_n representing s_n (or equivalently, $\phi_r(A_n)$); these (rational) entries will remain unchanged when the model is converted to rational seminormal form, and hence cannot be improved by any subsequent diagonal rescalings. An added advantage of this strategy is that the diagonal entries corresponding to each W_{n-1}-component of multiplicity two depend only on the parameter for that component, so the parameters may be optimized independently of each other.

With these considerations in mind, for each W_{n-1}-component of multiplicity two with associated parameter t, we use the least common denominator of the corresponding diagonal entries of $\phi_r(A_n)$ as our objective function when optimizing the choice of t.

In theory, finding a value for t that optimizes this objective function is a difficult number-theoretic problem. However in practice, we found that all of the instances of this problem that occur among the nine non-free representations of $W(\mathcal{E}_8)$ are amenable to a brute force search that finds "good" (but not provably optimal) solutions. More explicitly, our optimization algorithm proceeds by first making the change of variable $t \to (a/q)^{1/2}t$, where a denotes the unique square-free integer such that $(a/q)^{1/2}$ is rational; equivalently, this amounts to replacing q with a in (4.1). We then evaluate the objective function at $t = 0$ and at each rational $t > 0$ whose numerator and denominator sum to $2, 3, \ldots$, stopping at some predetermined maximum sum, such as 100 or 1000.

We applied the above algorithm to the initial orthogonal models for the nine non-free representations of $W(\mathcal{E}_8)$ produced by the methods described in the previous subsection, and obtained least common denominators of 320 for $R_{3240}(\pm 594)$, 80 for $R_{4536}(\pm 378)$, 72 for $R_{5600}(\pm 280)$, 224 for $R_{6075}(\pm 405)$, and 315 for both doubleton components of R_{7168}. The corresponding t-values used to produce these quasi-optimal denominators all involved rationals with single-digit numerators and denominators.

Finally, once a suitably optimal orthogonal hereditary matrix model of rational type has been identified, one may convert it to an optimal rational seminormal form via the algorithms of Sections 3E and 3F.

References

[BM] D. Barbasch and A. Moy, *A unitarity criterion for p-adic groups*, Invent. Math. **98** (1989), 19–37. MR 1010153 (90m:22038)

[B] D. Barbasch, *Unitary spherical spectrum for split classical groups*, preprint.

[Be] M. Benard, *On the Schur indices of characters of the exceptional Weyl groups*, Ann. of Math. **94** (1971), 89–107. MR 0297887 (45:6939)

[F1] J. S. Frame, *Orthogonal group matrices of hyperoctahedral groups*, Nagoya Math. J. **27** (1966), 585–590. MR 0197583 (33:5748)

[F2] J. S. Frame, *The classes and representations of the groups of 27 lines and 28 bitangents*, Ann. Mat. Pura Appl. **32** (1951), 83–119. MR 0047038 (13:817i)

[F3] J. S. Frame, *The characters of the Weyl group E_8*, Computational Problems in Abstract Algebra (Proc. Conf., Oxford, 1967), Pergamon, Oxford, 1970, pp. 111–130. MR 0269751 (42:4646)

[GP] M. Geck and G. Pfeiffer, *Characters of finite Coxeter groups and Iwahori-Hecke algebras*, Oxford Univ. Press, New York, 2000. MR 1778802 (2002k:20017)

[G] C. Greene, *A rational-function identity related to the Murnaghan-Nakayama formula for the characters of S_n*, J. Algebraic Combin. **1** (1992), 235–255. MR 0936289 (89k:94005)

[Gy] A. Gyoja, *On the existence of a W-graph for an irreducible representation of a Coxeter group*, J. Algebra **86** (1984), 422–438. MR R0732258 (85k:20144b)

[JK] G. James and A. Kerber, *The Representation Theory of the Symmetric Group*, Addison-Wesley, Reading, MA, 1981. MR 0644144 (83k:20003)

[KL1] D. Kazhdan and G. Lusztig, *Representations of Coxeter groups and Hecke algebras*, Invent. Math. **53** (1979), 165–184. MR 0560412 (81j:20066)

[KL2] D. Kazhdan and G. Lusztig, *A topological approach to Springer's representations*, Adv. in Math. **38** (1980), 222–228. MR 597198 (82f:20076)

[K] T. Kondo, *The characters of the Weyl group of type F_4*, J. Fac. Sci. Univ. Tokyo Sect. I **11** (1965), 145–153. MR 0185018 (32:2488)

[R] A. Ram, *Seminormal representations of Weyl groups and Iwahori-Hecke algebras*, Proc. London Math. Soc. (3) **75** (1997), 99–133. MR 1444315 (98d:20007)

[Ru] D. E. Rutherford, *Substitutional Analysis*, University Press, Edinburgh, 1948. MR 0027272 (10:280i)

[Sp] T. A. Springer, *A construction of representations of Weyl groups*, Invent. Math. **44** (1978), 279–293. MR 0491988 (58:11154)

[S1] J. R. Stembridge, *On the eigenvalues of representations of reflection groups and wreath products*, Pacific J. Math. **140** (1989), 353–396. MR 1023791 (91a:20022)

[S2] J. R. Stembridge, *A Maple package for root systems and finite Coxeter groups*, available electronically at `www.math.lsa.umich.edu/~jrs/maple.html`.

[OV] A. Okounkov and A. Vershik, *A new approach to representation theory of symmetric groups*, Selecta Math.(N.S.) **2** (1996), 581–605. MR 1443185 (99g:20024)

[Y] A. Young, *The collected papers of Alfred Young*, University of Toronto Press, Toronto, 1977. 0439548 (55:12438)

Department of Mathematics, University of Michigan, Ann Arbor, Michigan 48109–1109

E-mail address: `jrs@umich.edu`