Four-dimensional reflection groups and electrostatics

Maxim Olshanii [a,*], Yuri Styrkas [a,c], Dmitry Yampolsky [a],
Vanja Dunjko [a], Steven G. Jackson [b]

[a] *Department of Physics, University of Massachusetts Boston, Boston, MA 02125, USA*
[b] *Department of Mathematics, University of Massachusetts Boston, Boston, MA 02125, USA*
[c] *Princeton High School, Princeton, NJ 08540, USA*

ARTICLE INFO

ABSTRACT

We present a new class of electrostatics problems that are exactly solvable by adding finitely many image charges. Given a charge at some location inside a cavity bounded by up to four conducting grounded segments of spheres: if the spheres have a symmetry derived via a stereographic projection from a 4D finite reflection group, then this is a solvable generalization of the familiar problem of a charge inside a spherical cavity. There are 19 three-parametric families of finite groups formed by inversions relative to at most four spheres, each member of each family giving a solvable problem. We solve a sample problem that derives from the reflection group $\mathbf{D}_4$ and requires 191 image charges.

1. Introduction

In electrostatics, an important category of exactly solvable problems consists of those that can be solved by the method of images [1,2]. We will restrict ourselves to the subcategory where only finitely many image charges are needed and where the boundary conditions correspond to grounded conductors. So far, to the best of our knowledge [3–6], the list of problems known to belong to this fundamental subcategory has been unchanged since at least Maxwell's 1873 treatise [7], where it appears as follows (Art. 167, p. 206):

---

 * Corresponding author.
   E-mail addresses: Maxim.Olchanyi@umb.edu (M. Olshanii), ystyrkas@gmail.com (Y. Styrkas),
Dmitry.Yampolsky@umb.edu (D. Yampolsky), vanja.dunjko@umb.edu (V. Dunjko), Steven.Jackson@umb.edu (S.G. Jackson).

The method of electrical images may be applied to any space bounded by plane or spherical surfaces all of which cut one another in angles which are submultiples of two right angles. In order that such a system of spherical surfaces may exist, every solid angle of the figure must be trihedral, and two of its angles must be right angles, and the third either a right angle or a submultiple of two right angles. Hence the cases in which the number of images is finite are—(1) A single spherical surface or a plane. (2) Two planes, a sphere and a plane, or two spheres intersecting at an angle π/n. (3) These two surfaces with a third, which may be either plane or spherical, cutting both orthogonally. (4) These three surfaces with a fourth cutting the first two orthogonally and the third at an angle π/n'. Of these four surfaces at least one must be spherical.

According to Ref. [4] (p. 302), some of the best-known expositions of the method of images are the textbooks by Jeans [5] and Smythe [6]. Remarkably, we were unable to find any post-Maxwell sources that treat cases (3) and (4), with the sole exception of Jeans, who only mentions the case of three planes all intersecting at right angles (Art. 211, p. 188).

In the present work, we will substantially enlarge Maxwell's list. Our workhorse is the well-developed mathematical theory of *finite reflection groups* [8]. For our purposes, we will need only a few basic facts from this theory. Nevertheless, since the theory is perhaps unfamiliar to many readers, we will explain these relevant facts in detail. Reflection groups also underlie the exact solvability, via Bethe ansatz, of certain models in many-body 1D quantum mechanics [9–11].

Finite reflection groups concern planar mirror reflections. The connection to spherical surfaces is provided by the stereographic projection, through the following fact: if we stereographically project a planar mirror reflection in N dimensions, we obtain a sphere inversion in $N-1$ dimensions.

We will show that the method of images can be used to find the field induced by a point charge (placed at an arbitrary location inside a grounded conducting cavity) if the shape of the cavity has the following properties. First, the walls of the cavity should be segments of up to four spheres (which may have an infinite radius, i.e. may be planes). Second, the spheres in question should be (and we will devote a good portion of the text explaining what all of this means) stereographic projections, from 4D to 3D, of great hypercircles that lie on the planes that are the generating mirrors of a 4D finite reflection group.

The locations and values of the image charges are then derived, in effect, by a repeated application of the familiar formula used to find the image charge that solves the problem of a point charge next to a conducting sphere. The properties of reflection groups (and of the stereographic projection) guarantee that there will be only finitely many image charges and that the charge values can be unambiguously assigned.

2. Outline

The construction we have just briefly described includes as special cases many (but not all) elements from Maxwell's list. We will introduce the basic concepts of finite reflection groups in Section 3, illustrating them in Section 3.1 with the help of a case that is on the Maxwell's list. We will proceed in Section 3.2 to explain some further properties of reflection groups, including their direct products (Section 3.3) and Coxeter diagrams (Section 3.4). We will then be ready to present, in Section 4, the first example of a solvable problem that is not on Maxwell's list. This will use nontrivial properties of reflection groups, but it will be entirely in 3D, so the stereographic projection will not be used yet. The specific example is treated in Sections 4.1 and 4.2. In order to advance to 4D finite reflection group, we will need to discuss the stereographic projection, which we do in Section 5. At that point, we will be ready to present, in Section 6, an outline of our full 4D procedure. This will also be the point, in Section 6.1, where we contrast our method to that of the Kelvin transform. In Section 7, we review the sphere inversion for a single spherical cavity, and proceed to explain, in Section 8, how to construct solvable problems using piecewise-spherical conducting cavities. Section 9 gives computational details. In Section 10, we provide a worked nontrivial example of our procedure. Finally, we conclude with a summary and outlook in Section 11.

3. A crash course in finite reflection groups (and what they have to do with the method of images)

We assume the reader is familiar with the two simplest examples of the method of images, each of which requires just one image charge: a point charge next to a grounded conducting plane, and a point charge either inside or outside a grounded conducting sphere. In the first example, the location of the image charge is the *mirror reflection* (with respect to the conducting plane) of the original charge; in the second, it is the *spherical inversion* (relative to the conducting sphere) of the original charge.

For now, we will be considering a single point charge next to a *system of planar conductors*, a class of problems that uses only mirror reflections; we will return to spherical inversions in Section 6, after we have discussed the stereographic projection in Section 5.

When it comes to systems of planar conductors, the key question is under what circumstances we get *finitely* many images after we have reflected the original charge – and then, iteratively, also reflected each newly obtained image charge – relative to each of the planes of the conductors. The definite answer to this question comes from the mathematical theory of finite reflection groups. All mathematical facts concerning that theory will come from Ref. [8].

When a line or a plane (or, in higher dimensions, a hyperplane) is used to reflect objects, we will refer to it as a 'mirror'. It will be convenient to identify the mirror by a unit vector $\hat{\alpha}$ that is normal to it. Thus, when we reflect a vector $\vec{\lambda}$ with respect to the plane characterized by $\hat{\alpha}$, the result is given by $\vec{\lambda} - 2(\vec{\lambda} \cdot \hat{\alpha})\hat{\alpha}$. This same formula holds in any number of spatial dimensions.

Let us now consider a simple example in 2D.

3.1. A simple example: group $\mathbf{A}_2$ (a.k.a. $\mathbf{I}_2(3)$)

Consider the following problem, which belongs to case (2) from Maxwell's list: two planes intersecting at an angle π/n, with $n = 3$; see Fig. 1. Our region of interest is the orange wedge, which is bounded by two conducting half-planes. A charge (the solid black dot) is placed in this wedge, at some arbitrary, generic location. This location defines the radius-vector $\vec{r}_0$. The full planes that contain the conducting half-planes will be called the *generating mirrors*, shown in solid green and dashed red in the Figure. (The term 'generating mirrors' is a technical term from the theory of finite reflection groups. In this section, we will introduce several additional technical terms, just to see to what sorts of objects they refer. Then, in the following sections, we will explain in more detail what these terms really mean.) The unit normal vectors $\hat{\alpha}_1$ and $\hat{\alpha}_2$ are chosen so that a vector $\vec{r}$ (whose initial point is the origin) lies in the orange wedge if and only if the projections $\vec{r} \cdot \hat{\alpha}_1$ and $\vec{r} \cdot \hat{\alpha}_2$ are both positive. When the unit normal vectors are chosen this way, they are called the *simple roots* of the reflection group.

Now we start iteratively reflecting $\vec{r}_0$ and its images with respect to the generating mirrors. More precisely, let S_{old} be a set in which we will be collecting the distinct image charges. Let us initially place just the vector $\vec{r}_0$ in that set. Let S_{new} be another set, which will be a temporary storage for the outcomes of reflections during each iteration; let it initially be empty. Now we iterate the following process: for each element of S_{old}, we reflect that element once through each generating mirror, adding the outcome of each reflection to S_{new}. Once this is done, we check to see if S_{new} contains any vectors that are not already in S_{old}. If yes, we add those vectors to S_{old} (i.e. we set $S_{\text{old}} = S_{\text{old}} \cup S_{\text{new}}$), make the set S_{new} empty again, and repeat the procedure. If no, we stop. In the present case, we will find that once we generate the five image charges, so that S_{old} contains six charges total, the process stops: a reflection of any one of these charges with respect to either generating mirror produces one of the charges that is already in S_{old}. One can prove that if the angle between the generating mirrors were something other than π/n, with n a positive integer, the process would never stop: there would be an infinity of image charges.

The six images are called the *orbit* of the initial charge under the *action* of the reflection group. In other words, every image charge is connected to the original charge by the action of an element of the group. Thus, this particular group has six elements (including the identity, which maps the initial charge to itself); we say that the *order* of the group is six.
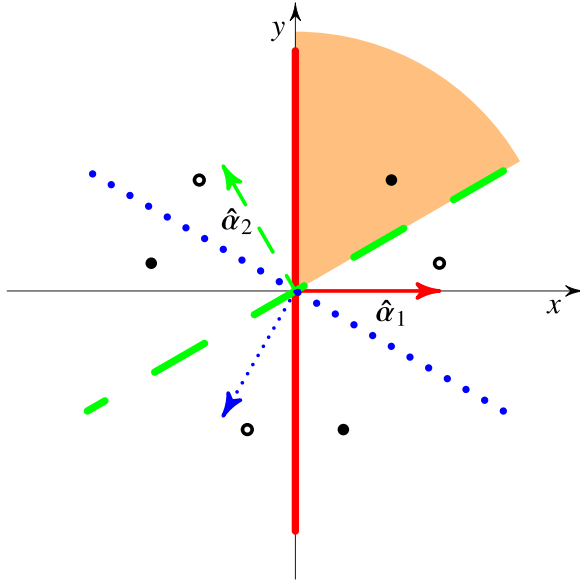
Fig. 1. Two planes intersecting at an angle of $\pi/3$ and the associated system of images. In the terminology of finite reflection groups, this is a group that belongs to two families and can be equivalently designated as either $\mathbf{A}_2$ or $\mathbf{I}_2(3)$. The solid green and dashed red lines are the *generating mirrors* of the group. The vectors $\hat{\boldsymbol{\alpha}}_1$ and $\hat{\boldsymbol{\alpha}}_2$ are the corresponding normal vectors, whose orientation is chosen so that they constitute a system of *simple roots* this reflection group. The orange wedge is the *principal chamber* of the group, which is where the original charge is located, denoted by a solid black dot; the final solution will give the electric field in this wedge, assuming that the portions of the red and green lines that bound it are grounded conductors. The group has another mirror, in dotted blue, whose action can be replicated by sequences of reflections through the two generating mirrors. The set of all mirrors partitions the space into six chambers, where six is the order of the reflection group. There is one image charge in each non-principal chamber. An image charge is either the same as the original or the negative of it, and is respectively denoted either by a solid dot or by an open circle. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

Note that there is a third plane, shown in dotted blue, that can also be considered a mirror. Its location can be obtained by reflecting one of the two generating mirrors with respect to the other. (Once we have the set of these three mirrors, reflecting any one of them with respect to any of the other ones produces a mirror that is already in the set; we will call this set of three mirrors a *full set* of mirrors.) As far as generating images, this third mirror does not add anything new: a reflection of any point through this third mirror is equivalent to some sequence of reflections with respect to the two generating mirrors. In particular, adding this extra mirror does not result in any extra image charges. However, it is useful to consider the full system of mirrors, because it leads to the concept of *chambers*. Note that the three mirrors in the Figure partition the ambient (2D) space into six congruent regions, which are the chambers. It turns out there is one charge (image or original) per chamber. Thus the order of the group is equal to the number of chambers. Note that all three mirrors intersect at a point, which also means that all the chambers meet at a point.

There are several additional facts to notice. First, not every element of the reflection group is a reflection—in other words, some sequences of reflections with respect to the mirrors are not equivalent to any single reflection with respect to some mirror in the full set. Consider, for example, the solid dot in the lower right chamber bounded by the solid green and dotted blue mirrors. It cannot be obtained from the original charge by any single reflection; in fact, what links these two is a rotation by $2\pi/3$. Second, all the charges are at the same distance from the origin, and so lie on a circle. Third, if we rotate the whole system of mirrors, none of the facts noted above change. Therefore, as far as reflection groups are concerned, we will think of all the rotated versions of this system of mirrors as being equivalent.

3.2. A general description of reflection groups

The preceding discussion suggests a procedure for how one could, at least in principle if not in practice, identify finite reflection groups in any number of spatial dimensions. We will now describe that procedure in general terms. Then, in Section 4.1, we will go through it again on the example of a specific 3D group.

Pick a dimensionality D of the ambient space ($D = 2, 3, 4, 5, \ldots$). Pick some finite number of hyperplanes in that space to serve as mirrors. Let us call these the 'seed' mirrors (which is *not* a technical term; these seed mirrors may or may not turn out to be what we will later call the *generating* mirrors, which *is* a technical term). Finally, pick a vector $\vec{r}_0$. Introduce the sets S_{old} and S_{new}. Initialize them so that S_{old} contains just $\vec{r}_0$, while S_{new} is empty. Now iterate the following procedure. First, reflect each element of S_{old} with respect to each of the seed mirrors, adding the outcome to S_{new}. After this has been done, check if S_{new} contains any elements that are not in S_{old} already. If yes, add all such elements to S_{old} (i.e. set $S_{\text{old}} = S_{\text{old}} \cup S_{\text{new}}$), reset S_{new} so that it is again empty, and repeat; if no, stop. For almost all choices of the seed mirrors, the procedure will never stop and the size of S_{old} will grow without bound (for almost all choices of $\vec{r}_0$). But for some special choices of the seed mirrors, eventually the procedure stops no matter how we pick $\vec{r}_0$. At that point, S_{old} will contain a finite number of vectors such that, if we reflect any of them with respect to any seed mirror, the result will already be in S_{old}. We say that such a system of seed mirrors produces a *finite reflection group*. The *order* of the group is the maximal size that S_{old} can have no matter how $\vec{r}_0$ is picked. (In fact, a generic choice of $\vec{r}_0$ results in the maximal size of S_{old}.)

Once we have identified a system of seed mirrors that produces a finite reflection group, it may turn out that this reflection group has extra mirrors, beyond the seed ones with which we started. The *full* set of mirrors of the reflection group can be obtained by initializing S_{old} with our seed mirrors (S_{new} again starts as an empty set), and then iterating as follows. In each iteration, we first reflect each element of S_{old} (which is a mirror) with respect to all the remaining elements of S_{old} (which are also mirrors). The result is another mirror, which we add to S_{new}. After this has been done, we check if there are any mirrors in S_{new} that are not already in S_{old}. If yes, add those mirrors to S_{old}, reset S_{new} so that it is again empty, and repeat; if no, stop. If the seed mirrors produce a finite reflection group, this process is guaranteed to stop, at which point S_{old} contains the full set of mirrors of the reflection group: if we reflect any element of S_{old} with respect to any of the other elements, the result is a mirror that is itself already in S_{old}. Note that the number of mirrors in the full set of mirrors of a reflection group is, in general, smaller than the order of the group—sometimes vastly so. The mirrors of the full set all intersect at a single point.

A key property of the full set of mirrors is that it partitions the ambient space into disjoint congruent regions called *chambers*. The original vector $\vec{r}_0$ is in one chamber, and all the other chambers contain exactly one image. Thus, there are as many images as there are chambers, and their number is called the *order* of the group.

Note that if a system of mirrors of a group is rotated as a whole by any amount, the result still constitutes the same finite reflection group. This property will be important below.

When the originally chosen seed mirrors do generate a finite reflection group, it may be possible to discard some of the seed mirrors and yet have the iterative procedure (using just the retained mirrors) still generate the same set of images as before. More generally, we may start with the full set of mirrors, and do the following procedure. We start discarding the mirrors one by one; whenever we discard one, we test to see if the remaining ones still generate all the images. If not, we reinstate the mirror; if yes, we leave it out. Then we continue discarding mirrors, but any mirror we had to reinstate at any previous step we do not attempt to discard again. Eventually we will end up only with mirrors that we had to reinstate at some point, and we stop. The result is a system of *generating mirrors*. Depending on the order in which we proceeded to discard the mirrors from the full set, we will in general end up with different systems of generating mirrors. However, all them are geometrically congruent to each other. In particular, given a full set of mirrors of a reflection group, all systems of generating mirrors have the same number of mirrors, which is called the *rank* of the reflection group; we will denote it by ℓ. Of course, given a system of generating mirrors, all the other mirrors of the group can be obtained by initializing S_{old} with just the generating mirrors, and iterating as before.

For each mirror in a full set of mirrors, there are two unit normal vectors, which are negatives of each other. These vectors are called the *roots*, and we see that there are twice as many roots as there are mirrors. (Actually, in the full mathematical theory, it is useful to consider normal vectors which are normalized to certain special values that are not necessarily unity. For our purposes, however, we may assume that all roots are normalized to one.)

One consequence of what we have said so far is that in order to completely characterize a reflection group, it is enough to know (at least one set of) its generating mirrors; thus, it is enough to know the roots (unit normal vectors) of these generating mirrors. There are 2ℓ of them, two per generating mirror. We will now show a useful way to pick just one root per generating mirror, giving a set of roots we will call the *simple roots*. For any system of ℓ generating mirrors, at least one chamber will be enclosed entirely by the generating mirrors in that set (if there is more than one such chamber, just pick one of them); this will be the *principal chamber*. It is always possible to pick, from the 2ℓ roots characterizing the generating mirrors, a set of ℓ of them, $\hat{\boldsymbol{\alpha}}_i$, $i = 1, \ldots, \ell$, one for each generating mirror, in such a way that a vector $\vec{r}$ belongs to the principal chamber if and only if $\vec{r} \cdot \hat{\boldsymbol{\alpha}}_i > 0$ for all $i = 1, \ldots, \ell$. In that case, the set of vectors $\hat{\boldsymbol{\alpha}}_i$ is called a set of *simple roots* of the reflection group, usually denoted by Δ. It is always a linearly independent set, but generally not an orthogonal one.

It follows that the rank ℓ of a reflection group is the lowest dimension that the ambient space must have in order for that reflection group to be realizable in it; the simple roots are a basis (though not an orthonormal one) for that minimally-dimensional ambient space. Since we will ultimately be dealing with the 4D space, we can use any reflection group of rank 4 or lower.

All finite reflection groups, in all spatial dimensions, have been classified and their properties are very well understood. In particular, we know all finite reflection groups of rank 4 or lower. It turns out there are 19 of them, and we list all of them in Tables 1 and 2. As is clear from the previous discussion, each reflection group is completely defined by its simple roots. There are standard choices for how to write the simple roots of all possible reflection groups. They can be looked up in any number of sources, and in particular in Ref. [8]. In Tables 1 and 2, for each reflection group we give a set of its simple roots, but we should note that the way we present them is not the standard way of writing them. Why we did this will become clear below. In addition, for each group we also give its Coxeter diagram (a concept we will explain below), the number of mirrors, and the order.

Referring back to Maxwell's list, the reflection groups corresponding to two planes meeting at an angle of π/n are called $\mathbf{I}_2(n)$. Four of them have aliases, because they also count as members of other families of groups; see the caption of Table 1. In the group notation, the subscript always denotes the rank. Note that the mirrors of $\mathbf{I}_2(n)$ are the planes of symmetry of a regular n-gon.

3.3. Direct products of reflection groups; reducible and irreducible groups

When we consider the problem in Fig. 1 (or indeed any problem where two planes meet at an angle π/n), it is clear that the problem remains solvable if we add another conducting plane parallel to the xy-plane: all we have to do is reflect all the charges about that plane and invert their signs. The resulting system of planes is also a reflection group, denoted by $\mathbf{I}_2(n) \times \mathbf{A}_1$. Here $\mathbf{A}_1$ is the usual notation for the rank-1 group consisting of a single mirror, and '$\times$' denotes the *direct product* of groups. We will now explain, through an explicit construction, what that is in general. Namely, given two arbitrary reflection groups, for which we assume we know their simple roots, we will show how to construct a system of simple roots for their direct product.

Note that a group is called *irreducible* if it cannot be written as a direct product of two other groups; if it can, then it is *reducible*. In Tables 1 and 2, all the groups that are *not* written as direct products are in fact irreducible.

Here is how to construct a direct product of reflection groups. Let g_1 be a finite reflection group of rank r_1, order o_1, and with a total of m_1 mirrors, and let g_2 be a reflection group whose corresponding numbers are r_2, o_2, and m_2. To construct the system of mirrors, simple roots, and chambers for $g_1 \times g_2$, we will need a space of at least $r_1 + r_2$ dimensions; for definiteness, let us choose $\mathbb{R}^{r_1+r_2}$. The simple roots of g_1 define an r_1-dimensional hyperplane in that space, while the simple roots of

Table 1

The 19 finite reflection groups of rank 4 or less, part 1 of the list. The subscript to the (bold capital) letter name of the group denotes the rank. Thus, the rank of the group is the sum of the subscripts. For example, both $\mathbf{H}_3$ and $\mathbf{I}_2(m) \times \mathbf{A}_1$ have rank 3, and so they each have three simple roots. The variables m_1, m_2, and m take integer values of 3 or more. In groups that contain $\mathbf{I}_2(m)$, $c_m = \cos\frac{\pi}{m}$. In $\mathbf{H}_4$ and groups that contain $\mathbf{H}_3$, $a_1 = -\frac{1}{4}\left(1+\sqrt{5}\right)$, $b_1 = \frac{1}{2}\sqrt{\frac{1}{2}\left(5-\sqrt{5}\right)}$, $a_2 = -\sqrt{\frac{1}{10}\left(5+\sqrt{5}\right)}$, $b_2 = \frac{1}{2}\sqrt{1+\frac{2}{\sqrt{5}}}\left(3-\sqrt{5}\right)$. Further, in $\mathbf{H}_4$, $c = \frac{1}{4}\sqrt{1+\frac{2}{\sqrt{5}}}\left(\sqrt{5}-5\right)$ and $d = \frac{1}{2}\sqrt{\frac{1}{2}\left(3-\sqrt{5}\right)}$. Note the following: i. $\mathbf{A}_2$, $\mathbf{B}_2$, $\mathbf{H}_2$, and $\mathbf{G}_2$ are included respectively as $\mathbf{I}_2(3)$ through $\mathbf{I}_2(6)$; ii. $\mathbf{D}_n$ only exists for $n \geqslant 4$; iii. $\mathbf{F}_n$ only exists for $n = 4$; and, iv. $\mathbf{H}_n$ only exists for $n = 2, 3, 4$.

group	Coxeter diagram	Num. of mirrors	order	simple roots	
$\mathbf{A}_4$	●—●—●—●	10	120	$(1,0,0,0)$　$\left(0,-\frac{1}{\sqrt{3}},\sqrt{\frac{2}{3}},0\right)$	$\left(-\frac{1}{2},\frac{\sqrt{3}}{2},0,0\right)$　$\left(0,0,-\frac{\sqrt{3/2}}{2},\frac{\sqrt{5/2}}{2}\right)$
$\mathbf{B}_4(=\mathbf{C}_4)$	●—●—●—⁴—●	16	384	$(1,0,0,0)$　$\left(0,-\frac{1}{\sqrt{3}},\sqrt{\frac{2}{3}},0\right)$	$\left(-\frac{1}{2},\frac{\sqrt{3}}{2},0,0\right)$　$\left(0,0,-\frac{\sqrt{3}}{2},\frac{1}{2}\right)$
$\mathbf{H}_4$	●—⁵—●—●—●	60	14 400	$(1,0,0,0)$　$(0,a_2,b_2,0)$	$(a_1,b_1,0,0)$　$(0,0,c,d)$
$\mathbf{F}_4$	●—●—⁴—●—●	24	1152	$(1,0,0,0)$　$\left(0,-\sqrt{\frac{2}{3}},\frac{1}{\sqrt{3}},0\right)$	$\left(-\frac{1}{2},\frac{\sqrt{3}}{2},0,0\right)$　$\left(0,0,-\frac{\sqrt{3}}{2},\frac{1}{2}\right)$
$\mathbf{D}_4$	(branching diagram)	12	192	$(1,0,0,0)$　$\left(0,-\frac{1}{\sqrt{3}},\sqrt{\frac{2}{3}},0\right)$	$\left(-\frac{1}{2},\frac{\sqrt{3}}{2},0,0\right)$　$\left(0,-\frac{1}{\sqrt{3}},-\frac{1}{\sqrt{6}},\frac{1}{\sqrt{2}}\right)$
$\mathbf{A}_3 \times \mathbf{A}_1$	●—●—● ●	7	48	$(1,0,0,0)$　$\left(0,-\frac{1}{\sqrt{3}},\sqrt{\frac{2}{3}},0\right)$	$\left(-\frac{1}{2},\frac{\sqrt{3}}{2},0,0\right)$　$(0,0,0,1)$
$\mathbf{B}_3(=\mathbf{C}_3) \times \mathbf{A}_1$	●—●—⁴—● ●	10	96	$(1,0,0,0)$　$\left(0,-\sqrt{\frac{2}{3}},\frac{1}{\sqrt{3}},0\right)$	$\left(-\frac{1}{2},\frac{\sqrt{3}}{2},0,0\right)$　$(0,0,0,1)$
$\mathbf{H}_3 \times \mathbf{A}_1$	●—⁵—●—● ●	16	240	$(1,0,0,0)$　$(0,a_2,b_2,0)$	$(a_1,b_1,0,0)$　$(0,0,0,1)$
$\mathbf{I}_2(m_1) \times \mathbf{I}_2(m_2)$	●—m_1—● ●—m_2—●	$m_1 + m_2$	$4\,m_1 m_2$	$(1,0,0,0)$　$(0,0,1,0)$	$\left(-c_{m_1},\sqrt{1-c_{m_1}^2},0,0\right)$　$\left(0,0,-c_{m_2},\sqrt{1-c_{m_2}^2}\right)$
$\mathbf{I}_2(m) \times \mathbf{A}_1 \times \mathbf{A}_1$	●—m—● ● ●	$m + 2$	$8m$	$(1,0,0,0)$　$(0,0,1,0)$	$\left(-c_m,\sqrt{1-c_m^2},0,0\right)$　$(0,0,0,1)$

g_2 can always be chosen in $\mathbb{R}^{r_1+r_2}$ so that they are all orthogonal to all simple roots of g_1. An obvious way to do this is to take the simple roots of g_1 (which can always be taken as r_1-dimensional row vectors) and append r_2 zeros to these row vectors; then we similarly *pre*-pend r_1 zeros to the r_2-dimensional row vectors that are the simple roots of g_2. Once such a choice is made, collectively the simple roots of g_1 and g_2 will be the simple roots of $g_1 \times g_2$, which thus has rank $r_1 + r_2$, and whose total number of mirrors is $m_1 + m_2$. The order of the direct product turns out to be the product of the orders, $o_1 o_2$. To construct the direct product of more than two (but still finitely many) reflection groups, proceed iteratively: first construct the direct product of two of them, then the direct product of the result with the third, and so on. It turns out that the order in which we do this does not matter ('up to isomorphism' of groups, as mathematicians say; for us, we may say that changing the order in which we construct the product of groups is equivalent to changing one's mind about which spatial dimension is called dimension 1, which one is called dimension 2, etc.).

Table 2

The 19 finite reflection groups of rank 4 or less, part 2 of the list. As before, in groups that contain $\mathbf{I}_2(m)$, $c_m = \cos\frac{\pi}{m}$. The final entry in the table, $\mathbf{A}_1$, corresponds to the classic problem with one spherical cavity—the very inspiration for this paper.

group	Coxeter diagram	Num. of mirrors	order	simple roots
$\mathbf{A}_1 \times \mathbf{A}_1 \times \mathbf{A}_1 \times \mathbf{A}_1$	• • • •	4	16	$(1,0,0,0)$ $(0,1,0,0)$ $(0,0,1,0)$ $(0,0,0,1)$
$\mathbf{A}_3$	•—•—•	6	24	$(1,0,0,0)$ $\left(-\frac{1}{2}, \frac{\sqrt{3}}{2}, 0, 0\right)$ $\left(0, -\frac{1}{\sqrt{3}}, \sqrt{\frac{2}{3}}, 0\right)$
$\mathbf{B}_3(=\mathbf{C}_3)$	•—•—⁴—•	9	48	$(1,0,0,0)$ $\left(-\frac{1}{2}, \frac{\sqrt{3}}{2}, 0, 0\right)$ $\left(0, -\sqrt{\frac{2}{3}}, \frac{1}{\sqrt{3}}, 0\right)$
$\mathbf{H}_3$	•—⁵—•—•	15	120	$(1,0,0,0)$ $(a_1, b_1, 0, 0)$ $(0, a_2, b_2, 0)$
$\mathbf{I}_2(m) \times \mathbf{A}_1$	•—ᵐ—• •	$m+1$	$4m$	$(1,0,0,0)$ $\left(-c_m, \sqrt{1-c_m^2}, 0, 0\right)$ $(0,0,1,0)$
$\mathbf{A}_1 \times \mathbf{A}_1 \times \mathbf{A}_1$	• • •	3	8	$(1,0,0,0)$ $(0,1,0,0)$ $(0,0,1,0)$
$\mathbf{I}_2(m)$	•—ᵐ—•	m	$2m$	$(1,0,0,0)$ $\left(-c_m, \sqrt{1-c_m^2}, 0, 0\right)$
$\mathbf{A}_1 \times \mathbf{A}_1$	• •	2	4	$(1,0,0,0)$ $(0,1,0,0)$
$\mathbf{A}_1$	•	2	2	$(1,0,0,0)$

In the example of $\mathbf{I}_2(n) \times \mathbf{A}_1$, the 'hyperplane' was chosen to be the *xy*-plane. Group $\mathbf{I}_2(n)$ is of rank 2 with simple roots $(1, 0)$ and $(-1/2, \sqrt{3}/2)$, while $\mathbf{A}_1$ is of rank 1 with simple root (1). So the rank of $\mathbf{I}_2(n) \times \mathbf{A}_1$ is $2 + 1 = 3$. It is of order $4n$ and has a total of $n + 1$ mirrors. We obtain its simple roots by appending $3 - 2 = 1$ zeros to the simple roots of $\mathbf{I}_2(n)$, obtaining $(1, 0, 0)$ and $(-1/2, \sqrt{3}/2, 0)$, and by prepending $3 - 1 = 2$ zeros to the simple root of $\mathbf{A}_1$, obtaining $(0, 0, 1)$.

3.4. Coxeter diagrams

At this point we can introduce *Coxeter diagrams*, which are a basic tool in the theory of finite reflection groups. In brief: since the relative orientation of simple roots is all that matters, simple roots (and thus the finite reflection group as a whole) can be efficiently represented in a diagrammatic form, as follows. Each simple root is represented by a node (i.e. a dot). In most finite reflections groups, an overwhelming majority of simple roots are mutually orthogonal, and so we connect two nodes by an edge only if the corresponding simple roots are *not* mutually orthogonal. Each edge therefore corresponds to two mirrors intersecting at a dihedral angle which is not $\pi/2$. The angle is, however, always of the form π/m, with m an integer. We place m above the edge, signifying that the angle between the mirrors is π/m. It turns out that $m = 3$ is by far the most common label, and so by convention we leave the edge unlabeled if $m = 3$ (whereas, as we said, we omit drawing the edge altogether if $m = 2$). The result is called a *Coxeter diagram*.

Using simple roots, the label m (which we place above the edge connecting the nodes that represent two simple roots $\vec{\alpha}_i$ and $\vec{\alpha}_j$) is given by $m = \pi/\arccos\left(-\vec{\alpha}_i \cdot \vec{\alpha}_j\right)$, where we assumed the simple roots are normalized.

A group is irreducible just in case its Coxeter diagram is connected. It should now be easy to verify that the Coxeter diagrams in the two tables do correspond to the simple roots of the group. For example, for $\mathbf{A}_1$, $\bullet$; for $\mathbf{I}_2(m)$, $\bullet\overset{m}{-}\bullet$; for $\mathbf{I}_2(n) \times \mathbf{A}_1$, $\bullet\overset{m}{-}\bullet\ \bullet$; for $\mathbf{A}_3$, $\bullet\!-\!\bullet\!-\!\bullet$; etc. Let us look at the last example in detail. Note the inner products of the simple roots of $\mathbf{A}_3$: $\vec{\alpha}_1 \cdot \vec{\alpha}_2 = \vec{\alpha}_2 \cdot \vec{\alpha}_3 = -1/2$ and $\vec{\alpha}_1 \cdot \vec{\alpha}_3 = 0$. This corresponds to $m = 3,\ 3,\ 2$. We see that there are indeed two unlabeled edges corresponding to the two instances of $m = 3$, and that the first and last nodes are not connected by an edge, corresponding to $m = 2$.

4. First extension of Maxwell's list of solvable problems: irreducible finite reflection groups of rank 3

Considering how we solved the electrostatics problem associated with $\mathbf{I}_2(n) \times \mathbf{A}_1$, we might suspect that all finite reflection groups of rank 3 will generate a similar solvable electrostatics problem for a point charge in a cavity bounded by three grounded conducting planar walls. This is in fact correct, as we will show momentarily. Looking at Tables 1 and 2, we see that, besides the reducible groups of rank 3 (namely, $\mathbf{A}_1 \times \mathbf{A}_1 \times \mathbf{A}_1$ and the family $\mathbf{I}_2(n) \times \mathbf{A}_1$), there are also three irreducible ones: $\mathbf{A}_3$, $\mathbf{B}_3$ (alias $\mathbf{C}_3$), and $\mathbf{H}_3$. It is these that give solvable problems that were not on Maxwell's list. To construct the generating mirrors in 3D, we need the 3D simple roots of these groups. These can be obtained from Table 2. True, they are given there as 4D vectors, but we actually originally obtained them as 3D vectors, and then embedded them in 4D space by appending a zero as the fourth coordinate. Thus, now we are just reversing this, discarding the zero fourth coordinate.

We should emphasize that all of the solvable problems we are presently discussing, which are related to reflection groups in 3D, can also be obtained as special cases of our general 4D procedure, which we will be explaining shortly. Nevertheless, it may be helpful to outline how one constructs a solvable problem in this simpler 3D case, where we do not use a stereographic projection.

After that, we will proceed to use these 3D reflection groups to illustrate the method of stereographic projection in the simpler case of mapping 3D to 2D. Once we have done this, we will proceed to our final goal, which is the method of stereographic projection from 4D to 3D.

4.1. An example: group $\mathbf{A}_3$ (without stereographic projection)

For definiteness, let us pick the group $\mathbf{A}_3$, which is the symmetry group of the regular tetrahedron. For the groups $\mathbf{A}_n$, the standard form of the simple roots is given as embedded in the space of $n+1$ dimensions, where they define an n-dimensional hyperplane. The three simple roots of $\mathbf{A}_3$ are given in the literature as $(1, -1, 0, 0)$, $(0, 1, -1, 0)$, and $(0, 0, 1, -1)$. To find what these vectors are in 3D space, we apply the Gram–Schmidt process to them to generate an orthonormal basis for the 3D space that they span, and then find their coordinates relative to that basis. The result (after normalization) is $\hat{\alpha}_1 = (1, 0, 0)$, $\hat{\alpha}_2 = (-1/2, \sqrt{3}/2, 0)$, and $\hat{\alpha}_3 = (0, -1/\sqrt{3}, \sqrt{2/3})$. This is exactly what appears in Table 2, after we discard the fourth (zero) coordinate. It is clear that none of the simple roots is orthogonal to both of the other ones, so this case indeed is not on Maxwell's list.

As we said, this group is a symmetry group of the regular tetrahedron. We can start from the three simple roots we have just written down; see Fig. 2(a).

They define three generating mirrors. We can iteratively (as we described above) reflect the generating mirrors about each other until no new mirrors are produced. In this case, we will get three extra mirrors, for a total of six. They are depicted in Fig. 2(b). The six mirrors will be the six planes of symmetry of the regular tetrahedron, and so may be described as follows. Consider a tetrahedron. Pick a pair of vertices, and consider the edge that connects them. Now look at the remaining pair of vertices, and consider the midpoint of the edge that connects *them*. The former edge and the latter midpoint define a plane. It should be clear from the construction that this is a plane of symmetry of the tetrahedron. There are $\binom{4}{2} = 6$ ways to pick a pair of vertices, and
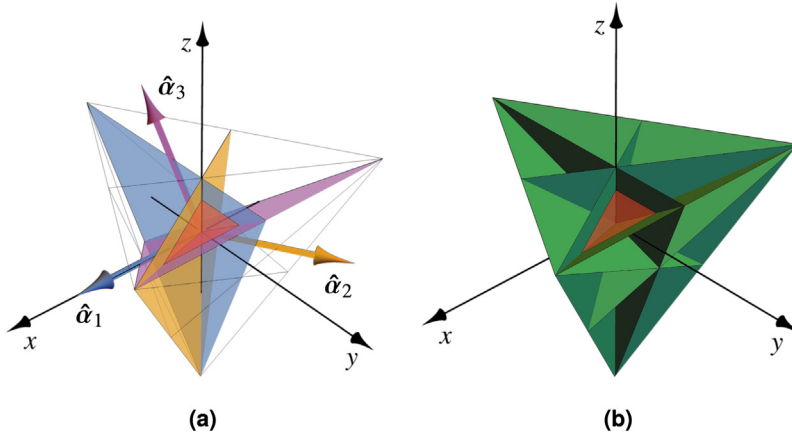
Fig. 2. Geometry of $\mathbf{A}_3$. (a) The simple roots and the corresponding generating mirrors. A part of the principal chamber is shown in red. (b) The full system of mirrors, with part of the principal chamber in red. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

thus there are six planes of symmetry. Any three can be taken as generating mirrors, and the three simple roots we have written down will indeed give a triplet of such planes. All the mirrors meet at a point (this is the case with all finite reflection groups, in all dimensions), and we will set the origin of our coordinate system at that point. Note that we still have the freedom to rotate our system of mirrors about the origin. The full set of mirrors partitions the space into chambers, which are visible in Fig. 2(b). To determine their number, note that a tetrahedron has 4 faces, and on each face we see into six chambers. So there are 24 chambers total, which is also the order of $\mathbf{A}_3$. One of the 24 chambers is the principal chamber, defined as the set of points $\vec{r}$ such that all of $\vec{r} \cdot \hat{\boldsymbol{\alpha}}_i$ is positive for $i = 1,\ 2,\ 3$. In Fig. 2(a) and (b), a part of the principal chamber is shown in red. All other chambers are the images of the principal chamber 'under the action' of the reflection group. This means that if we start iteratively reflecting the entire principal chamber about the generating mirrors, we will eventually generate the remaining 23 chambers, and after that we will generate no further ones.

4.2. Constructing a solvable electrostatics problem

To construct a solvable electrostatic problem, consider the principal chamber. It is enclosed by parts of the generating mirrors. Suppose these parts are grounded conducting surfaces. Place a point charge inside the principal chamber: our electrostatics problem is to find the electric field inside this chamber. To solve it by the method of images, iteratively reflect that point charge through the generating mirrors until 23 distinct images are created, one in each non-principal chamber; further reflections will not generate any new ones. For future reference, note that all charges are the same distance from the origin, and so lie on the surface of a sphere. The magnitudes of all the charges are the same; the only nonobvious question is whether it is possible to assign the charge *signs* in a consistent manner. What is required is that any pair of charges related by a single reflection have opposite signs. That way, their contribution to the potential at the mirror between them will be zero. Furthermore, our iterative construction guarantees that for every charge and every mirror, there is a charge that is the reflection of the first charge relative to that mirror. Thus, at every mirror, the potentials due to all the charges will cancel out.

Here is how to assign the charge signs. Note that every image charge is the result of some sequence of reflections from the original charge. Let the image charge have the same sign as the original if it takes an even number of reflections to generate it from the original, and let it have the opposite sign if it takes an odd number. This guarantees any pair of charges related by a single

reflection have opposite signs: after all, they will be reachable from the original charge through reflection sequences whose lengths differ by one. Thus the sequence lengths will have opposite parities, and so the image charges will have opposite signs.

The only possible problem is this: note that each image charge is reachable from the original one through a variety of reflection sequences, which generally have unequal lengths. And if some of these sequence lengths are odd numbers while others are even, we are in trouble. Luckily, a general theorem of reflection groups guarantees that no matter which sequence of reflections we use to get from the original charge to any given image, the parity of the sequence length will always be the same. Thus the signs of the image charges can be unambiguously assigned, and the full system of charges will solve our electrostatics problem in the principal chamber.

5. Stereographic projection

Our main result concerns reflection groups in 4D, and the way we will go from the abstract 4D space to our physical 3D one will be via a stereographic projection. Here we will first describe our procedure in the more familiar 3D setting. For an account of the stereographic projection from 3D to 2D, we refer to Ref. [12], which contains basic facts and proofs using Euclidean geometry.

Choose a system of generating mirrors in 3D, e.g. corresponding to $\mathbf{A}_3$. The generating mirrors all meet at a point. Choose a sphere whose center is at that point, of any radius R'. Form the intersections of this sphere and the generating mirrors. The result is three great circles, as shown in Fig. 3(a). The region corresponding to the principal chamber is a spherical triangle. Now choose some point on the sphere to be the 'South Pole'; the point diametrically opposite is then the 'North Pole'. Consider the 2D plane S tangent to the sphere at the South Pole. The stereographic projection of points on the sphere to this plane is then found as follows. Pick a point of interest, P', on the sphere. Consider the line L that passes through both the North Pole and the point of interest. The stereographic projection of the point of interest is the point P on the tangent plane S at which the line L crosses the that plane. This is shown in Fig. 3(b). The 'North Pole' is marked by the red point on the topmost end of the black line. This line shows the projection of a point on the blue great circle. Note that the system of mirrors was tilted so that the spherical triangle corresponding to the principal chamber is near the South Pole, but also so that no great circle passes through the South Pole. The resulting projection is shown in (c). If the system of mirrors were tilted so that two of the generating mirrors passed through the South Pole, then two of the arcs become straight lines. The third one still must be a circular ark, because the three great circles corresponding to the generating mirrors never all meet at one point.

If we pick a coordinate system such that its origin is at the center of the sphere, so that the South Pole is at $(0, 0, -R')$, it is easy to show that the stereographic projection (x, y) of a point (x', y', z') on the sphere (i.e. such that $(x')^2 + (y')^2 + (z')^2 = (R')^2$) is given by

$$(x, y) = \frac{2R'}{R' - z'} \left(x', y' \right).$$

If the point on the sphere has spherical coordinates (R', θ', ϕ') so that $(x', y', z') = R' (\sin\theta' \cos\phi', \sin\theta' \sin\phi', \cos\theta')$, then the polar coordinates (ρ, ϕ) of the projected point are $\rho = 2R' \left| \cot\frac{\theta'}{2} \right|$ and $\phi = \phi'$ (and of course the z-coordinate, if we choose to include it, is $-R'$). Relevant properties of the stereographic projections are as follows.

First, a great circle on the sphere that passes through the poles is mapped to a line in the plane (but it will be convenient to consider this a limiting case of a circle of as its radius goes to infinity). If the great circle does not pass through the poles and lies in the plane whose unit normal is $\boldsymbol{n} = (n_x, n_y, n_z) = (\sin\theta_{\boldsymbol{n}} \cos\phi_{\boldsymbol{n}}, \sin\theta_{\boldsymbol{n}} \sin\phi_{\boldsymbol{n}}, \cos\theta_{\boldsymbol{n}})$, then it is mapped to a circle C in the plane, whose radius is $R = 2R' |\sec\theta_{\boldsymbol{n}}|$ and whose center is $\vec{\boldsymbol{O}} = -2R'(\sec\theta_{\boldsymbol{n}})(n_x, n_y, 0) + (0, 0, -R')$.

Second, recall that the great circle is an intersection of a mirror and the chosen sphere. Pick a point $\vec{\boldsymbol{p}}'_1$ on the sphere and reflect it through this mirror. The reflection $\vec{\boldsymbol{p}}'_2$ is also on the surface of the sphere. Now consider the stereographic projections $\vec{\boldsymbol{p}}_1$ and $\vec{\boldsymbol{p}}_2$ of the two points. It turns out that
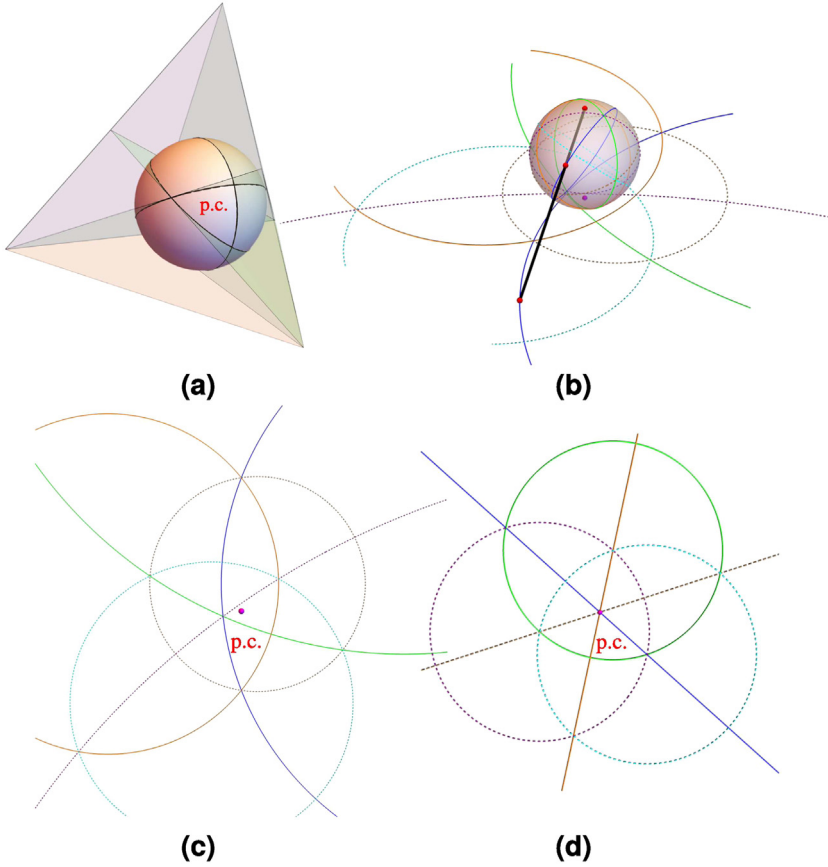
Fig. 3. Stereographic projection of reflection group geometry from 3D to 2D, on the example of **A**₃. The regions corresponding to the principal chamber are marked by the red letters 'p.c.' (a) The great circles on a sphere that are the intersections of the sphere and the generating mirrors. (b) The stereographic projection of the great circles from (a) onto a plane tangent to the sphere at the sphere's 'South Pole', which is marked by the point in magenta. The solid lines correspond the generating mirrors, while the dotted ones correspond to the other mirrors (c) The resulting 2D system of circles. The 'South Pole' is marked by the point in magenta. The region corresponding to the principal chamber is bounded by circular arcs. (d) If the system of mirrors were tilted so that two of the generating mirrors passed through the South Pole, then two of solid arcs become straight lines. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

one of the points is inside the circle C, one outside, and that they are *circle inversions* of each other, with C as the inversion circle. This means that $\vec{p}_1$, $\vec{O}$, and $\vec{p}_2$ lie on a line, and that the distances from the points to the center of the circle satisfy $d(\vec{p}_1, \vec{O}) \, d(\vec{p}_2, \vec{O}) = R^2$, where $d(\vec{p}, \vec{O}) = |\vec{p} - \vec{O}|$. Here for any vector $\vec{v}$, $|\vec{v}|$ is its norm, defined as $|\vec{v}| = \sqrt{\vec{v} \cdot \vec{v}}$.

Third, stereographic projection is *conformal*, meaning that if two curves on the surface of the sphere intersect at some angle, the curves that are their projections will intersect at the same angle.

Fourth, the stereographic projection is reversible; it defines a one-to-one correspondence between the sphere (without the North Pole) and the plane.

We should mention that we were unable to find an *explicit* mention of the second property in the literature, except for the special case where the great circle is parallel to the plane to which we are projecting. In particular, it is not mentioned in Ref. [12]. Our explicit proofs (for 3D→2D and 4D→3D cases) are, so far, brute-force computations which we will not reproduce.

6. An outline of our procedure in 4D

We will now have to repeat the procedure of the previous section with all numbers of spatial dimensions raised by one. Pick a reflection group g; let ℓ be its rank and $|g|$ its order. Place its ℓ generating mirrors in the 4D space. This means that ℓ must be 4 or less. There are 19 finite reflection groups of rank 4 or less, and they are all listed in Tables 1 and 2. The generating mirrors of g always meet at a point, and we choose a coordinate system whose origin is at that point. We are free to rotate our system of mirrors about the origin; below, we will briefly discuss how to parametrize 4D rotations.

Now pick a radius $\mathcal{R}$ and consider a 4D hypersphere (a 3D object), centered at the origin, of that radius. The intersections of the hypersphere and the generating mirrors are great hypercircles S_i', $i = 1, \ldots, \ell$. The intersection of the sphere and the principal chamber, which we will denote A', is a 3D region on the hypersphere. Choose a point P_1' in A', and find all its images P_j', $j = 2, \ldots, |g|$, relative to the generating mirrors.

Now we stereographically project to the 3D hyperplane tangent to the hypersphere at its South Pole, $(0, 0, 0, -\mathcal{R})$. The ℓ projections of the great hypercircles are S_i, each of which is either a plane (if it comes from a great hypercircle that passes through the poles) or else a sphere. Note that by 4D-rotating our system of mirrors we can make up to three great hypercircles pass through the poles.

The projections of all the points are P_i, $i = 1, \ldots, |g|$. The projection of A' is a 3D region bounded by segments of S_i's, and it contains P_1. Note that although P_i, $i = 2, \ldots, |g|$ were obtained via the stereographic projection of P_i''s from 4D, we could obtain them in a different way, entirely within 3D. Namely, we just iteratively sphere-invert P_1 relative to the spheres S_i (for any S_i that is a plane, we just reflect as usual). Below, we will in fact refer to this procedure to figure out how to assign the values to the image charges.

Indeed, we need not have started with 4D points (such as P_1'). We may first use the stereographic projection to generate the 3D region bounded by segments of the S_i's. Now we pick a point P_1 in that region (this will be the location of our charge) and iteratively sphere-invert it to obtain the other points P_i (these will be the locations of our image charges). Note that since the stereographic projection is reversible, all those points are images of some points P_i', $i = 1, \ldots, |g|$, on the 4D hypersphere; and all of these P_i''s are images of each other under the reflection through the mirrors of the 4D finite reflection group.

Note that the above implies that we can use the stereographic projection to generate a finite group of sphere inversions (possibly supplemented by plane reflections). In fact, our original motivation for this project was to find out what concept plays the same role for sphere inversions as finite reflection groups do for plane reflections; it just turned out that sphere inversions are simply related to plane reflections (in one spatial dimension higher) through the stereographic projection.

6.1. Comparison to the Kelvin transform

At this point, we should mention the well-known Kelvin transform [13], which, given any problem solvable by images, can be used to generate a different solvable problem. Physics sources tend to refer to the Kelvin transform as the 'method of inversion' [1,5,6], but since sphere inversions appear in our procedure in a different way, we will not use that name. The Kelvin transform involves sphere-inverting the geometry of the original problem relative to an arbitrary inversion sphere, and adjusting the resulting charges in a particular way. Sphere inversion transforms planes and spheres to planes and spheres, and in particular can transform planes into spheres. Indeed, the Kelvin transform is the origin of the spherical surfaces in Maxwell's list of solvable problems.

These properties of spherical inversions sound remarkably similar to those of stereographic projections. This is not an accident: the stereographic projection of a sphere s is equivalent to a spherical inversion of s relative to an inversion sphere S, whose radius is the diameter of s, centered at the North Pole of s. In general, under spherical inversion, a sphere that passes through the origin of the inversion sphere gets mapped to a plane. If it is further true that the radius of S is twice

the radius of *s*, simple geometry shows that the spherical inversion is equivalent to a stereographic projection.

Since our 4D procedure also generates solvable problems involving intersecting spheres and planes, one might suspect that our construction and the construction using the Kelvin transform are closely related. However, we will now show that neither is reducible to the other.

To see that the Kelvin transform construction cannot treat all the cases that our 4D construction can, note that Kelvin transform requires a solvable 3D problem as input. If we do not use our 4D construction, all solvable problems must ultimately come from reflection groups of rank 3 or lower (via the methods described in Section 4). The geometries that result from applying the Kelvin transform to those cannot be equivalent to those that are obtained (via our 4D procedure) from *irreducible* reflection groups of rank 4. To see this, recall that both the spherical inversion and the stereographic projection are *conformal*, i.e. angle-preserving. On the other hand, reflection groups are completely determined by the dihedral angles between their generating mirrors (as encoded in their Coxeter diagrams). After transformation (be it a spherical inversion or a stereographic projection), whatever spheres and planes we obtain will still cross each other at those same angles. So a spherical inversion of a 3D system of generating mirrors cannot produce the same angles as a stereographic projection of an *irreducible* rank-4 system.

And to see, conversely, that our 4D construction cannot treat all the cases that the Kelvin transform can, note that when our construction maps the principal chamber to an area completely enclosed by spherical or planar surfaces, the original charge should be placed within this area. On the other hand, the Kelvin transform can be used to produce problems where the charge is on the outside: just start with a π/n wedge (with a point charge inside of it) and choose an inversion circle inside this wedge. After a spherical inversion, the half-planes that define the wedge will become intersecting spheres, with the appropriate parts of them missing so that the result is a single 'double bubble' surface, which completely encloses a finite region of space. In this case, it is the image charges (of the original solvable-wedge problem) whose Kelvin transforms will end up on the inside, while the Kelvin transform of the original charge will be on the outside.

In our 4D construction, after we stereographically project the great hypercircles to ordinary 3D spheres and planes, we have the locations of all the conducting surfaces and all the charges, the original one and the image charges. However, we still do not know which values we need to assign to the image charges. To determine these values, we will need to use sphere inversions. Note that these sphere inversions are not Kelvin transforms, because neither the bounding walls not the original charge are transformed. Let us describe this type of spherical inversion quantitatively.

7. Sphere inversion for a single spherical conducting cavity

Consider a point charge q_1 at a position $\boldsymbol{p}_1$ inside an empty spherical cavity of radius *R* centered at a point $\boldsymbol{O}$, surrounded by a grounded conductor (Fig. 4). Consider a point $\boldsymbol{p}_2$ related to the point $\boldsymbol{p}_1$ through a sphere inversion, with the cavity wall playing the role of the inversion sphere:

$$\boldsymbol{p}_2 = \boldsymbol{O} + \left(\frac{R}{|\boldsymbol{p}_1 - \boldsymbol{O}|}\right)^2 (\boldsymbol{p}_1 - \boldsymbol{O}) . \tag{1}$$

Let r_1 and r_2 be the distances between a point on the cavity surface and the points $\boldsymbol{p}_1$ and $\boldsymbol{p}_2$, respectively. For any point on the cavity surface, it can be proven that

$$\frac{r_1^2}{|\boldsymbol{p}_1 - \boldsymbol{O}|} = \frac{r_2^2}{|\boldsymbol{p}_2 - \boldsymbol{O}|} . \tag{2}$$

From this property, it follows immediately that if the value of the image charge is assigned as

$$q_2 = -\sqrt{\frac{|\boldsymbol{p}_2 - \boldsymbol{O}|}{|\boldsymbol{p}_1 - \boldsymbol{O}|}} q_1 , \tag{3}$$

then the electrostatic potential created by the physical charge and its image on any point on the wall will vanish: $q_1/r_1 + q_2/r_2 = 0$ The resulting field will be the correct solution of the Poisson equation with (zero) Dirichlet boundary conditions, thus solving the problem of the field induced by a charge.
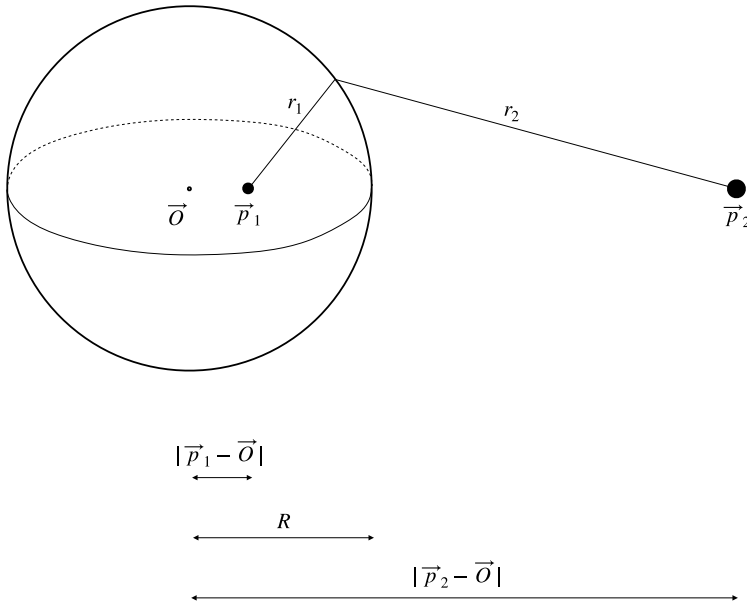
Fig. 4. A charge inside a conducting spherical cavity: notations and definitions. $\boldsymbol{O}$ is the center of the cavity. The vector $\boldsymbol{p}_1$ is the position of the physical charge and r_1 is the distance between a particular point on the cavity wall and the physical charge, while $\boldsymbol{p}_2$ and r_2 are the corresponding quantities for the image charge. The distances $|\boldsymbol{p}_1 - \boldsymbol{O}|$ and $|\boldsymbol{p}_2 - \boldsymbol{O}|$ govern the assignment of the value of the image charge.

8. Piecewise-spherical conducting cavities solvable with the method of images

As we have been saying, our solvable problems will involve regions (cavities) bounded by segments of spheres and planes, where the latter coincide with the 4D→3D stereographic projections of great hypercircles associated with a finite reflection group in 4D. Let us, however, first explicitly write down the conditions that must be satisfied in order for any problem involving regions bounded by grounded conducting segments of spheres and planes to be solvable by the method of (finitely many) images.

Imagine an empty cavity surrounded by a grounded conductor. Assume that its walls are formed by segments of spherical surfaces. The field induced by a point charge placed inside the cavity can be constructed using the method of images if the system satisfies the following conditions, which are spherical-geometry analogs of the requirements described in Section 4.2. Note that if a surface is a plane rather than a sphere, in what follows, a sphere inversion should be replaced by a simple reflection.

Solvability conditions:

I. Consider the set of image charge locations produced via chains of inversions from Eq. (1) with respect all the spheres, one sphere at each link of the chain. For any starting location, there should be a finite number of image locations such that once they are all in the set, further inversions of its elements will only produce locations that are already in the set.

II. One should be able to unambiguously assign the values of the image charges via a sequential application of the rule in Eq. (3). The assigned charge value should not depend on the particular sequence of inversions that produced its location.

III. The sphere inversions should not produce any image charges within the cavity of interest.

Let us show that these conditions are indeed sufficient. Consider one of the spherical surfaces defining the cavity. According to condition I., any charge – regardless of whether it is the physical

charge or an image charge – will have a counterpart linked to it, in both directions, by an inversion with respect to the surface of interest. According to condition II., the two will create a zero potential on the cavity surface. The rest of the charges will form similar pairs, with similar results. Finally, condition III. guarantees that no 'ghost' charges are required by the solution obtained, i.e. that inside the cavity there is only the charge that is present in the statement of the problem.

9. Generation of piecewise-spherical conducting cavities using a 4D reflection group and a 4D→3D stereographic projection

9.1. The setup

As we have been suggesting, a fruitful way to construct a set of spherical surfaces that satisfy the solvability conditions outlined in the previous section is to use the construction described in Section 6. That construction involves a 4D→3D stereographic projection of great hypercircles on the surface of a 4D hypersphere, so let us describe this transformation quantitatively; see Fig. 5. A 4D→3D stereographic projection takes a point $\boldsymbol{p}'$ on a hypersphere centered at $\mathfrak{O}$ and of radius $\mathfrak{R}$,

$$
\boldsymbol{p}' \equiv \begin{pmatrix} p'_x \\ p'_y \\ p'_z \\ p'_w \end{pmatrix} = \begin{pmatrix} \sin(\Theta_{\boldsymbol{p}'})\sin(\theta_{\boldsymbol{p}'})\cos(\phi_{\boldsymbol{p}'})\mathfrak{R} \\ \sin(\Theta_{\boldsymbol{p}'})\sin(\theta_{\boldsymbol{p}'})\sin(\phi_{\boldsymbol{p}'})\mathfrak{R} \\ \sin(\Theta_{\boldsymbol{p}'})\cos(\theta_{\boldsymbol{p}'})\mathfrak{R} \\ \cos(\Theta_{\boldsymbol{p}'})\mathfrak{R} \end{pmatrix} . \tag{4}
$$

and converts it to a point $\boldsymbol{p}$ on the tangential 'horizontal' 4D hyperplane – identified with the physical 3D space – that is tangent to the hypersphere at the 'South Pole' $(x, y, z, w) = (0, 0, 0, -\mathfrak{R})$. In Cartesian coordinates, (p'_x, p'_y, p'_z, p'_w) on the surface of the hypersphere gets mapped to $(p_x, p_y, p_z, p_w) = \frac{2\mathfrak{R}}{\mathfrak{R}-p'_w}(p'_x, p'_y, p'_z, 0) + (0, 0, 0, -\mathfrak{R})$. In spherical coordinates,

$$
\boldsymbol{p} \equiv \begin{pmatrix} p_x \\ p_y \\ p_z \\ p_w \end{pmatrix} = \begin{pmatrix} 2\cot(\frac{\Theta_{\boldsymbol{p}'}}{2})\sin(\theta_{\boldsymbol{p}'})\cos(\phi_{\boldsymbol{p}'})\mathfrak{R} \\ 2\cot(\frac{\Theta_{\boldsymbol{p}'}}{2})\sin(\theta_{\boldsymbol{p}'})\sin(\phi_{\boldsymbol{p}'})\mathfrak{R} \\ 2\cot(\frac{\Theta_{\boldsymbol{p}'}}{2})\cos(\theta_{\boldsymbol{p}'})\mathfrak{R} \\ -\mathfrak{R} \end{pmatrix} . \tag{5}
$$

The final step is an orthogonal projection Pr to the 3D space, which is accomplished by simply deleting the final (w) coordinate: $\mathrm{Pr}(x, y, z, w) = (x, y, z)$. Since it is unlikely to cause confusion, we will not always carefully distinguish the points before and after the orthogonal projection.

The stereographic projection is invertible: the 3D point (p_x, p_y, p_z) has the inverse image

$$
\begin{pmatrix} p'_x \\ p'_y \\ p'_z \\ p'_w \end{pmatrix} = \frac{4\mathfrak{R}^2}{4\mathfrak{R}^2 + p_x^2 + p_y^2 + p_z^2} \begin{pmatrix} p_x \\ p_y \\ p_z \\ -2\mathfrak{R} \end{pmatrix} + \begin{pmatrix} 0 \\ 0 \\ 0 \\ \mathfrak{R} \end{pmatrix} .
$$

A great hypercircle on our hypersphere is the set of 4D points that satisfy

$$
(\boldsymbol{p}')^2 = \mathfrak{R}^2 \tag{6}
$$
$$
\boldsymbol{p}' \cdot \boldsymbol{n} = 0 , \tag{7}
$$

where

$$
\boldsymbol{n} \equiv \begin{pmatrix} n_x \\ n_y \\ n_z \\ n_w \end{pmatrix} = \begin{pmatrix} \sin(\Theta_{\boldsymbol{n}})\sin(\theta_{\boldsymbol{n}})\cos(\phi_{\boldsymbol{n}}) \\ \sin(\Theta_{\boldsymbol{n}})\sin(\theta_{\boldsymbol{n}})\sin(\phi_{\boldsymbol{n}}) \\ \sin(\Theta_{\boldsymbol{n}})\cos(\theta_{\boldsymbol{n}}) \\ \cos(\Theta_{\boldsymbol{n}}) \end{pmatrix} . \tag{8}
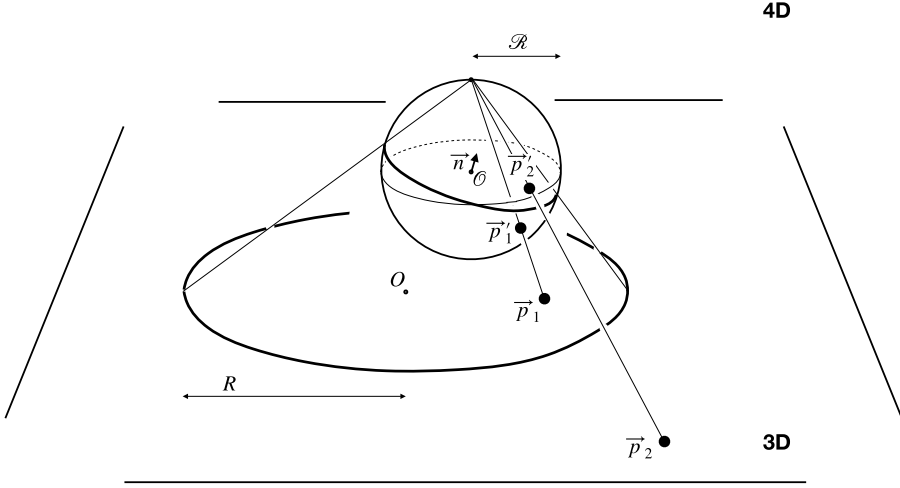$$

Fig. 5. An artistic rendering of the relationship between a 4D reflection on a hypersphere and a 3D sphere inversion, through a 4D stereographic projection. The points $\boldsymbol{p}_1'$ and $\boldsymbol{p}_2'$ are on a hypersphere of radius $\mathscr{R}$ whose center is at $\mathfrak{O}$. These two points are related by a 4D reflection with respect to a hyperplane through $\mathfrak{O}$ whose unit normal vector is $\boldsymbol{n}$. The 3D points $\boldsymbol{p}_1$ and $\boldsymbol{p}_2$ are the images of $\boldsymbol{p}_1'$ and $\boldsymbol{p}_2'$ under the stereographic projection from the hypersphere to a 'horizontal' hyperplane. The latter hyperplane is identified with the physical 3D space. It can be shown that $\boldsymbol{p}_1$ and $\boldsymbol{p}_2$ are then related via a 3D sphere inversion with respect to a sphere whose center is $\boldsymbol{O}$ and whose radius is R. This sphere is the stereographic image of the great hypercircle formed at the intersection of the hyperplane characterized by $\boldsymbol{n}$ and the hypersphere.

We may interpret $\boldsymbol{n}$ as the unit vector normal to a hyperplane that passes through $\mathfrak{O}$, the center of the hypersphere. The intersection of this hyperplane and the hypersphere is the great hypercircle of interest.

Under the stereographic projection described by Eqs. (4) and (5), the great hypercircle defined by Eqs. (6)–(8) transforms to a 3D sphere of radius

$$R = 2|\sec(\Theta_{\boldsymbol{n}})|\mathscr{R} \tag{9}$$

centered at

$$\mathfrak{O} \equiv \begin{pmatrix} O_x \\ O_y \\ O_z \\ O_w \end{pmatrix} = \begin{pmatrix} -2\tan(\Theta_{\boldsymbol{n}})\sin(\theta_{\boldsymbol{n}})\cos(\phi_{\boldsymbol{n}})\mathscr{R} \\ -2\tan(\Theta_{\boldsymbol{n}})\sin(\theta_{\boldsymbol{n}})\sin(\phi_{\boldsymbol{n}})\mathscr{R} \\ -2\tan(\Theta_{\boldsymbol{n}})\cos(\theta_{\boldsymbol{n}})\mathscr{R} \\ -\mathscr{R} \end{pmatrix} . \tag{10}$$

Notice that all the spheres produced this way satisfy the relation

$$R^2 = \boldsymbol{O}^2 + 4\mathscr{R}^2 , \tag{11}$$

with $\boldsymbol{O}$ being the 3D version of the corresponding 4D vector $\mathfrak{O}$, i.e. its orthogonal projection Pr to the xyz space:

$$\boldsymbol{O} \equiv \mathrm{Pr}\,\mathfrak{O} = (O_x, O_y, O_z) . \tag{12}$$

Let $\boldsymbol{p}_{1,\,4D}'$ and $\boldsymbol{p}_{2,\,4D}'$ be two points on the hypersphere. Suppose they are related by a reflection via a 4D mirror defined by the unit normal vector $\boldsymbol{n}$, in other words that

$$\boldsymbol{p}_{2,\,4D}' = \boldsymbol{p}_{1,\,4D}' - 2(\boldsymbol{n} \cdot \boldsymbol{p}_{1,\,4D}')\boldsymbol{n} .$$

As we already mentioned, it can be explicitly shown that the stereographic images of $\boldsymbol{p}_{1,\,4D}'$ and $\boldsymbol{p}_{2,\,4D}'$ are related by a sphere inversion. To apply Eq. (1), we take $\boldsymbol{p}_1 = \mathrm{Pr}\,\boldsymbol{p}_{1,\,4D} = (p_{1,\,x},\ p_{1,\,y},\ p_{1,\,z})$ and similarly $\boldsymbol{p}_2 = \mathrm{Pr}\,\boldsymbol{p}_{2,\,4D}$. The inversion sphere parameters R and $\boldsymbol{O}$ are as in Eqs. (9) and (12).

Let us finally consider the 3D spheres forming the cavity of interest. Let us assume that these 3D spheres stereographically originate from great hypercircles that are, in turn, produced by the generating mirrors of a 4D reflection group, as explained in Section 6.

9.2. Solvability condition I: finite number of image charges

Let us place a charge q at the location $\boldsymbol{p}$ inside the cavity described above. Let $\boldsymbol{p}'$ be the stereographic inverse image of $\boldsymbol{p}$ (so that $\boldsymbol{p}$ is the stereographic image of $\boldsymbol{p}'$). Start iteratively reflecting $\boldsymbol{p}'$ through the generating mirrors of the 4D finite reflection group. For each pair of 4D positions related by a reflection, their stereographic projections are related by a sphere inversion, and vice versa. Since we know there are only finitely many images produced by the 4D reflections, it follows that the iterative sphere inversions can produce only a finite number of image charges (the same number as the 4D reflections produce).

9.3. Solvability condition II: consistency of charge assignment

This is a sphere-inversion analog of the discussion in Section 6 about how to assign charge values to the image charges in the case of mirror reflections. But this is one place where the sphere inversion is noticeably more complicated than its 4D mirror-reflection counterpart, because while the reflections only affect the signs of the image charges, sphere inversions also affect their magnitudes.

If an inversion sphere is a stereographic image of a great hypercircle on a hypersphere of radius $\mathscr{R}$, then any two points $\boldsymbol{p}_1$ and $\boldsymbol{p}_2$ related by the inversion obey

$$\frac{|\boldsymbol{p}_1 - \boldsymbol{O}|}{|\boldsymbol{p}_2 - \boldsymbol{O}|} = \frac{F(\boldsymbol{p}_1)}{F(\boldsymbol{p}_2)} , \tag{13}$$

with

$$F(\boldsymbol{p}) = \boldsymbol{p}^2 + 4\mathscr{R}^2 . \tag{14}$$

Notice that $F(\boldsymbol{p}_1)$ depends neither on the other point $\boldsymbol{p}_2$, nor on the center $\boldsymbol{O}$ of the inversion sphere, nor on the radius R of the inversion sphere. The same holds for $F(\boldsymbol{p}_2)$, upon a $\boldsymbol{p}_1 \leftrightarrow \boldsymbol{p}_2$ substitution. The charge assignment rule in Eq. (3) becomes

$$q_2 = -\sqrt{\frac{F(\boldsymbol{p}_2)}{F(\boldsymbol{p}_1)}} q_1 . \tag{15}$$

It is easy to show that this relationship will become

$$q_{1+m} = (-1)^m \sqrt{\frac{F(\boldsymbol{p}_2)}{F(\boldsymbol{p}_1)}} q_1 \tag{16}$$

if q_{1+m} and q_1 are linked by a chain of m inversions (instead of a single inversion). In particular, any image charge q_{image} will be related to the original physical charge q_{physical} as

$$q_{\text{image}} = (-1)^{m_{\text{image}}} \sqrt{\frac{F(\boldsymbol{p}_{\text{image}})}{F(\boldsymbol{p}_{\text{physical}})}} q_{\text{physical}} , \tag{17}$$

where $(-1)^{m_{\text{image}}}$ is the parity of the number of inversions linking the physical charge and the image charge in question. As we already mentioned in Section 6, it is known that for reflection groups, the parity of the number of reflections leading to a particular member of the group is the same for any chain of reflections. Hence, the same will be true for sphere inversions that are stereographically linked to the members of a 4D reflection group, and the charge value assigned using Eq. (17) is indeed independent of the sequence of inversions used to get from the original charge to the image charge.

9.4. Solvability condition III: no image charges outside of the conductor

Consider again $\mathbf{p}'$, the stereographic inverse image of $\mathbf{p}$, where the latter is the location of the physical charge. We have already seen above that there will be one image charge per non-principal chamber, and so in particular all the 4D mirror-reflection images of $\mathbf{p}'$ will lie outside of the principal chamber. The intersection of the principal chamber with the hypersphere will, under the stereographic projection, become the 3D cavity of interest. Likewise, the 4D reflection images will become the locations of the 3D image charges, and none of them will be located inside the cavity of interest.

10. An example: a solvable cavity associated with the reflection group $\mathbf{D}_4$

According to Table 1, the 4D reflection group $\mathbf{D}_4$ has the Coxeter diagram ●━━< . Thus it has four simple roots (and so four generating mirrors), three of which are mutually orthogonal, while the fourth one makes an angle of $\pi/3$ with each of the first three. The total number of group elements is 192, of which 12 are pure reflections. Once the mutual orientation of the generating mirrors is fixed (basically by the Coxeter diagram), we have to choose the orientation of the mirrors as a group. The point where all the generating mirrors meet should be fixed at the origin, so the allowed transformations are 4D rotations. In four dimensions, rotations are parametrized by 6 real parameters. Indeed, the full reflection group can be parametrized as a product six rotation matrices [14], each describing a rotation in one of six pairs of Cartesian planes: xy, xz, yz, xw, yw, and zw. We will refer to these as *elementary rotations*. For example, the matrix describing an elementary rotation of a 4D vector (x, y, z, w) that takes place in the yw-plane is

$$R_{yw}(\theta_{yw}) = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & \cos\theta_{yw} & 0 & \sin\theta_{yw} \\ 0 & 0 & 1 & 0 \\ 0 & -\sin\theta_{yw} & 0 & \cos\theta_{yw} \end{pmatrix} .$$

The other elementary rotations work similarly, where each one gets its own angle (e.g. θ_{xy}, θ_{xz}, etc.). All six angles can be set independently.

Note that the concept of a rotation axis is not useful in dimensions higher than three, where the dimensionality of what is left fixed by a rotation in general depends on the rotation—unlike in 3D, where every rotation leaves fixed a line through the origin. In 4D there are rotations that, like our elementary ones, leave a whole 2D plane fixed. But the most general 4D rotation is a double rotation, a simultaneous rotation in two perpendicular planes, which leaves only the origin fixed.

Recall that our stereographic projection maps into the xyz space. Let us parametrize an arbitrary 4D rotation by first performing elementary rotations in planes involving the fourth dimension w, and then performing the rotations that take place wholly within the xyz space. For example, we can use the matrix $R_{yz}R_{xz}R_{xy}R_{zw}R_{yw}R_{xw}$, where R_{xw} is a function of θ_{xw}, R_{yw} of θ_{yw}, etc. In this case, the final three rotations ($R_{yz}R_{xz}R_{xy}$) will have the effect of simply 3D-rotate the final system of spheres and planes. The 3D→2D analog is a rotation of the sphere in Fig. 3(b) about the z-axis. In Fig. 3(c) and (d), this would have the effect of rigidly rotating the whole plane about the South Pole.

The first three rotations, however, can be used to control which great hypercircles stereographically map into planes and which into spheres, and what are the radii of the spheres and the locations of spheres and planes. The 3D→2D analog is rotating the sphere in Fig. 3(b) about the x or y axis. This will generally deform the patterns we get in the plane, changing the radii of the spheres and possibly even transforming them into planes.

As an example, we choose the following set of the generating mirror normals:

$$
\begin{aligned}
&(1, 0, 0, 0) \\
&(0, 1, 0, 0) \\
&(0, 0, 1, 0)
\end{aligned}
\tag{18}
$$

$$\left(\frac{1}{2}, \frac{1}{2}, \frac{1}{2}, \frac{1}{2}\right) . \tag{19}$$

The corresponding angles are then given by Eq. (8), which gives

$$
\begin{aligned}
&\Theta = \tfrac{\pi}{2};\ \theta = \tfrac{\pi}{2};\ \phi = 0 \\
&\Theta = \tfrac{\pi}{2};\ \theta = \tfrac{\pi}{2};\ \phi = \tfrac{\pi}{2} \\
&\Theta = \tfrac{\pi}{2};\ \theta = 0;\ \phi = \text{any}
\end{aligned}
\tag{20}
$$

$$
\Theta = \tfrac{\pi}{3};\ \theta = \arg\left(\tfrac{1+i\sqrt{2}}{\sqrt{3}}\right);\ \phi = \tfrac{\pi}{4}\ .
\tag{21}
$$

Under a stereographic projection in Eqs. (4) and (5), Eq. (9) says that the first three great hypercircles transform to 3D spheres of infinite radius, i.e. to 3D planes. These planes will cross the origin. The fourth hypercircle becomes a sphere. The resulting cavity is given by the following set of inequalities:

$$
x > 0,\ y > 0,\ z > 0\ ,
\tag{22}
$$

and

$$
\left(x + \frac{R}{2}\right)^2 + \left(y + \frac{R}{2}\right)^2 + \left(z + \frac{R}{2}\right)^2 < R^2\ ,
\tag{23}
$$

with

$$
R = 4\mathcal{R}\ ,
\tag{24}
$$

where $\mathcal{R}$ is the radius of the hypersphere the stereographic projection originates from; see Eqs. (12), (10), and (9) for the formulae for the center and the radius of the ball defined by Eq. (23). Notice that since the stereographic projection is a conformal transformation, the angles between the resulting 3D surfaces are the same as the angles between the 4D hyperplanes from Eqs. (18) and (19) that generate them. In particular, the spherical segment of the cavity surface crosses each of the three planar segments at an angle of 60°.

We can use the method of images to find the electrostatic potential created by a point charge placed anywhere inside the cavity bounded by the surfaces defined in Eqs. (22) and (23), where the surfaces are grounded conducting walls. The image charge locations will be given by sequential applications of the reflections about the planar cavity boundaries and sphere inversions in Eq. (1). The number of image charges, all situated outside the cavity, is 191. The charge values can be unambiguously assigned using the rule in Eqs. (17) and (14), indeed the original rule in Eq. (3). For the reflections, when an inversion sphere degenerates into a plane, the rules in Eqs. (17), (14) and (3) become $q_{\text{image}} = (-1)^m q_{\text{physical}}$ and $q_2 = -q_1$, respectively, with m still being the total number of the inversions and reflections linking q_{image} and q_{physical}.

In Fig. 6(a), we show a cubic segment of space, centered at the origin, of a linear size $0.6R$. The tetrahedron-like shape in the middle of the subfigure represents the shape of an electrostatic cavity – with grounded conducting walls – solvable using the method of images. A sample charge (red) is placed at the grand diagonal, at a distance $0.05R$ away from the spherical segment of the surface of the cavity. (Solvability persists for any position of the charge inside the cavity.) There are 191 image charges, but only a subset of them are visible (white) in the subfigure. The total number of charges, corresponding to the one physical charge supplemented by all the image charges, is 192, which is also the total number of group elements – both reflections and rotations – in the reflection group $\mathbf{D}_4$. Together with the sample (physical) charge, the image charges generate a potential inside the cavity that vanishes at the cavity walls.

Fig. 6(b) presents the surfaces of zero potential for the electrostatic potential induced by the sample charge of the previous subfigure. Shown there is a cubic segment of space, centered at the origin, of a linear size $0.36R$. A part of this zero-potential surface coincides with the surface of the cavity. The view of that part is partially obscured by the eight additional spheres (additional to the four surfaces forming the cavity) on which the potential vanishes as well. There must be a total of 12 surfaces, each either a sphere or a plane, because $\mathbf{D}_4$ has a total of 12 mirrors, and the symmetry of the problem guarantees they all must have the same potential. The potential also vanishes on the continuations of the surface fragments forming the cavity of interest.
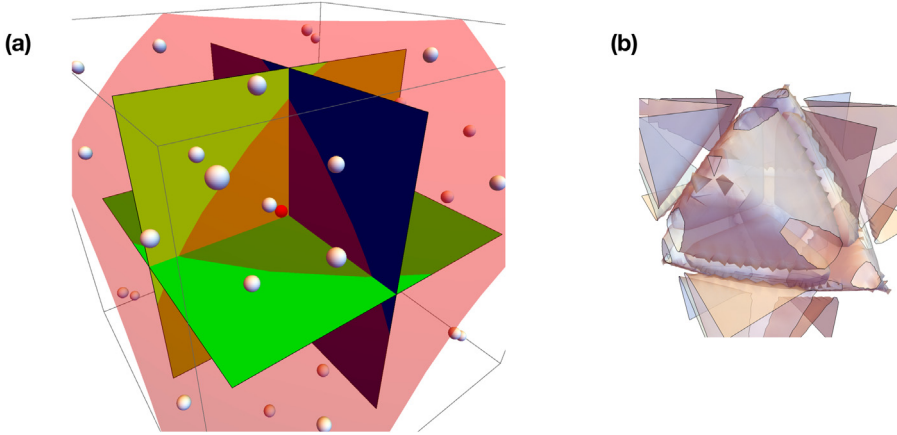
Fig. 6. A solvable electrostatic problem associated with the four-dimensional reflection group $\mathbf{D}_4$ (a subgroup of the full symmetry group of the tesseract). (a) A cubic segment of space, centered at the origin, of a linear size $0.6R$. (b) The surfaces of zero potential for the electrostatic potential induced by the sample charge of the previous subfigure. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

11. Summary and outlook

In all branches of physical science, problems whose solution is expressible in closed form are invaluable for the development of physical insight, while also serving as benchmarks on which numerical or approximate (but more generally applicable) methods of solution can be tested [3,11, 15,16]. In some dynamical cases, exact solvability is itself the mechanism underlying important physical effects (such as lack of thermalization [17]) and distinct physical phenomena (such as solitons [16]). This is due to the fact that exactly solvable (also known as *integrable*) time-dependent systems have a great (often infinite) number of dynamical conservation laws.

We have shown that in electrostatics, the class of problems solvable by the method of images is considerably larger than what was previously known. In particular, we showed that the solution of a classic electrostatic problem – finding the field induced by a point charge placed inside a spherical cavity with grounded conducting walls – can be generalized to any cavity whose walls are represented by 4D stereographic projections of great hypercircles formed by the intersections of a 4D hypersphere and the mirrors of any known finite 4D reflection group of rank 4 or lower. We used the group $\mathbf{D}_4$ as a worked example.

In the process, we pointed out that every known problem that is solvable by adding finitely image charges is based on some finite reflection group. Starting with such a group, one may use its mirrors as they are or apply further transformations to them: either a 4D→3D stereographic projection, which we introduced in this work, or the well-known Kelvin transform. In this way we obtain all presently known solvable problems.

The scope of problems solvable by our construction is as follows:

- Any finite reflection group of rank 4 or less can be used, including groups that are reducible and those that are lower-dimensional but embedded in 4D; see Tables 1 and 2.
- For each of the reflection groups listed in Tables 1 and 2, there will be a three-parametric family of orientations of the mirrors in the 4D space, leading to a three-parametric family of 3D sphere radii and positions. The remaining three parameters of the six-parametric family of the 4D rotations will control the trivial 3D rotation of the resulting 3D cavity.

A 2D generalization of our scheme (a 3D stereographic projection from a sphere to a plane) is worth considering. There, one obtains a family of solvable problems involving cylindrical conducting cavities and charged wires. The problem of assigning the values of charges is expected to disappear (as one trivially gets a set of sign-alternating charges). However, this property alone only guarantees that the electrostatic potential will be a non-zero constant on each of the cylindrical segments; it is still in principle possible that the potential will not be given by the same constant on each of the cylinders. However, we expect that a relationship analogous to Eqs. (13) and (14) can be proven in the 2D case as well. If so, then it will be possible to prove that the potential in fact has the same value at every point on the cavity surface.

The property in Eqs. (13) and (14) – that so far seems completely accidental – may prove to be a consequence of a deeper connection. One may conjecture that a map exists between the solutions of the Poisson equation on the hypersphere and the solutions of the Poisson equation on a plane stereographically connected to that hypersphere. In this case, the solutions for the class of problems we considered will become images of the hyperspherical solutions which were obtained via *pure reflections*, obviating the need to go through sphere inversion procedure in 3D. The charges on the hypersphere will have the same magnitude and an alternating sign. The relationship between a hyperspherical charge q' and the 3D charge q will be then given by $q = \text{const} \times \sqrt{F(\boldsymbol{p})}\, q'$, where $\boldsymbol{p}$ is the location of the charge q, and the function F is given by Eq. (14). The proof of this conjecture may involve (a) the fact that the 3D sphere-, and, potentially, 4D hypersphere inversions convert solutions of the Poisson equation to solutions of the Poisson equation and (b) the fact that a 4D stereographic projection can be reinterpreted as a 4D hypersphere inversion.

Finally, it is hard to believe that our construction is not in some way connected to the Kelvin transform after all. A progress regarding the program from the previous paragraph might help elucidate this connection, and perhaps the other way around as well.

Declaration of competing interest

Acknowledgments

References

[1] L.D. Landau, L.P. Pitaevskii, Electrodynamics of Continuous Media: V. 8, second ed., Pergamon, New York, 1984, http://dx.doi.org/10.1016/B978-0-08-030275-1.50025-4.
[2] J.D. Jackson, Classical Electrodynamics, third ed., Wiley, Hoboken, NJ, 2009.
[3] D. Palaniappan, Electr. Eng. 94 (2011) 107–116, http://dx.doi.org/10.1007/s00202-011-0221-7.
[4] A. Garg, Classical Electromagnetism in a Nutshell, Princeton University Press, Princeton, NJ, 2012.
[5] J.H. Jeans, The Mathematical Theory of Electricity and Magnetism, fifth ed., Cambridge University Press, Cambridge, UK, 2009, http://dx.doi.org/10.1017/CBO9780511694356, reprinted.
[6] W.R. Smythe, Static and Dynamic Electricity, third ed., Hemisphere Pub. Corp., New York, 1989.
[7] J.C. Maxwell, A Treatise on Electricity and Magnetism, Vol. 1, Macmillan, London, UK, 1873, Ch. XI.
[8] J.E. Humphreys, Reflection Groups and Coxeter Groups, Cambridge University Press, Cambridge, U.K., 2012, http://dx.doi.org/10.1017/CBO9780511623646.
[9] M. Olshanii, S.G. Jackson, New J. Phys. 17 (2015) 105005, http://dx.doi.org/10.1088/1367-2630/17/10/105005.
[10] M. Gaudin, J.-S. Caux, The Bethe Wavefunction, Cambridge University Press, Cambridge, UK, 2014.
[11] B. Sutherland, Beautiful Models: 70 Years of Exactly Solved Quantum Many-Body Problems, World Scientific, Singapore, 2004, http://dx.doi.org/10.1142/5552.
[12] B.A. Rosenfeld, N.D. Sergeeva, Stereographic Projection, Mir Publishers, Moscow, Russia, 1977, translated from the Russian by Vitaly Kisin.
[13] O.D. Kellogg, Foundations of Potential Theory, Springer, Berlin, 1967, http://dx.doi.org/10.1007/978-3-642-86748-4.
[14] A.J. Hanson, in: P.S. Heckbert (Ed.), Graphics Gems IV, Academic Press, San Diego, CA, USA, 1994, pp. 149–170, http://dx.doi.org/10.1016/B978-0-12-336156-1.50024-0, Ch. II.6.

[15] S. Albeverio, F. Gesztesy, R. Høegh-Krohn, H. Holden, Solvable Models in Quantum Mechanics, second ed., AMS Chelsea Pub., Providence, RI, 2005, with an appendix by Pavel Exner.

[16] M.J. Ablowitz, H. Segur, Solitons and the Inverse Scattering Transform, SIAM, Philadelphia, PA, 1981, http://dx.doi.org/10.1137/1.9781611970883.

[17] T. Kinoshita, T. Wenger, D.S. Weiss, Nature 440 (2006) 900–903, http://dx.doi.org/10.1038/nature04693.