Irreducible Circuits and Coxeter Arrangements

John R. Stembridge

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

25 May 2006
minor revisions, 16 January 2007

Contents

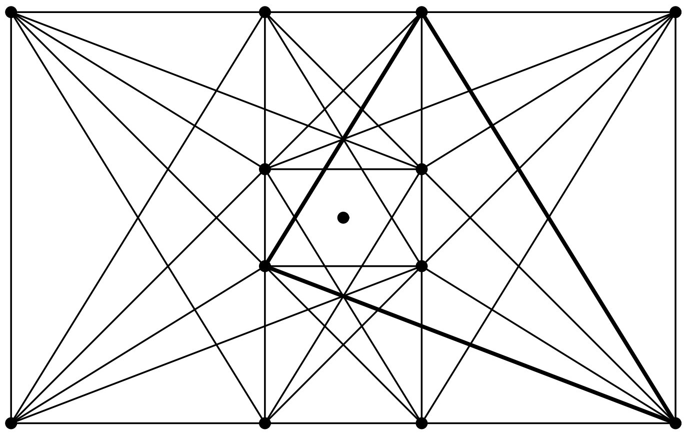

This work was supported by a grant from the NSA.

0. Introduction

Among the most basic combinatorial questions in the study of real hyperplane and vector arrangements are those involving positivity or feasibility: which vectors are in the nonnegative linear span of other vectors? Or the dual question: which sets of half-spaces (corresponding to hyperplane sides) have nonempty intersections?

Of course these questions are completely trivial, and can be considered "solved," either in the computational sense by linear programming, or in the conceptual sense by the theory of oriented matroids [**2**] [**4**].

However, we would like to point out that although the circuits of oriented matroids provide the means to answer such questions, there are many cases of interest in which they are an excessively redundant structure for this purpose. To give a rough idea of this redundancy, suppose that $\alpha_1, \alpha_2, \alpha_3$ are linearly independent members of a vector arrangement Φ. If $\alpha \in \Phi$ is in the strictly positive span of $\{\alpha_1, \alpha_2\}$, then $\{-\alpha, \alpha_1, \alpha_2\}$ is a positive circuit of Φ; i.e., a minimal subset of Φ with a positive dependence relation. If $\{-\beta, \alpha, \alpha_2, \alpha_3\}$ is another positive circuit of Φ, then it follows immediately that β is in the positive linear span of $\{\alpha_1, \alpha_2, \alpha_3\}$ and thus $\{-\beta, \alpha_1, \alpha_2, \alpha_3\}$ is also a positive circuit, but one whose existence is a purely formal consequence of other circuits.

In this paper, we introduce the concept of "irreducible" circuits for centrally symmetric vector arrangements (see Section 1 for the definition). These circuits have the property that they generate every (positive) circuit of the oriented matroid in the above sense. Furthermore, the irreducible circuits may be partitioned into equivalence classes so that every minimal set of generating circuits consists of a system of representatives from these classes (see Remark 1.3 below).

In the generic case, every circuit is irreducible. On the other hand, in a degenerate arrangement, irreducible circuits may be much rarer and simpler than general circuits. This is illustrated by our main result, an analysis of the irreducible circuits in the highly degenerate vector arrangements arising from Coxeter arrangements; i.e., root systems. What we find is that, up to symmetry, there are very few irreducible circuits in these cases, the main exception being $\mathcal{H}_4$. The proof is distributed over three separate sections, in part because the simply-laced case (Section 2) has significant additional features, including an interesting connection with spectral graph theory (see Remark 2.4 and Corollary 2.5).

Our motivation for this work came from attempts to understand in a uniform way the combinatorial structure of the regions of various subarrangements of Coxeter arrangements, as well as affine arrangements involving translates of root hyperplanes, such as the Shi arrangement, the Catalan arrangement, and others (e.g., see [**1**], [**5**], and [**6**]). However in this paper, we have restricted our attention to central hyperplanes for simplicity.

Finally, it is noteworthy that although irreducible circuits make sense in any oriented matroid, we do not know whether the basic results of Section 1 are valid in the setting of abstract (as opposed to realizable) oriented matroids.

We thank Lou Billera, Vic Reiner, and Tom Zaslavsky for helpful comments.

1. Irreducible circuits

Let V be a real Euclidean space with inner product $\langle \cdot, \cdot \rangle$, and $\Phi \subset V$ a finite subset such that for all $\alpha \in \Phi$, we have $c\alpha \in \Phi$ if and only if $c = \pm 1$. One should view Φ as a collection of normal vectors chosen from each side of a set of hyperplanes $\mathcal{H}$. A useful example to have in mind for the sequel is to take Φ to be a root system, in which case $\mathcal{H}$ is a Coxeter arrangement.

By exchanging vectors with their negatives (if necessary), every dependence relation involving pairwise independent members of Φ may be rewritten in the form

$$\alpha = c_1\alpha_1 + \cdots + c_n\alpha_n, \quad \text{where } c_i > 0 \text{ for } 1 \leqslant i \leqslant n. \tag{1.1}$$

We say that (1.1) is a *positively pointed* dependence relation with *base* $\{\alpha_1, \ldots, \alpha_n\}$ and *apex* α. If the base is linearly independent, then we call (1.1) a positively pointed *circuit* of *rank* n. Since the coefficients c_i are uniquely determined in this case, we may identify positively pointed circuits with apex-base pairs (α, Δ).

If (α, Δ) is a positively pointed circuit involving members of Φ, then $\Delta \cup \{-\alpha\}$ is a circuit of the matroid associated to Φ. However, there is extra information implicit in this set; namely, that the essentially unique dependence relation among its members has positive coefficients. Moreover, the oriented matroid associated to Φ may be viewed as the abstract incidence structure that records (a) the supports of all circuits with positive coefficients, and (b) the bijection that associates α and $-\alpha$ for each $\alpha \in \Phi$.

Although this is a non-traditional way to describe the oriented matroid of a centrally symmetric vector configuration, let us note that similar techniques *are* traditional in the subject (e.g., see the discussion of the "Lawrence construction" in Section 9.3 of [**2**]).

We define a positively pointed circuit in Φ with base Δ to be *irreducible* (with respect to Φ) if there is no positively pointed circuit in Φ whose base is a proper subset of Δ. In geometric terms, this means that no proper face of the simplicial cone generated by Δ contains any members of $\Phi - \Delta$. Note that every positively pointed circuit has rank at least 2, so every positively pointed circuit of rank exactly 2 is necessarily irreducible.

For example, in the root system $\Phi = \mathcal{A}_{\ell-1} = \{\varepsilon_i - \varepsilon_j : 1 \leqslant i \neq j \leqslant \ell\}$ (a set of normal vectors for the braid arrangement), the dependence relation

$$\varepsilon_1 - \varepsilon_6 = (\varepsilon_1 - \varepsilon_4) + (\varepsilon_4 - \varepsilon_2) + (\varepsilon_2 - \varepsilon_3) + (\varepsilon_3 - \varepsilon_6)$$

is a typical positively pointed circuit. However, it is not an irreducible circuit. Indeed, any set of consecutive terms of the base of this circuit is the base of another positively pointed circuit; e.g.,

$$\varepsilon_1 - \varepsilon_2 = (\varepsilon_1 - \varepsilon_4) + (\varepsilon_4 - \varepsilon_2).$$

Furthermore, it is not hard to show that all of the irreducible circuits in $\mathcal{A}_{\ell-1}$ are of the form $\varepsilon_i - \varepsilon_k = (\varepsilon_i - \varepsilon_j) + (\varepsilon_j - \varepsilon_k)$ for distinct i, j, k.

A subset $\Psi \subseteq \Phi$ is defined to be *positively closed* (with respect to Φ) if for every positively pointed dependence relation in Φ with base Δ and apex α,

$$\Delta \subseteq \Psi \quad \text{implies} \quad \alpha \in \Psi. \tag{1.2}$$

The *positive closure* of $\Theta \subseteq \Phi$, denoted $\mathrm{cl}(\Theta)$, is the smallest positively closed subset of Φ that contains Θ; i.e., the intersection of Φ with the nonnegative linear span of Θ.

In the braid arrangement (i.e., $\Phi = \mathcal{A}_{\ell-1}$), the subsets of Φ may be identified with reflexive binary relations on the set $[\ell] := \{1, \dots, \ell\}$, the positively closed sets are the transitive relations, and $\Psi \mapsto \mathrm{cl}(\Psi)$ is the operation of transitive closure.

Proposition 1.1. *A subset $\Psi \subseteq \Phi$ is positively closed if and only if (1.2) holds for every irreducible circuit of Φ with base Δ and apex α.*

Before proving this result, it will be convenient to define the *height* of a positively pointed dependence relation $\alpha = \sum_i c_i\alpha_i$ to be $\sum_i c_i\|\alpha_i\|/\|\alpha\|$, where $\|\alpha\| = \langle \alpha, \alpha\rangle^{1/2}$ denotes the usual Euclidean norm. By the Cauchy-Schwarz inequality, one knows that $\|\beta + \gamma\| < \|\beta\| + \|\gamma\|$ unless β and γ are linearly dependent, so $\|\alpha\| < \sum_i c_i\|\alpha_i\|$ unless the summands are pairwise dependent. The latter is forbidden in a positively pointed dependence relation, so we have

Lemma 1.2. *Every positively pointed dependence relation has height > 1.*

Proof of Proposition 1.1. The stated condition is clearly necessary. For the converse, suppose that (1.2) holds for every irreducible circuit, but that there is some positively pointed dependence relation $\alpha = c_1\alpha_1 + \cdots + c_n\alpha_n$ such that $\alpha_1, \dots, \alpha_n \in \Psi$ and $\alpha \notin \Psi$. Among all such relations with n minimal, choose one of minimum height. (The space of positively pointed dependence relations is polyhedral, so it is clear that height-minimizers exist.) If there is a dependence relation $\sum a_i\alpha_i = 0$ with at least one positive coefficient, one may take t to be the smallest positive member of $\{c_i/a_i : 1 \leqslant i \leqslant n\}$, so that the coefficients of the dependence relation

$$\alpha = (c_1 - ta_1)\alpha_1 + \cdots + (c_n - ta_n)\alpha_n.$$

are nonnegative and include at least one zero, contradicting the minimality of n.

The chosen dependence relation must therefore be a positively pointed circuit, so there is an irreducible circuit of the form $\beta = \sum_{i \in I} b_i\alpha_i$ for some $I \subseteq [n]$ and $\beta \in \Phi$. Furthermore, we must have $\beta \in \Psi$ (and hence $\beta \neq \alpha$), since (1.2) is assumed to hold for all such circuits. It follows that

$$\alpha = t\beta + \sum_{i \in I}(c_i - tb_i)\alpha_i + \sum_{i \notin I} c_i\alpha_i$$

is a positively pointed dependence relation for small $t > 0$, and if t is the minimum of

$\{c_i/b_i : i \in I\}$, it will have n positive coefficients (the minimum possible) and height

$$\begin{aligned} t\frac{\|\beta\|}{\|\alpha\|} &+ \sum_{i\in I}(c_i - tb_i)\frac{\|\alpha_i\|}{\|\alpha\|} + \sum_{i\notin I} c_i \frac{\|\alpha_i\|}{\|\alpha\|} \\ &= h + t\frac{\|\beta\|}{\|\alpha\|}\Big(1 - \sum_{i\in I} b_i \frac{\|\alpha_i\|}{\|\beta\|}\Big) = h - t\frac{\|\beta\|}{\|\alpha\|}(h'-1), \end{aligned}$$

where h and h' denote the heights of the original circuit and the irreducible circuit. By Lemma 1.2, the height of the new circuit is strictly less than h, a contradiction. $\square$

REMARK 1.3. (a) The pair (α, Δ) is a positively pointed circuit of Φ if and only if the positive closure of Δ includes α, and no proper subset of Δ has this property. Since Proposition 1.1 shows that the closure operator is completely determined by irreducible circuits, it follows that the irreducible circuits of Φ also determine the corresponding oriented matroid.

(b) It is possible for several irreducible circuits to share the same base. For example, if $\dim V = 2$, this occurs whenever Φ includes vectors normal to 4 or more hyperplanes.

(c) Regard two irreducible circuits as equivalent if they share the same base. Minor modifications of the above proof show that Proposition 1.1 may be strengthened as follows: Given any fixed set $\mathcal{C}$ of equivalence class representatives of the irreducible circuits of Φ, a subset Ψ is positively closed if and only if (1.2) holds for every $(\alpha, \Delta) \in \mathcal{C}$. Conversely, it is not hard to see that this fails if $\mathcal{C}$ does not contain at least one member from each equivalence class of irreducible circuits.

(d) The concept of an irreducible circuit makes sense in any oriented matroid; what is not so clear is whether Proposition 1.1 is valid in the non-realizable cases. Certainly the above proof is not.

Let $\alpha^\perp = \{\lambda \in V : \langle \lambda, \alpha \rangle = 0\}$ denote the central hyperplane corresponding to any nonzero $\alpha \in V$, and let

$$H^+(\alpha) := \{\lambda \in V : \langle \lambda, \alpha \rangle > 0\}$$

denote the positive side of $\alpha^\perp$. The connected regions of the complement of a union of hyperplanes $\beta_1^\perp \cup \cdots \cup \beta_n^\perp$ are the nonempty half-space intersections of the form

$$H^+(\pm\beta_1) \cap \cdots \cap H^+(\pm\beta_n).$$

Thus in order to catalog the connected regions defined by various subarrangements of $\mathcal{H}(\Phi) = \{\alpha^\perp : \alpha \in \Phi\}$, one needs to characterize when

$$H^+(\alpha_1) \cap \cdots \cap H^+(\alpha_n) = \varnothing$$

for various subsets $\{\alpha_1, \ldots, \alpha_n\} \subseteq \Phi$.

Example 1.4. A graphical arrangement $\mathcal{G}$ is a subarrangement of the braid arrangement $\mathcal{H}(\mathcal{A}_{\ell-1})$ (for some ℓ). Such an arrangement may be identified with a simple graph Γ on the vertex set $[\ell]$, and selecting a half-space for each member of $\mathcal{G}$ amounts to selecting an orientation of Γ. Moreover, it is easy to see that the intersection of these half-spaces is a non-empty region of the complement of $\mathcal{G}$ if and only if the orientation is acyclic.

Proposition 1.5. *If $\Psi = \{\alpha_1, \dots, \alpha_n\} \subseteq \Phi$, then*

$$H^+(\alpha_1) \cap \cdots \cap H^+(\alpha_n) \neq \varnothing \text{ if and only if } (-\Psi) \cap \mathrm{cl}(\Psi) = \varnothing.$$

In other words (generalizing Example 1.4), one may test whether the intersection of positive half spaces corresponding to some $\Psi \subseteq \Phi$ is a nonempty region of the complement of $\mathcal{H}(\Psi)$ by using the irreducible circuits of Φ to compute the positive closure of Ψ, and observe whether this closure includes members of $-\Psi$. In principle, one could use the irreducible circuits of Ψ for this purpose, instead of Φ, but it may happen (as in the case of graphical arrangements) that these circuits are simpler in Φ. Also, once the irreducible circuits of Φ are known, one may prefer to avoid the expense of computing the irreducible circuits of its subarrangements.

Proof. Let $C_i = H^+(\alpha_1) \cap \cdots \cap H^+(\alpha_i)$ for $1 \leqslant i \leqslant n$.

Every $\alpha \in \mathrm{cl}(\Psi)$ is in the nonnegative linear span of Ψ and hence $C_n \subseteq H^+(\alpha)$ for such α. Therefore if $\mathrm{cl}(\Psi)$ includes both α and $-\alpha$, then $C_n \subseteq H^+(\alpha) \cap H^+(-\alpha) = \varnothing$. Conversely, if $C_n = \varnothing$, then there is an index i such that $C_{i-1} \neq \varnothing$ and $C_i = \varnothing$; i.e.,

$$\langle \lambda, \alpha_1 \rangle > 0, \dots, \langle \lambda, \alpha_{i-1} \rangle > 0 \quad \Rightarrow \quad \langle \lambda, \alpha_i \rangle < 0. \tag{1.3}$$

In particular, the fact that C_{i-1} is open forces the last inequality to be strict. It follows that $-\alpha_i$ must be in the nonnegative linear span of $\alpha_1, \dots, \alpha_{i-1}$ (and hence $-\alpha_i \in \mathrm{cl}(\Psi)$); otherwise, there would be a hyperplane in V separating $-\alpha_i$ from the cone generated by $\alpha_1, \dots, \alpha_{i-1}$. That is, there would exist $\lambda \in V$ such that $\langle \lambda, \alpha_1 \rangle > 0, \dots, \langle \lambda, \alpha_i \rangle > 0$, contradicting (1.3). Thus both $-\alpha_i$ and α_i are in $\mathrm{cl}(\Psi)$. □

2. Simply-laced root systems

Now suppose that Φ is a simply-laced root system. Renormalizing if necessary, we may assume that $\langle \alpha, \alpha \rangle = 2$ for all $\alpha \in \Phi$, in which case the defining features of Φ are

(2.1) $\langle \alpha, \beta \rangle \in \mathbf{Z}$ for all $\alpha, \beta \in \Phi$, and

(2.2) $\alpha - \langle \alpha, \beta \rangle \beta \in \Phi$ for all $\alpha, \beta \in \Phi$.

Since the roots have length $\sqrt{2}$, property (2.1) implies

(2.3) $\langle \alpha, \beta \rangle \in \{1, 0, -1\}$ or $\alpha = \pm\beta$ for all $\alpha, \beta \in \Phi$.

In other words, the angle between every pair of independent roots is 60°, 90°, or 120°.

Remark 2.1. Although it does not seem to be well-known, it is noteworthy that any set Ψ of vectors of length $\sqrt{2}$ satisfying (2.3) is necessarily a subset of some simply-laced root system. Indeed, (2.1) holds trivially. If (2.2) fails, say $\gamma = \alpha - \langle \alpha, \beta \rangle \beta \notin \Psi$ for some $\alpha, \beta \in \Psi$, then an easy calculation shows that $\langle \gamma, \gamma \rangle = 2$. Moreover, γ is in the integer span of α and β, so (2.1) (and therefore also (2.3)) continues to hold for $\Psi' := \Psi \cup \{\gamma\}$. If Ψ' also violates (2.2), then we may continue to add vectors, but eventually we must reach a configuration that satisfies (2.2). Otherwise, we could pack infinitely many points on a sphere so that every pair of points is at least 60° apart.

THEOREM 2.2. *If $\Delta = \{\alpha_1, \ldots, \alpha_n\}$ is the base and α is the apex of a positively pointed circuit in a simply-laced root system Φ, then (α, Δ) is irreducible if and only if*

(a) $\langle \alpha, \alpha_j \rangle = 1$ *for all j, and*
(b) *either $\langle \alpha_i, \alpha_j \rangle \in \{0, 1\}$ for all $i \neq j$, or $n = 2$.*

Proof. We have $\alpha = \sum c_i \alpha_i$ for certain positive coefficients c_i $(1 \leqslant i \leqslant n)$.

If (α, Δ) is irreducible and $\langle \alpha_i, \alpha_j \rangle \in \{0, 1\}$ for all $i \neq j$, then $\langle \alpha, \alpha_j \rangle = \sum_i c_i \langle \alpha_i, \alpha_j \rangle \geqslant 2c_j > 0$, so (2.3) implies $\langle \alpha, \alpha_j \rangle = 1$ and (a) follows. If $\langle \alpha_i, \alpha_j \rangle \notin \{0, 1\}$ for some $i \neq j$, then $\langle \alpha_i, \alpha_j \rangle = -1$ (again by (2.3)) and (2.2) implies that $\alpha_i + \alpha_j$ is a root. Since no proper subset of Δ has a root in its positive linear span (given that the circuit is irreducible), this forces $\Delta = \{\alpha_i, \alpha_j\}$ and $\alpha = \alpha_i + \alpha_j$, whence (a) and (b) clearly hold.

Conversely, if (α, Δ) fails to be irreducible, then $n > 2$ and there is a proper $I \subset [n]$ and a root $\beta = \sum_{i \in I} b_i \alpha_i \in \Phi - \Delta$ with $b_i > 0$ for all $i \in I$. Given that (b) holds, we must have $\langle \alpha_i, \alpha_j \rangle \geqslant 0$ for all i and j. By the same reasoning as above, it follows that $\langle \beta, \alpha_i \rangle = 1$ for all $i \in I$, and hence $\sum_{i \in I} b_i = \sum_{i \in I} b_i \langle \beta, \alpha_i \rangle = \langle \beta, \beta \rangle = 2$. Given that (a) holds, we thus obtain $\langle \beta, \alpha \rangle = \sum_{i \in I} b_i \langle \alpha_i, \alpha \rangle = \sum_{i \in I} b_i = 2$. Hence $\beta = \alpha$, a contradiction. □

REMARK 2.3. Recall (see Remark 1.3(b)) that a vector configuration may have more than one irreducible circuit with the same base. However, in a simply-laced root system, the above result shows that the projections of the apex onto each base vector are uniquely determined, and hence there is at most one irreducible circuit with a given base.

Let α_0 and $\Delta = \{\alpha_1, \ldots, \alpha_n\}$ be the apex and base of some irreducible circuit in a simply-laced root system Φ. We define the *type* of (α_0, Δ) to be the isomorphism class of the smallest root system containing $\Delta \cup \{\alpha_0\}$. Theorem 2.2 shows that the apex is not orthogonal to any base vector, so the type is irreducible; i.e., one of $\mathcal{A}_n$, $\mathcal{D}_n$, $\mathcal{E}_6$, $\mathcal{E}_7$, or $\mathcal{E}_8$.

In the case $n = 2$, the only possible type is $\mathcal{A}_2$, and it is clear that the dependence relation must take the form $\alpha_0 = \alpha_1 + \alpha_2$, where $\langle \alpha_1, \alpha_2 \rangle = -1$. Otherwise, we have $n > 2$ and Theorem 2.2 implies that $\langle \alpha_i, \alpha_j \rangle \in \{0, 1\}$ for all $i \neq j$, so we may encode the Gram matrix $G = [\langle \alpha_i, \alpha_j \rangle]_{0 \leqslant i,j \leqslant n}$ by means of a graph Γ with vertex set $\{0, 1, \ldots, n\}$ and an edge between i and j if $\langle \alpha_i, \alpha_j \rangle = 1$. We call $\Gamma = \Gamma(\alpha, \Delta)$ the *diagram* of (α, Δ).

REMARK 2.4. We may similarly associate to any simple graph Γ a symmetric matrix $G(\Gamma)$ with 2's on the diagonal, but $G(\Gamma)$ will be realizable as the Gram matrix of a set of vectors in a Euclidean space only if $G(\Gamma)$ is positive semidefinite. In such cases, the vectors necessarily have length $\sqrt{2}$ and satisfy (2.3), and thus (Remark 2.1) are roots in

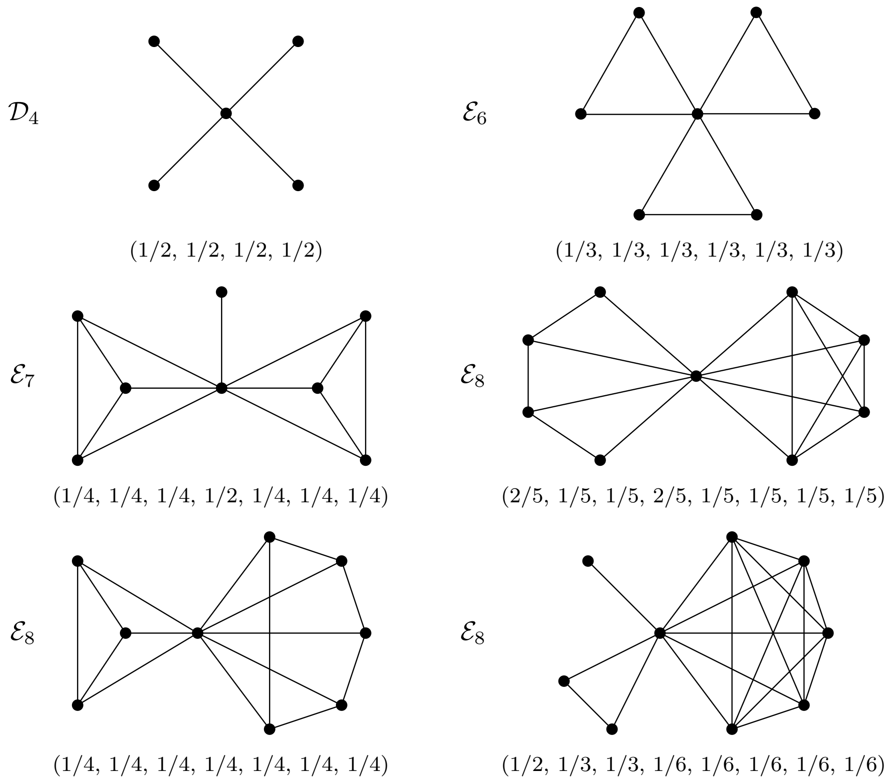

FIGURE 1: Diagrams of irreducible circuits and their apex coordinates.

a simply-laced root system; this observation enabled the classification of such graphs by Cameron, Goethals, Seidel, and Shult [**3**]. Similarly, there is also a classification of the graphs Γ for which the eigenvalues of $G(\Gamma)$ are $\leqslant 4$—these are the simply-laced Dynkin diagrams of finite or affine type.

COROLLARY 2.5. *A graph Γ with $n+1$ vertices is the diagram of an irreducible circuit of rank $n > 2$ in a simply-laced root system if and only if*

(a) *the matrix $G(\Gamma)$ is positive semidefinite of rank n,*
(b) *the kernel of $G(\Gamma)$ has a (necessarily unique) member with n positive coordinates and one coordinate equal to -1, and*
(c) *the vertex whose coordinate is -1 in (b) is adjacent to every other vertex in Γ.*

Proof. If $\alpha_0 = \sum_i c_i \alpha_i$, then the column vector $(-1, c_1, \ldots, c_n)$ is in the kernel of the Gram matrix $G = [\langle \alpha_i, \alpha_j \rangle]$, so it is clear that the stated conditions are necessary. Conversely, if $G(\Gamma)$ is positive semidefinite, then (see Remark 2.4) it is the Gram matrix of a set of roots in some simply-laced root system Φ, and Theorem 2.2 shows that the given conditions force these roots to form an irreducible circuit in Φ. $\square$

Theorem 2.6. *Up to isometry, there are seven different irreducible circuits that can occur in a simply-laced root system. Their types are $\mathcal{A}_2$, $\mathcal{D}_4$, $\mathcal{E}_6$, $\mathcal{E}_7$, and $\mathcal{E}_8$ (the last occurring in three distinct ways), and their diagrams for the cases of rank > 2 are illustrated in Figure 1.*

As a convenience, along with the diagrams in Figure 1 we have provided the coordinates of the apex in terms of the base. This information is redundant at least in principle, since these coordinates are completely determined by Γ (more specifically, the kernel of $G(\Gamma)$). For example, one sees that the $\mathcal{D}_4$-circuit amounts to a base configuration of four mutually orthogonal roots $\alpha_1, \dots, \alpha_4$ and an apex of the form $(\alpha_1 + \cdots + \alpha_4)/2$. We have left the identification between nodes and coordinates to the reader, but symmetry considerations alone are enough to make this determination, with one exception. In the exceptional case, the coordinate $2/5$ is associated to the two nodes of degree 2.

Proof. Let α be the apex and $\Delta = \{\alpha_1, \dots, \alpha_n\}$ the base of an irreducible circuit in a simply-laced root system Φ. Theorem 2.2 shows that (α, Δ) is also an irreducible circuit in any simply-laced root system that contains $\Delta \cup \{\alpha\}$, so (using the classification of root systems) we may assume that $\Phi = \mathcal{D}_\ell$ ($\ell \geqslant 4$) or $\Phi = \mathcal{E}_8$.

Property (2.2) amounts to the statement that the reflections through hyperplanes orthogonal to roots generate a (Coxeter) group W of automorphisms of Φ. One knows that the roots in an irreducible, simply-laced root system form a single W-orbit, so we may fix a single choice for α (depending on Φ).

We may also assume that $n > 2$ and $\langle \alpha_i, \alpha_j \rangle \geqslant 0$ for all i, j; otherwise, Theorem 2.2 implies $\alpha = \alpha_1 + \alpha_2$ and that (α, Δ) is of type $\mathcal{A}_2$.

In the case $\Phi = \mathcal{D}_\ell = \{\pm\varepsilon_i \pm \varepsilon_j : 1 \leqslant i < j \leqslant \ell\}$ (where $\varepsilon_1, \dots, \varepsilon_\ell$ are orthonormal), we may select $\alpha = \varepsilon_1 + \varepsilon_2$, whence the constraint $\langle \alpha, \alpha_j \rangle = 1$ (Theorem 2.2) forces

$$\Delta \subset \{\varepsilon_1 \pm \varepsilon_i : 2 < i \leqslant \ell\} \cup \{\varepsilon_2 \pm \varepsilon_i : 2 < i \leqslant \ell\}.$$

If a base vector has a positive ε_i-coordinate ($i > 2$), there must also be a base vector with a negative ε_i-coordinate, and vice-versa; otherwise, α could not be in the span of Δ. Furthermore, $\varepsilon_1 + \varepsilon_i$ and $\varepsilon_2 - \varepsilon_i$ cannot both occur in Δ (they have inner product -1), so Δ must be a union of pairs of the form $\varepsilon_1 \pm \varepsilon_i$ or $\varepsilon_2 \pm \varepsilon_i$. In addition, we must have at least one pair of each type; again, the span of Δ could not otherwise include α. Thus Δ includes four roots of the form $\varepsilon_1 \pm \varepsilon_i$ and $\varepsilon_2 \pm \varepsilon_j$ for some $i \neq j$. The sum of these four roots is 2α, whence Theorem 2.2 implies that they form the base of an irreducible circuit with apex α; i.e., $\Delta = \{\varepsilon_1 \pm \varepsilon_i, \varepsilon_2 \pm \varepsilon_j\}$ and (α, Δ) is of type $\mathcal{D}_4$. Note that the diagram of (α, Δ) appears in Figure 1.

For the remainder of the proof, we may assume $\Phi = \mathcal{E}_8$.

In particular, since $\mathcal{D}_\ell$ has been analyzed completely, the only remaining possibility for the type of (α, Δ) is $\mathcal{E}_n$, so as soon as the remaining diagrams in Figure 1 are confirmed, we know that the types claimed for these diagrams must be correct.

In the usual presentation, $\mathcal{E}_8$ consists of the roots of $\mathcal{D}_8$ as above, together with the 128 vectors of the form $(\pm\varepsilon_1 \pm \cdots \pm \varepsilon_8)/2$ with an even number of negative signs. If we fix the choice of $\alpha = (\varepsilon_1 + \cdots + \varepsilon_8)/2$, then

$$\{\beta \in \Phi : \langle \alpha, \beta \rangle = 1\} = \{\varepsilon_i + \varepsilon_j : 1 \leqslant i < j \leqslant 8\} \cup \{\alpha - (\varepsilon_i + \varepsilon_j) : 1 \leqslant i < j \leqslant 8\},$$

and Theorem 2.2 implies that Δ must be a subset of these roots.

For all such roots β, let us define $\bar\beta = \varepsilon_i + \varepsilon_j$ if $\beta = \varepsilon_i + \varepsilon_j$ or $\beta = \alpha - (\varepsilon_i + \varepsilon_j)$. Notice that $\varepsilon_i + \varepsilon_j$ and $\alpha - (\varepsilon_i + \varepsilon_j)$ have inner product -1, and thus Δ cannot include both; all other inner products are $\geqslant 0$. It follows that $\bar\Delta := \{\bar\alpha_1, \ldots, \bar\alpha_n\}$ has n distinct members.

LEMMA 2.7. *If $n > 2$ and $\Theta = \{\beta_1, \ldots, \beta_n\} \subset \{\varepsilon_i + \varepsilon_j : 1 \leqslant i < j \leqslant 8\}$, then there is an irreducible circuit (α, Δ) with $\Theta = \bar\Delta$ if and only if*

(a) *the span of $\Theta \cup \{\alpha\}$ is n-dimensional, and*
(b) *there exist nonzero scalars $c_1, \ldots, c_n$ such that $c_1\beta_1 + \cdots + c_n\beta_n \in \mathbf{R}\alpha$.*

Moreover, under these conditions, nodes i and j are adjacent in the diagram $\Gamma(\alpha, \Delta)$ if and only if $\langle \beta_i, \beta_j \rangle = 1$ and $c_ic_j > 0$ or $\langle \beta_i, \beta_j \rangle = 0$ and $c_ic_j < 0$ for $1 \leqslant i < j \leqslant n$.

Recall that there are $n + 1$ nodes in $\Gamma(\alpha, \Delta)$, with the extra node (the apex) being adjacent to every other node. Thus the above result shows that $\bar\Delta$ determines $\Gamma(\alpha, \Delta)$.

Proof. Suppose Θ satisfies (a) and (b). By renormalization, there must be a dependence relation of the form $b\alpha = \sum c_i\beta_i$, where $b = 0$ or 1 and each c_i is nonzero. Hence,

$$(b - c)\alpha = \sum_{i \in I} c_i\beta_i - \sum_{j \in J} c_j(\alpha - \beta_j),$$

where $I = \{i : c_i > 0\}$, $J = I^c = \{j : c_j < 0\}$, and $c = \sum_{j \in J} c_j < 0$. Thus α is in the strictly positive span of $\Delta = \{\beta_i : i \in I\} \cup \{\alpha - \beta_j : j \in J\}$ and Theorem 2.2 implies that (α, Δ) is an irreducible circuit with $\bar\Delta = \Theta$. (In particular, note that $\operatorname{Span}\Delta = \operatorname{Span}\Delta \cup \{\alpha\} = \operatorname{Span}\Theta \cup \{\alpha\}$ is n-dimensional by (a), so Δ is linearly independent.)

Conversely, if $\Delta = \{\alpha_1, \ldots, \alpha_n\}$ is the base of an irreducible circuit with apex α, one knows that the unique relation of the form $\alpha = \sum c_i\alpha_i$ has positive coefficients, so

$$\alpha = \sum_{i \in I} c_i\bar\alpha_i + \sum_{j \in J} c_j(\alpha - \bar\alpha_j),$$

where $I = \{i : \bar\alpha_i = \alpha_i\}$ and $J = I^c = \{j : \bar\alpha_j = \alpha - \alpha_j\}$. It follows that

$$(1 - c)\alpha = \sum_{i \in I} c_i\bar\alpha_i - \sum_{j \in J} c_j\bar\alpha_j,$$

where $c = \sum_{j \in J} c_j$. Thus $\Theta := \bar\Delta$ satisfies (b), and $\operatorname{Span}\bar\Delta \cup \{\alpha\} = \operatorname{Span}\Delta \cup \{\alpha\} = \operatorname{Span}\Delta$ is n-dimensional, so (a) holds as well.

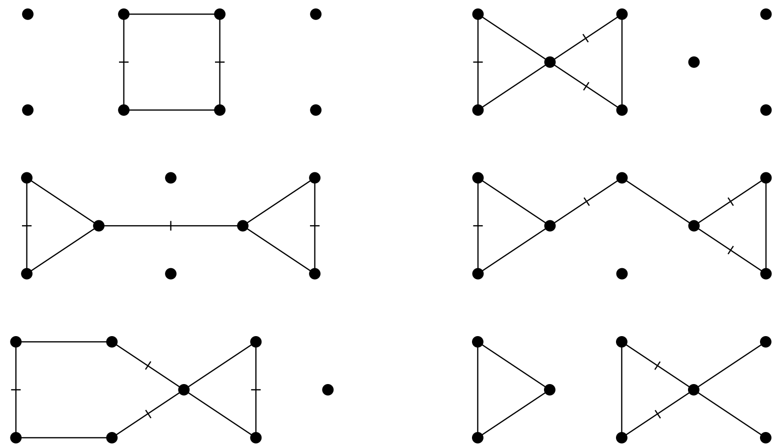

FIGURE 2: Examples of the graphs $\Sigma(\Delta)$.

The fact that $\operatorname{Span}\bar{\Delta} \cup \{\alpha\}$ is n-dimensional implies that the space of dependence relations $b\alpha = \sum b_i\bar{\alpha}_i$ is one-dimensional. Thus for any such (nontrivial) relation, I as defined above must either be the set of indices i such that $b_i > 0$, or the set of indices i such that $b_i < 0$, and J must be I^c. It follows that if $b_ib_j > 0$, then either $\bar{\alpha}_i = \alpha_i$ and $\bar{\alpha}_j = \alpha_j$, or $\bar{\alpha}_i = \alpha - \alpha_i$ and $\bar{\alpha}_j = \alpha - \alpha_j$. Either way, we have $\langle\bar{\alpha}_i, \bar{\alpha}_j\rangle = \langle\alpha_i, \alpha_j\rangle$, and thus i and j are adjacent in $\Gamma(\alpha, \Delta)$ if and only if $\langle\bar{\alpha}_i, \bar{\alpha}_j\rangle = 1$. Similarly, if $b_ib_j < 0$, then either $\bar{\alpha}_i = \alpha_i$ and $\bar{\alpha}_j = \alpha - \alpha_j$, or $\bar{\alpha}_i = \alpha - \alpha_i$ and $\bar{\alpha}_j = \alpha_j$, in which case $\langle\bar{\alpha}_i, \bar{\alpha}_j\rangle = 1 - \langle\alpha_i, \alpha_j\rangle$, and i and j are adjacent in $\Gamma(\alpha, \Delta)$ if and only if $\langle\bar{\alpha}_i, \bar{\alpha}_j\rangle = 0$. $\square$

We may encode $\bar{\Delta}$ by an n-edge simple graph $\Sigma(\Delta)$ on the vertex set $[8]$ by declaring i adjacent to j if $\varepsilon_i + \varepsilon_j \in \bar{\Delta}$, or equivalently, if $\varepsilon_i + \varepsilon_j$ or $\alpha - (\varepsilon_i + \varepsilon_j)$ is in Δ.

In Figure 2 we list the graph encodings of $\bar{\Delta}$ for several irreducible circuits (α, Δ). For example, the vertices of the 7-edge graph in Figure 2 may be numbered so that the set of encoded vectors is $\Theta = \{\beta_1, \ldots, \beta_7\}$, where

$$\beta_1 = \varepsilon_1 + \varepsilon_2,\ \beta_2 = \varepsilon_1 + \varepsilon_3,\ \beta_3 = \varepsilon_2 + \varepsilon_3,\ \beta_4 = \varepsilon_3 + \varepsilon_6,$$
$$\beta_5 = \varepsilon_6 + \varepsilon_7,\ \beta_6 = \varepsilon_6 + \varepsilon_8,\ \beta_7 = \varepsilon_7 + \varepsilon_8.$$

One sees that there is a dependence relation $\beta_2 + \beta_3 + \beta_5 + \beta_6 = \beta_1 + 2\beta_4 + \beta_7$ and that $\operatorname{Span}\Theta$ is 6-dimensional and omits α, so by Lemma 2.7, there is an irreducible circuit (α, Δ) such that $\Theta = \bar{\Delta}$. The validity of the other graphs may be confirmed similarly.

To simplify the reconstruction of diagrams, we have placed marks on some of the edges in Figure 2 to indicate the distribution of positive and negative coordinates relevant to Lemma 2.7(b). Thus two edges of $\Sigma(\Delta)$ correspond to an adjacent pair of nodes in $\Gamma(\alpha, \Delta)$ if they either have a common endpoint and are both marked or both unmarked, or if they are disjoint and one is marked and the other is unmarked. In this way, one may easily check that the diagrams corresponding to the graphs in Figure 2 are those listed in Figure 1.

To complete the proof of Theorem 2.6, we must show that there are no other possible diagrams; we will accomplish this by analyzing the graphs $\Sigma(\Delta)$.

We say that a connected graph is a *unicycle* if it has a unique cycle (or equivalently, an equal number of vertices and edges); it is *odd* or *even* according to the length of this cycle. We say that a tree (or more generally, any connected bipartite graph) is *balanced* if the unique 2-coloring of its vertices has an equal number of vertices of each color.

LEMMA 2.8. *If $\bar{\Delta}$ is linearly independent, then every connected component of $\Sigma(\Delta)$ is an odd unicycle or a balanced tree.*

Proof. If a connected component of $\Sigma(\Delta)$ has vertex set $I \subset [8]$, then the corresponding subset of $\bar{\Delta}$ belongs to a subspace of dimension $|I|$; namely, $\operatorname{Span}\{\varepsilon_i : i \in I\}$. Thus any such component cannot have more than $|I|$ edges; otherwise, $\bar{\Delta}$ would be dependent. If the component has exactly $|I|$ edges it is a unicycle. If it has fewer edges, it is a tree.

If the component has a (necessarily unique) cycle of length k, then up to renumbering there are vectors in $\bar{\Delta}$ of the form $\beta_i = \varepsilon_i + \varepsilon_{i+1}$ for $1 \leqslant i \leqslant k$ (subscripts mod k). If k is even, these vectors satisfy the dependence relation $\sum(-1)^i\beta_i = 0$, contradicting the independence of $\bar{\Delta}$. Thus all unicycles must be odd.

If the component is a tree, let $I = A \cup B$ be the unique partition of the vertices into the color classes of a 2-coloring of the tree. Each edge of $\Sigma(\Delta)$ has either an endpoint in A and an endpoint in B, or both endpoints are disjoint from I, so $\bar{\Delta}$ belongs to the subspace of vectors $\sum b_i\varepsilon_i$ such that $\sum_{i\in A} b_i = \sum_{j\in B} b_j$. Lemma 2.7 implies $\alpha \in \operatorname{Span}\bar{\Delta}$, so α must also belong to this subspace; i.e., $|A| = |B|$ and the tree is balanced. $\square$

LEMMA 2.9. *If $\bar{\Delta}$ is linearly dependent, then $\Sigma(\Delta)$ has at least one isolated vertex, and the edges of $\Sigma(\Delta)$ form an even cycle or a pair of odd cycles connected by a path.*

Proof. Lemma 2.7 implies that $\operatorname{Span}\bar{\Delta}$ has dimension $n-1$ and that the unique dependence relation involving members of $\bar{\Delta}$ has full support. It follows that every proper subset of $\bar{\Delta}$ is linearly independent, so by the reasoning in the proof of Lemma 2.8, every proper subgraph of $\Sigma := \Sigma(\Delta)$ must be a union of trees and odd unicycles.

Let $\bar{\Delta}_1, \dots, \bar{\Delta}_k$ be the subsets of $\bar{\Delta}$ corresponding to the nontrivial connected components of Σ. We have $\operatorname{Span}\bar{\Delta} = \operatorname{Span}\bar{\Delta}_1 \oplus \cdots \oplus \operatorname{Span}\bar{\Delta}_k$, so the fact that $\bar{\Delta}$ is linearly dependent forces some block $\bar{\Delta}_i$ to be dependent. However, the unique dependence relation involving $\bar{\Delta}$ has full support, so this is possible only if $k = 1$; i.e., Σ is connected, except possibly for isolated vertices.

If Σ has an even-length cycle, then (as noted in the proof of Lemma 2.8), an alternating sum of the corresponding members of $\bar{\Delta}$ vanishes, so this must be the unique dependence relation in $\bar{\Delta}$. Since this relation must have full support, the cycle necessarily uses all n edges of Σ. Also, we cannot have $n = 8$ in this case (the maximum possible); otherwise, $\operatorname{Span}\bar{\Delta}$ would include α and $\operatorname{Span}\bar{\Delta} \cup \{\alpha\}$ would be only 7-dimensional, contradicting Lemma 2.7(a). In particular, Σ must have isolated vertices.

The remaining possibility is that all cycles of Σ have odd length. Moreover, Σ must have at least two such cycles; otherwise, $\bar{\Delta}$ would be linearly independent. These cycles must be pairwise edge-disjoint, for if there is an edge between vertices i and j that participates in an odd cycle, then $\varepsilon_i - \varepsilon_j$ may be expressed as an alternating sum over the members of $\bar{\Delta}$ corresponding to the remaining members of the cycle. It follows that if this edge participates in two or more such cycles, then there is a dependence relation in $\bar{\Delta}$ that does not involve $\varepsilon_i + \varepsilon_j$, a contradiction.

Thus Σ has at least two edge-disjoint cycles of odd length. All edges of Σ belong to the same connected component, so there must be a path (possibly of length 0) connecting two odd cycles. However, the subgraph formed by the two cycles and the path has more edges than vertices, so the corresponding subset of $\bar{\Delta}$ is linearly dependent, and hence this subgraph must use all n edges of Σ. We must also have $n \leqslant 8$ (the dimension of the ambient space), so one or more of the vertices must be isolated. $\square$

Given a root β, let $r_\beta \in W$ denote the reflection through the hyperplane orthogonal to β, and W_α the subgroup of W generated by the reflections that fix α (i.e., $\{r_\beta : \langle \alpha, \beta \rangle = 0\}$). This group isometrically permutes the irreducible circuits with apex α, so it suffices to restrict our attention to representatives from each W_α-orbit.

The group W_α also permutes the 28 unordered pairs $\{\varepsilon_i + \varepsilon_j, \alpha - (\varepsilon_i + \varepsilon_j)\}$; indeed, these are the roots in $\{\beta \in \Phi : \langle \alpha, \beta \rangle = 1\}$. This induces an action of W_α on the graphs with vertex set [8] that is compatible with the action of W_α on irreducible circuits in the sense that $w\Sigma(\Delta) = \Sigma(w\Delta)$ for all $w \in W_\alpha$.

To describe the action of W_α on a graph Σ more explicitly, observe that the roots orthogonal to α (a copy of the root system $\mathcal{E}_7$) consist of all vectors of the form $\varepsilon_i - \varepsilon_j$ (i.e., the roots of $\mathcal{A}_7$), together with all 70 vectors of the form $(\pm\varepsilon_1 \pm \cdots \pm \varepsilon_8)/2$ with an equal number of positive and negative coordinates. The reflections corresponding to the roots of the first type generate the symmetric group of all permutations of the vertex set [8], whereas to describe the action of a reflection r_β of the second type, let $B \subset [8]$ denote the four indices corresponding to positive coordinates in β. If $i \in B$ and $j \notin B$, then r_β fixes $\varepsilon_i + \varepsilon_j$ and $\alpha - (\varepsilon_i + \varepsilon_j)$, so i and j are adjacent in Σ if and only if they are adjacent in $r_\beta\Sigma$. On the other hand, if $i, j \in B$ or $i, j \notin B$, then r_β interchanges $\varepsilon_i + \varepsilon_j$ and $\alpha - (\varepsilon_k + \varepsilon_\ell)$, where k and ℓ denote the remaining two members of B or B^c (respectively), so i and j are adjacent in Σ if and only if k and ℓ are adjacent in $r_\beta\Sigma$.

Given B and β as above, we refer to $r_\beta\Sigma$ as the *B-reflection* of Σ.

The following result completes the proof of Theorem 2.6.

Lemma 2.10. *Every W_α-orbit of irreducible circuits (α, Δ) of rank $n > 2$ includes at least one whose graph $\Sigma(\Delta)$ is listed in Figure 2.*

Proof. The only graph satisfying the conditions described in Lemma 2.9 that is *not* listed in Figure 2 is a 6-cycle. However, by taking B to be the leftmost four vertices in the graphs below, one sees that a 6-cycle is the B-reflection of one of the listed graphs.

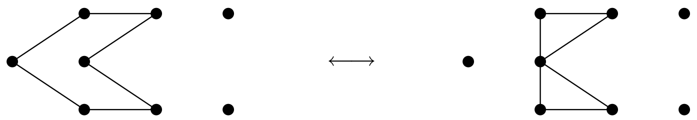

The remaining possibility is that $\bar{\Delta}$ is linearly independent, and hence (Lemma 2.8) every connected component of $\Sigma = \Sigma(\Delta)$ is a balanced tree or an odd unicycle.

If there is a set of four vertices $B = \{i_1, i_2, i_3, i_4\}$ such that i_2, i_3 and i_4 are mutually non-adjacent in Σ, and the neighborhood of i_1 is a subset of B, then i_1 is an isolated vertex of the B-reflection of Σ and we say that B *isolates* i_1. An isolated vertex is an unbalanced tree, so the corresponding irreducible circuit must satisfy the conditions of Lemma 2.9 and its graph must belong to the same W_α-orbit as one of the graphs listed in Figure 2.

If Σ has no vertex of degree $\leqslant 1$, then it must be a union of disjoint (odd) cycles, one of length 3 and one of length 5. If we take i_1, i_2, i_3 to be three consecutive vertices of the 5-cycle and i_4 any vertex of the 3-cycle, then $\{i_1, i_2, i_3, i_4\}$ isolates i_2.

A vertex of degree 0 is an unbalanced tree, so we may assume henceforth that Σ has a vertex i_1 of degree 1, and that i_2 is the unique vertex adjacent to i_1.

Suppose Σ is disconnected. If there is a pair of non-adjacent vertices i_3 and i_4 among the components not involving i_1, then $\{i_1, i_2, i_3, i_4\}$ isolates i_1, so we may assume that Σ has exactly two components: the one containing i_1, and a clique. If this clique is a balanced tree, it must have two vertices, each of degree 1. However, the above argument applies equally well to any component with a vertex of degree 1, so the component containing i_1 must also be a clique. Since i_1 has degree 1, this allows for only four vertices in Σ, a contradiction. Thus, the clique must be an odd unicycle; i.e., a 3-cycle. The component containing i_1 therefore has 5 vertices and cannot be a balanced tree, so it must be a unicycle containing a 3-cycle. If a vertex i_3 of this component is not adjacent to i_2, then we may take i_4 to be a vertex of the other component and use $\{i_1, i_2, i_3, i_4\}$ to isolate i_1. Thus we may assume that i_2 is adjacent to all vertices of its component, in which case Σ is the last graph listed in Figure 2.

The remaining possibility is that Σ is connected. In case Σ is a balanced tree, then a 2-coloring of Σ has four vertices of each color, and we may take i_3 and i_4 to be two vertices with the same color as i_2 in order for $\{i_1, i_2, i_3, i_4\}$ to isolate i_1.

If Σ is an odd unicycle, we may assume i_1 is chosen to be at maximum distance d from the vertices of the unique cycle. If the cycle edges are deleted, the result may be viewed as a forest of rooted trees, the roots being the vertices that participate in the cycle.

- If i_2 has at least two other children in addition to i_1, say i_3 and i_4, then both have degree 1 and the $\{i_1, i_2, i_3, i_4\}$-reflection of Σ is disconnected, one of the components being a 3-cycle formed by i_1, i_3, and i_4. Thus the W_α-orbit of Σ belongs to one of the previous cases.

- If i_2 has two children and is not a cycle vertex (i.e., $d > 1$), then the neighborhood of i_2 consists of three independent vertices: i_1, the other child (say) i_3, and the parent (say) i_4 of i_2, whence $\{i_1, i_2, i_3, i_4\}$ isolates i_2.
- If i_2 has two children (i_1, and say i_3) and $d = 1$, then parity considerations imply that there is another vertex i_4 not in the (odd) cycle. It cannot be in the tree rooted at i_2, so it has a parent i_5 in another tree, and $\{i_1, i_3, i_4, i_5\}$ isolates i_4.
- If i_2 has only one child (namely, i_1) and $d > 1$, then parity considerations imply that there is another vertex i_3 that is not in the (odd) cycle, and a vertex i_4 in the cycle that is not adjacent to i_2 or i_3, in which case $\{i_1, i_2, i_3, i_4\}$ isolates i_2 or i_1, depending on whether i_2 is a child of i_3.
- If i_2 has only one child, $d = 1$, and the length of the cycle is at least 5, then the neighborhood of i_2 has exactly three independent vertices; namely, i_1 and the two vertices on the cycle adjacent to i_2, so again we may isolate i_2.

The above considerations leave out only those unicycles in which the tree rooted at each cycle vertex has at most two vertices (in particular, $d = 1$), and there are only three vertices in the cycle. However, such graphs could have at most six vertices. $\square$

REMARK 2.11. In an irreducible, simply-laced root system Φ, we claim that isometric irreducible circuits are necessarily in the same W-orbit.

Suppose that we have two irreducible circuits of type $\Psi \subseteq \Phi$. By Theorem 2.6, we know that $\Psi = \mathcal{A}_2$, $\mathcal{D}_4$, $\mathcal{E}_6$, $\mathcal{E}_7$, or $\mathcal{E}_8$; a special feature of these root systems is that they are maximal within their ranks among simply-laced root systems. It follows that $\Psi = \Phi \cap \operatorname{Span} \Psi$, so a generic vector λ in the orthogonal complement of $\operatorname{Span} \Psi$ will have the property that $\langle \lambda, \alpha \rangle = 0$ for $\alpha \in \Psi$ and $\langle \lambda, \alpha \rangle \neq 0$ for $\alpha \in \Phi - \Psi$. This property is preserved if we replace (λ, Ψ) with $(w\lambda, w\Psi)$ for any $w \in W$. In particular, by choosing w so that $w\lambda$ is dominant, one obtains that $w\Psi$ is necessarily parabolic; i.e., generated by simple roots. Thus, every copy of Ψ in Φ is in the W-orbit of a parabolic copy.

A second feature of the root systems $\mathcal{D}_4$, $\mathcal{E}_6$, $\mathcal{E}_7$, and $\mathcal{E}_8$ is that their Dynkin diagrams each occur at most once as a subdiagram of any (connected) Dynkin diagram, so there is at most one parabolic copy of each in Φ. (This fails for $\mathcal{A}_2$, but it is still easy to see that all parabolic copies of $\mathcal{A}_2$ in Φ are W-conjugate.)

Up to W-symmetry, we may therefore assume that the two irreducible circuits generate the same root system Ψ. An isometry between the two circuits thus induces an automorphism of Ψ. One knows that the automorphism group of any root system is the semi-direct product of its reflection group and its group of Dynkin diagram automorphisms. However, the diagrams of $\mathcal{E}_7$ and $\mathcal{E}_8$ have only trivial automorphisms, and it is not hard to check that for the remaining isometry classes (of types $\mathcal{A}_2$, $\mathcal{D}_4$, and $\mathcal{E}_6$), there is an irreducible circuit fixed (setwise) by all Dynkin diagram automorphisms, so the claim follows.

3. Crystallographic root systems

Now assume that Φ is a crystallographic root system. Here, the defining features are

(3.1) $\langle \alpha, \beta^\vee \rangle \in \mathbf{Z}$ for all $\alpha, \beta \in \Phi$, and

(3.2) $\alpha - \langle \alpha, \beta^\vee \rangle \beta \in \Phi$ for all $\alpha, \beta \in \Phi$,

where $\beta^\vee := 2\beta / \langle \beta, \beta \rangle$ denotes the co-root corresponding to β.

The roots within a given irreducible component of Φ necessarily have at most two possible lengths; if only one length occurs, then the component is simply-laced. If two lengths occur ("long" and "short"), then their ratio is either $\sqrt{2}$ or $\sqrt{3}$, the latter occurring only if the component is isomorphic to $\mathcal{G}_2$.

The following result generalizes part of Theorem 2.2.

Proposition 3.1. *If $\Delta = \{\alpha_1, \dots, \alpha_n\}$ is the base and α is the apex of an irreducible circuit of rank $n > 2$ in a crystallographic root system Φ, then*

(a) $\langle \alpha, \alpha_j^\vee \rangle > 0$ *for all j, and*

(b) $\langle \alpha_i, \alpha_j^\vee \rangle \geqslant 0$ *for all i, j.*

Proof. If (b) fails, then $\langle \alpha_i, \alpha_j^\vee \rangle = -c < 0$ for some $i \neq j$, and (3.2) implies that $\alpha_i + c\alpha_j$ is a root. Given that (α, Δ) is irreducible, no proper subset of Δ has a root in its positive linear span, so this forces $\Delta = \{\alpha_i, \alpha_j\}$ and $n = 2$, a contradiction. Also, the expansion $\alpha = \sum c_i \alpha_i$ has positive coefficients, so (b) implies $\langle \alpha, \alpha_j^\vee \rangle = \sum_i c_i \langle \alpha_i, \alpha_j^\vee \rangle \geqslant 2c_j > 0$, and thus (a) follows. $\square$

Remark 3.2. In view of Theorem 2.2, one might expect that the above properties characterize the irreducible circuits of rank > 2 among all positively pointed circuits. However, the geometry of a positively pointed circuit alone does not generally determine whether it is irreducible in a given (crystallographic) root system. For example, in

$$\mathcal{B}_\ell = \{\pm\varepsilon_i : 1 \leqslant i \leqslant \ell\} \cup \{\pm\varepsilon_i \pm \varepsilon_j : 1 \leqslant i < j \leqslant \ell\},$$

the four roots $\varepsilon_1 \pm \varepsilon_2$, $\varepsilon_3 \pm \varepsilon_4$ form the base of a positively pointed circuit with apex $\varepsilon_1 + \varepsilon_3$ (provided $\ell \geqslant 4$). The simplicial cone generated by this base has roots of $\mathcal{B}_\ell$ on its boundary (namely, ε_1 and ε_3), so it is not irreducible in $\mathcal{B}_\ell$. On the other hand, this circuit *is* irreducible in $\mathcal{D}_\ell$.

Extending the terminology we introduced for simply-laced root systems, define the *type* of an irreducible circuit (α, Δ) in a crystallographic root system to be the isomorphism class of the smallest root system containing $\Delta \cup \{\alpha\}$. Proposition 3.1(a) shows that the type is necessarily irreducible (the rank 2 case being trivial).

In an irreducible circuit of rank 2 there may be roots β, γ such that $\langle \beta, \gamma^\vee \rangle < 0$. In all other crystallographic cases, Proposition 3.1(b) implies $\langle \beta, \gamma^\vee \rangle \geqslant 0$, and hence either

(3.3) $\langle \beta, \gamma^\vee \rangle = \langle \gamma, \beta^\vee \rangle = 0$, or

(3.4) $\langle \beta, \gamma^\vee \rangle = \langle \gamma, \beta^\vee \rangle = 1$, or

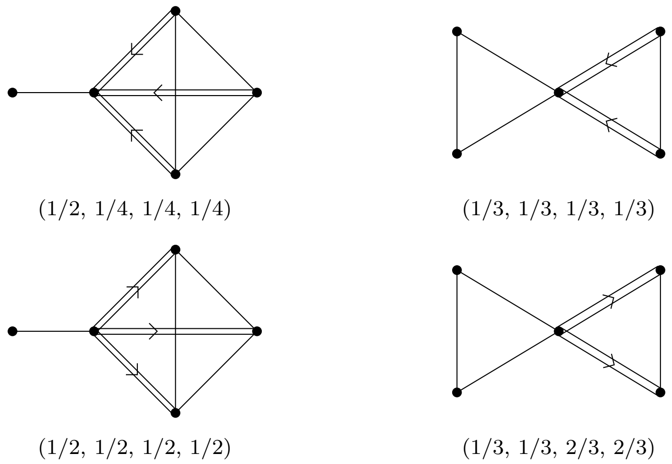

FIGURE 3: The irreducible circuits of type $\mathcal{F}_4$.

(3.5) $\langle \beta, \gamma^\vee \rangle = 1$, $\langle \gamma, \beta^\vee \rangle = 2$, or
(3.6) $\langle \beta, \gamma^\vee \rangle = 1$, $\langle \gamma, \beta^\vee \rangle = 3$,

possibly after exchanging β and γ. Adapting the conventions of Dynkin diagrams, we define the *diagram* of an irreducible circuit (α, Δ) with nonnegative scalar products to be the graph $\Gamma = \Gamma(\alpha, \Delta)$ with vertex set $\Delta \cup \{\alpha\}$ and an unoriented simple edge between β and γ, or a double or triple edge oriented from γ to β, when (3.4), (3.5), or (3.6) hold, respectively. With these conventions, the diagram completely determines the geometry of (α, Δ) up to an overall scalar factor.

THEOREM 3.3. *If (α, Δ) is an irreducible circuit of rank > 2 in an irreducible, multiply-laced root system Φ, then $\Phi = \mathcal{F}_4$ and the type of (α, Δ) is $\mathcal{F}_4$. Moreover, the diagrams of all such circuits and their apex coordinates are provided in Figure 3.*

Proof. The Coxeter group W generated by reflections through hyperplanes orthogonal to roots acts as a group of (isometric) automorphisms of Φ, so it suffices to consider one apex α from each W-orbit of roots. Once α is fixed, Proposition 3.1(a) implies

$$\Delta \subset \Phi_1(\alpha) := \{\beta \in \Phi : \langle \alpha, \beta^\vee \rangle \geqslant 1\} - \{\alpha\}.$$

Now consider the case $\Phi = \mathcal{B}_\ell$, using the roots ε_1 and $\varepsilon_1 + \varepsilon_2$ as orbit representatives. For the apex $\alpha = \varepsilon_1$, we have

$$\Phi_1(\alpha) = \{\varepsilon_1 \pm \varepsilon_i : 2 \leqslant i \leqslant \ell\}.$$

The roots $\varepsilon_1 + \varepsilon_i$ and $\varepsilon_1 - \varepsilon_i$ cannot both be in Δ; otherwise, since ε_1 is in their positive span, they would exhaust all of Δ and contradict the hypothesis that $|\Delta| > 2$. Thus Δ consists of at most one root of the form $\varepsilon_1 \pm \varepsilon_i$ for $2 \leqslant i \leqslant \ell$. However, ε_1 is independent of any such set of roots, so there is no irreducible circuit of rank > 2 with apex ε_1.

Similarly, if $\alpha = \varepsilon_1 + \varepsilon_2$, then

$$\Phi_1(\alpha) = \{\varepsilon_1, \varepsilon_2\} \cup \{\varepsilon_1 \pm \varepsilon_i : 3 \leqslant i \leqslant \ell\} \cup \{\varepsilon_2 \pm \varepsilon_i : 3 \leqslant i \leqslant \ell\}.$$

There can be at most one root of each of the forms $\varepsilon_1 \pm \varepsilon_i$ and $\varepsilon_2 \pm \varepsilon_i$ in Δ (otherwise, ε_1 or ε_2 would be roots in the positive span of a proper subset of Δ). Also, Δ cannot contain both $\varepsilon_1 + \varepsilon_i$ and $\varepsilon_2 - \varepsilon_i$ (say), or we contradict Proposition 3.1, nor can it contain both $\varepsilon_1 + \varepsilon_i$ and $\varepsilon_2 + \varepsilon_i$ (say), since every point in the strictly positive span of Δ would have a positive ε_i-coordinate, whereas the ε_i-coordinate of α vanishes. Thus, Δ consists of at most one root of the form $\varepsilon_1 \pm \varepsilon_i$ or $\varepsilon_2 \pm \varepsilon_i$ for $3 \leqslant i \leqslant \ell$, along with possibly ε_1 or ε_2. All such sets of roots are linearly independent, so the only such sets that have α in their linear span are those that include ε_1 and ε_2. However, since α is in the positive span of ε_1 and ε_2 alone, this would force (α, Δ) to have rank 2, a contradiction.

Thus the claimed result holds for $\mathcal{B}_\ell$, and hence also for $\mathcal{C}_\ell$, since (α, Δ) is an irreducible circuit in Φ if and only if $(\alpha^\vee, \{\beta^\vee : \beta \in \Delta\})$ is irreducible in the dual root system.

The only remaining possibility (using the classification of root systems) is $\Phi = \mathcal{F}_4$.

In the usual realization, $\mathcal{F}_4$ consists of the 32 roots of $\mathcal{B}_4$, together with the 16 vectors of the form $(\pm\varepsilon_1 \pm \cdots \pm \varepsilon_4)/2$. For the orbit of short roots, we may select $\alpha = (\varepsilon_1 + \cdots + \varepsilon_4)/2$ as a representative apex, in which case

$$\Phi_1(\alpha) = \{\varepsilon_i : 1 \leqslant i \leqslant 4\} \cup \{\alpha - \varepsilon_i : 1 \leqslant i \leqslant 4\} \cup \{\varepsilon_i + \varepsilon_j : 1 \leqslant i < j \leqslant 4\}.$$

The base Δ cannot contain two of $\varepsilon_1, \ldots, \varepsilon_4$, or else a proper subset of Δ would have a root of the form $\varepsilon_i + \varepsilon_j$ in its positive span. Similarly, $\alpha - \varepsilon_1$ and $\alpha - \varepsilon_2$ have the root $\varepsilon_3 + \varepsilon_4$ in their positive span, so Δ may contain at most one of $\{\alpha - \varepsilon_i : 1 \leqslant i \leqslant 4\}$. Also, ε_i and $\alpha - \varepsilon_i$ have α in their positive span, so Δ cannot include both. Up to W-symmetry, we may therefore assume that the short roots in Δ are a subset of $\{\varepsilon_1, \alpha - \varepsilon_2\}$.

If there are no short roots in Δ (i.e., $\Delta \subset \{\varepsilon_i + \varepsilon_j : 1 \leqslant i < j \leqslant 4\}$), then we may assume (say) $\varepsilon_1 + \varepsilon_2 \in \Delta$. Note that α is on the hyperplane orthogonal to $\beta := \varepsilon_1 + \varepsilon_2 - \varepsilon_3 - \varepsilon_4$, whereas $\varepsilon_1 + \varepsilon_2$ is on the positive side of $\beta^\perp$, and $\varepsilon_3 + \varepsilon_4$ is the only long root in $\Phi_1(\alpha)$ on the negative side, so Δ must include $\varepsilon_3 + \varepsilon_4$ as well. However, α is in the positive span of these two roots, a contradiction.

Reflecting through the root hyperplane $\beta^\perp$ fixes α and interchanges ε_1 and $\alpha - \varepsilon_2$, so if there is only one short root in Δ, we may assume that it is ε_1. However, ε_1 is on the positive side of $\beta^\perp$, so Δ must include the only long root on the negative side: $\varepsilon_3 + \varepsilon_4$. By permuting the coordinates $\varepsilon_2, \varepsilon_3, \varepsilon_4$, it follows similarly that $\varepsilon_2 + \varepsilon_3, \varepsilon_2 + \varepsilon_4 \in \Delta$. The maximum possible rank is 4, so this forces $\Delta = \{\varepsilon_1, \varepsilon_2 + \varepsilon_3, \varepsilon_3 + \varepsilon_4, \varepsilon_4 + \varepsilon_2\}$. Conversely,

$$\alpha = (1/2)\varepsilon_1 + (1/4)(\varepsilon_2 + \varepsilon_3) + (1/4)(\varepsilon_3 + \varepsilon_4) + (1/4)(\varepsilon_4 + \varepsilon_2),$$

so these roots form the base of a positively pointed circuit with apex α; its diagram is the first one listed in Figure 3. It is necessarily irreducible, otherwise some proper subset of

Δ would be the base of an irreducible circuit of rank 2 or 3. However, Theorem 2.6 and the preceding argument establish that there are no irreducible circuits of rank 3 in any crystallographic root system, and it is easy to see that there is no root in the positive span of any two members of Δ. The type of this circuit is necessarily $\mathcal{F}_4$; if it were of type $\mathcal{B}_4$ or $\mathcal{C}_4$, it would have occurred as an irreducible circuit in those root systems as well.

If ε_1 and $\alpha-\varepsilon_2$ are the two short roots in Δ, then since both of these are on the positive sides of (and α is on) the hyperplanes orthogonal to $\varepsilon_1+\varepsilon_4-\varepsilon_2-\varepsilon_3$ and $\varepsilon_1+\varepsilon_3-\varepsilon_2-\varepsilon_4$, Δ must include the unique long roots in $\Phi_1(\alpha)$ on the negative sides: $\varepsilon_2+\varepsilon_3$ and $\varepsilon_2+\varepsilon_4$. Hence $\Delta=\{\varepsilon_1,\alpha-\varepsilon_2,\varepsilon_2+\varepsilon_3,\varepsilon_2+\varepsilon_4\}$. Conversely,

$$\alpha=(1/3)\varepsilon_1+(1/3)(\alpha-\varepsilon_2)+(1/3)(\varepsilon_2+\varepsilon_3)+(1/3)(\varepsilon_2+\varepsilon_4),$$

so these roots form the base of a positively pointed circuit with apex α; its diagram is the second one listed in Figure 3. By reasoning similar to the previous case, it follows that (α,Δ) is irreducible of type $\mathcal{F}_4$.

Finally, the irreducible circuits in $\mathcal{F}_4$ for which the apex is a long root may be obtained by dualization. In the dual root system, the roles of long and short are interchanged, but $\mathcal{F}_4$ is self-dual, so the diagrams of such circuits are obtained by reversing the orientations of the diagrams corresponding to irreducible circuits with a short apex. $\square$

Remark 3.4. The above argument proves slightly more: two irreducible circuits of type $\mathcal{F}_4$ belong to the same W-orbit if and only if they have the same diagram. Alternatively, noting that the Dynkin diagram of $\mathcal{F}_4$ has only trivial automorphisms, this fact follows by reasoning similar to Remark 2.11.

4. Non-crystallographic root systems

Aside from the trivial rank 2 cases, the only irreducible root systems that are not crystallographic (i.e., (3.2) holds, but not (3.1)) are $\mathcal{H}_3$ and $\mathcal{H}_4$. If the roots are normalized to have length $\sqrt{2}$, then for every pair α,β in $\mathcal{H}_3$ or $\mathcal{H}_4$, we have

$$\langle\alpha,\beta\rangle\in\{0,\,\pm1/g,\,\pm1,\,\pm g\},\ \text{or}\ \alpha=\pm\beta, \tag{4.1}$$

where $g=1.618...$ denotes the golden ratio.

Now suppose (α,Δ) is an irreducible circuit of rank >2 in $\mathcal{H}_3$ or $\mathcal{H}_4$. By inspecting the rank 2 root systems, one sees that there is a root in the positive linear span of two roots β and γ if and only if $\langle\beta,\gamma\rangle<0$ (cf. (3.2)) or $\langle\beta,\gamma\rangle=1/g$. Therefore,

$$\langle\beta,\gamma\rangle\in\{0,1,g\}\ \text{for all distinct}\ \beta,\gamma\in\Delta. \tag{4.2}$$

Conversely, a positively pointed circuit of rank 3 with a base that satisfies this condition is necessarily irreducible. Moreover, Theorems 2.6 and 3.3 establish that there are no other irreducible circuits of rank 3, so any such circuit is necessarily of type $\mathcal{H}_3$.

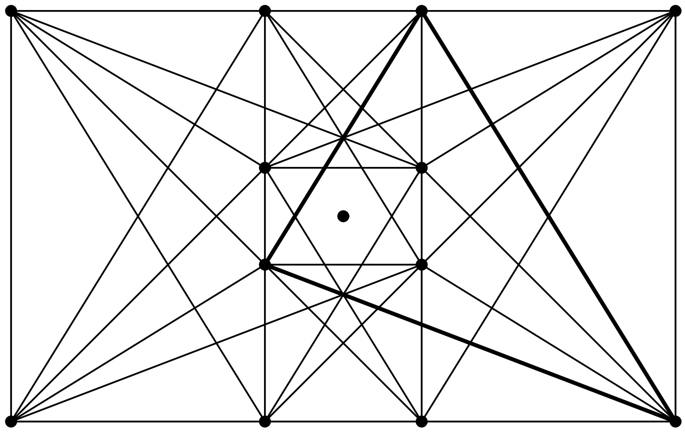

FIGURE 4: A hyperplane section of $\mathcal{H}_3$.

Proposition 3.1 and its proof are valid in the non-crystallographic case, so a further necessary condition on (α, Δ) is that $\langle \alpha, \beta \rangle > 0$ for all $\beta \in \Delta$. For any fixed choice of apex α in $\mathcal{H}_3$, there are 12 roots β with this property (not including α), and the intersections of the rays $\{c\beta : c > 0\}$ with the affine hyperplane $\langle \cdot, \alpha \rangle = 2$ are the 12 non-central points in Figure 4. We have drawn an edge between each pair of points whose corresponding roots satisfy (4.2), so that an irreducible circuit of rank 3 with apex α corresponds to a triangle with the center point in its interior and none of the other 12 points on its boundary edges. The thick edges in Figure 4 illustrate one such triangle; it is not hard to see that there are exactly 4 of these triangles, and that they are all equivalent up to symmetry. In fact, there are two roots in $\mathcal{H}_3$ that are orthogonal to α and each other; the reflections corresponding to these two roots generate the four symmetries in Figure 4.

Thus we conclude that there do exist irreducible circuits of type $\mathcal{H}_3$ in $\mathcal{H}_3$ and that they form a single orbit relative to the group W generated by reflection symmetries. The geometry of these circuits is such that the base roots are mutually orthogonal, and it is easy to check that (unordered) orthogonal triples of roots in $\mathcal{H}_3$ form a single W-orbit, so every triple of orthogonal roots in $\mathcal{H}_3$ is the base of an irreducible circuit. (In fact, every triple of orthogonal roots is the base of three distinct but isometric irreducible circuits.)

The root system $\mathcal{H}_4$ contains $\mathcal{H}_3$ as a root subsystem, so there exist positively pointed circuits in $\mathcal{H}_4$ with the same geometry as the irreducible $\mathcal{H}_3$-circuits. Since $\mathcal{H}_3$ is the largest rank 3 root system, these circuits are necessarily irreducible in $\mathcal{H}_4$ as well. Moreover, one can check that unordered orthogonal triples form a single reflection-orbit in $\mathcal{H}_4$, so every triple of orthogonal roots in $\mathcal{H}_4$ is the base of an irreducible circuit of type $\mathcal{H}_3$.

A consequence of the previous observation is that a positively pointed circuit (α, Δ) of rank 4 in $\mathcal{H}_4$ is irreducible if and only if (4.2) holds and there is no triple of orthogonal roots in Δ. Note that although $\mathcal{D}_4$ is a root subsystem of $\mathcal{H}_4$, the only irreducible circuit of this type has a base with four mutually orthogonal roots (see Figure 1), and thus cannot

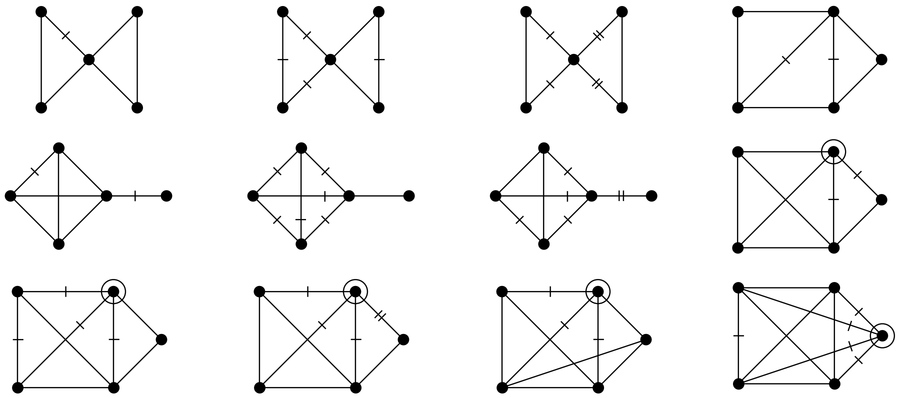

FIGURE 5: Diagrams of irreducible circuits of type $\mathcal{H}_4$.

be irreducible in $\mathcal{H}_4$. Geometry alone eliminates $\mathcal{F}_4$ as a possibility, so any irreducible circuit of rank 4 in $\mathcal{H}_4$ is necessarily of type $\mathcal{H}_4$.

In a machine computation, we fixed a root $\alpha \in \mathcal{H}_4$ and generated orbit representatives for all 4-subsets of $\{\beta \in \mathcal{H}_4 : \langle \beta, \alpha \rangle > 0\}$ satisfying (4.2). After rejecting subsets containing orthogonal triples, we found that 12 of these were bases of positively pointed circuits with apex α; by the preceding discussion, each is necessarily irreducible of type $\mathcal{H}_4$. The diagrams of these 12 circuits are displayed in Figure 5, using the convention that pairs of roots with inner product g (resp., $1/g$) are encoded by edges with a single (resp., double) mark. In diagrams with more than one node of degree 4, the apex is circled.

References

[**1**] C. Athanasiadis, Generalized Catalan numbers, Weyl groups and arrangements of hyperplanes, *Bull. London Math. Soc.* **36** (2004), 294–302.

[**2**] A. Björner, M. Las Vergnas, B. Sturmfels, N. White and G. M. Ziegler, "Oriented matroids" (2nd ed.), Cambridge Univ. Press, Cambridge, 1999.

[**3**] P. J. Cameron, J.-M. Goethals, J. J. Seidel, and E. E. Shult, Line graphs, root systems, and elliptic geometry, *J. Algebra* **43** (1976), 305–327.

[**4**] P. Orlik and H. Terao, "Arrangements of hyperplanes," Springer, Berlin 1992.

[**5**] A. Postnikov and R. Stanley, Deformations of Coxeter hyperplane arrangements, *J. Combin. Theory Ser. A* **91** (2000), 544–597.

[**6**] J.-Y. Shi, Sign types corresponding to an affine Weyl group, *J. London Math. Soc. (2)* **35** (1987), 56–74.