TRANSACTIONS OF THE
AMERICAN MATHEMATICAL SOCIETY
Volume 00, Number 0, Pages 000–000
S 0002-9947(XX)0000-0

TRINITY SYMMETRY AND KALEIDOSCOPIC REGULAR MAPS

DAN ARCHDEACON, MARSTON CONDER, AND JOZEF ŠIRÁŇ

Abstract. A cellular embedding of a connected graph (also known as a map) on an orientable surface has trinity symmetry if it is isomorphic to both its dual and its Petrie dual. A map is regular if for any two incident vertex-edge pairs there is an automorphism of the map sending the first pair onto the second. Given a map M with all vertices of the same degree d, for any e relatively prime to d the power map M^e is formed from M by replacing the cyclic rotation of edges at each vertex on the surface with the eth power of the rotation. A map is kaleidoscopic if all of its power maps are pairwise isomorphic. In this paper, we present a covering construction that gives infinite families of kaleidoscopic regular maps with trinity symmetry.

1. Introduction and basic concepts

The Platonic solids have fascinated people for thousands of years. They are very symmetric — in some sense, the most symmetric shapes possible. More specifically, they are the only spherical shapes with arc-transitive automorphism group. So attention turns naturally to very symmetric nonspherical shapes — that is, to discrete structures with large automorphism groups that lie on other surfaces, such as regular maps. Regular maps have been highly studied and catalogued for small orders and small genera. In this paper we aim to go even further.

There are important operations on maps, such as taking the dual, or taking the Petrie dual (whose faces are the left-right paths of the original map), or forming a new map by taking an integer e relatively prime to all vertex degrees and replacing at every vertex the cyclic order of edges emanating from the vertex by the eth power of that order. Any case in which such an operation yields a map isomorphic to the original one may be regarded as an 'external symmetry' of the map. From this point of view, one may say that a map has the 'ultimate level of symmetry' if the map is regular and has all possible external symmetries. Our aim is to construct such 'super-symmetric' maps.

We will make the above concepts more precise in the two subsections of this Introduction. Our two main results are presented in Section 2. Techniques for

2010 *Mathematics Subject Classification.* Primary 05E18, 20B25, 57M15.

The first author thanks the Open University which hosted him while much of this research was conducted.

The second author's work was supported by the Marsden Fund of New Zealand, project UOA-1015, and assisted by use of the Magma system [5].

The third author acknowledges support from the VEGA Research Grant 1/0781/11, the APVV Research Grant 0223-10, and the APVV support as part of the EUROCORES Programme EUROGIGA, project GREGAS, ESF-EC-0009-10, financed by the European Science Foundation.

proving our results are developed in Sections 3 and 4. The actual proofs are in Sections 5 and 6, followed by concluding remarks in Section 7.

1.1. **Maps and automorphisms.** A map M is a cellular embedding of a connected graph or multigraph G on a surface. For the most part these surfaces will be connected and orientable, although we will make occasional mention of the disconnected or non-orientable cases. We can describe the embedding combinatorially in terms of rotations. To begin with, we will do this for embeddings on orientable surfaces. Fix an orientation on the surface, clockwise or anticlockwise, making the surface *oriented*. This orientation induces a cyclic permutation of the edge-ends incident with a vertex v, which we call the *local rotation* at v. A *rotation* is then any product of local rotations (over any/all vertices), which is then a permutation of the set $\mathcal{E}$ of all edge-ends of the underlying graph of the map. Let R be the product of all local rotations. This rotation contains all the information needed to recover the embedding. To see this, let I be the involutory permutation of the set $\mathcal{E}$ that swaps the two ends of each edge. Then the pair (R, I) of permutations of $\mathcal{E}$ completely determines the map. Indeed, orbits of the permutations I, R, and RI can be identified with edges, vertices, and face boundary walks of the map, and their mutual incidence is given by non-empty intersection of the orbits. We may thus identify an oriented map M with the corresponding permutation pair (R, I) acting transitively on the set of half-edges $\mathcal{E}$, and write $M = (\mathcal{E}; R, I)$.

We have seen that embeddings on oriented surfaces correspond to certain pairs of permutations, and vice versa. To understand how this may be extended to arbitrary surfaces, observe that every edge end of an embedded graph has two 'sides' on the supporting surface; these sides are usually called *flags* (or *blades*, see [6]). This way, for every edge of a graph we may associate a set of four flags. Let $\mathcal{F}$ be the set of all flags of the embedded graph; observe that $|\mathcal{F}| = 2|\mathcal{E}|$. Let T be the involutory permutation of $\mathcal{F}$ that interchanges the two flags associated with each edge-end, and let L be the involutory permutation of $\mathcal{F}$ that interchanges the two flags appearing at the same side of each edge. Observe the important relation $LT = TL$. Finally, let C be the involutory permutation of $\mathcal{F}$ that interchanges every two flags forming a 'corner' (two edge-ends meeting at a vertex on the boundary of a face).

If $M = (\mathcal{E}; R, I)$ is an oriented map, and C, L, T are the permutations of $\mathcal{F}$ defined as above, then the two products CT and TL are permutations of $\mathcal{F}$ that represent in a natural sense the effect of the two permutations R and I of edge-ends, respectively. Furthermore, the permutation group generated by CT and TL has two orbits on $\mathcal{F}$, with the two flags associated with each edge-end always lying in different orbits. We note that this description of a map by three involutions is suitable also for maps on non-orientable surfaces, corresponding to the situation where the permutation group generated by CT and TL has a single orbit on $\mathcal{F}$. In any case, we sometimes use the notation $M = (\mathcal{F}; C, L, T)$ if a representation of the map M is necessary in terms of the three involutions C, L, T acting on the flag set $\mathcal{F}$. For more background on algebraic theory of maps, see the survey-type papers [10, 6] and the monograph [9].

A *map isomorphism* $\theta\colon M \to M'$ between two oriented maps M and M' is an isomorphism of the underlying graphs that extends to a homeomorphism of the corresponding surfaces and preserves the set of face boundary walks. In algebraic terms, if $M = (\mathcal{E}; R, I)$ and $M' = (\mathcal{E}'; R', I')$, then a map isomorphism $\theta\colon M \to M'$ will be identified with a bijection $\mathcal{E} \to \mathcal{E}'$ between the corresponding sets of edge-ends, such that $\theta(x^R) = (\theta(x))^{R'}$ and $\theta(x^I) = (\theta(x))^{I'}$ for every edge-end $x \in \mathcal{E}$. Note that by connectedness, a map isomorphism $\theta\colon M \to M'$ is completely determined by the image of any particular edge-end.

Let θ be a map isomorphism from $M = (\mathcal{E}; R, I)$ to itself. Then θ commutes with R and I in the sense explained before, and hence preserves the orientation of the supporting surface. We call any such θ an *orientation-preserving* automorphism of M. The family of all orientation-preserving automorphisms of M forms a group under composition of mappings, called the *orientation-preserving automorphism group* of M, and denoted by $\mathrm{Aut}^+(M)$. By the remark at the end of the previous paragraph we have $|\mathrm{Aut}^+(M)| \le |\mathcal{E}|$, or, equivalently, $|\mathrm{Aut}^+(M)| \le 2|E|$ where E denotes the edge set of the underlying graph of the map. We therefore have an upper bound on the number of orientation-preserving 'symmetries' of an oriented map. If the equality $|\mathrm{Aut}^+(M)| = 2|E|$ is achieved, the map M is called *orientably-regular*. In that case, $\mathrm{Aut}^+(M)$ acts regularly on the edge-ends of M, and M has as much orientation-preserving symmetry as possible.

Given any oriented map $M = (\mathcal{E}; R, I)$, we can form its *oriented mate* $M^{-1} = (\mathcal{E}; R^{-1}, I)$. In general, M and M^{-1} need not be isomorphic, but if they are, then the map M is called *reflexible*. In such a case, an isomorphism $\theta\colon M \to M^{-1}$ with the property that $\theta(x^R) = (\theta x)^{R^{-1}}$ and $\theta(x^I) = (\theta x)^I$ for every $x \in \mathcal{E}$ is called an *orientation-reversing* automorphism of M. From this point on, orientation-preserving and orientation-reversing automorphisms (if any) will be simply called *automorphisms*, and the group of all automorphisms of a map M will be denoted by $\mathrm{Aut}(M)$.

If an oriented map M is reflexible, then $\mathrm{Aut}^+(M)$ is a subgroup of index two in $\mathrm{Aut}(M)$, while if M is not reflexible, then $\mathrm{Aut}^+(M) = \mathrm{Aut}(M)$, and M is called *chiral*. In either case we have $|\mathrm{Aut}(M)| \le 2|\mathcal{E}| = 4|E|$. Oriented maps achieving the equality $|\mathrm{Aut}(M)| = 4|E|$ are called *regular*, since in that case $\mathrm{Aut}(M)$ acts regularly on the sides (or flags) of M. Such maps have as much symmetry as possible.

In the general setting, when a map is represented as $M = (\mathcal{F}; C, L, T)$, a permutation f of $\mathcal{F}$ is an automorphism of M if and only if for every flag $x \in \mathcal{F}$ one has $(x^C)^f = (x^f)^C$, $(x^L)^f = (x^f)^L$, and $(x^T)^f = (x^f)^T$ — that is, if and only if f commutes with all of the three involutions in the algebraic description of M.

In what follows we will extend the concept of symmetry of a map in several ways, depending on certain operations on maps, which we will introduce next.

1.2. **New maps from old.** In this section we describe a number of methods of forming (possibly) new maps from a given *primal map* M. We will first consider operations that do not change the underlying graph of M.

Powers of maps: Taking the oriented mate of a map has a natural generalization, obtained by replacing R^{-1} by any integral power R^e of the rotation R, for e relatively prime to the degree of every vertex of the map. For an oriented map M, the *degree* of M is defined as the least common multiple of all vertex degrees of M. If $M = (\mathcal{E}; R, I)$ is an oriented map of degree d, and e is relatively prime to d, then the map $M^e = (\mathcal{E}; R^e, I)$ is called the *e th power* of M. In terms of flags, if $M = (\mathcal{F}; C, L, T)$, then $M^e = (\mathcal{F}; C_e, L, T)$ where $C_e = (CT)^{e-1}C$. The e th power of M has the same underlying graph as M, but, in general, the supporting surfaces of M and M^e may be different.

One may form up to $\varphi(d)$ powers of an oriented map M of degree d, where φ is the Euler totient function. Some of these powers may be isomorphic to the original map, and we will address this situation in Section 2. We note that the operation of taking a power appears to have been first considered by Wilson in [20].

The Petrie dual: Suppose we begin a walk along some edge e, and when we first encounter a vertex we continue the walk along the edge immediately to our left, then at the next vertex we continue along the edge immediately to our right, and so on, in an alternating left-right manner. Eventually a directed edge is repeated in the same left-right sense, thereafter the walk is periodic. Each period is a closed walk called a *Petrie polygon.* (It might also be called a *left-right walk*, more accurately since it is not strictly a polygon, but we prefer the historic term.) Note that the concept of 'left-right' needs the map M, so this is not just a graph theoretic concept.

Observe that the set of all Petrie polygons cover the set of all edges of the embedded graph, with each edge lying in exactly two polygons. Now consider another embedding of this graph on a surface, such that the Petrie polygons of M are the faces. (It can be shown that such an embedding exists — that is, you never get pinch points — and we will come back to this in a short while.) The resulting map is called the *Petrie dual* (or sometimes the *Petrial*) of M, and denoted by M^P. The Petrie dual might not be on the same surface as the primal map M, and it need not even be orientable; for example, the reader is invited to prove that that the Petrie dual of the tetrahedral map on the sphere is a quadrangulation of the projective plane.

If a map M is described in terms of three involutory permutations in the form $M = (\mathcal{F}; C, L, T)$, then its Petrie dual is $M^P = (\mathcal{F}; C, LT, T)$. This description is valid for all maps.

Powers and the Petrie operation represent ways of modification of an embedding without changing the underlying graph. We will now discuss two more operations on maps that may change the underlying graph as well.

Geometric duals: As is commonly known, the *geometric dual* M^* of a map M can be formed from the embedding of the underlying graph by letting the faces of M become the vertices of M^*, and letting each edge e of M become a dual edge e^* joining the faces of M on either side of e, so that the vertices of M become the faces of M^*. If the supporting surface for M is oriented, then the rotation on the dual is given by the order of edges around the primal faces. In algebraic terms, if M and M^* are represented by rotations R and R^* and involutions I and I^*, then $R^* = RI$ and $I^* = I$. More generally, if the primal M is $(\mathcal{F}; C, L, T)$ then the dual M^* is $(\mathcal{F}; C, T, L)$ — that is, with $\mathcal{F}^* = \mathcal{F}$, $C^* = C$, $L^* = T$, and $T^* = L$. In other words, the flag sets are the same, and the involutions describing the corners are the same, but the concept of moving 'longitudinally' by L along an edge is swapped with that of moving 'transversally' by T across an edge. Under either point of view, we easily see that the process of taking the dual is involutory, so we may say "The dual of the dual is the primal".

It is interesting to compare algebraic descriptions of the primal map $M = (\mathcal{F}; C, L, T)$ with its dual M^* and Petrie dual M^P. These are obtainable by replacing the ordered pair (L, T) by (T, L) and (LT, T) respectively. Since L and T are commuting involutions, the group $\langle L, T\rangle$ is isomorphic to the Klein four-group $\mathbb{Z}_2 \times \mathbb{Z}_2$, with automorphism group $\mathrm{Aut}(\mathbb{Z}_2 \times \mathbb{Z}_2) \cong S_3$, and the assignments

$(L,T) \mapsto (T,L)$ and $(L,T) \mapsto (LT,T)$ give two of its six automorphisms. The third involutory automorphism comes from the assignment $(L,T) \mapsto (L,LT)$, which corresponds to what is known as the *opposite* of the map M. The other two non-trivial automorphisms (each of order 3) may be called *trialities*. Disregarding the identity, we obtain a total of five non-trivial operators on maps in this way; these were introduced by Wilson in [20] and further studied in [21, 11].

The primal map, the geometric dual map, and the Petrie dual map can be described in a succinct way through the concept of a 'medial' map, as follows.

Medial maps: Let M be a map on a surface S based on an embedded graph G. We form the *medial graph* of M, denoted by $\mathrm{med}(M)$, by taking as vertices the edges of M, and then joining two vertices by an edge in $\mathrm{med}(M)$ if they represent consecutive edges in a face of M (or, equivalently, if they represent two consecutive edges in the rotation around a vertex of M^*). The medial graph embeds in S in a natural way which we call the *medial map of* M, and we also denote by $\mathrm{med}(M)$. The medial graph is 4-regular and face-2-colorable: one class of faces corresponds to the vertices of M, say the *black* faces, and the other class corresponds to the faces of M, say the *white faces*. Observe that both the black faces and the white faces partition the edges of the medial. Also $\mathrm{med}(M)$ is isomorphic to $\mathrm{med}(M^*)$, with the only difference between them being that the role of black faces and white faces are interchanged. Moreover, any 4-regular face-2-colorable map on a surface is the medial map of a pair (M, M^*).

The Petrie dual of M can also be described in terms of $\mathrm{med}(M)$. Recall that the medial map is 4-regular. A *straight-ahead walk*, or *SAW*, in the medial map is formed by walking along a edge, and whenever a vertex is encountered, continuing not along the left or right edge, but directly along the opposite edge. The set of SAWs partition the edges of the medial graph. We may call the set of SAWs the *grey faces* of the medial. The medial map of the Petrie dual can be formed by taking the embedding of the medial graph using the white faces and the grey faces. Hence the medial graphs of M and M^P are the same, but are embedded differently.

2. Our main results

We will be interested in maps that have the 'absolutely highest level of symmetry'. What should this be? To begin with, one should require such a map to be regular — that is, to have the largest possible number of map automorphisms. Recall from Section 1.1 that a finite map $M = (\mathcal{F}; C, L, T)$ is regular if and only if it has $|\mathcal{F}|$ automorphisms. But there are two more ways of thinking about symmetries of maps.

In Section 1.2 we introduced operations that form the geometric dual and the Petrie dual of a map, and mentioned the six Wilson operators [20] generated by these two types of duality. A natural way to proceed is to consider maps that are invariant with respect to Wilson's operations. We say that the map M is *self-dual* or *self-Petrie* if M is isomorphic to M^* or M^P, respectively. The map M is said to have *trinity symmetry* if it is both self-dual and self-Petrie, that is, if the maps M, M^* and M^P are pairwise isomorphic.

From the algebraic description, we immediately obtain the observation that if M has trinity symmetry, then all the six maps obtained by Wilson's operations are isomorphic to each other [20]. Also it is obvious from the definition of regularity that if M is regular, then so are the maps M^* and M^P, as well as the remaining three maps obtained by Wilson's operations. Regular maps with trinity symmetry could therefore be considered to be the 'most symmetric' maps with respect to automorphisms and dualities. Infinite families of such regular maps were constructed in [15].

But can one look for even more symmetries in regular maps with trinity symmetry? In Section 1.2 we discussed powers of maps. If a power of a map is isomorphic to the original map, this isomorphism can be viewed as an additional 'external symmetry' of the map.

Following [14], we call a non-zero integer e an *exponent* of the map M if the e th power map M^e is isomorphic to M. Since a product of two exponents is again an exponent, we can speak of the *exponent group* of a map. If d is the degree of M, then the exponent group of M is isomorphic to a subgroup of $\mathbb{Z}_d^*$, the multiplicative group of units in the ring $\mathbb{Z}_d$.

Maps that have 'all possible exponents' are of special interest. We call a map M *kaleidoscopic* if every integer e relatively prime to the degree of M is an exponent of M. Thus a kaleidoscopic map M of degree d has $\varphi(d)$ exponents, and its exponent group is isomorphic to $\mathbb{Z}_d^*$. In particular, since $-1(\equiv d-1 \bmod d)$ is relatively prime to d, a kaleidoscopic map has to be reflexible. A 'residual finiteness' construction for kaleidoscopic regular maps was presented in [18].

Maps that are regular, kaleidoscopic, and have trinity symmetry, may therefore be designated as the 'absolutely most symmetric maps'. But are there any such maps at all? Well, yes: a trivial example is a cycle of length two embedded on the sphere. More complex examples are not easy to find. Using our main results we give many such examples in Section 7.

We now state our first main result.

Theorem 2.1. *If there is an oriented regular kaleidoscopic map M of degree d with trinity symmetry, then for any integer $n \geq 2$, there is an oriented regular kaleidoscopic map of degree dn with trinity symmetry and automorphism group isomorphic to $(\mathbb{Z}_n)^{1+\varepsilon} \rtimes \mathrm{Aut}(M)$, where ε is the number of edges of M.*

Applying Theorem 2.1 to the spherical embedding of a cycle of length two, we obtain a family of examples in which we have, additionally, a full description of the automorphism groups of the maps. We state this as a separate result.

Theorem 2.2. *For every integer $n \geq 1$, there is a map M_n of degree $2n$ with $2n^2$ vertices on an orientable surface of genus $n^3 - 2n^2 + 1$ such that:*

(i) *M_n is regular, kaleidoscopic, and has trinity symmetry, and*
(ii) *the automorphism group of M_n has order $8n^3$ and defining presentation*

$$\mathrm{Aut}(M_n) = \langle\, a, b, c, z \mid a^2,\ b^2,\ c^2,\ z^2,\ abc,\ (az)^{2n},\ (bz)^{2n},\ (cz)^{2n},\ (azbzcz)^2 \,\rangle.$$

We note that it was suggested a long time ago by Wilson in the course of preparation of his doctoral dissertation [20] that the group with the above presentation has order $8n^3$ and is the automorphism group of a regular map with trinity symmetry, and that this was checked by the same author by computer for all $n \leq 50$ [22]. We prove it (and the fact that each M_n is kaleidoscopic) for all n.

The rest of the paper is concerned with the proof of both of the two theorems above, along with development of corresponding theory.

We conclude this section with a note on the history of the problem. Wilson [20] appears to be the first to have considered, back in the 1970's, constructions of regular maps with extra symmetry properties – namely, orientably-regular maps that are self-dual and self-Petrie, and orientably regular maps isomorphic to their powers. Although significant advances in the theory of orientably-regular and regular maps were made in the 1980s and 1990s (see e.g. [13, 17] for surveys), including new methods of constructing spherical self-dual maps [2] and a detailed treatment of exponents [14], a non-trivial example of a kaleidoscopic regular map with trinity symmetry was still beyond reach. This can be explained by the results of this article: the first non-trivial (that is, admitting exponents other than ± 1) example of a kaleidoscopic regular map with trinity symmetry has genus 9 and a census covering regular maps of such a genus was not available until 2001 [7]. Constructions of infinite families of kaleidoscopic orientably-regular maps and orientably-regular maps with trinity symmetry, based on residual finiteness of groups, were given in [18] and [15], but despite their algebraic similarity, it appears impossible to unify them to yield the 'absolutely most symmetric maps' as furnished by the main results of this paper.

3. Voltages, lifts, and automorphisms

The proofs of our main results involve regular coverings of maps, with branch points at vertices as well as face centers. These can be constructed by means of corner voltage assignments, as introduced in [3] and studied in extended generality in [1]. We first briefly describe the techniques involved, including lifts of automorphisms.

Let $M = (\mathcal{F}; C, L, T)$ be a map and let H be a group, called the *voltage group*. A *corner voltage assignment* on M in H is any mapping $\alpha : \mathcal{F} \to H$ such that $\alpha(x^C) = (\alpha(x))^{-1}$ for every flag $x \in \mathcal{F}$. The pair (M, α) gives rise to a *lift* M^α of the map M, defined as follows. Let $\mathcal{F}' = \mathcal{F} \times H$ and let C', L' and T' be permutations of the set $\mathcal{F}'$ defined by

$$(x, g)^{C'} = (x^C, g\alpha(x)),\ (x, g)^{L'} = (x^L, g) \text{ and } (x, g)^{T'} = (x^T, g) \text{ for all } (x, g) \in \mathcal{F}'.$$

The action of the permutation group $\langle C', L', T' \rangle$ on the set $\mathcal{F}'$ need not be transitive. Every connected component of this action determines a map, and the connected components are pairwise isomorphic maps. Hence we may denote by M^α any such connected component, and still write $M^\alpha = (\mathcal{F}'; C', L', T')$, with the understanding that the action of the group $\langle C', L', T' \rangle$ refers to it action on a connected component — that is, to some subset of the flag set $\mathcal{F}'$ on which the group acts transitively.

To describe the connected components we first need to introduce one further concept. A *flag-walk* in the map $M = (\mathcal{F}; C, L, T)$ is a sequence $W = (x_0, x_1, \ldots, x_k)$ of flags with the property that for $1 \le i \le k$ we have $x_i = {x_{i-1}}^{X_i}$ for some $X_i \in \{C, L, T\}$. The flag-walk W is *closed* if $x_k = x_0$, and *based at* x if x is the initial flag x_0. The *voltage* $\alpha(W)$ of W is defined as the product $\alpha(W) = a_1 a_2 \ldots a_k$ of elements of the voltage group H, where $a_i = 1$ if $x_i = {x_{i-1}}^L$ or ${x_{i-1}}^T$, and $a_i = \alpha(x_{i-1})$ if $x_i = {x_{i-1}}^C$, for $1 \le i \le k$. Now for any flag $x \in \mathcal{F}$, let H_x be the set of all voltages $\alpha(W)$ of *closed* flag-walks W based at x. Then H_x is a subgroup of H, sometimes called the *local group at* x. All such local groups are conjugate to each other in H, and the index of each H_x in H is equal to number of (pairwise isomorphic) connected components to which the notation M^α refers. The reason for this (see [3]) is that flags (x, h) and (x, h') lie in the same component of the lift if and only if $h' = h\alpha(W)$ for some close walk at x, and this happens if and only if h' and h are in the same coset of the local group.

Next, let $\pi_\alpha : M^\alpha \to M$ be the *covering projection* associated with the lift α, that takes each $(x, g) \in \mathcal{F} \times H$ to $x \in \mathcal{F}$. This is a map homomorphism in the sense that $\pi_\alpha X = X'\pi_\alpha$ for each $X \in \{C, L, T\}$, and gives a *regular covering* of M by M^α (and a regular covering of the corresponding surfaces) in terms of algebraic topology.

For completeness, let us mention that corner voltage assignments are equivalent to ordinary voltage assignments on the underlying graph of the medial map of M. Indeed, the edges of $\mathrm{med}(M)$ correspond to 2-element flag-sets of the form $\{x, x^C\}$, effectively consisting of the ends of an edge. Note that it is difficult to work with exponents of lifted maps in terms of medial maps, since the operation of taking a power of a map destroys the medial; hence we prefer working with flags.

We now turn to the question of 'lifting' map isomorphisms and automorphisms. The reason for this is that our construction will employ corner voltage assignments on a map with trinity symmetry and all possible exponents, in order to lift to a covering map that has the same properties. To do this, we need to ensure that isomorphisms between the various duals and powers will lift to isomorphisms of the covering maps. We also want the covering maps to be regular, so we need to ensure that the automorphisms in the regular base map will lift to automorphisms in the covering map.

To make the concept of lifting of a map isomorphism more precise, let M and N be maps and let $f : M \to N$ be a map isomorphism. Let α and β be corner voltage assignments on M and N in the same group H, and let π_α and π_β be the respective covering projections. Then a mapping $\tilde{f} : M^\alpha \to M^\beta$ is said to be a *lift* of f if $f\pi_\alpha = \pi_\beta \tilde{f}$.

The special case where $M = N$ and $\alpha = \beta$ is particularly important. For any $h \in H$, the bijection i_h defined on flags of M^α by $i_h(x, g) = (x, hg)$ is an automorphism of M^α, and is a lift of the identity automorphism of M. Such automorphisms are known as *deck transformations*, and these form a group isomorphic to the voltage group. Also as a general consequence of the theory of lifts [3] we know that if a map isomorphism $f : M \to N$ lifts onto an isomorphism $\tilde{f} : M^\alpha \to M^\beta$, then *all* such lifts of f have the form $i_h\tilde{f}$ where h ranges over all elements of the voltage group. In particular, if an isomorphism has a lift, then the number of its lifts is equal to the order of the voltage group H.

We note that the concept of a lift can be defined in greater generality for map homomorphisms, but this is not of concern to us here.

Tools for lifting map isomorphisms and automorphisms are prevalent in the literature. The best one for us comes from the following result, which is a slight modification of Theorem 9 of [3]; see also Propositions 6 and 7 in [1].

Theorem 3.1. *Let M and N be maps with corner voltage assignments α and β in some group H. Then a map isomorphism $f : M \to N$ lifts to a map isomorphism $\tilde{f} : M^\alpha \to N^\beta$ if and only if for any closed flag-walk W with origin at a fixed flag the following condition is satisfied: $\alpha(W) = 1_H$ if and only if $\beta(f(W)) = 1_H$.* □

In particular, when $M = N$ and $\alpha = \beta$ this gives a criterion for lifting automorphisms from M to M^α, namely that the automorphism of M has to preserve the set of walks with net voltage 1_H. It also has nice implications for lifting regular maps. If the map M is regular, with flag set $\mathcal{F}$, and *all* of the $|\mathcal{F}|$ automorphisms of M lift to automorphisms of the map M^α (obtained from M by a corner voltage assignment $\alpha\colon \mathcal{F} \to H$), then every automorphism f of M lifts onto $|H|$ automorphisms of M^α of the form $i_h\tilde{f}$; it then follows (with the help of some more general theory from [3]) that the collection of all such lifted automorphisms of M^α forms a group that acts transitively on the flags of M^α, so M^α is also a regular map. We state this as a separate result; see [3] for details.

Theorem 3.2. *Let M be a regular map with a corner voltage assignment α in a group H. If all automorphisms of M lift to automorphisms of M^α, then M^α is a regular map.* □

4. Lifts and exponents

Let $M = (\mathcal{F}; C, L, T)$ be a map and let M^α be a regular lift of M, that is, a regular covering of M. If e is an integer relatively prime to both the degree of M and the degree of M^α, then $(M^\alpha)^e$ is a regular covering of M^e, and it follows that there is a voltage assignment β on M^e such that $(M^e)^\beta$ is isomorphic to $(M^\alpha)^e$. In this section we work out a 'canonical' form for β, which will be useful later.

Let α be a corner voltage assignment on M in some group H. Recall that this means that for any flag-walk $W = (x_0, x_1, \dots x_k)$ in M we have $\alpha(W) = a_1 a_2 \dots a_k$ where $a_i = 1$ if $x_i = {x_{i-1}}^L$ or ${x_{i-1}}^T$, and $a_i = \alpha(x_{i-1})$ if $x_i = {x_{i-1}}^C$. Also let D denote the degree of M — that is, the least common multiple of the degrees of all vertices of M.

To compute the degree of the lift M^α, we first need to determine the degrees of individual vertices of the lift. By general theory, this can be done as follows. Let v be a vertex of M of degree d, let x be any flag at v, and let o be the order of the element $\alpha(v; x) = \alpha(x)\alpha(x^{CT}) \dots \alpha(x^{(CT)^{d-1}})$ in the group H; this order may depend on v but not on the choice of the flag x at v. Then any vertex of M^α that is a lift of v has degree od in M^α. It follows that the degree D^α of M^α is the least common multiple of such products od taken over all vertices of M. Hence, in particular, D^α is a multiple of D.

We are now ready to discuss powers of M. Since our interest is in powers that work for both M and M^α, consider an *arbitrary* positive integer e relatively prime to D^α; then also e is relatively prime to D. As noted earlier, the eth power M^e of M is the map $(\mathcal{F}; C_e, L, T)$ where $C_e = (CT)^{e-1}C$. Define a corner voltage assignment $\beta = \alpha_e$ on M^e by setting

$$\beta(x) = \alpha_e(x) = \alpha(x)\alpha(x^{CT}) \cdots \alpha(x^{(CT)^{e-1}}) \ \text{ for each flag } \ x \in \mathcal{F}.$$

This effectively assigns the pair (M^e, α_e) to the pair (M, α). Note that the integer e depends on both M and α. We would like to have an assignment of the form $(M, \alpha, e) \mapsto (M^e, \alpha_e, f)$, for some suitably defined integer f. An extra bonus would accrue if this correspondence was involutory, meaning that one could interchange the roles of M and M^e and have $(M^e)^f = M$ and $(\alpha_e)_f = \alpha$. We will show that this can be done.

Note first that the power e was taken as an arbitrary integer relatively prime to D^α, the degree of M^α. Hence a natural first step is to find the degree D^{α_e} of $(M^e)^{\alpha_e}$. Again by taking a flag x at a vertex v of degree d in M^e, we see that the degree of any lift of v in $(M^e)^{\alpha_e}$ is d times the order of the element $\alpha_e(v; x) = \alpha_e(x)\alpha_e(x^{C_eT}) \dots \alpha_e(x^{(C_eT)^{d-1}})$ in G.

Lemma 4.1. *The degrees of M^α and $(M^e)^{\alpha_e}$ are the same, that is, $D^\alpha = D^{\alpha_e}$.*

Proof. Let $y_j = x^{(CT)^{ej}}$ for $0 \le j \le d-1$. Since $C_e = (CT)^{e-1}C$, we have $C_eT = (CT)^e$. We evaluate $\alpha_e(v;x)$ with the help of the fact that $\alpha(x^{(CT)^d}) = \alpha(x)$ as follows:

$$\begin{aligned}
\alpha_e(v;x) &= \alpha_e(x)\alpha_e(x^{C_eT})\cdots\alpha_e(x^{(C_eT)^{d-1}}) \\
&= \alpha_e(x)\alpha_e(x^{(CT)^e})\cdots\alpha_e(x^{((CT)^e)^{d-1}}) \\
&= \alpha_e(y_0)\alpha_e(y_1)\cdots\alpha_e(y_{d-1}) \\
&= \textstyle\prod_i \alpha(y_0{}^{(CT)^i})\prod_i \alpha(y_1{}^{(CT)^i})\cdots\prod_i \alpha(y_{d-1}{}^{(CT)^i}) \\
&= \textstyle\prod_i \alpha(x^{(CT)^i})\prod_i \alpha(x^{(CT)^{e+i}})\cdots\prod_i \alpha(x^{(CT)^{e(d-1)+i}}) \\
&= \textstyle\prod_j \alpha(x^{(CT)^j})\prod_j \alpha(x^{(CT)^{d+j}})\cdots\prod_j \alpha(x^{(CT)^{d(e-1)+j}}) \\
&= [\alpha(x)\alpha(x^{CT})\ldots\alpha(x^{(CT)^{d-1}})]^e \\
&= [\alpha(v;x)]^e,
\end{aligned}$$

where the products $\prod_i$ and $\prod_j$ are taken over the ranges $0 \le i \le e-1$ and $0 \le j \le d-1$, respectively. Since e is relatively prime to d, the elements $\alpha(v;x)$ and $\alpha_e(v;x) = [\alpha(v;x)]^e$ have the same order in H. Therefore by the theory explained earlier, the lifted maps M^α and $(M^e)^{\alpha_e}$ have the same degree. □

We may now carry through our plan as indicated. Let f be an *arbitrary* positive integer such that $ef \equiv 1 \bmod D^\alpha$, that is, such that $ef = tD^\alpha + 1 = tD^{\alpha_e} + 1$ for some positive integer t. By the definition of powers of maps, it follows that $(M^e)^f$ is isomorphic to M, with an isomorphism provided by the identity mapping on the common set of flags $\mathcal{F}$. Next, given the voltage assignment β on M^e as above, let us by analogy introduce a corner voltage assignment $\gamma = \beta_f$ on $M = (M^e)^f$ by letting

$$\gamma(x) = \beta_f(x) = \beta(x)\beta(x^{C_eT})\cdots\beta(x^{(C_eT)^{f-1}})$$

for any $x \in \mathcal{F}$. Since $(M^e)^f$ is isomorphic to M, successive application of the exponent e to M followed by f to M^e does not change the map M. We show that a similar thing holds for the voltage assignments just introduced.

Lemma 4.2. *With the above notation, we have $\alpha = \gamma$, or equivalently, $(\alpha_e)_f = \alpha$.*

Proof. We carry out a computation similar to the one in the proof of Lemma 4.1. Let $x \in \mathcal{F}$ be a flag at a vertex v of degree d in M, and let o be the order of the element $\alpha(v;x) = \alpha(x)\alpha(x^{CT})\ldots\alpha(x^{(CT)^{d-1}})$ in G. Then since od divides D^α, we see that $ef = tD^\alpha + 1 = ods + 1$ for some integer s. Now let $y_j = x^{(CT)^{ej}}$ for $0 \le j \le f-1$. Then using our description of the voltage assignments β and γ and the fact that $C_e = (CT)^{e-1}C$, we successively obtain the following:

$$\begin{aligned}
\gamma(x) &= \beta(x)\beta(x^{C_eT})\cdots\beta(x^{(C_eT)^{f-1}}) \\
&= \alpha_e(x)\alpha_e(x^{C_eT})\cdots\alpha_e(x^{(C_eT)^{f-1}}) \\
&= \alpha_e(x)\alpha_e(x^{(CT)^e})\cdots\alpha_e(x^{((CT)^e)^{f-1}}) \\
&= \alpha_e(y_0)\alpha_e(y_1)\cdots\alpha_e(y_{f-1}) \\
&= \textstyle\prod_i \alpha(y_0{}^{(CT)^i})\prod_i \alpha(y_1{}^{(CT)^i})\cdots\prod_i \alpha(y_{f-1}{}^{(CT)^i}) \\
&= \textstyle\prod_i \alpha(x^{(CT)^i})\prod_i \alpha(x^{(CT)^{e+i}})\cdots\prod_i \alpha(x^{(CT)^{e(f-1)+i}}),
\end{aligned}$$

where the products $\prod_i$ are taken over the range $0 \le i \le e-1$. The resulting expression has $ef = ods+1$ constituents, with $os+1$ terms of the form $\alpha(x^{(CT)^{d\ell}})$ for $0 \le \ell \le os$, and for $1 \le m \le d-1$ it contains os terms of the form $\alpha(x(CT)^{d\ell+m})$

for $0 \le \ell \le os-1$. Note that the very last term $\alpha(x^{(CT)^{e(f-1)+(e-1)}})$ is simply equal to $\alpha(x)$. Also using $x^{(CT)^d} = x$ we see that for any fixed m, the os values $\alpha(x^{(CT)^{d\ell+m}})$ for $0 \le \ell \le os-1$ are equal to each other. Accordingly, we find that

$$\gamma(x) = [\alpha(x)\alpha(x^{CT})\ldots\alpha(x^{(CT)^{d-1}})]^{os} \cdot \alpha(x) = [\alpha(v;x)]^{os}\alpha(x)\ .$$

But o is the order of the element $\alpha(v;x)$, therefore $\gamma(x) = \alpha(x)$, and then since x was an arbitrary flag of M, we deduce that $(\alpha_e)_f = \gamma = \alpha$, as claimed. □

This enables us to prove the following basic result on exponents and lifts. In both its statement and proof we use the notation introduced earlier.

Proposition 4.3. *The maps $(M^\alpha)^e$ and $(M^e)^{\alpha_e}$ are isomorphic.*

Proof. We have $M = (\mathcal{F}; C, L, T)$ and $M^e = (\mathcal{F}; C_e, L, T)$, where $C_e = (CT)^{e-1}C$. Both lifts via α have the same flag set $\mathcal{F}' = \mathcal{F} \times H$, and $M^\alpha = (\mathcal{F}'; C', L', T')$ where

$$(x,g)^{C'} = (x^C, g\alpha(x)),\ (x,g)^{L'} = (x^L, g) \text{ and } (x,g)^{T'} = (x^T, g)$$

for each flag $(x,g) \in \mathcal{F}'$. It follows that $(M^\alpha)^e = (\mathcal{F}'; (C')_e, L', T')$ where $(C')_e = (C'T')^{e-1}C'$. On the other hand, we have $(M^e)^{\alpha_e} = (\mathcal{F}'; (C_e)', L', T')$, with the same L' and T' as for M^α and with $(C_e)'$ taking a flag (x,g) to the flag $(xC_e, g\alpha_e(x))$.

We show that the identity mapping on $\mathcal{F}'$ provides an isomorphism $(M^\alpha)^e \to (M^e)^{\alpha_e}$. Since L' and T' are common to both maps, all we need to do is show that $(C')_e = (C_e)'$. But this happens simply because for any flag $(x,g) \in \mathcal{F}'$ we have

$$(x,g)^{(C')_e} = (x,g)^{(C'T')^{e-1}C'} = (x^{(CT)^{e-1}C}, g\alpha(x)\alpha(x^{CT})\ldots\alpha(x^{(CT)^{e-1}})),$$

which gives $(x,g)^{(C')_e} = (x^{(CT)^{e-1}C}, g\alpha_e(x)) = (x^{C_e}, g\alpha_e(x)) = (x,g)^{(C_e)'}$ for all (x,g). □

Thus, applying a power to a lift of a map is equivalent to applying the lifting construction to the same power of the map, although under a different corner voltage assignment (the explicit form of which will be helpful in proving our main results).

5. Proof of Theorem 2.1

Theorem 2.1 asserts that if there is an oriented, regular, kaleidoscopic map of degree d with trinity symmetry, then for any integer $n \ge 2$ there exists an oriented, regular, kaleidoscopic map of degree dn with trinity symmetry. We now prove this fact.

Let M be a kaleidoscopic regular map with trinity symmetry on an orientable surface. This means that the regular map has all the possible additional external symmetries: it is self-dual and self-Petrie, and admits every e relatively prime to the degree d of the map as an exponent. We will work with the algebraic representation $M = (\mathcal{F}; C, L, T)$.

Let n be any integer greater than 1, and let $H = \mathbb{Z}_n \times \ldots \times \mathbb{Z}_n$ be the direct product of $|\mathcal{F}|/2$ copies of $\mathbb{Z}_n$, under componentwise addition.

We will represent elements of H as follows. First, partition the set of flags $\mathcal{F}$ into $|\mathcal{F}|/2$ two-element subsets of the form $\{x, x^C\}$ (the pair associated with a corner of M) for $x \in \mathcal{F}$, and let $\mathcal{F}_2$ denote the set of all such two-element subsets. Note that $\mathcal{F}_2$ corresponds precisely to the edge set of the medial map of M. We will now consider elements of H as $|\mathcal{F}|/2$-tuples with entries in $\mathbb{Z}_n$, indexed by the set $\mathcal{F}_2$. Any element of H can then be written in the form $(g_z)_{z\in\mathcal{F}_2}$. Also for any particular $z'\in\mathcal{F}_2$ we define the *unit vector* $[z']$ to be the element $(g_z)_{z\in\mathcal{F}_2}$ of H given by $g_z=1$ if $z=z'$ and $g_z=0$ for $z\neq z'$.

Let us now define a corner voltage assignment α on M with values in H, by setting $\alpha(x)=[z]$ and $\alpha(x^C)=-[z]$ whenever $z=\{x,x^C\}\in\mathcal{F}_2$. Note that α is injective on $\mathcal{F}$; and moreover, $\alpha(x)$ and $\alpha(x')$ are independent unit vectors whenever x and x' are flags associated with different corners.

Next, consider the lift M^α of M determined by α. By equivalence of corner voltage assignments on the map M with ordinary voltage assignments on the medial map $\mathrm{med}(M)$, our voltage assignment α on M is equivalent to an ordinary voltage assignment on the edges of $\mathrm{med}(M)$ having zero voltages on a spanning tree. Since $\mathrm{med}(M)$ has $|\mathcal{F}|/2$ edges and $|\mathcal{F}|/4$ vertices, the equivalent voltage assignment on edges of $\mathrm{med}(M)$ generates a subgroup of index $n^{|\mathcal{F}|/4-1}$ in H. Thus M^α has $n^{|\mathcal{F}|/4-1}$ (pairwise isomorphic) connected components. From now on, let $\tilde{M}$ be a fixed connected component of M^α. Then the number of flags of $\tilde{M}$ is $|\mathcal{F}|\cdot|H|/n^{|\mathcal{F}|/4-1}=|\mathcal{F}|\cdot n^{|\mathcal{F}|/2-|\mathcal{F}|/4+1}$.

Observe that $\tilde{M}$ is an orientable map — which can be seen (without considering orbits) by noting that all coverings of an orientable map have orientable supporting surfaces. The degree of $\tilde{M}$ can be determined as in Section 4. Since every vertex of M has degree d, the degree of every vertex of $\tilde{M}$ is od where o the order of the element $\alpha(x)\alpha(x^{CT})\ldots\alpha(x^{(CT)^{d-1}})$ in the group H. Since the values of $\alpha(x^{(CT)^i})$ for $0\le i\le d-1$ are linearly independent unit vectors, we find that $o=n$, so the degree of $\tilde{M}$ is dn.

In the remaining part of the proof, we show that $\tilde{M}$ is a kaleidoscopic regular map with trinity symmetry. To do this, we use Theorems 3.1 and 3.2 from Section 3 to prove that all the relevant map automorphisms and isomorphisms lift.

What we need to check are cases where walks with zero voltage are taken by the map automorphisms and isomorphisms to walks of zero voltage. In what follows, we let W be an arbitrary closed flag-walk of M, of the form $W=(x_0,x_1,\ldots,x_{k-1},x_k)$ where $x_i\in\mathcal{F}$ and $x_k=x_0$. Then the voltage of this given walk is $\alpha(W)=a_1+a_2+\ldots+a_{k-1}+a_k\in H$, where $a_i=0_H$ if $x_i={x_{i-1}}^L$ or ${x_{i-1}}^T$, while $a_i=\alpha(x_{i-1})$ if $x_i={x_{i-1}}^C$.

We begin by considering regularity. Let f be any automorphism of M. Then we have $f(W)=(f(x_0),f(x_1),\ldots,f(x_{k-1}),f(x_k))$, and $\alpha(f(W))=a'_1+a'_2+\ldots+a'_{k-1}+a'_k$ where $a'_i=1$ if $f(x_i)=(f(x_{i-1}))^L$ or $(f(x_{i-1}))^T$, and $a'_i=\alpha(f(x_{i-1}))$ if $f(x_i)=(f(x_{i-1}))^C$. Let $a=a_i$ for some i (with $1\le i\le k$), and let x be the (unique) flag of M for which $\alpha(x)=a$. Also let $a'=a'_i$ be the corresponding value in the voltage sum for $\alpha(f(W))$, and let $x'=f(x)$ be the flag of M for which $\alpha(x')=\alpha(f(x))=a'$. Next let m_1 and m_2 be the number of occurrences of a and $-a$, respectively, in the sum $\alpha(W)=a_1+a_2+\ldots+a_{k-1}+a_k$. Now assume that $\alpha(W)=0$. Then by the definition of α on individual flags (and in particular, the linear independence of α-images of flags from different corners), we find that $m_1-m_2\equiv 0 \bmod n$. But m_1 and m_2 are also the number of times that a' and $-a'$ occur in the sum $\alpha(f(W))=a'_1+a'_2+\ldots+a'_{k-1}+a'_k$, and the occurrences of $\pm a'$ in $\alpha(f(W))$ sum to 0. Since this argument can also be reversed (using the fact that

f is an automorphism), it shows that $\alpha(W) = 0$ if and only if $\alpha(f(W)) = 0$. But f was an arbitrary automorphism of M, so we conclude from Theorem 3.2 that $\tilde{M}$ is a regular map.

We now consider duality and Petrie duality. This time, let $f\colon M \to M^*$ be a map isomorphism, where $M^* = (\mathcal{F}; C^*, L^*, T^*)$ is the geometric dual of M, with $C^* = C$, $L^* = T$, and $T^* = L$. By the definition of the lift of a map, it is immediately obvious that the lift of the dual $(M^*)^\alpha$ can be identified with the dual of the lift $(M^\alpha)^*$. Since flag-walks in M and M^* are the same, and voltage on them depends only on adjacent pairs of the form x and $x^C = x^{C^*}$, we find that $\alpha(W) = 0$ if and only of $\alpha(f(W)) = 0$, by the same arguments as those developed in the previous paragraph. In an entirely similar manner, any isomorphism of M onto its Petrie dual $M^P = (\mathcal{F}; C^P, L^P, T^P) = (\mathcal{F}; C, LT, T)$ lifts onto an isomorphism of M^α onto its Petrie dual $(M^\alpha)^P$, which can be identified with the lift $(M^P)^\alpha$ of M^P. Hence if M is self-dual (respectively self-Petrie), then so is $\tilde{M}$.

It remains for us to show that $\tilde{M}$ has all arithmetically feasible exponents, whenever M does.

Let e be a positive integer relatively prime to dn, the degree of $\tilde{M}$. Then e is also relatively prime to d. We will assume that there exists an isomorphism $f\colon M \to M^e$ from M to its eth power $M^e = (\mathcal{F}; (CT)^{e-1}C, L, T)$. Let us define a corner voltage assignment $\beta = \alpha_e$ on M^e exactly as in Section 4, but in additive notation, so that $\beta(x) = \alpha(x) + \alpha(x^{CT}) + \cdots + \alpha(x^{(CT)^{e-1}})$ for every $x \in \mathcal{F}$. We show that f lifts onto an isomorphism from M^α to $(M^e)^\beta$.

So suppose $\alpha(W) = 0$ where W is a closed flag-walk as above. Again, consider a particular entry a in the sum $\alpha(W) = a_1 + a_2 + \ldots + a_{k-1} + a_k$, and let x be the flag in W for which $a = \alpha(x)$. If there are m_1 occurrences of the flags x and x^C (in that order) in consecutive positions in the flag-walk W, and m_2 such occurrences of the same two flags in reverse order, then in the sum $\alpha(W) = a_1 + a_2 + \ldots + a_{k-1} + a_k$ we have m_1 of the a_i equal to a and m_2 of the a_i equal to $-a$. Then since $\alpha(W) = 0$, and since α-images of flags from different corners are independent, this implies that $m_1 \equiv m_2 \bmod n$.

We now look at the effect that this has on $f(W)$ under the isomorphism $f\colon M \to M^e$. Let $f(W) = (y_0, y_1, \ldots, y_{k-1}, y_k)$. Then $\beta(f(W)) = b_1 + b_2 + \ldots + b_{k-1} + b_k$, where $b_i = 0$ if $y_i = {y_{i-1}}^L$ or ${y_{i-1}}^T$ while $b_i = \beta(y_{i-1})$ if $y_i = {y_{i-1}}^{(CT)^{e-1}C}$. Also let $y = f(x)$, where x is the flag considered in the previous paragraph, and let $b = \beta(y)$. By the definition of β, we have $\beta(y) = \alpha(y) + \alpha(y^{CT}) + \cdots + \alpha(y^{(CT)^{e-1}})$. Then since in $\alpha(W)$ we had m_1 summands equal to $\alpha(x)$ and m_2 summands equal to $-\alpha(x)$, it follows that in $\alpha(f(W))$ we have m_1 occurrences of $\alpha(y^{(CT)^j})$ and m_2 occurrences of $-\alpha(y^{(CT)^j})$, for $0 \le j \le i-1$, and since $m_1 \equiv m_2 \bmod n$, these occurrences sum to zero. Once again, because the values of α on flags from different corners are linearly independent, the above arguments show that if $\alpha(x)$ contributes net value zero to $\alpha(W)$, then $\beta(f(x))$ also contributes net value zero to $\beta(f(W))$, for all x. Hence if $\alpha(W) = 0$, then also $\beta(f(W)) = 0$.

Finally, note that we have defined β to be equal to α_e in the notation of Section 4. By Lemma 4.2, the above arguments can be carried out with the roles of M and M^e interchanged, and with f replaced by its inverse. It follows that $\alpha(W) = 0$ if and only if $\alpha_e(f(W)) = 0$. This means that the isomorphism $f\colon M \to M^e$ lifts onto an isomorphism $\tilde{f}\colon M^\alpha \to (M^e)^{\alpha_e}$, by Theorem 3.1. We can now apply Theorem

4.3, which tells us that the maps $(M^e)^{\alpha_e}$ and $(M^\alpha)^e$ are isomorphic. Thus we obtain an isomorphism between M^α and $(M^\alpha)^e$, and therefore e is an exponent of M^α. Since e was an arbitrary integer relatively prime to the degree of M^α, this implies that the map M^α is kaleidoscopic.

The lifting procedure used above is a reformulation of the 'homological lifting' method introduced in [4], and adopted in different terms elsewhere; for example, see [19, 1]. Further analysis in [16] and [12] shows that the automorphism group of the lift of a map M under such a voltage assignment is always a split extension of the direct product of β copies of $\mathbb{Z}_n$ by the group $\mathrm{Aut}(M)$, where β is the Betti number of the medial graph of M. By our previous calculations, we have $\beta = 1+|\mathrm{Aut}(M)|/4 = 1+\varepsilon$ where ε is the number of edges of M. Hence $\mathrm{Aut}(M^\alpha) \cong (\mathbb{Z}_n)^{1+\varepsilon} \rtimes \mathrm{Aut}(M)$, and this completes the proof. □

6. Proof of Theorem 2.2

The first part of Theorem 2.2 is a special case of Theorem 2.1, applied to a particularly simple map and its covers. Here we are able to completely determine the automorphism groups of the resulting kaleidoscopic maps with trinity symmetry, as described in the second part of Theorem 2.2. In this section we give a proof of both parts.

Let M be a map on the sphere whose underlying graph has two vertices u and v joined by a pair of parallel edges that appear dashed in Fig. 1. The vertices of the medial map $\mathrm{med}(M)$ are depicted as two squares in the middle of the two edges of M, and the four parallel edges of $\mathrm{med}(M)$ are depicted as solid lines. Using $\mathrm{med}(M)$, it is easy to see that M is isomorphic to its geometric dual as well as to its Petrie dual. Then since M is clearly both orientable and regular, and each vertex has degree 2, it follows that M is a kaleidoscopic regular map with trinity symmetry.

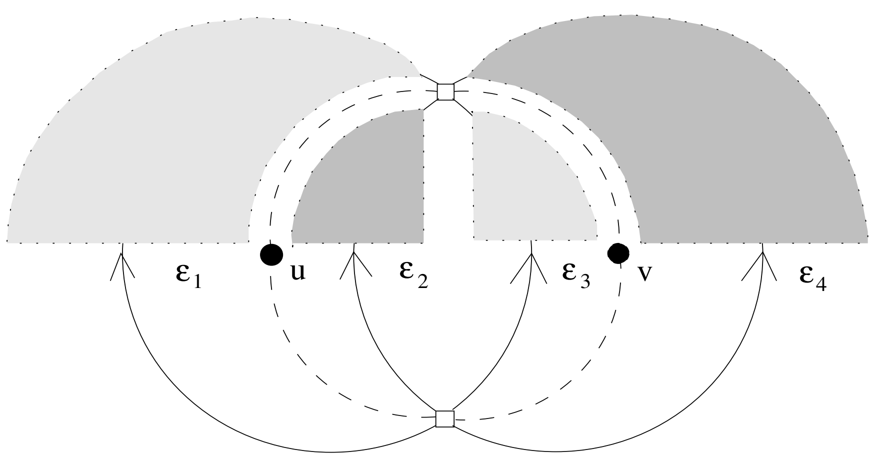

FIGURE 1. The two-vertex map M.

Now by Theorem 2.1, for every $n \geq 2$ there is a kaleidoscopic regular map of degree $2n$ with trinity symmetry on an orientable surface. This proves the first part of Theorem 2.2. A proof of the second part requires going into details of the lifting construction given in the proof of Theorem 2.1.

Constructing a voltage assignment on the corners of our two-vertex map M is equivalent to describing an ordinary voltage assignment on directed edges of $\mathrm{med}(M)$. As voltage group, take $H = \mathbb{Z}_n \times \mathbb{Z}_n \times \mathbb{Z}_n \times \mathbb{Z}_n$ (where $n \geq 2$), then direct the four medial edges as in Fig. 1, and label them with the (column) vectors $\varepsilon_1, \varepsilon_2, \varepsilon_3, \varepsilon_4$ from the standard basis of H (such that the j th entry of ε_i is δ_{ij}). Now define a voltage assignment α on the flags of M by $\alpha(\varepsilon_i) = \varepsilon_i$, that is, by letting the voltage of any directed edge ε_i be its label, and letting the reverse edge have voltage $-\varepsilon_i$ (using additive notation). In this way, the set of directed edges of $\mathrm{med}(M)$, and hence the set of all flags of M, can be taken as $\mathcal{F} = \{\pm\varepsilon_1, \pm\varepsilon_2, \pm\varepsilon_3, \pm\varepsilon_4\}$.

Next, define the two matrices

$$A = \begin{bmatrix} 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 \\ 0 & 1 & 0 & 0 \\ 1 & 0 & 0 & 0 \end{bmatrix} \quad \text{and} \quad B = \begin{bmatrix} 0 & 1 & 0 & 0 \\ 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 \end{bmatrix}$$

with entries in $\mathbb{Z}_n$, and let I_4 be the identity matrix in $\mathrm{GL}(4, \mathbb{Z}_n)$. Also define bijections f^C, f^L and f^T on $\mathcal{F}$ by

$$f^C(\varepsilon) = -I_4\varepsilon, \quad f^L(\varepsilon) = A\varepsilon \quad \text{and} \quad f^T(\varepsilon) = B\varepsilon \qquad \text{for all} \quad \varepsilon \in \mathcal{F}.$$

These are automorphisms of M that take the flag ε_1 to $-\varepsilon_1 = {\varepsilon_1}^C$, $\varepsilon_4 = {\varepsilon_1}^L$ and $\varepsilon_2 = {\varepsilon_1}^T$, respectively. In particular, these three automorphisms have the same effect on ε_1 as the three permutations C, L and T. By the general theory of maps on surfaces [6], it follows that the three automorphisms f^C, f^L and f^T generate the entire group $\mathrm{Aut}(M)$, which is isomorphic to $\langle -I_4, A, B\rangle \cong \mathbb{Z}_2 \times \mathbb{Z}_2 \times \mathbb{Z}_2$.

To determine the automorphism group of the lifted map, we need to see how automorphisms lift. By [3] we can do this by considering voltages of flag-walks and their images. We begin with a general observation of these for our map $M = (\mathcal{F}; C, L, T)$. Let $W = (x_0, x_1, \ldots x_k)$ be an arbitrary flag-walk in M. We know that $\alpha(W) = a_1 + \ldots + a_k$ where $a_j = 0$ (the identity element of H) if $x_j = {x_{j-1}}^L$ or ${x_{j-1}}^T$, and $a_j = \alpha(x_{j-1})$ if $x_j = {x_{j-1}}^C$. But in our notation, voltages on flags are simply their labels and therefore we have $a_j = \alpha(x_{j-1}) = x_{j-1}$ when $x_j = {x_{j-1}}^C$. Now let f be any automorphism of M. Then $f(W) = (y_0, y_1, \ldots, y_k)$ where $y_j = f(x_j)$, and $\alpha(f(W)) = b_1 + \ldots + b_k$ where $b_j = 0$ if $y_j = {y_{j-1}}^L$ or ${y_{j-1}}^T$ and $b_j = \alpha(y_{j-1}) = y_{j-1}$ if $y_j = {y_{j-1}}^C$. Since f is an automorphism, we have $y_j = {y_{j-1}}^C$ if and only if $x_j = {x_{j-1}}^C$, and similarly for L and T in place of C. But the action of f is given by multiplication of a matrix F from $\langle -I_4, A, B\rangle \cong \mathrm{Aut}(M)$, and so $y_{j-1} = Fx_{j-1}$ for all j for which $x_j = {x_{j-1}}^C$. By additivity, we may therefore conclude that $\alpha(f(W)) = F(\alpha(W))$.

For the next step, we focus on a connected component M^α of the lift containing the flag $(\varepsilon_1, 0) \in \mathcal{F} \times H$. By [3], the flag-set $\mathcal{F}'$ of this connected component contains precisely those flags (ε, g) for which g is the voltage of a closed walk in M based at ε_1. In the terminology of Section 3, $(\varepsilon, g) \in \mathcal{F}'$ if and only if $\varepsilon \in \mathcal{F}$ and g is an element of the local group $H_{\mathrm{loc}} = H_{\varepsilon_1}$. In order to determine this local group, observe that if a flag-walk W contains a flag ε_i followed by ${\varepsilon_i}^C$, then this pair adds ε_i to the sum $\alpha(W) = a_1 + \ldots + a_k$; this also corresponds to passing from one of the top four shaded flags in Fig. 1 to an opposite flag along the corresponding edge of the medial map. On the other hand, if W contains a flag ${\varepsilon_j}^C$ followed by ε_j, then this pair adds $-\varepsilon_i$ to the sum $\alpha(W)$, which corresponds to passing from one of the bottom four flags to an opposite top flag along the corresponding edge of the medial map. It follows that of W is any closed flag-walk in M then the sum of the entries of the vector $\alpha(W) \in H$ is zero. The converse is easily seen to be true as well. Hence the local group H_{loc} consists of all $h \in H = (\mathbb{Z}_n)^4$ orthogonal to the all 1s vector.

We have seen in the Proof of Theorem 2.1 that every automorphism of M lifts. In order to be more specific, we invoke a detailed version of Theorem 9 of [3]. This states that given any $h \in H$, every automorphism $f \in \mathrm{Aut}(M)$ lifts to an automorphism f_h of M^α described as follows. Let W be any flag-walk in M beginning at the flag ε_1 and terminating at an arbitrary flag $\varepsilon \in \mathcal{F}$. Then for any $h \in H_{\mathrm{loc}}$, the automorphism f lifts to the automorphism f_h of M^α given (here in additive notation) by

$$f_h(\varepsilon, \alpha(W)) = (f(\varepsilon), h + \alpha(f(W))) \ . \tag{6.1}$$

It is a consequence of Theorem 3.1 that this formula is independent of the choice of the walk W, in the sense that if W' is another walk starting at ε_1 and ending at ε such that $\alpha(W) = \alpha(W')$, then $\alpha(f(W)) = \alpha(f(W'))$.

To determine all the lifts of the automorphisms of M we will use the identification $\mathrm{Aut}(M) = \langle f^A, f^B, f^C \rangle \cong \langle -I_4, A, B \rangle$; the lifts will then be denoted by pairs (F, h) with $F \in \langle -I_4, A, B \rangle$ and $h \in H$. We will also use the fact established earlier that for any $F \in \langle -I_4, A, B \rangle$ we have $\alpha(FW) = F\alpha(W)$; that is, if $\alpha(W) = g \in H$ then $\alpha(FW) = Fg$. Accordingly, the formula (6.1) yields

$$(F, h)(\varepsilon, g) = (F\varepsilon, h + Fg) \tag{6.2}$$

for any flag $\varepsilon \in \mathcal{F}$, any $F \in \langle -I_4, A, B \rangle$, any $g \in H$ and any $h \in H_{\mathrm{loc}}$. For a composition of two lifts, we obtain

$$(F', h')(F, h)(\varepsilon, g) = (F', h')(F\varepsilon, h + Fg) = (F'F\varepsilon, h' + F'h + F'Fg) \tag{6.3}$$

which shows that $(F', h')(F, h) = (F'F, h' + F'h)$. This shows that the group $\mathrm{Aut}(M^\alpha)$ is isomorphic to a semidirect product $H_{\mathrm{loc}} \ltimes \langle -I_4, A, B \rangle \cong H_{\mathrm{loc}} \ltimes (\mathbb{Z}_2 \times \mathbb{Z}_2 \times \mathbb{Z}_2)$, with $\langle -I_4, A, B \rangle$ acting naturally on H_{loc} by left matrix multiplication.

We now proceed to determining the three generators of $\mathrm{Aut}(M^\alpha)$ that correspond to the action of C', L', T' on the flag $(\varepsilon_1, 0)$. Observe that $(\varepsilon_1, 0)^{C'} = ({\varepsilon_1}^C, 0 + \alpha(\varepsilon_1)) = (-\varepsilon_1, \varepsilon_1)$, $(\varepsilon_1, 0)^{L'} = (\varepsilon_4, 0)$, and $(\varepsilon_1, 0)^T = (\varepsilon_2, 0)$. Applying (6.2) to our generators $X \in \{-I_4, A, B\}$ and to the flag-walks $(\varepsilon_1, {\varepsilon_1}^X)$ for $X \in \{C, L, T\}$ yields the lifts $z = (-I_4, \varepsilon_1)$, $a = (A, 0)$ and $b = (B, 0)$, for which $z(\varepsilon_1, 0) = (-\varepsilon_1, \varepsilon_1)$, $a(\varepsilon_1, 0) = (\varepsilon_4, 0)$, and $b(\varepsilon_1, 0) = (\varepsilon_2, 0)$, respectively. Since these have the same effect as C', L', T' on the flag $(\varepsilon_1, 0)$, it follows that a, b, z generate the group $\mathrm{Aut}(M^\alpha)$. Direct calculations in $\mathrm{Aut}(M^\alpha) \cong H_{\mathrm{loc}} \ltimes \langle -I_4, A, B \rangle$, that is, using the formula (6.3), then give

$$a^2 = b^2 = z^2 = (ab)^2 = (bz)^{2n} = (az)^{2n} = (abz)^{2n} = (azbzabz)^2 = 1 \ .$$

So now define the abstract group G with presentation

$$G = \langle \, a, b, z \mid a^2, b^2, z^2, (ab)^2, (bz)^{2n}, (az)^{2n}, (abz)^{2n}, (azbzabz)^2 \, \rangle \ . \tag{6.4}$$

Then $\mathrm{Aut}(M^\alpha)$ is a quotient of G, and in particular, $|G| \geq |\mathrm{Aut}(M^\alpha)| = 8|H_{\mathrm{loc}}| = 8n^3$. To finish the proof of Theorem 2.2, we need to establish the reverse inequality.

All arguments that follow can be verified by easy calculations. Let $r = bz$, $s = za$, and $t = z$; this choice of notation has been chosen to agree with the one used in the accompanying tables of [8]. The presentation (6.4) is equivalent to

$$G = \langle\, r, s, t \mid r^{2n}, s^{2n}, t^2, (rs)^2, (rt)^2, (ts)^2, [r^2, s^2] \,\rangle\ . \tag{6.5}$$

Now let $u = r^2$, $v = s^2$, and $w = r^{-1}s^2r$. With the help of the definitions of r, s, t, u, v and w and the relations from (6.4) or (6.5), we find that conjugation by r fixes u and interchanges v with w, while conjugation by s takes u onto $(uwv)^{-1}$ and fixes both v and w. (For example, $uwv = r^2r^{-1}s^2rs^2 = (rs^2)^2 = (baza)^2 = bazabaza = abzbza = az(zb)^2za = s^{-1}u^{-1}s$.) Then conjugation of the commutator relation $[u, v] = 1$ by r and s gives $[u, w] = 1$ and $[(uwv)^{-1}, w] = 1$, from which it follows that $[v, w] = 1$ as well. Hence the subgroup $K = \langle u, v, w\rangle$ of G is abelian. Taking into account the realtions $u^n = v^n = w^n = 1$ gives $|K| \le n^3$. Next, the relations $(rt)^2 = (ts)^2 = 1$ imply that conjugation by t inverts both r and s and hence also both u and v, and furthermore, since $[s^2, r^2] = 1$ we find that $(r^{-1}s^2r)^t = rs^{-2}r^{-1} = r^{-1}r^2s^{-2}r^{-2}r = r^{-1}s^{-2}r$, so conjugation by t inverts w as well. Hence $K = \langle u, v, w\rangle$ is normal in G. The relations from (6.5) show that the factor group G/K is generated by three commuting involutions rK, sK, tK, and therefore $|G/K| \le 8$. Thus $|G| \le 8n^3$, which was the inequality we needed to prove that $|G| = 8n^3$.

Finally, we let $c = ab$ and rewrite (6.4) in the form

$$G = \langle\, a, b, c, z \mid a^2, b^2, c^2, z^2, abc, (az)^{2n}, (bz)^{2n}, (cz)^{2n}, (azbzcz)^2 \,\rangle$$

which is exactly the presentation from the last part of Theorem 2.2. □

7. Remarks

We have seen one building block for our construction of Theorem 2.1, namely the embedding of a cycle of length two in the sphere, which led to the family given in Theorem 2.2. Observe, however, that one can apply Theorem 2.1 to any of the new maps provided by Theorem 2.2. Indeed, for any $n \ge 2$, take the map M_n of degree $2n$ with $|\mathrm{Aut}(M_n)| = 8n^3$ from Theorem 2.2. The underlying graph of M_n has $\varepsilon_n = 2n^3$ edges, so by Theorem 2.1 we have a trinity symmetry kaleidoscopic regular map of degree $2nm$ with automorphism group of order $8n^3m^{1+2n^3}$, for any $m, n \ge 2$.

Are there any other 'small' ingredients to feed into Theorem 2.1? The answer is "Yes". An inspection of the census [8] of orientable regular maps of genus between 2 and 101 shows that there are 14 such maps; their labels (from the list accompanying the paper [8]), degrees, automorphism group orders, numbers of vertices, numbers of edges and genera are given in the table below.

Map	Degree	Group order	Vertices	Edges	Genus
R10.13	6	216	18	54	10
R21.14	6	480	40	120	21
R9.21	8	128	8	32	9
R9.23	8	128	8	32	9
R17.25	8	256	16	64	17
R33.39	8	512	32	128	33
R33.40	8	512	32	128	33
R33.41	8	512	32	128	33
R65.64	8	1024	64	256	65
R85.44	8	1344	84	336	85
R76.20	10	1000	50	250	76
R73.86	12	864	36	216	73
R81.127	12	960	40	240	81
R97.125	16	1024	32	256	97

Note that R10.13, R33.39 and R76.20 are the maps covered by Theorem 2.2 for $n = 3, 4, 5$. In combination with Theorem 2.1, the data for the remaining maps in the table yield the following extension of the first part of Theorem 2.2.

Corollary 7.1. *For any $n \geq 1$, there exist orientable kaleidoscopic regular maps with trinity symmetry having the following parameters:*

Degree	Group order	Vertices	Edges	Genus
$2n$	$8n^3$	$2n^2$	$2n^3$	$n^3 - 2n^2 + 1$
$6n$	$480n^{121}$	$40n^{120}$	$120n^{121}$	$60n^{121} - 40n^{120} + 1$
$8n$	$128n^{33}$	$8n^{32}$	$32n^{33}$	$16n^{33} - 8n^{32} + 1$
$8n$	$512n^{129}$	$32n^{128}$	$128n^{129}$	$64n^{129} - 32n^{128} + 1$
$8n$	$1024n^{257}$	$64n^{256}$	$256n^{257}$	$128n^{257} - 64n^{256} + 1$
$8n$	$1344n^{337}$	$84n^{336}$	$336n^{337}$	$168n^{337} - 84n^{336} + 1$
$12n$	$864n^{217}$	$36n^{216}$	$216n^{217}$	$108n^{217} - 36n^{216} + 1$
$12n$	$960n^{241}$	$40n^{240}$	$240n^{241}$	$120n^{241} - 40n^{240} + 1$

Of course one may re-apply the procedure described at the beginning of this section to any of the maps represented in this table, to produce further new infinite families of kaleidoscopic regular maps with trinity symmetry.

Also Theorem 2.1 extends automatically to non-orientable regular maps for any odd n; oddness of n is needed to ensure non-orientability of the lifts. The list of all non-orientable regular maps of genus up to 202 (associated with [8]) reveals examples of a kaleidoscopic maps with trinity symmetry of degree 6, 10 and 12, the smallest one being the map N12.3 of degree 6 with automorphism group of order 120. Thus for any odd n we have infinite families of non-orientable kaleidoscopic regular maps with trinity symmetry, for the following degrees and automorphism group orders:

Degree (n odd)	Group order
$6n$	$120n^{31}$
$10n$	$7200n^{1801}$
$12n$	$5040n^{1261}$

Before making two final remarks, we consider the automorphism group $\mathrm{Aut}(M)$ of a regular map $M = (\mathcal{F}; C, L, T)$. Since this acts regularly on the flags of M, we may identify the flag set $\mathcal{F}$ with the group $\mathrm{Aut}(M)$. Details have been worked out in the literature; see [6] for example. This enables us to establish a one-to-one correspondence between regular maps of type (d, ℓ) — that is, those with vertex degree d and face length ℓ — and groups with partial presentation of the form

$$\langle\, \omega, \lambda, \tau \mid \omega^2 = \lambda^2 = \tau^2 = (\lambda\tau)^2 = (\tau\omega)^d = (\omega\lambda)^\ell = \ldots = 1 \,\rangle \tag{7.1}$$

where 2, d and ℓ are true orders of the corresponding elements. One direction of this correspondence may be explained as follows. Let $M = (\mathcal{F}; C, L, T)$ be a regular map, let x be a flag of M, and let ω, λ and τ be the unique automorphisms of M taking x to x^C, x^L and x^T, respectively. Then these involutory automorphisms generate $G = \mathrm{Aut}(M)$, and satisfy a presentation of the form (7.1).

In group-theoretic terms, the map M is self-dual if and only if there exists an automorphism of the group G that fixes ω and interchanges λ with τ, and self-Petrie-dual if and only if there is an automorphism of G that fixes ω and interchanges λ with $\lambda\tau$. Trinity symmetry is then equivalent to extendability of any automorphism of the subgroup $\langle\lambda, \tau\rangle \cong \mathbb{Z}_2 \times \mathbb{Z}_2$ to the entire group. Finally, for the purpose of discussing our 'absolutely most symmetric maps', $e \in \mathbb{Z}_d^*$ will be considered to be an exponent of a regular map determined by the group presentation (7.1) if the group admits an automorphism fixing λ and τ and sending ω onto $(\omega\tau)^{e-1}\omega$.

Now observe that the examples and families of kaleidoscopic regular maps with trinity symmetry described in the above tables all have even degree. In general, the degree of an orientable kaleidoscopic regular map with trinity symmetry is always even, since the length of a Petrie polygon in an orientable map cannot be odd. But there do exist such 'super-symmetric' maps of odd degree in the non-orientable case.

For example, take the direct product $G = A_5 \times A_5 \times A_5$ of three copies of the alternating group A_5, and consider this as a permutation group on 15 points with three orbits of length 5. Define three involutions ω, λ, τ in G as follows:

$$\begin{aligned} \omega &= (1,2)(3,4)(6,7)(8,9)(11,12)(13,14), \\ \lambda &= (2,3)(4,5)(7,10)(8,9)(12,14)(13,15), \\ \tau &= (2,4)(3,5)(7,8)(9,10)(12,15)(13,14). \end{aligned}$$

It is not difficult to prove that ω, λ, τ together generate G. Also $\lambda\tau$ has order 2, while each of $\omega\lambda$, $\omega\tau$ and $\omega\lambda\tau$ has order 15. Hence G is the automorphism group of a non-orientable regular map of degree 15, with faces and Petrie polygons of length 15. Moreover, conjugation by the permutation $(3,4)(6,11)(7,12)(8,14)(9,13)(10,15)$ in S_{15} fixes ω and interchanges λ with τ, while conjugation by $(1,11)(2,12)(3,14)(4,13)$ $(5,15)(8,9)$ fixes ω and interchanges λ with $\lambda\tau$. Hence this map has trinity symmetry. Similarly, there exist permutations in S_{15} that centralize both λ and τ and conjugate ω to $(\omega\tau)^{e-1}\omega$ for each e in the group $\mathbb{Z}_{15}^* = \{1, 2, 4, 7, 8, 11, 13, 14\}$ of units mod 15, so the map is kaleidoscopic.

It would be interesting to know if there are any examples of odd *prime* degree. This is an open question.

Our final remark concerns the group presented in the second part of Theorem 2.2. We can now outline an alternative, independent proof of the fact that the group G with presentation (6.4), or, equivalently, (6.5), is the automorphism group of a regular kaleidoscopic map of degree $2n$ with trinity symmetry and with $|G| = 8n^3$. (Recall that this was suggested (without the 'kaleidoscopic' part) by Wilson some time ago [20].) We can prove it using group theory, without referring to coverings or voltages at all.

To allow translation between the presentation (6.4) and the above, let $a = \lambda$, $b = \tau$ and $z = \omega$, and let us use the notation of the proof of Theorem 2.2 in what follows. Calculations in the last part of the proof show that G contains an abelian normal subgroup K of index 8 generated by $u = r^2 = (bz)^2$, $v = s^2 = (za)^2$ and $w = r^{-1}s^2r = zb(za)^2bz$. Using Reidemeister-Schreier theory, it is not hard to show that K is isomorphic to $\mathbb{Z}_n \times \mathbb{Z}_n \times \mathbb{Z}_n$. (In fact, one can do this for the group without the relations $(bz)^{2n} = (az)^{2n} = (abz)^{2n} = 1$ and find that the pre-image of K is free abelian of rank 3, for example using the `Rewrite` command in Magma [5], and then factor out the normal subgroup generated by $(bz)^{2n}$, $(az)^{2n}$ and $(abz)^{2n}$.) Then since G/K is isomorphic to $\mathbb{Z}_2 \times \mathbb{Z}_2 \times \mathbb{Z}_2$, it follows that G is a split extension of $\mathbb{Z}_n \times \mathbb{Z}_n \times \mathbb{Z}_n$ by $\mathbb{Z}_2 \times \mathbb{Z}_2 \times \mathbb{Z}_2$, and hence $|G| = 8n^3$.

Moreover, it is easy to see that any permutation of the set $\{a, b, ab\}$ extends to an automorphism of G that fixes z, since all the defining relations in (6.4) for G are preserved. This shows that the map has trinity symmetry. It remains for us to deal with isomorphisms of the map and its powers.

Let e be any unit in $\mathbb{Z}_{2n}$, so that e is odd and relatively prime to n. Then we need to check that the assignment $(a, b, z) \mapsto (a, b, (zb)^{e-1}z)$ induces an automorphism of G. This time we use the presentation of G in the form (6.5). The above assignment takes $r = bz$ to $b(zb)^{e-1}z = (bz)^e = r^e$ (which has the same order as r), $s = za$ to $(zb)^{e-1}za = r^{1-e}s$, and $t = z$ to $(zb)^{e-1}z = z(bz)^{e-1} = tr^{e-1}$, and fixes $rs = ba$. In particular, it preserves the relations $r^{2n} = (rs)^2 = 1$. Next we recall that $s^{-1}us = azbzba = (uwv)^{-1}$. This gives $(r^{1-e}s)^2 = (u^{(1-e)/2}s)^2 = u^{(1-e)/2}s^2(uwv)^{(e-1)/2} = u^{(1-e)/2}v(uwv)^{(e-1)/2} = v^{(e+1)/2}w^{(e-1)/2}$, which lies in K, so has order dividing n and commutes with $u^e = r^{2e}$. In particular, this shows that the relators $s^{2n} = [r^2, s^2] = 1$ are preserved. Finally, the relators t^2 and $(rt)^2$ are taken to $(tr^{e-1})^2$ and $(r^e t)^2$, which are trivial since $trt = r^{-1}$, while $(st)^2$ is taken to $(r^{1-e}str^{e-1})^2 = r^{1-e}(st)^2r^{e-1}$, also trivial. Hence all the relations are preserved by the assignment, and since the images r^e, $r^{1-e}s$ and tr^{e-1} of r, s and t generate $\langle r, s, t\rangle = G$, we have an automorphism of G, as required for e to be an exponent. Thus G is the automorphism group of a kaleidoscopic regular map with trinity symmetry.

References

1. D.S. Archdeacon, P. Gvozdjak, and J. Širáň, Constructing and forbidding automorphisms in lifted maps, *Mathematica Slovaca* 47 (1997), 113–129.
2. D.S. Archdeacon and R.B. Richter, Construction and classification of self-dual polyhedra, *J. Combin. Theory Ser. B* 54 (1992), 37–63.
3. D.S. Archdeacon, R.B. Richter, J. Širáň, and M. Škoviera, Branched coverings of maps and lifts of map homomorphisms, *Australas. J. Combin.* 9 (1994), 109–121.
4. N.L. Biggs, *Algebraic Graph Theory*, Cambridge Tracts in Mathematics 67, Cambridge Univ. Press, 1974.
5. W. Bosma, J. Cannon and C. Playoust, The Magma Algebra System I: The User Language, *J. Symbolic Computation* **24** (1997), 235–265.
6. R.P. Bryant and D. Singerman, Foundations of the theory of maps on surfaces with boundary, *Quart. J. Math. Oxford Ser. (2)* 36 (1985), 17–41.

7. M. Conder and P. Dobcsányi, Determination of all regular maps of small genus, *J. Combinat. Theory Ser. B* 81 (2001), 224–242.
8. M.D.E. Conder, Regular maps and hypermaps of Euler characteristic -1 to -200, *J. Combin. Theory Ser. B* 99 (2009), 455–459, with data lists available at the website `http://www.math.auckland.ac.nz/~conder/hypermaps.html`.
9. J.L. Gross and T.W. Tucker, *Topological Graph Theory*, Wiley, 1987, and Dover, 2001.
10. G.A. Jones and D. Singerman, Theory of maps on orientable surfaces, *Proc. London Math. Soc.* 37 (1978), 273–307.
11. G.A. Jones and J.S. Thornton, Operations on maps, and outer automorphisms, *J. Combin. Theory Ser. B* 35 (1983), 93–103.
12. A. Malnić, R. Nedela and M. Škoviera, Lifting graph automorphisms by voltage assignments, *European J. Combin.* 21 (2000), 927–947.
13. R. Nedela, Regular maps – combinatorial objects relating different fields of mathematics, *J. Korean Math. Soc.* 38 (2001) 5, 1069-1105.
14. R. Nedela and M. Škoviera, Exponents of orientable maps, *Proc. London Math. Soc.* 75 (1997), 1–31.
15. R.B. Richter, J. Širáň and Y. Wang, Self-dual and self-Petrie-dual regular maps, *J. Graph Theory* 69 (2012) 2, 152-159.
16. J. Širáň, The 'walk calculus' of regular lifts of graph and map automorphisms, *Yokohama Math. Journal* 47 (1999), 113–128.
17. J. Širáň, Regular maps on a given surface – a survey, in: Topics in Discrete Mathematics, Algorithms Combin. 26, Springer, 2006, 591-609.
18. J. Širáň and Y. Wang, Maps with highest level of symmetry that are even more symmetric than other such maps: Regular maps with largest exponent groups, *AMS Contemporary Mathematics Series* 531 (2010), 95–102.
19. D. Surowski, Lifting map automorphisms and Macbeath's theorem, *J. Combin. Theory Ser. B* 50 (1990) 135–149.
20. S. Wilson, *New techniques for the construction of regular maps*, doctoral dissertation, University of Washington, Seattle, 1976.
21. S. Wilson, Operators over regular maps, *Pacific J. Math.* 81 (1979), 559–568.
22. S. Wilson, personal communication, 2010.

Department of Mathematics and Statistics, University of Vermont, Burlington, VT, USA 05405

E-mail address: dan.archdeacon@uvm.edu

Department of Mathematics, University of Auckland, Private Bag 92019, Auckland 1142, New Zealand

E-mail address: m.conder@auckland.ac.nz

Department of Mathematics and Statistics, Open University, Milton Keynes, MK7 6AA, UK, and Slovak University of Technology, 81368 Bratislava, Slovakia

E-mail address: j.siran@open.ac.uk