Applying unitarity tests in the Weyl group
Atlas of Lie Groups AIM Workshop II

John R. Stembridge ⟨jrs@umich.edu⟩

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

12–16 July 2004

Topics

I. Overview of the computation

Reviewing the setup from our lecture notes of last year [**S1**], we let R denote a crystallographic root system with V the ambient real Euclidean space, along with the usual choices (simple roots, positive roots,...), and W the Weyl group.

Given a reduced expression for the longest element of the Weyl group, say $w_0 = s_{i_1} \cdots s_{i_l}$, there is an induced ordering $\beta_1^\vee, \dots, \beta_l^\vee$ of the positive co-roots; namely,

$$\beta_l^\vee = \alpha_{i_l}^\vee,\ \beta_{l-1}^\vee = s_l \alpha_{i_{l-1}}^\vee,\ \dots,\ \beta_1^\vee = s_l \cdots s_2 \alpha_{i_1}^\vee.$$

Now for each $\nu \in V$, let $A(\nu)$ denote the following element of the group algebra $\mathbb{R}W$:

$$A(\nu) := (1 + \langle \nu, \beta_1^\vee \rangle s_{i_1}) \cdots (1 + \langle \nu, \beta_l^\vee \rangle s_{i_l}).$$

Elementary Facts 1–3.

(1) *$A(\nu)$ is independent of the choice of reduced expression.*

(2) *If $w_0\nu = -\nu$, then $A(\nu)$ is Hermitian in every unitary representation of W.*

(3) *$A(\nu)$ is invertible if and only if $\langle \nu, \beta^\vee \rangle \neq \pm 1$ for all roots β.*

Let V_0 denote the subspace of V fixed by $-w_0$.

There is a spherical representation of a split p-adic group corresponding to each dominant $\nu \in V_0$, and thanks to work of Barbasch-Moy, we know that it is unitary if and only if $A(\nu)$ is positive semidefinite (psd) as a (Hermitian) operator on $\mathbb{R}W$.

Similarly, we know that unitarity for spherical representations of real groups requires $\sigma(A(\nu))$ to be positive semidefinite for *some* irreducible representations σ of W. Thus,

Main Problem. *For every unitary irrep $\sigma \in \widehat{W}$, and all $\nu \in V_0^+$, determine when $\sigma(A(\nu))$ is psd. More ambitiously, compute the signature of every $\sigma(A(\nu))$ as well.*

It is natural to carve up V into cells according to where a given point ν may sit relative to each hyperplane $\langle \nu, \beta^\vee \rangle = \pm 1$ (i.e., where the singularities occur in each factor of $A(\nu)$). Given that we are interested only in the case of dominant ν, we will discard the hyperplanes $\langle \nu, \beta^\vee \rangle = -1$ (for $\beta^\vee$ positive). With this in mind, we define a *cell* to be any of the (non-empty) sets obtained by selecting one of

$$\{\nu \in V : \langle \nu, \beta^\vee \rangle > 1\},\ \{\nu \in V : \langle \nu, \beta^\vee \rangle = 1\},\ \{\nu \in V : \langle \nu, \beta^\vee \rangle < 1\}$$

for each positive co-root $\beta^\vee$, and taking the intersection. Note that each cell is polyhedral and open relative to its affine span.

It is important to note that cells will often include points that are not in the dominant chamber; there even exist cells that are entirely disjoint from the dominant chamber. A second complication is that we are only interested in points where $w_0\nu = -\nu$. We say that a cell C is a *dominant cell* if $C \cap V_0^+$ is non-empty.

A critical ingredient we need to make the Main Problem feasible is

Non-Elementary Fact 4. *For each unitary W-rep σ, the signature of $\sigma(A(\nu))$ is constant within each (dominant) cell.*

This follows from the non-elementary fact, told to me by Dan Barbasch, that the rank of $A(\nu)$ is constant within each dominant cell.

II. New developments

In last year's lecture, we identified four subproblems needed to solve the Main Problem:

Subproblem A. *Generate explicit matrices for each simple reflection, as well as a matrix for the invariant form, for each desired irrep σ of W.*

This is now completely solved—there is a paper discussing the solution in [**S2**], as well as files and documentation available online at

⟨`http://atlas.math.umd.edu/unitarity/weyl/hereditary`⟩

These models have rational matrix entries, are sparse, and have a W-invariant bilinear form that is diagonal.

Note that $W(E_7)$ has 60 irreps, the largest of which have degree 512, and $W(E_8)$ has 112 irreps, the largest of which have degree 7168.

It would be nice to have $\mathbb{Z}$-models for these representations, but not if it means sacrificing the sparsity of the models. Also, the rationality of the parameter ν might destroy any advantage of having an integral model.

Subproblem B. *Devise an efficient way to visit each dominant cell C and select a representative point ν from C.*

This problem is essentially solved. The cells have nice explicit descriptions (see the discussion in [**S1**]), and we have fast, efficient code for visiting each one. To select the point ν, we currently use linear programming; alternatively, each cell C contains a unique face $F = F(C)$ in the cells of the associated affine Weyl group that is minimal in the Bruhat order and of the same dimension. Selecting the barycenter of F provides a canonical way to choose a representative from the cell C. Given the tendency for messy coordinates in ν to create monstrous matrices for $\sigma(A(\nu))$, having a canonical choice may be preferable to linear programming. (But this is just a guess.)

To provide some perspective on the scale of this problem, note that in E_7, there are $113,100$ dominant cells. Using our current Maple code on an Athlon MP 2800+, it takes about 7.5 minutes to visit each dominant cell, and about 1 hour to both visit and select a point ν (via linear programming methods) from these cells.

Similarly, E_8 has $1,070,716$ dominant cells, it takes about 3 hours to visit each one, and about 1 day (estimate) to both visit and select a representative point from each cell.

Subproblem C. *Evaluate the matrix $\sigma(A(\nu))$.*

This is theoretically easy, and thanks to the code we designed to take advantage of the sparsity of the matrices for σ, also practical. In effect, the matrix for $\sigma(A(\nu))$ is computed one row at a time. Note that the average number of nonzero entries in each row of the matrix $\sigma(1 + \langle \nu, \beta^\vee \rangle s_i)$ is never more than 10.5, and usually much less.

On the other hand, note that if we have 1GB of RAM available, then the average number of bits we can afford to allocate to each entry of a (dense) matrix of size $(7168)^2$ is about 20.

Subproblem D. *Test whether $\sigma(A(\nu))$ is psd, and (possibly) compute the signature.*

Note that the standard algorithms for testing positive semi-definiteness of an $N \times N$ matrix are based on the LU-factorization, and are roughly equivalent to computing one determinant of degree N. In particular, they use about $O(N^3)$ scalar operations.

This is the main computational bottleneck. A year ago, this looked barely feasible for E_7 on a commodity PC, and for E_8 perhaps only on some high-end hardware or a cluster.

However, the main news in this department is that we have some new tricks that greatly accelerate the computation of the matrices $\sigma(A(\nu))$, and allow us to reduce testing the positivity of $\sigma(A(\nu))$ to a space whose dimension is (on average) roughly $1/4$ of the degree of σ. In particular, this made it easy for us to determine all of the psd cells for E_7 (i.e., the spherical unitary dual for p-adic E_7), and now we are optimistic that the same techniques may make E_8 solvable on ordinary PC hardware.

The tricks will be discussed in more detail in Section IV.

III. Computing the SpUD for p-adic E_7 and (perhaps) E_8

In this section, we describe in more detail the algorithm we used to determine the spherical unitary dual (SpUD) for p-adic E_7; i.e., finding all dominant cells C such that $\sigma(A(\nu))$ is psd for all $\nu \in C$ and all irreps σ of $W(E_7)$.

We intend to apply this method to E_8 as well.

Since we are now concerned only with these two Weyl groups, we will add the simplifying hypothesis that $w_0 = -1$.

A. Adaptive search.

In the first phase of the algorithm, we make a careful selection of a subset of the irreps of W, say $\sigma_1, \ldots, \sigma_r$, and determine a set of representative points $\nu_1, \ldots, \nu_k$ from the dominant cells for which $\sigma_i(A(\nu))$ is psd for all i. Since positive semidefiniteness is a closed condition (and expensive to test), it is enough to produce representative points ν_j from cells that are *maximal* with respect to the property of being psd. (The cell containing ν_j is "maximal" if it is not in the closure of any of the cells containing the other ν_i's.)

Temporarily avoiding the question of which representations σ_i to choose, the list of points ν_j is generated by a search that is iterated r times, once for each irrep σ_i. At all stages of the search, two lists are maintained:

- The *GoodList* contains representative points ν_j such that $\sigma(A(\nu_j))$ is psd in all irreps σ tested up to this point in the computation.
- The *BadList* contains representative points ν_j such that $\sigma(A(\nu_j))$ is not psd in some irrep σ.

Furthermore, since we prefer to avoid maintaining lists of potentially 1 million representative points (as would happen for E_8), both lists should be minimal in the sense that no point in the *GoodList* should be in the closure of the cell containing some other point in the *GoodList*, and the same should be true of the *BadList*.

Starting with both lists initially empty, we visit every dominant cell C (see Subproblem B) and test the irrep σ_1. We use the data describing C (hyperplanes and half spaces containing the cell) to decide if C is in the closure of a cell in the *GoodList*, or if some cell in the *BadList* is in the closure of C. If either of these conditions hold, then we skip C. Otherwise, we generate a representative point ν from C, compute the matrix $\sigma_1(A(\nu))$, and test whether it is psd. If yes, then we add ν and C to the *GoodList* and delete any existing members of *GoodList* that are now redundant. If no, then we add ν and C to the *BadList* and delete any existing members of *BadList* that are now redundant.

Once the processing of σ_1 has finished, we reset *GoodList* to empty (but not *BadList*) and process $\sigma_2, \sigma_3, \dots$ in the same way. Once the results appear to stabilize, (i.e., the *GoodList* and *BadList* are unchanged after processing an irrep σ), we have a "probable SpUD"; i.e., the *GoodList* contains a list of representative points that may possibly be psd in *all* irreps of W. In any case, all such cells must be in the closure of the *GoodList* cells.

B. Ranking the irreps of W.

Now consider the problem of deciding the order in which to test the irreps of W. Note that the operator $A(\nu)$ is a product of factors of the form $1 + \langle \nu, \beta^\vee \rangle s_i$ for various positive co-roots $\beta^\vee$. Most cells are on one or more of the hyperplanes $\langle \nu, \beta^\vee \rangle = 1$, and these give rise to singular factors $1 + s_i$ of rank $(N + t)/2$, where $N = \mathrm{tr}(\sigma(1))$ and $t = \mathrm{tr}(\sigma(s_1))$ denote the dimension and trace of a reflection in the chosen irrep σ. The lower the rank of the operator, the harder it is to find negative parts of the associated quadratic form, so (heuristically) it is better to use irreps in which the eigenvalues of a reflection have a relatively high percentage of $+1$'s compared to -1's.

With this in mind, we define the *value* of the irrep σ to be $(N+t)/2N$, a quantity that achieves a maximum of 1 for the trivial representation and a minimum of 0 for the sign representation. We test the (non-trivial) irreps of W in order of decreasing value.

For reference, we list the irreps of W of highest value for E_6 E_7 and E_8 in Table 1. We label a given irrep σ by the triple $[N, t, \varepsilon]$, where ε denotes the sign of $\mathrm{tr}(\sigma(w_0))$; these triples uniquely distinguish the irreps of each Weyl group of type E up to isomorphism.

In the cases of E_6 and E_7, the irreps marked by (*)'s indicate a minimal set that suffice to determine the cells that are psd in all irreps. In the case of E_7, this list is absolutely

E_6		E_7		E_8	
Irrep	Value	Irrep	Value	Irrep	Value
$[1, 1, +]$	1.000	$[1, 1, +]$	1.000	$[1, 1, +]$	1.000
$[6, 4, -]$	$0.833(*)$	$[7, 5, -]$	0.857	$[8, 6, -]$	0.875
$[20, 10, +]$	$0.750(*)$	$[27, 15, +]$	$0.778(*)$	$[35, 21, +]$	0.800
$[15, 5, +]$	$0.667(*)$	$[21, 11, -]$	$0.762(*)$	$[28, 14, +]$	0.750
$[15, 5, -]$	0.667	$[21, 9, +]$	0.714	$[84, 42, +]$	0.750
$[30, 10, -]$	$0.667(*)$	$[35, 15, +]$	$0.714(*)$	$[112, 56, -]$	0.750
$[64, 16, 0]$	0.625	$[56, 24, -]$	$0.714(*)$	$[50, 20, +]$	0.700
		$[15, 5, -]$	0.667	$[160, 64, -]$	0.700
		$[105, 35, -]$	0.667	$[210, 84, +]$	0.700
		$[120, 40, +]$	0.667	$[560, 196, -]$	0.675
		$[189, 51, -]$	0.635	$[567, 189, +]$	0.667
		$[105, 25, +]$	$0.619(*)$	$[300, 90, +]$	0.650

TABLE 1: Irreps of W of high value.

minimal in the sense that for each irrep σ marked by a (*), there is a point ν such that $\sigma(A(\nu))$ is not psd, but $\sigma'(A(\nu))$ *is* psd in every other irrep σ'.

C. Verification.

Once a probable-SpUD has been determined, the algorithm enters the verification phase. We take the list of representative points $\nu_1, \ldots, \nu_k$ from the *GoodList*, and for every j and every irrep σ of W, test whether $\sigma(A(\nu_j))$ is psd. Of course if this fails, then we must re-run the adaptive search with the offending irrep σ.

Note that this hardest part of the problem, but the good news is that it is embarrassingly parallel—each matrix may be tested for psd-ness independently, and thus may easily be farmed out to separate CPUs.

However in the worst cases, these problems are still challenging. It might have taken at least a CPU-month to complete this phase of the computation for the 107 maximal psd cells of E_7. However, by taking advantage of the dimension-reducing trick below, we were able to carry it out in about one CPU-day.

IV. A dimension-reducing trick

Suppose $\nu_1, \ldots, \nu_k$ are representatives of the maximal cells that are suspected of being psd in every irrep of W. Thanks to work of J.-K. Yu from last year [**Y**], we know an *a priori* description of the big cells (i.e., co-dimension 0) of this type. These are the cells such that $A(\nu)$ is positive definite in every irrep of W. For example, in E_7 there are 8 such cells, and in E_8 there are 16. We may omit these cells from the verification test—which is fortunate, since we have no dimension-reducing tricks for these cases.

The remaining maximal cells (99 in the case of E_7) each lie on one or more of the hyperplanes $\langle \nu, \beta^\vee \rangle = 1$. In fact, it is surprising that each of the 99 maximal psd cells is of co-dimension at least 3, and thus lies on at least three of these hyperplanes. We expect that the psd cells for E_8 will have similar features.

For each hyperplane of this type, there is a corresponding factor of $A(\nu)$ of the form $1 + s_i$ for some i. The rank of this factor is $r = \operatorname{tr} \sigma((1 + s_i)/2)$ in each irrep σ. If we can find a cheap way to change coordinates, then this allows for the possibility of testing the psd-ness of $\sigma(A(\nu))$ in an r-dimensional space, which cuts the dimension of the computation by an average factor of 2 (as σ varies over the irreps of W).

Similarly, if s_i and s_j commute, and there are two *orthogonal* co-roots $\beta_1^\vee$ and $\beta_2^\vee$ such that $\langle \nu, \beta_1^\vee \rangle = \langle \nu, \beta_2^\vee \rangle = 1$, then one can show that there is a reduced expression for w_0 so that $A(\nu)$ has a pair of consecutive factors of the form $(1 + s_i)(1 + s_j)$. In such cases, the rank of $\sigma(A(\nu))$ can be at most

$$\operatorname{tr} \sigma\left(\frac{1 + s_i}{2} \cdot \frac{1 + s_j}{2}\right),$$

so this provides the potential for cutting the dimension by a factor of 4, on average.

Remarkably, each of the 99 non-maximal psd cells in E_7 have a pair of orthogonal co-roots with this property. The key question is thus whether we can find a cheap way to change coordinates so that these dimension-reductions can be exploited.

Turning to the matrix models we constructed in [**S2**], it should be noted at this point that the simple reflections for $W(E_7)$ and $W(E_8)$ are ordered so that s_1 and s_2 commute. Moreover, the models are "hereditary" in the sense that the coordinates may be partitioned into bases for irreps for each of the parabolic subgroups generated by $s_1, \ldots, s_i$, for all i. Since s_1 and s_2 commute, this means that the matrices representing s_1 and s_2 are always diagonal, and the matrices representing $1+s_1$ and $(1+s_1)(1+s_2)$ are also diagonal, with nonzero entries equal to 2 or 4 (respectively).

Now if the factors $1+s_1$ or $(1+s_1)(1+s_2)$ appear somewhere in the middle of an expression for $A(\nu)$, there is no reason to expect $\sigma(A(\nu))$ to be sparse. However, if we can arrange for these factors to appear first (or last), then the Hermitian-ness of $A(\nu)$ will force the matrix of $\sigma(A(\nu))$ to vanish outside of an $r \times r$ principal submatrix, where r denotes the rank of $\sigma(1+s_1)$ or $\sigma((1+s_1)(1+s_2))$.

The following is a straightforward calculation.

Fact. *If $w_0 = -1$ and $c = \langle \mu, \alpha_j^\vee \rangle \neq \pm 1$, then*

$$A(s_j\mu) = \frac{1}{1-c^2}(1-cs_j)A(\mu)(1-cs_j).$$

In particular, if $|c| < 1$, then $A(s_j\mu)$ and $A(\mu)$ have the same signature in every W-irrep, and if $|c| > 1$, then $A(s_j\mu)$ and $A(\mu)$ have opposite signatures.

The above calculation suggests that we define the *broken W-orbit* of a point ν in V to consist of the smallest set $BW(\nu)$ such that $\nu \in BW(\nu)$ and

$$s_j\mu \in BW(\nu) \text{ for all } \mu \in BW(\nu) \text{ and all } j \text{ such that } \langle \mu, \alpha_j^\vee \rangle \neq \pm 1.$$

For every point $\mu \in BW(\nu)$, either $A(\nu)$ and $A(\mu)$ have the same signature or they have opposite signatures in every irrep of W, and the alternative is easily predictable.

Proposition. *If some $\mu \in BW(\nu)$ satisfies $\langle \mu, \alpha_1^\vee \rangle = \langle \mu, \alpha_2^\vee \rangle = 1$, then in every hereditary representation σ of W, the matrix of $\sigma(A(\mu))$ vanishes outside of a principal submatrix of order r, where r denotes the rank of $\sigma((1+s_1)(1+s_2))$.*

Proof. Express $A(\mu)$ in terms of a reduced word of the form $w_0 = s_1 s_2 \cdots$. $\square$

For 97 of the 99 maximal psd cells in E_7 of positive co-dimension, we found a suitable μ satisfying the above proposition by trial and error. More specifically, for each orthogonal pair of co-roots $\beta_1^\vee$ and $\beta_2^\vee$ such that $\langle \nu, \beta_1^\vee \rangle = \langle \nu, \beta_2^\vee \rangle = 1$, we found $w \in W$ such that $w\alpha_i^\vee = \beta_i^\vee$, and tested to see whether the lexicographically first reduced expression for w proved that $w\nu \in BW(\nu)$. For each ν, we found at least one w with this property in 97 of the 99 cases.

In the remaining two cases, we did not systematically search $BW(\nu)$, but we did find $\mu \in BW(\nu)$ such $\langle \mu, \alpha_1^\vee \rangle = 1$ (thereby cutting the dimension by an average factor of 2). Moreover, in these two cases, the number of co-roots such that $\langle \nu, \beta^\vee \rangle = 1$ is exceptionally large, the ranks of $\sigma(A(\mu))$ small, and the computations easy.

References

[**S1**] J. Stembridge, Lecture notes for AIM/Atlas of Lie groups workshop, July 2003.

[**S2**] J. Stembridge, Explicit matrices for irreducible representations of Weyl groups, *Represent. Theory* **8** (2004), 267–289.

[**Y**] J.-K. Yu, Generic unitary spherical parameters, July 2003. Lecture notes from Atlas Workshop I.