2016

On Edge-Minimization of Vertex Minimal Graphs with Dicyclic Automorphism Group

Peter E. Huston
Otterbein University, peter.huston@otterbein.edu

On Edge Minimization of Vertex-Minimal Graphs with Dicyclic Automorphism Group

Peter Huston

April 1, 2016

Submitted in partial fulfillment of the requirements for graduation with honors

Jeremy Moore, PhD Project Advisor	Advisor's Signature
Ryan Berndt, PhD Second Reader	Second Reader's Signature
Jonathan DeCoster, PhD Honors Representative	Honors Representative's Signature

Acknowledgements

Parts of this work were preformed in collaboration with Lindsay Cadwallader (University of Connecticut), Stephen Graves (University of Texas), and Cesar Meza (Pomona College). I would also like to thank Jeremy Moore and Ryan Berndt for their extensive guidance and support during the writing and editing process. Without their gracious assistance, this work would probably not have been completed, and it certainly would not be intelligible!

This material is based upon work supported by the National Science Foundation under Grant No. DMS-1359101. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation.

Abstract

We find the smallest degree of a graph with automorphism group isomorphic to the dicyclic group with $4n$ elements, denoted $\alpha(\mathrm{Dic}_n)$. We also find the fewest edges a minimum-order graph with dicyclic automorphism group and a minimal number of vertices can have. For n not a power of 2, the value of $\alpha(\mathrm{Dic}_n)$ is significantly less than the best previously known upper bound, $8n$. Such an edge-minimized vertex-minimal graph is constructed and shown to have automorphism group Dic_n. That the exhibited graph is minimal is verified using a combination of techniques similar to those developed previously for Abelian groups and earlier results for the special case Dic_{2^n}.

Table of Contents

0 Definitions and Conventions

These basic definitions are presented for reference:

Definition 0.1. *A **graph** is an ordered pair (V, E) where V is a nonempty set and the elements of E are pairs of elements of V. For a graph Γ, we adopt the notational convention that $\Gamma = (V(\Gamma), E(\Gamma))$.*

One may think of the elements of V as a set of points or "vertices" in space, and an element $e \in E$ as representing a curve or "edge" connecting the elements of e.

Definition 0.2. *The **automorphism group** of a graph Γ, denoted $\mathrm{Aut}(\Gamma)$, is the group of permutations ϕ of the members of $V(\Gamma)$ such that $\{\phi(v), \phi(w)\} \in E(\Gamma)$ if and only if $\{v, w\}$ in $E(\Gamma)$, with the group operation being function composition.*

In other words, the members of $\mathrm{Aut}(\Gamma)$ are those permutations of the vertices of Γ that map edges onto edges, or "preserve" edges.

The next few definitions and notations pertain specifically to our problem and have been commonly used in the literature:

Definition 0.3. *For a group G, let*

$$\alpha(G) = \min\{|V| : \Gamma = (V, E) \text{ is a graph and } \mathrm{Aut}((V, E)) = G\}.$$

Definition 0.4. *For a group G and a positive integer n, let*

$$e(G, n) = \min\{|E| : \Gamma = (V, E) \text{ is a graph, } \mathrm{Aut}((V, E)) = G, \text{ and } |V| = n\}$$

if it exists, or ∞ if the set is empty.

Definition 0.5. *For an integer $n \geq 2$, the **dicyclic group** with $4n$ elements has the presentation*

$$\mathrm{Dic}_n = \langle \sigma, \tau | \sigma^{2n} = \tau^4 = 1, \sigma^n = \tau^2, \sigma^\tau = \tau\sigma^{-1} \rangle.$$

Using this notation, we can more concisely write our goals as determining $\alpha(\mathrm{Dic}_n)$ and $e(\mathrm{Dic}_n, \alpha(\mathrm{Dic}_n))$.

Usually, we write elements of Dic_n in the form $\sigma^k\tau^b$ where $0 \leq k < 2n$ and $b \in \{0, 1\}$, with the exception that we write τ^2 rather than σ^n.

Definition 0.6. *Suppose $\phi : G \to S_k$ is an injective homomorphism such that $\phi(G)$ is the automorphism group of some graph. Then we say ϕ is a **realizable embedding** of G on k vertices.*

When two groups G and H are isomorphic, we write $G \cong H$; we reserve the notation $G = H$ for *equal* groups; that is, where the two groups have the same elements and group operation. This distinction eliminates ambiguity in the case where multiple permutation groups are isomorphic to Dic_n, but differ in other respects of interest.

Definition 0.7. *Suppose a group G acts on $\{1, 2, \ldots k\}$. Then for any $a \in \{1, 2, \ldots k\}$, the **orbit** or **vertex orbit** of a under G is $\mathcal{O}_G\{a\} = \{ga : g \in G\}$.*

In each case, the subscript G may be omitted when the group in question is clear. We also write $\mathcal{O}_g\{a\}$ instead of $\mathcal{O}_{\langle g \rangle}\{a\}$.

Definition 0.8. *Suppose a group G acts on $\{1, 2, \ldots k\}$. Then for any $\{a, b\} \subseteq \{1, 2, \ldots k\}$, the **edge orbit** of $\{a, b\}$ under G is $\mathcal{O}_G\{a, b\} = \{\{ga, gb\} : g \in G\}$.*

Definition 0.9. *Suppose a group G acts on a set X. For any $Y \subseteq X$, we define $\mathrm{fix}\, Y = \bigcap_{y \in Y} \mathrm{stab}(y)$.*

Definition 0.10. *Suppose Γ is a graph and G is a group acting on the vertices of Γ. Then we define the **group of G respecting automorphisms of Γ** to be*

$$A(\Gamma, G) = \{\gamma \in \mathrm{Aut}(\Gamma) : \gamma v \in \mathcal{O}_G\{v\} \forall v \in V(\Gamma)\}.$$

Using a somewhat unusual notation for the prime factorization of the parameter n in the definition of the dicyclic group will simplify much of the following work, for reasons that will become apparent later. We adopt the convention that

$$n = 2^b \prod_{i=1}^{I} p_i \prod_{j=1}^{J} q_j^{m_j}$$

where $\{p_i\}$ are the prime factors of n of multiplicity 1 which are greater than 5, $\{q_j\}$ are other odd prime factors of n, and m_j is the multiplicity of q_j as a divisor of n. To make our structured factorization unique, we also demand that $p_a < p_b$ and $q_a < q_b$ whenever $a < b$ and both exist.

Suppose $(x+1, x+2, x+3, \dots, x+k)$ is a cycle in the disjoint cycle decomposition of σ. Then the vertex denoted by $x+1+r$ is $x+1+(r \mod k)$; similarly, the vertex denoted by $x+k-r$ is $x+k-(r \mod k)$.

If a sum or product has a fractional upper bound, the variable being summed over takes on every value less than the bound. For example $\sum_{k=0}^{\frac{5}{2}} 1 = 3$.

1 Introduction

The first major result regarding the automorphism groups of graphs was obtained by Frucht, who showed that every group was the automorphism group of some graph [3]. Subsequently, he further showed that Γ could always be chosen to be a 3-regular graph [2]. The search for $\alpha(G)$ began in earnest when Sabidussi, building on Frucht's work, showed [5] that $\alpha(G) = \mathcal{O}(|G| \log(\log(|G|)))$. Babai improved this when he showed [1] that $\alpha(G) \leq 2|G|$ so long as G is not $\mathbb{Z}_3$, $\mathbb{Z}_4$, or $\mathbb{Z}_5$; these special cases have been addressed elsewhere. The precise value of $\alpha(G)$ is known for a few specific cases, most significantly for Abelian groups; Arlinghaus rederived previous unpublished work and extended it to all Abelian groups [6]. The techniques in Arlinghaus are particular relevant to the present work, because many details in Arlinghaus regarding the ways different graphs may be combined and regarding special cases involving small prime numbers carry over to the case of dicyclic groups. For a more detailed summary of results on the values of $\alpha(G)$ and $e(G,k)$ for various G and k, refer to [4].

By 1990, progress had largely stopped. Many who were interested in the problem found it difficult to better the bound Babai had set forth. Babai's bound is sharp for many groups, and the techniques Babai and Sabidussi employed in proving their bounds seem to have been optimized. Babai's approach, in short, is to create two copies of the Cayley graph of G and add edges connecting them in such a way as to ensure there are no automorphisms outside of those given by applying the same element of G to each copy. The constructions he and Sabidussi employed rely on obtaining an arbitrary minimal generating set of G, without paying particular attention to the relations between the generators. In contrast, by using techniques designed for the specific groups under consideration, Graves, Graves, and Lauderdale showed that $\alpha(\mathrm{Dic}_{2^m}) = 2^{m+3} = 2|\mathrm{Dic}_n|$; in other words, they showed that in this case, Babai's bound is sharp [4]. We will extend their techniques to discover $\alpha(\mathrm{Dic}_n)$ for all n.

2 Properties of Dicyclic Groups

If G is a group and $\Gamma = (V, E)$ is a graph with $G \cong \mathrm{Aut}(\Gamma)$, the group action of G on the vertices of Γ specifies an injective homomorphism $\phi : G \to S_V$. More explicitly, ϕ is the composition of the isomorphism between G and $\mathrm{Aut}(\Gamma)$ and the inclusion homomorphism from $\mathrm{Aut}(\Gamma)$ into S_V. On the other hand, although any group is, up to isomorphism, the automorphism group of some graph (see [3]), it is not the case that any particular group of permutations is a graph automorphism group. The minimal counterexample is $\mathbb{Z}_3 \cong \langle (123) \rangle$. In order to determine the minimum number of vertices of a graph Γ with $\mathrm{Aut}(\Gamma) \cong \mathrm{Dic}_n$, we will proceed in two steps. First, we will examine Dic_n algebraically and classify all injective homomorphisms, or embeddings, from Dic_n to S_N (where $N \in \mathbb{N}$). Then, we will exploit our classification to identify those realizable embeddings $\phi : \mathrm{Dic}_n \to S_N$ for which N is minimal.

Aside from serving as an introduction to working with Dic_n, the following fact is of interest to us because it gives an easy way to show that some group is not Dic_n: find an order two element other than τ^2. With a few exceptions, this will be our approach.

Lemma 2.1. *Suppose* Dic_n *is defined as in Definition 0.5. Then* $\sigma^n = \tau^2$ *is the only element of order two, and every element in the set* Dic_n *this is not in* $\langle \sigma \rangle$ *has order four.*

Proof. Elements in the set $\mathrm{Dic}_n \setminus \langle \sigma \rangle$ have the form $\sigma^k \tau$ for some $k \in \{0, 1, \dots, 2n-1\}$. We know that $|\sigma^k \tau| = 4$ because $(\sigma^k \tau)^2 = \sigma^k \tau \sigma^k \tau = \sigma^k \sigma^{-k} \tau^2 = \tau^2$ and $|\tau^2| = 2$.

Because σ has order $2n$, we know that $\sigma^n = \tau^2$ is the only element of order two in $\langle\sigma\rangle$. Since every element not in $\langle\sigma\rangle$ has order four, $\sigma^n = \tau^2$ is the only element of order two in all of Dic_n. $\square$

To classify embeddings of Dic_n into S_N, we will first determine how Dic_n may act on a single point. It turns out that the restriction of σ and τ to the orbit of a single vertex must take on one of three forms. These forms are the foundation of our work and are referenced very frequently, so we will give them concise names.

Let σ_0 and τ_0 be the restriction of σ and τ to $\mathcal{O}\{1\}$. Then:

Definition 2.2. *If $\sigma_0 = (1, 2, \dots k)(k+1, k+2, \dots 2k)$ and $\tau_0 = \prod_{x=1}^{\frac{k}{2}-1} (x, 2k+1-x, x+\frac{k}{2}, 2k+1-(x+\frac{k}{2}))$, then we say that $\mathcal{O}\{1\}$ is a **twisted** vertex orbit, and that the cycles $(1, 2, \dots k)$ and $(k+1, k+2, \dots 2k)$ are **twisted together** by τ.*

Definition 2.3. *If $\sigma_0 = (1, 2, \dots k)$ and $\tau_0 = \prod_{x=0}^{\frac{k}{2}} (x, k+1-x)$ or $\tau_0 = \sigma_0 \prod_{x=0}^{\frac{k}{2}} (x, k+1-x)$, then we say that $\mathcal{O}\{1\}$ is a **self-reversed** vertex orbit, and that the cycle $(1, 2, \dots k)$ is **reversed onto itself** or **self-reversed** by τ.*

Definition 2.4. *If $\sigma_0 = (1, 2, \dots k)(k+1, k+2, \dots 2k)$ and $\tau_0 = \prod_{x=0}^{k} (x, 2k+1-x)$, then we say that $\mathcal{O}\{1\}$ is a **pair-reversed** vertex orbit, and that the cycles $(1, 2, \dots k)$ and $(k+1, k+2, \dots 2k)$ are **reversed onto each other** or **pair-reversed together** by τ.*

The following lemma verifies that these three forms describe all possible restrictions of σ and τ to a single vertex orbit and determines what values the parameter k may possess.

Lemma 2.5. *Suppose σ and τ are functions in S_N such that $\langle\sigma, \tau\rangle \cong \mathrm{Dic}_n$. Then the permutations σ_0 and τ_0 take one of the following forms, up to a renaming of the elements of $\mathcal{O}\{1\}$:*

1. *$\mathcal{O}\{1\}$ is a twisted vertex orbit, and σ_0 is the product of two k-cycles for some k which divides $2n$ but not n.*
2. *$\mathcal{O}\{1\}$ is a pair-reversed vertex orbit, and σ_0 is the product of two k-cycles for some k which divides n.*
3. *$\mathcal{O}\{1\}$ is a self-reversed vertex orbit, and σ_0 is a k-cycle for some k which divides n.*

Proof. Let k be the length of the cycle containing 1 in the disjoint cycle decomposition of σ. Because $|\sigma| = 2n$, we know that k must be a divisor of $2n$. We can also get something from the group presentation. Because $\sigma\tau = \tau\sigma^{-1}$, for any $x \in \{1, 2, \dots k\}$ we have $\sigma\tau(x) = \tau\sigma^{-1}(x) = \tau(x-1)$ (modulo k). More generally, $\sigma^r\tau(x) = \tau\sigma^{-r}(x) = \tau(x-r)$. By iterated application of this identity, we can see that $(\tau(k), \tau(k-1), \tau(k-2), \dots \tau(1))$ is a k-cycle in the decomposition of σ. The individual cases are:

$$\begin{aligned} \sigma(\tau(k)) &= \tau(k-1) \\ \sigma(\tau(k-1) &= \tau(k-2) \\ &\dots \\ \sigma(\tau(2)) &= \tau(1) \\ \sigma(\tau(1)) &= \tau(k) \end{aligned}$$

Of course, the image of a cycle under τ may just be the same cycle.

Suppose τ maps $\{1, 2, \dots k\}$ onto itself. Then some manipulations allow us to discover the order of τ_0:

$$\begin{aligned} \tau\sigma^{-r}(k) &= \sigma^r\tau(k) \\ \tau(k-r) &= \tau(k) + r \\ \tau(r) &= \tau(k) - r \\ \tau(r) &= k - (k + r - \tau(k)) \\ \tau^2(r) &= \tau(k - (k + r\tau(k)) \\ \tau^2(r) &= \tau(k) + (k + r - \tau(k) \\ \tau^2(r) &= r \end{aligned}$$

So in this case, $|\tau_0| = 2$.

With this fact, we are ready to determine the potential forms of τ_0 in several cases. First, consider the case where k is not also a divisor of n. Since k must divide $2n$, we may write $k = 2q$ where q divides n. (Note that this means k is even, so the product defined in Definition 2.2 is not degenerate.) Since q is the greatest common divisor of n and k, it follows that $q = an + bk$ for some $a, b \in \mathbb{Z}$. Because $2n$ is a multiple of k, we know a is odd (or else we would have q a multiple of k), and $q = n$ modulo k. Therefore $\sigma_0^n = \sigma_0^q$, and (as $q = \frac{k}{2}$) we know that $|\sigma_0^n| = 2$. By our group relations, $\sigma_0^n = \tau_0^2$, so τ_0^2 is a product of 2-cycles, and τ_0 must be a product of 4-cycles. If $|\tau_0| = 4$, then τ must map the elements of the cycle $(1, 2, \ldots k)$ onto a different k-cycle in σ, because otherwise (by the above) $|\tau_0| = 2$. If we arbitrarily denote this cycle $(k+1, k+2, \ldots 2k)$ and select $\tau(1) = 2k$, then the fact that $\tau_0^2(x) = \sigma_0^q(x) = x + q$ forces τ_0 to take the form given in Definition 2.2.

On the other hand, suppose k is a divisor of n. In the case where τ maps $(1, 2, \ldots k)$ onto itself, the value of $\tau_0(k)$ and the fact that $\tau_0(r) = \tau_0(k) - r$ determine the form of τ_0. If k is odd, a renaming of the vertices allows τ_0 to be written as the first product given in Definition 2.3. If k is even and $\tau(k)$ is odd, the same product arises, but if $\tau(k)$ is even, the second product given in Definition 2.3 arises instead. The only difference between these two maps is that the former has no fixed points and the latter has two fixed points. As it will later turn out, these self-reversed cycles of even length will not be of much importance to us, and if the length of the cycle is odd, the two products are the same.

Finally, in the case where k is a divisor of n but τ does not map $(1, 2, \ldots k)$ onto itself, we may again assume (up to renaming) that $\tau(1) = 2k$. From this assumption and the fact that τ reverses the cycle, τ_0 must satisfy Definition 2.4. □

Given the last lemma and a list of factors of $2n$, one may enumerate all possible forms of a single vertex orbit in an embedding of Dic_n.

Now, we turn to the question of which combinations of vertex orbits generate all of Dic_n, rather than some strict subgroup. We will arrive at a result weaker than the converse to Lemma 2.5 which gives the conditions under which a combination of vertex orbits produces Dic_n. It turns out that these restrictions are not subtle; any combination suffices, so long as $|\sigma| = 2n$ and $|\tau|$ are correct:

Theorem 2.6. *Suppose σ and τ are elements of a symmetric group S_k such that $|\sigma| = 2n$ and the restrictions of σ and τ to each vertex orbit are consistent with Lemma 2.5. Then $\langle \sigma, \tau \rangle \cong Dic_n$.*

Proof. We proceed similar to the previous proof. For each x in $\{1, 2, \ldots k\}$, define σ_x to be the product of the cycles in the disjoint cycle decomposition of σ which permute elements of $\mathcal{O}_{\langle \sigma, \tau \rangle}\{x\}$; define τ_x analogously. Let R be a set containing an single member of each orbit of vertices. Observe that σ_x commutes with σ_y and τ_y whenever $y \notin \mathcal{O}\{x\}$, because they permute disjoint sets of symbols. Similarly, τ_x commutes with σ_y and τ_y if $y \notin \mathcal{O}\{x\}$. Therefore, if the group relations for Dic_n are satisfied by (σ_x, τ_x) for each x, they are satisfied by $\sigma = \prod_{x \in R} \sigma_x, \tau = \prod_{x \in R} \tau_x$. We know that $|\sigma| = 2n$ by hypothesis, so the lcm of the lengths of the cycles in the disjoint cycle decomposition of σ must be $2n$. By the above observation, as long as $|\sigma| = 2n$, any σ and τ permitted by Lemma 2.5 will satisfy the correct group relations.

Now, we must show that no additional relations are satisfied. Call the lengths of the cycles in the decomposition of σ by the names $\{k_1, k_2, \ldots k_n\}$. Because $\mathrm{lcm}\{k_1, k_2, \ldots k_n\} = 2n$, there must be some j such that 2^{b+1} divides k_j. But this means that k_j does not divide n, so by the lemma two cycles of length k_j must be twisted together by τ, meaning that $|\tau| = 4$ exactly. Since σ and τ satisfy the appropriate group relations, $\langle \sigma, \tau \rangle \leq \mathrm{Dic}_n$, we may write any element of $\langle \sigma, \tau \rangle$ in the form $\sigma^k \tau^a$, and any additional relation may be written in the form $\sigma^r \tau^a = \sigma^s \tau^b$ where $(r, a), (s, b) \in \mathbb{Z}_{2n} \times \mathbb{Z}_4, (k, a) \neq (r, b)$. An additional nontrivial relation of this form where $a = b$ would mean that $|\sigma| < 2n$, contradicting our assumption. Therefore, our relation must be of the form $\sigma^r \tau = \sigma^s$, or by conjugation, $\tau = \sigma^k$ where $k = s - r$. Up to rotation, we may assume the two twisted cycles of length k_j are written on the first $2k_j$ symbols; so again, we find that $\tau(1) \notin \mathcal{O}_{\langle \sigma \rangle}\{1\}$, contradicting any relation of the form $\tau = \sigma^k$. Either way, we have reached a contradiction. □

With Theorem 2.6 in hand, we may greatly simplify the statements and proofs of future results by establishing a canonical way of labeling embeddings, or in other words, a way of choosing a particular embedding of $\mathrm{Dic}_n \to S_k$ from the many embeddings which differ only by an inner automorphism of S_k, which we view as merely relabeling the vertices of corresponding graphs. Our scheme is as follows:

- The vertices of each cycle in σ are to be consecutive and in ascending order. For example, we may have $\sigma = (1, 2, 3, 4)(5, 6, 7, 8)$, but not $\sigma = (1, 2, 3, 4)(6, 7, 8, 9)$ or $\sigma = (1, 2, 3, 4)(5, 7, 6, 8)$.
- Vertex orbits should be ordered by the smallest prime factor of the length of a cycle in σ. For example, a pair-revered pair of 9-cycles come before a pair-revered pair of 5 cycles.

- Among vertex orbits with the same smallest prime factor, twisted vertex orbits come first, than pair-reversed, then self-reversed.
- Vertex orbits with the same smallest prime factor and type should be ordered so that smaller orbits come before larger ones.

As a first example, the embedding of Dic_{30} which has a twisted pair of 4-cycles, pair-reversed pairs of 3-cycles and 5-cycles, and a self-reversed 3-cycle may be written as follows:

$$\sigma = (1,2,3,4)(5,6,7,8)(9,10,11)(12,13,14)(15,16,17)(18,19,20,21,22)(23,24,25,26,27)$$
$$\tau = (1,8,3,6)(2,7,4,5)(9,14)(10,13)(11,12)(15,17)(18,27)(19,26)(20,25)(21,24)(22,23)$$

3 Restrictions on Graphs with Dicyclic Automorphism Group

So far, we have classified all injective homomorphisms from Dic_n to S_N, including those defined by group actions of graph automorphism groups. Now, we turn to the question of which ones correspond to graphs of minimal order.

Suppose we have some embedding $\phi : \mathrm{Dic}_n \to S_k$, and we wish to know if it is realized by a graph Γ on the vertices $\{1, 2, \ldots k\}$ with $\mathrm{Aut}(\Gamma) = \phi(\mathrm{Dic}_n)$; that is, the automorphisms of Γ are literally the same permutations, not merely an isomorphic permutation group. Since a graph automorphism must preserve adjacency, it is immediate that $\phi(\mathrm{Dic}_n) \leq \mathrm{Aut}(\Gamma)$ if and only if $\{a, b\} \in E(\Gamma)$ precisely when $\mathcal{O}_{\phi(\mathrm{Dic}_n)}\{a, b\} \subset E(\Gamma)$. Therefore, a graph $\Gamma = (V, E)$ with $\mathrm{Dic}_n \leq \mathrm{Aut}(\Gamma)$ is uniquely identified by the embedding $\phi : \mathrm{Dic}_n \to S_V$ and the set of edge orbits of V under $\phi(\mathrm{Dic}_n)$ which are included in E. For ease of notation, we will often omit the embedding ϕ when there is only one embedding being considered, and we will write $\mathcal{O}\{a, b\}$ where $\mathcal{O}_{\mathrm{Dic}_n}\{a, b\}$ is intended, reserving the use of subscripts for denoting edge orbits under other groups.

Using this method of describing graphs associated with a particular embedding, we will now obtain several results concerning requirements that realizable embeddings must satisfy. Doing this serves two purposes. First, when we later exhibit an edge-minimal graph corresponding to the smallest realizable embedding of Dic_n, the work of this section will provide an intuitive explanation of much of the graph's structure. Second, we will use these lemmas to show which embedding of Dic_n is vertex-minimal by ruling out all smaller embeddings.

Our first result is the generalization of Lemma 7 from [4] to Dic_n.

Lemma 3.1. *Embeddings of* Dic_n *into symmetric groups with less than two twisted vertex orbits are not realizable.*

Proof. By Theorem 2.6, any embedding Dic_n must contain at least one twisted vertex orbit. So, suppose that Γ is a graph corresponding to an embedding ϕ of Dic_n with only one twisted vertex orbit, made up of $2k$ vertices. We claim that the permutation given by

$$\gamma(v) = \begin{cases} \tau^2(v) & v \in \mathcal{O}_\sigma\{1\} \\ v & v \notin \mathcal{O}_\sigma\{1\} \end{cases}$$

is a member of $\mathrm{Aut}(\Gamma)$ but not of Dic_n; more concisely, we call γ an extra automorphism.

Clearly, γ preserves all edges between vertices not in $\mathcal{O}\{1\}$, because all such vertices are left in place. For the same reason, all edges between vertices in $\mathcal{O}_\sigma\{\ell + 1\}$ and vertices not in $\mathcal{O}\{1\}$ are preserved. Suppose $v \in \mathcal{O}_\sigma\{1\}$ and $w \notin \mathcal{O}\{1\}$. By Lemma 2.5, we know that all non-twisted vertex orbits have cycles of length a factor of n in σ, so $\tau^2 = \sigma^n$ fixes all such vertices. Since $\tau^2 \in \mathrm{Aut}(\Gamma)$, whenever $\{v, w\}$ is an edge, $\tau^2\{v, w\} = \{\tau^2(v), w\}$ must also be an edge. But $\gamma(v) = \tau^2(v)$, so $\gamma\{v, w\} = \tau^2\{v, w\}$ is an edge, and all edges between vertices in $\mathcal{O}\{1\}$ and the rest of the graph are preserved be γ.

The only remaining edges are within $\mathcal{O}\{1\}$. If both endpoints of an edge are in $\mathcal{O}_\sigma\{1\}$, we know that γ acts the same as τ^2 on that edge, preserving it; similarly, γ preserves edges between vertices in $\mathcal{O}_\sigma\{\ell + 1\}$. This leaves only edges with one vertex in $\mathcal{O}_\sigma\{1\}$ and one vertex in $\mathcal{O}_\sigma\{k + 1\}$. If $v \in \mathcal{O}_\sigma\{1\}$ then $\tau(v) \in \mathcal{O}_\sigma\{k + 1\}$, so for every edge between these two orbits under σ is of the form $\{v, \sigma^k\tau(v)\}$ for some k. But $(\sigma^k\tau)^2 = \tau^2$, so $\sigma^k\tau\{v, \sigma^k\tau(v)\} = \{\sigma^k\tau(v), \tau^2(v)\} = \gamma(\{v, \sigma^k\tau(v)\})$, and γ also preserves these edges.

Proofs that show that γ preserves an edge $e \in E(\Gamma)$ by giving another edge $f \in E(\Gamma)$ and an automorphism $g \in \mathrm{Dic}_n$ where $g(f) = \gamma(e)$ are common, so to abbreviate and clarify them, we write in charts. One summarizing this proof is given below.

$\mathcal{O}\{v_1, v_2\}$	$e \in \mathcal{O}\{v_1, v_2\}$	$\gamma(e)$	$\pi \in \langle\sigma, \tau\rangle : \pi(e*) = \gamma(e)$
$\mathcal{O}\{1, 1 + k\}$	$\{1 + x, 1 + k + x\}$ $\{2\ell - x, 2\ell - (k + x)\}$	$\tau^2\{1 + x, 1 + k + x\}$ $\{2\ell - x, 2\ell - (k + x)\}$	$\tau^2\sigma^k\{1 + x, 1 + k + x\})$ $\sigma^k\{2\ell - x, 2\ell - (k + x)\}$
$\mathcal{O}\{1, \ell + 1 + k\}$	$\{1 + x, \ell + 1 + (k + x - 1)\}$ $\{2\ell - x, \tau^2(1) - (k + x)\}$	$\{\tau^2(1) + x, \ell + 1 + (k + x - 1)\}$ $\{2\ell - x, 1 - (k + x)\}$	$\sigma^x\tau\sigma^{-k}\{1, \ell + 1 + (k - 1)\}$ $\sigma^{-x}\tau^{-1}\sigma^k\{2\ell, \tau^2(1) - k\}$

∎

Before we continue, we will need a technical lemma that will be applied in proving many of the following results. It is well known that a graph Γ and its complement have the same automorphism group, because automorphisms must preserve non-edges as well as edges in order to preserve adjacency in the graph. In the following lemma, we will obtain an analogous but stronger result about the subgroup $A(\Gamma, G) \leq \mathrm{Aut}(\Gamma)$, which is defined in Definition 0.10. It turns out that replacing all edges between a pair of vertex orbits with non-edges and vice versa does not change $A(\Gamma, G)$.

Lemma 3.2. *Suppose $G \leq \mathrm{Aut}(\Gamma)$. For any $v, w \in V(\Gamma)$, define $C(v, w)$ to be $\{\{a, b\} : a \in \mathcal{O}_G\{v\}, b \in \mathcal{O}_G\{w\}, a \neq b\}$, and define $\Gamma' = (V(\Gamma), [E(\Gamma) \setminus C(v, w)] \bigcup [C(v, w) \setminus E(\Gamma)])$. Then $A_G(\Gamma) = A_G(\Gamma')$.*

Proof. Suppose that ϕ is some automorphism in $A_G(\Gamma)$. Consider some arbitrary edge (nonedge) $e = \{a, b\}$. If $\mathcal{O}\{e\}$ is not a subset of $C(v, w)$, then e is an edge (nonedge) in both Γ and Γ', so ϕ maps e to an edge (nonedge) in either case. On the other hand, if $\mathcal{O}\{e\} \subseteq C(v, w)$, then $\mathcal{O}\{\phi(e)\} \subseteq C(v, w)$, by the definition of A_G; that is $A_G(\Gamma)$ leaves $C(v, w)$ invariant. Therefore, both e and $\phi(e)$ are edges in $C(v, w) \bigcap E(\Gamma)$ (nonedges in $C(v, w) \setminus E(\Gamma)$), and by definition of Γ', both are nonedges (edges) in Γ'.

Since in every case, ϕ maps edges and nonedges of Γ' to edges and nonedges respectively, $\phi \in A_G(\Gamma')$. Because ϕ was chosen arbitrarily, we have $A_G(\Gamma) \leq A_G(\Gamma')$; to prove that $A_G(\Gamma') \leq A_G(\Gamma)$, switch Γ and Γ' and apply the above logic again. □

Arlinghaus previously obtained a slightly weaker form of this result; see Lemma 3.3 of [6]. As we will soon see, this result has many applications to vertex minimization. Most of the remaining work will be based on a number of lemmas giving conditions under which a graph Γ has $\mathrm{Aut}(\Gamma) > \mathrm{Dic}_n$, which are proven by exhibiting an element $\gamma \in \mathrm{Aut}(\Gamma)$ such that $\gamma \notin \mathrm{Dic}_n$. Without exception, it will be clear that $\gamma \in A(\Gamma, \mathrm{Dic}_n)$, so Lemma 3.2 will greatly broaden the sets of graphs that each lemma rules out. For example, if a result implies that a graph Γ with $\mathrm{Aut}(\Gamma) = \mathrm{Dic}_n$ must include k edge orbits between a pair of vertex orbits, Lemma 3.2 also allows us to demand that least $2k$ such orbits exist.

Showing that self-reversed 1-cycles will not appear in the vertex minimal embedding of Dic_n is a relatively simple application of the result we have just obtained.

Lemma 3.3. *Suppose $\phi : \mathrm{Dic}_n \to S_k$ is an isomorphism such that the embedding $\phi' : \mathrm{Dic}_n \to S_{k+1}$ given by adding the symbol $k + 1$ fixed by σ and τ has a corresponding graph Γ' with $\mathrm{Aut}(\Gamma') = \mathrm{Dic}_n$. Let Γ be the induced subgraph of Γ' on the first k vertices. Then $A_{\mathrm{Dic}_n}(\Gamma) = \mathrm{Dic}_n$.*

Proof. Let Γ^+ be the graph formed by adding the vertex $k+1$ to Γ, but no edges. By examination of the permutations involved, we can see that $A_{\mathrm{Dic}_n}(\Gamma^+) = A_{\mathrm{Dic}_n}(\Gamma)$. Any possible Γ' can be formed by adding some edge orbits between other vertices and $k + 1$ to Γ^+. There is exactly one edge orbit between every other vertex orbit and $\{k + 1\}$, as $k + 1$ is fixed by the members of $\phi'(\mathrm{Dic}_n)$. Therefore, by Lemma 3.2, $A_{\mathrm{Dic}_n}(\Gamma') = A_{\mathrm{Dic}_n}(\Gamma^+) = A_{\mathrm{Dic}_n}(\Gamma)$. □

The next two restrictions on realizable embeddings we will obtain are the most useful in explaining the structure of the vertex-minimal embedding of Dic_n, so after we have them, we will exhibit the minimal embedding.

The embedding of Dic_n with the fewest vertices simply includes a twisted vertex orbit of minimum size and a self-reversed p^m cycle for each multiplicity m prime divisor p of n. However, if a self-reversed p^m cycle is the only cycle with length a multiple of p in some embedding, the resulting embedding will not be realizable. The following lemma gives a much more general result, albeit one that is harder to state.

Lemma 3.4. *Suppose $\phi : \mathrm{Dic}_n \to S_N$ is an embedding where there exists a self-reversed ℓ_0-cycle of σ such that there is no sequence $\ell_0, \ell_1, \ldots \ell_m$ of lengths of cycles where $\gcd(\ell_{k-1}, \ell_k) > 1$ whenever all $1 \leq k \leq m$ and ℓ_m is a non-self-reversed cycle of σ. Then ϕ is not realizable.*

Proof. Let W be the set of all vertices of Γ in cycles of σ in all sequences of cycles of lengths $\ell_0, \ell_1, \ldots \ell_m$ with the property that $\gcd(\ell_{k-1}, \ell_k) > 1$ for all k. By hypothesis, all of these cycles are self-reversed cycles, and none of them have length with a common factor with the length of a non-self reversed cycle. Consider $\gamma : V(\Gamma) \to V(\Gamma)$ given by

$$\gamma(v) = \begin{cases} \tau(v) & v \in W \\ v & v \notin W \end{cases}$$

Clearly, γ is an automorphism of induced by W and $V(\Gamma) \setminus W$; this leaves only the edges between W and $V(\Gamma) \setminus W$. Let Γ' be the graph formed by removing all edges between W and $V(\Gamma) \setminus W$; clearly, $\gamma \in \mathrm{Aut}(\Gamma')$, and because γ is equal to some member of Dic_n at every point, we know that $\gamma \in A(\Gamma', \mathrm{Dic}_n)$. We also know that vertices of W come from self-reversed cycles and lengths of cycles within W have no common factors with lengths of cycles within $V(\Gamma) \setminus W$. Therefore, there is only one orbit of edges between $\mathcal{O}\{w\}$ and $\mathcal{O}\{v\}$ for any $w \in W$, $v \in V(\Gamma) \setminus W$. By Lemma 3.2, this means that $\gamma \in A(\Gamma, \mathrm{Dic}_n) \leq \mathrm{Aut}(\Gamma)$ as well. □

The next few results combine to rule out embeddings of Dic_{pk} where $p \in \{3, 5\}$ is coprime to k, and the only cycles with length a multiple of p are two pair-reversed p-cycles. A result about what edge orbits must be included in the graph is also obtained, which will be revisited later.

Lemma 3.5. *Suppose $i : \mathbb{Z}_p \to \{0, 1\}$ is some function with either $|i^{-1}(0)| \leq 2$ or $|i^{-1}(1)| \leq 2$ where $p \geq 7$ is a prime. Then there exists $x \in Z_p$ such that the function $\gamma_x(k) = x - k$ satisfies $i(k) = i(\gamma_x(k))$.*

Proof. Suppose without loss of generality that $|i^{-1}(0)| \leq 2$. Consider the case $|i^{-1}(0)| = 2$. We may assume without loss of generality that $|i^{-1}(0)| = \{0, r\}$. Then we may choose $x = r$. In the case $|i^{-1}(0)| = 1$, we may assume $i^{-1}(0) = \{0\}$, in which case we choose $x = 0$. In the case $|i^{-1}(0)|$, we may choose any x. ☐

Lemma 3.6. *Suppose p is a prime, with k coprime to p, and $\phi : \mathrm{Dic}_{pk} \to S_N$ is an embedding containing a pair of pair-reversed p cycles on the vertices $\{x + 1, x + 2, \ldots x + 2p\}$. Let Γ be a corresponding graph with no edge orbits between $\mathcal{O}\{x + 1\}$ and other vertex orbits with size not coprime to p. If Γ does not contain edge orbits of the form $\{x + 1, x + p + 1 + r\}$ where the partition $\{\{r\}, \mathbb{Z}_p \setminus \{r\}\}$ forms an asymmetric 2-coloring of a p-cycle, then $\mathrm{Aut}(\Gamma) > \mathrm{Dic}_{pk}$.*

Proof. Any orbit of edges between $\mathcal{O}\{x + 1\}$ and another vertex orbit is of the form $\mathcal{O}\{x + 1, v\}$. Since all vertex orbits other than $\mathcal{O}\{x + 1\}$ with vertices adjacent to $\mathcal{O}\{x + 1\}$ have size coprime to $\mathcal{O}\{x + 1\}$, the orbit $\mathcal{O}\{x + 1, v\}$ contains all edges between vertices of $\mathcal{O}_\sigma\{x + 1\}$ and vertices of $\mathcal{O}_\sigma\{v\}$, as well as all edges between $\mathcal{O}_\sigma\{x + 2p\}$ and $\mathcal{O}_\sigma\{\tau(v)\}$. Consequently, any permutation of vertices γ that preserves orbits under σ will preserve edges between $\mathcal{O}\{x + 1\}$ and the rest of the graph.

Let $R = \{r : \mathcal{O}\{x + 1, x + p + 1 + r\} \subseteq E(\Gamma)\}$. Let $i : \mathbb{Z}_p \to \{0, 1\}$ be given by

$$i(r) = \begin{cases} 1 \text{ if } r \in R \\ 0 \text{ otherwise} \end{cases}$$

By assumption, i does not give an asymmetric 2-coloring of a p-cycle. Let z be the element of $\mathbb{Z}_p$ corresponding to i given by Lemma 3.5. Let τ_r reverse the p cycles onto themselves; explicitly,

$$\tau_r(v) = \begin{cases} v \text{ if } v \notin \mathcal{O}\{x + 1\} \\ x + p - k \text{ where } v = x + 1 + k \in \mathcal{O}_\sigma\{x + 1\} \\ x + 2p - k \text{ where } v = x + p + 1 + k \in \mathcal{O}_\sigma\{x + p + 1\} \end{cases}$$

We claim that the function

$$\gamma(v) = \begin{cases} v & v \notin \mathcal{O}\{x + 1\} \\ \tau_r(v) & v \in \mathcal{O}_\sigma\{x + 1\} \\ \sigma^z \tau_r(v) & v \in \mathcal{O}_\sigma\{x + p + 1\} \end{cases}$$

is an order two member of $\mathrm{Aut}(\Gamma)$ other than τ^2.

Routine calculations show that γ is of order 2. Because $\gamma(v) \in \mathcal{O}_\sigma\{v\}$ for all v, the above reasoning holds, and we immediately know that γ preserves all edges involving vertices outside of $\mathcal{O}\{x + 1\}$. The calculations exhibited in the following chart show that γ preserves the remaining edges. For each edge e in the remaining orbits, we give another edge e' and an automorphism $g \in \mathrm{Dic}_{pk}$ such that $\gamma(e) = g(e')$. Since g is an automorphism and e' is an edge, $\gamma(e)$ must be an edge, as desired.

As shown in the following chart, γ sends $\mathcal{O}\{x + 1, x + p + 1 + r\}$ to $\mathcal{O}\{x + 1, x + p + 1 + \gamma_z(r)\}$ where γ_z is defined as in Lemma 3.5, permuting the orbits of the form $\mathcal{O}\{x + 1, +p + 1 + r\}$ amongst themselves.

$\mathcal{O}\{e\}$	e	$\gamma(e)$	$g(e')$
$\mathcal{O}\{x + 1, x + 1 + r\}$	$\{x + 1 + k, x + 1 + r + k\}$ $\{x + p + 1 + k, x + p + 1 + r + k\}$	$\{x + p - k, x + p - (r + k)\}$ $\{x + 2p - k + z, x + 2p - (r + k) + z\}$	$\sigma^{-(k+r)}(\{x + 1, x + 1 + r\})$ $\sigma^{z-(k+r)}(\{x + p + 1, x + p + 1 + r + k\})$
$\mathcal{O}\{x + 1, x + p + 1 + r\}$	$\{x + 1 + k, x + p + 1 + r + k\}$	$\{x + p - k, x + 2p - (r + k) + z\}$	$\sigma^{-(k+r+z)}(\{x + 1, x + 1 + i(r)\})$

☐

Corollary 3.7. *Suppose that $p \in \{3, 5\}$, that k is coprime to p, and $\phi : \mathrm{Dic}_{pk} \to S_N$ is an embedding containing a pair of pair-reversed p cycles on the vertices $\{x + 1, x + 2, \ldots x + 2p\}$. Let Γ be a corresponding graph with no edge orbits between $\mathcal{O}\{x + 1\}$ and other vertex orbits with size not coprime to p. Then $\mathrm{Aut}(\Gamma) > \mathrm{Dic}_{pk}$.*

Proof. There are $p < 2 \cdot 3$ orbits of edges of the form $\mathcal{O}\{x + 1, x + p + 1 + r\}$. If fewer than 3 edge orbits of the form $\mathcal{O}\{x + 1, x + p + 1 + r\}$ were included in $E(\Gamma)$, the previous lemma tells us that $\mathrm{Aut}(\Gamma) > \mathrm{Dic}_n$. Otherwise, fewer than 3 edge orbits of the form $\mathcal{O}\{x + 1, x + p + 1 + r\}$ were excluded. By applying Lemma 3.2, we know that Γ has the same automorphism group as some graph Γ' including only the edge orbits between vertices of $\mathcal{O}\{x + 1\}$ that Γ excluded, and hence $\mathrm{Aut}(\Gamma) > \mathrm{Dic}_n$ again. ☐

4 An Edge-Minimal Vertex-Minimal Graph with Dicyclic Automorphism Group

We will now exhibit an edge-minimal graph Γ with $\mathrm{Aut}(\Gamma) = \mathrm{Dic}_n$ and $|V(\Gamma)| = \alpha(\mathrm{Dic}_n)$, and verify that it has the correct automorphism group. We will only consider n not a power of 2, because the case $n = 2^b$ has been previously addressed. [4]

It will be useful to determine the stabilizers of different vertex orbits and vertices. The following lemma summarizes the calculations.

Lemma 4.1. *Suppose* Γ *is a graph with* $\mathrm{Aut}(\mathrm{Dic}_n)$*, with* $v \in V(\Gamma)$ *and* $k = |\mathcal{O}_\sigma\{v\}|$*. Then for some* $r \in \mathbb{Z}_k$*, the values given in the following chart are correct:*

Type of $\mathcal{O}\{v\}$	$\mathrm{fix}(\mathcal{O}\{v\})$	$\mathrm{stab}(v)$
Twisted	$\langle\sigma^k\rangle$	$\langle\sigma^k\rangle$
Pair-reversed	$\langle\sigma^k\rangle$	$\langle\sigma^k\rangle$
Self-reversed	$\langle\sigma^k\rangle$	$\langle\sigma^k\rangle \bigcup \sigma^r\tau + \langle\sigma^k\rangle$

Proof. No matter what type of vertex orbit $\mathcal{O}\{v\}$ is, each cycle in σ on the vertices of $\mathcal{O}\{v\}$ has length k, so in each case σ^k fixes all the vertices of $\mathcal{O}\{v\}$. Because $|\sigma| = 2n$, the size of $\langle\sigma^k\rangle$ is $\frac{2n}{k}$. For twisted and pair-reversed $\mathcal{O}\{v\}$, there are $2k$ vertices in $\mathcal{O}\{v\}$, and by the orbit-stabilizer theorem, $|\operatorname{stab} v| = \frac{|\mathrm{Dic}_n|}{|\mathcal{O}\{v\}} = \frac{4n}{2k} = \frac{2n}{k}$, meaning that $\langle\sigma^k\rangle$ is all of $\operatorname{stab} v$. On the other hand, suppose $\mathcal{O}\{v\}$ is self-reversed. Then there are only k vertices in $\mathcal{O}\{v\}$, so $|\operatorname{stab} v| = 2 \cdot \frac{2n}{k}$, twice the size of $\langle\sigma^k\rangle$. Because $\tau(v) \in \mathcal{O}_\sigma\{v\}$, there exists an r such that $\sigma^r\tau(v) = v$, so $\sigma^r\tau + \langle\sigma^k\rangle \subseteq \operatorname{stab} v$. Combined, these two cosets have the required $2 \cdot \frac{2n}{k}$ elements. In general, we see that

$$\begin{aligned} v &= \sigma^r\tau(v) \\ \sigma^s(v) &= \sigma^{r+s}\tau(v) \\ \sigma^s(v) &= \sigma^{r+2s}\tau\sigma^s(v), \end{aligned}$$

meaning that $\sigma^{r+2s}\tau + \langle\sigma^k\rangle$ fixes $\sigma^s(v)$. Since this coset depends on s, its elements are not in $\mathrm{fix}(\mathcal{O}\{v\})$. ☐

In the graph we will show, there are no edge orbits between vertex orbits with coprime odd lengths, and with one exception, cycles will only have prime power length. The graph may be constructed by taking a particular configuration of edge orbits and vertex orbits, or 'chunk', corresponding to each maximal prime power divisor of $2n$, or 15 in the special case where 3 and 5 are maximal prime power divisors of n, and connecting all other chunks to the chunk corresponding to 2^{b+1} in a certain way. That the graph has the correct automorphism group will essentially follow from the fact that the subgraphs induced by the the chunk corresponding to $2^{b+1} + 1$ and any other chunk have the correct automorphism group, and that powers of different primes are coprime. The next 8 lemmas exhibit the minimal graph under for certain classes of n with few prime factors; the constructions they use will then be combined to give the result in general.

While the following lemma is stated for any odd number z, we will make use of two special cases: $z = p^m$ where $m > 1$ or $p \in \{3, 5\}$, and $(p, z) = (3, 15)$. These are the cases in which the graph given in the lemma is vertex minimal. The techniques used in proving the lemma are essentially borrowed from Graves et all (see [4] Theorem 18).

Lemma 4.2. *Suppose* $n = 2^b z$*, where* $b > 0$ *and* $z > 1$ *is odd, is a positive even integer; let* p *be a prime factor of* z*. Consider the embedding of* $\phi : \mathrm{Dic}_n \to S_N$ *where* σ *contains* 4 2^{b+1}*-cycles twisted by* τ*, 2* p*-cycles pair-reversed together by* τ*, and a single* z*-cycle self-reversed by* τ*. Then the graph* Γ *on* $2^{b+3} + 2p + z$ *vertices with edges orbits generated by* $\{\{1, 2^{b+1}+1\}, \{1, 2^{b+2}+1\}, \{1, 2^{b+2}+2\}, \{1, 3\cdot 2^{b+1}+1\}, \{1, 2^{b+3}+1\}, \{2^{b+3}+1, 2^{b+3}+p+1\}, \{2^{b+3}+1, 2^{b+3}+2p+2\}, \{2^{b+3}+2p+1, 2^{b+3}+2p+2\}\}$ *under the action of* $\phi(\mathrm{Dic}_n)$ *has* $\mathrm{Aut}(\Gamma) \cong \mathrm{Dic}_n$*.*

Proof. The techniques used in proving this lemma will be illustrative of future proofs. First, we prove that $\mathrm{Aut}(\Gamma) = A(\Gamma, \mathrm{Dic}_n)$; that is, that the orbit of a vertex under $\mathrm{Aut}(\Gamma)$ is the same as its orbit under $\phi(\mathrm{Dic}_n)$. Equivalently, this means that any automorphism $\gamma \in \mathrm{Aut}(\Gamma)$ is, at any particular vertex, the restriction of an element of Dic_n. Then we will only need to show that γ is the restriction of the same element of Dic_n at all vertices; or in other words, that different vertices 'agree' about the identity of γ.

For similarity with future proofs, let $x = 2^{b+3}$. There are four orbits of vertices in Γ: $\mathcal{O}\{1\}$, $\mathcal{O}\{2^{b+2}+1\}$, $\mathcal{O}\{x+1\}$, and $\mathcal{O}\{x+2p+1\}$. Examination of the orbits comprising Γ shows that the degrees of vertices in different orbits are as follows:

$\mathcal{O}\{v\}$	$\deg(v)$
$\mathcal{O}\{1\}$	$5+p$
$\mathcal{O}\{2^{b+2}+1\}$	3
$\mathcal{O}\{x+1\}$	$2^{b+1}+1+\frac{z}{p}$
$\mathcal{O}\{x+2p+1\}$	4

Since $2^{b+1} \geq 2$ and $p \geq 2$, it is clear that only vertices in $\mathcal{O}\{x+2p+1\}$ have degree 4; therefore, this vertex orbit is invariant under $\mathrm{Aut}(\Gamma)$. Since the only neighbors of $\mathcal{O}\{x+2p+1\}$ outside of that orbit are the members of $\mathcal{O}\{x+1\}$, that orbit is also invariant under $\mathrm{Aut}(\Gamma)$. The two remaining twisted vertex orbits have vertices of different degree, so their vertices may not be interchanged, and they are also invariant under $\mathrm{Aut}(\Gamma)$. Hence $\mathrm{Aut}(\Gamma) = A(\Gamma, \mathrm{Dic}_n)$.

Let γ be a member of $\mathrm{Aut}(\Gamma)$. We claim that there is a unique element g in Dic_n such that $\gamma\{1, x+2p+2\} = g\{1, x+2p+2\}$. First, consider the case that $\gamma(1) \in \mathcal{O}_\sigma\{1\}$, say $\gamma(1) = \sigma^a(1)$. Since $\mathcal{O}\{x+2p+2\}$ is self-reversed, $\gamma(x+2p+2) = \sigma^b(x+2p+2)$ for some b. Then, any element g of $\sigma^{-b} + \mathrm{stab}\,\mathcal{O}\{x+2p+2\} \bigcap \sigma^{-a} + \mathrm{stab}\,\mathcal{O}\{1\}$ is equal to γ^{-1} at both 1 and $x+2p+2$. By Lemma 4.1, $\mathrm{stab}\,\mathcal{O}\{1\} = \langle \sigma^{2^{b+1}} \rangle$ and $\mathrm{stab}\,\mathcal{O}\{x+2p+1\} = \langle \sigma^z \rangle$, and because z and 2^{b+1} are coprime, the intersection is nonempty, so such a g exists. Since $2^{b+1}z = n$, the intersection has precisely one element, and g is unique. On the other hand, if $\gamma(1) = \sigma^a\tau(1)$, we know that $\gamma(x+2p+2) = \sigma^b\tau(x+2p+2)$ for some b, and we may obtain the desired g similarly.

Let $\alpha = g^{-1}\gamma$. If we can show that α is the identity, then we will know that $\gamma = g$, and hence $\gamma \in \mathrm{Dic}_n$. We already know that α fixes 1 and $x+2p+2$, so we will argue that an automorphism that fixes these two vertices must fix all of $V(\Gamma)$.

From the two vertices that α is known to fix and the fact that $\mathrm{Aut}(\Gamma) = A(\Gamma, \mathrm{Dic}_n)$, we may sequentially deduce that:

- α leaves $\mathcal{O}_\sigma\{x+1\}$ invariant, because these vertices are the neighbors of 1 in $\mathcal{O}\{x+1\}$.
- Similarly, α leaves $\mathcal{O}_\sigma\{x+p+1\}$, the non-neighbors of 1 in $\mathcal{O}\{x+1\}$, invariant.
- α also leaves $\mathcal{O}_\sigma\{1\}$ and $\mathcal{O}_\sigma\{2^{b+1}+1\}$ invariant, as the former is the set of neighbors of $\mathcal{O}_\sigma\{x+1\}$ in $\mathcal{O}\{1\}$, and the latter is the set of non-neighbors.
- Finally, α leaves $\mathcal{O}_\sigma\{2^{b+2}+1\}$ and $\mathcal{O}_\sigma\{3 \cdot 2^{b+1}+1\}$ invariant, because the members of the former orbit have 2 neighbors in $\mathcal{O}_\sigma\{1\}$ while the latter only have 1 such neighbor.
- α fixes $x+1$, the only common neighbor of fixed vertices 1 and $x+2p+2$.
- α fixes $x+p+3$, the only unfixed neighbor of $x+2p+2$ in $\mathcal{O}\{x+1\}$.
- α fixes $x+3$, the only common neighbor of fixed vertices 1 and $x+p+3$.
- α fixes $2^{b+2}+2$, the only neighbor of 1 in $\mathcal{O}_\sigma\{2^{b+2}+1\}$ that is not adjacent to another neighbor of 1.
- α fixes $2^{b+1}+2$, the remaining neighbor of 1 in $\mathcal{O}_\sigma\{2^{b+2}+1\}$.
- α fixes the remaining neighbors of 1, because they are the only neighbors of 1 in the invariant orbits $\mathcal{O}_\sigma\{2^{b+1}+1\}$ and $\mathcal{O}_\sigma\{3 \cdot 2^{b+1}+1\}$.
- α fixes 2, the only neighbor of $2^{b+2}+2$ in $\mathcal{O}_\sigma\{1\}$.

Notice that the edges of $\mathcal{O}\{x+2p+1, x+2p+2\}$ form a z-cycle; since the edges of this cycle are the only edges between vertices in $\mathcal{O}\{x+2p+1\}$, a set of vertices left invariant by $\mathrm{Aut}(\Gamma)$, this z-cycle must be preserved by α. Because α fixes one of the vertices of the cycle, namely $x+2p+2$, we know that α either fixes all the vertices of $\mathcal{O}\{x+2p+1\}$ or else flips the cycle about $x+2p+2$. Suppose that α indeed flips the z-cycle about $x+2p+2$. There is only one orbit of edges connecting $\mathcal{O}\{x+1\}$ and $\mathcal{O}\{x+2p+1\}$, namely $\mathcal{O}\{x+1, x+2p+2\}$, so α would also have to reverse the order of $\mathcal{O}_\sigma\{x+1\}$. However, we already know that two vertices of $\mathcal{O}_\sigma\{x+1\}$ are fixed, namely $x+1$ and $x+3$, and that $|\mathcal{O}_\sigma\{x+1\}| = p$ is odd, so α does not reverse the order of $\mathcal{O}_\sigma\{x+1\}$, a contradiction. Hence α fixes every vertex in $\mathcal{O}\{x+2p+1\}$. Since each vertex in $\mathcal{O}\{x+2p+1\}$ has one neighbor in $\mathcal{O}_\sigma\{x+1\}$, all members of $\mathcal{O}_\sigma\{x+1\}$ are also fixed; the same goes for the members of $\mathcal{O}_\sigma\{x+p+1\}$.

To see that the remainder of the graph is fixed, recall that every neighbor of 1 is fixed, and that $2 = \sigma(1)$ and $x+2p+3 = \sigma(x+2p+2)$. Applying the arguments we have just used and replacing each vertex v with $\sigma(v)$ suffices to show that every neighbor of 2 is fixed, as well as 3 and $x+2p+4$. By induction, every vertex in $\mathcal{O}_\sigma\{1\}$ is fixed, as are all their neighbors, which include all the vertices of $\mathcal{O}\{1\}$ and $\mathcal{O}\{2^{b+2}+1\}$. Therefore, α fixes every vertex of Γ, as desired. □

When z is prime, the orbit of edges between the vertices of the self-reversed p-cycle are not needed, as shown below.

Lemma 4.3. *Suppose $n = 2^b p$, where $b > 0$ and p is an odd prime. Consider the embedding of $\phi : \mathrm{Dic}_n \to S_N$ where σ contains 4 2^{b+1}-cycles twisted by τ, 2 p-cycles pair-reversed together by τ, and a single p-cycle self-reversed by τ. Then the graph Γ on $2^{b+3}+3p$ vertices with edges orbits generated by $\{\{1, 2^{b+1}+1\}, \{1, 2^{b+2}+1\}, \{1, 2^{b+2}+2\}, \{1, 3\cdot 2^{b+1}+1\}, \{1, 2^{b+3}+1\}, \{2^{b+3}+1, 2^{b+3}+p+1\}, \{2^{b+3}+1, 2^{b+3}+2p+2\}\}$ under the action of $\phi(\mathrm{Dic}_n)$ has $\mathrm{Aut}(\Gamma) \cong \mathrm{Dic}_n$.*

Proof. We proceed as in the previous proof. Let $x = 2^{b+3}$. Calculation determines that vertices in the given orbits have the following degrees:

$\mathcal{O}\{v\}$	$\deg(v)$
$\mathcal{O}\{1\}$	$5+p$
$\mathcal{O}\{2^{b+2}+1\}$	3
$\mathcal{O}\{x+1\}$	$2^{b+1}+1+\frac{z}{p}$
$\mathcal{O}\{x+2p+1\}$	2

A similar argument to the one given in Lemma 4.2 shows that each vertex orbit is invariant under $\mathrm{Aut}(\Gamma)$.

Define γ, g, and ρ as in the previous proof.

All of the bulleted deductions in the previous proof still hold, so we know that $x+3$ is fixed by ρ. Hence $x+2p+3$, the unique neighbor of $x+3$ in $\mathcal{O}\{x+2p+1\}$, is also fixed. By the Chinese remainder theorem, there exists some k such that σ^k is equal to σ^2 on the vertices of $\mathcal{O}\{x+1\}$ and $\mathcal{O}\{x+2p+1\}$ and equal to the identity on vertices of $\mathcal{O}\{1\}$ and $\mathcal{O}\{2^{b+2}+1\}$. By applying the argument previously given with each vertex v replaced with $\sigma^k(v)$, we may show that all vertices in $\mathcal{O}\{x+1\}$ are fixed. This means that all vertices in $\mathcal{O}\{x+2p+1\}$ are fixed as well, because each has a unique neighbor in $\mathcal{O}_\sigma\{x+1\}$. We may show that the remainder of the graph is fixed in the same way as in the previous lemma. □

If a prime factor p is sufficiently large and has multiplicity 1, the self-reversed p-cycle can be omitted entirely.

Lemma 4.4. *Suppose $n = 2^b p$ where $b > 0$ and $p > 7$ is a prime number. Let $\phi : \mathrm{Dic}_n \to S_N$ be the embedding where σ contains 4 2^{b+1}-cycles twisted by τ and 2 p-cycles pair reversed by τ. Then the graph Γ on $2^{b+3}+2p$ vertices with edges orbits generated by $\{\{1, 2^{b+1}+1\}, \{1, 2^{b+2}+1\}, \{1, 2^{b+2}+2\}, \{1, 3\cdot 2^{b+1}+1\}, \{1, 2^{b+3}+1\}, \{2^{b+3}+1, 2^{b+3}+p+1\}, \{2^{b+3}+1, 2^{b+3}+p+3\}, \{2^{b+3}+1, 2^{b+3}+p+4\}\}$ under the action of $\phi(\mathrm{Dic}_n)$ has $\mathrm{Aut}(\Gamma) \cong \mathrm{Dic}_n$.*

Proof. We proceed as in the previous proof. Let $x = 2^{b+3}$. Calculation determines that vertices in the given orbits have the following degrees:

$\mathcal{O}\{v\}$	$\deg(v)$
$\mathcal{O}\{1\}$	$5+p$
$\mathcal{O}\{2^{b+2}+1\}$	3
$\mathcal{O}\{x+1\}$	$2^{b+1}+3$

Observe that vertices in $\mathcal{O}\{1\}$ and $\mathcal{O}\{x+1\}$ have degree greater than 3, so $\mathcal{O}\{2^{b+2}+1\}$ is invariant. As for the remaining orbits, vertices in $\mathcal{O}\{1\}$ have even degree (since p is odd), while those in $\mathcal{O}\{x+1\}$ have odd degree, so $\mathcal{O}\{1\}$ and $\mathcal{O}\{x+1\}$ are also invariant. Therefore, $\mathrm{Aut}(\Gamma) = A(\Gamma, \mathrm{Dic}_n)$.

Let γ be a member of $\mathrm{Aut}(\gamma)$. By the same reasoning as in Lemma 4.2, there exists $g \in \mathrm{Dic}_n$ such that $g(1) = \gamma(1)$ and $g(x+1) = \gamma(x+1)$. Define $\alpha = g^{-1}\gamma$, and we will again show that α is the identity.

From the two vertices that α is known to fix and the fact that $\mathrm{Aut}(\Gamma) = A(\Gamma, \mathrm{Dic}_n)$, we may sequentially deduce that:

- α fixes $x+p+5$, because $x+p+5$ is a common neighbor of the other two neighbors of $x+1$, which are not adjacent to each other.
- α leaves $\{x+p+2, x+p+4\}$ invariant, because these two vertices are the remaining unfixed neighbors of $x+1$.
- α fixes $x+3$, the only unfixed common neighbor of $x+p+2$ and $x+p+4$.

Observe that $x+3=\sigma^2(x+1)$. With $x+3$ is fixed, we may apply the above argument replacing each vertex v with $\sigma^2(v)$, and by induction, all of $\mathcal{O}_{\sigma^2}\{x+1\}$ is fixed. But since $\operatorname{stab}\mathcal{O}\{x+1\}=\langle\sigma^p\rangle$ and p is coprime to 2, we know that $\mathcal{O}_{\sigma^2}\{x+1\}=\mathcal{O}_\sigma\{x+1\}$. This also means that the vertices of $\mathcal{O}_\sigma\{x+p+1\}$ are fixed, since each has a different set of three neighbors in $\mathcal{O}_\sigma\{x+1\}$. Therefore, all of $\mathcal{O}\{x+1\}$ is fixed. From here, the same reasoning used in the proof of Lemma 4.2 will show that the remaining vertices of Γ are fixed. □

In the special case $p=7$, the above construction can be used if an extra orbit of edges is included.

Lemma 4.5. *Suppose $n=2^bp$ where $b>0$ and $p\geq 7$ is a prime number. Let $\phi:\mathrm{Dic}_n\to S_N$ be the embedding where σ contains 4 2^{b+1}-cycles twisted by τ and 2 p-cycles pair reversed by τ. Then the graph Γ on $2^{b+3}+2p$ vertices with edges orbits generated by $\{\{1,2^{b+1}+1\},\{1,2^{b+2}+1\},\{1,2^{b+2}+2\},\{1,3\cdot 2^{b+1}+1\},\{1,2^{b+3}+1\},\{2^{b+3}+1,2^{b+3}+2\},\{2^{b+3}+1,2^{b+3}+p+1\},\{2^{b+3}+1,2^{b+3}+p+3\},\{2^{b+3}+1,2^{b+3}+p+4\}\}$ under the action of $\phi(\mathrm{Dic}_n)$ has $\mathrm{Aut}(\Gamma)\cong\mathrm{Dic}_n$.*

Proof. We proceed as in the previous proof. Let $x=2^{b+3}$. Calculation determines that vertices in the given orbits have the following degrees:

$\mathcal{O}\{v\}$	$\deg(v)$
$\mathcal{O}\{1\}$	$5+p$
$\mathcal{O}\{2^{b+2}+1\}$	3
$\mathcal{O}\{x+1\}$	$2^{b+1}+5$

Again, observe that vertices in $\mathcal{O}\{1\}$ and $\mathcal{O}\{x+1\}$ have degree greater than 3, so $\mathcal{O}\{2^{b+2}+1\}$ is invariant. As for the remaining orbits, vertices in $\mathcal{O}\{1\}$ have even degree because p is odd, while those in $\mathcal{O}\{x+1\}$ have odd degree, so $\mathcal{O}\{1\}$ and $\mathcal{O}\{x+1\}$ are also invariant. Therefore, $\mathrm{Aut}(\Gamma)=A(\Gamma,\mathrm{Dic}_n)$.

Let γ, g, and α be defined as in Lemma 4.4. As in the proof of Lemma 4.2, we note that the sets $\mathcal{O}_\sigma\{x+1\}$ and $\mathcal{O}_\sigma\{x+p+1\}$ are invariant under α because they are the sets of neighbors and non-neighbors (respectively) of 1, a vertex fixed by α. The edges connecting the vertices of $\mathcal{O}_\sigma\{x+1\}$ form a p-cycle which must be preserved by α; the same goes for $\mathcal{O}_\sigma\{x+p+1\}$. Because $x+1$ is fixed, γ must either fix the p-cycle of the vertices of $\mathcal{O}_\sigma\{x+1\}$ or flip it about $x+1$. Therefore,

- α leaves $\{x+p,x+2\}$ invariant, because they are the neighbors of $x+1$ in $\mathcal{O}_\sigma\{x+1\}$.
- α leaves $\{x+p+3,x+p+4\}$ invariant, because they are the neighbors of $x+1$ in $\mathcal{O}_\sigma\{x+p+1\}$ which are adjacent to members of $\{x+p,x+2\}$.
- α fixes $x+p+1$, the remaining neighbor of $x+1$ in $\mathcal{O}_\sigma\{x+p+1\}$.
- α either fixes $\mathcal{O}_\sigma\{x+p+1\}$ or flips it about $x+p+1$, because $x+p+1$ is fixed and α must preserve the p-cycle of edges between the vertices of $\mathcal{O}_\sigma\{x+p+1\}$.
- α fixes $\mathcal{O}_\sigma\{x+p+1\}$, because flipping it about $x+p+1$ does not leave $\{x+p+3,x+p+4\}$ invariant.
- α fixes $\mathcal{O}_\sigma\{x+1\}$, because each vertex has a different set of neighbors in $\mathcal{O}_\sigma\{x+p+1\}$.

From here, the same reasoning used in the proof of Lemma 4.2 will show that the remaining vertices of Γ are fixed. □

When n is odd similar constructions may be used as for even n, with some modification to take into the account the fact that there are less potential orbits of edges between the 8 vertices of the twisted cycles.

Lemma 4.6. *Suppose $n>1$, is a positive odd integer, and let p be a prime factor of n. Consider the embedding of $\phi:\mathrm{Dic}_n\to S_N$ where σ contains 4 2-cycles twisted by τ, 2 p-cycles pair-reversed together by τ, and a single n cycle self-reversed by τ. Then the graph Γ on $8+2p+n$ vertices with edges orbits generated by $\{\{1,3\},\{1,5\},\{1,7\},\{1,8+1\},\{5,8+1\},\{8+1,8+p+1\},\{8+1,8+2p+2\},\{8+2p+1,8+2p+2\}\}$ under the action of $\phi(\mathrm{Dic}_n)$ has $\mathrm{Aut}(\Gamma)\cong\mathrm{Dic}_n$.*

Proof. Again, let $x=8$. The degrees of vertices in Γ are as follows:

$\mathcal{O}\{v\}$	$\deg(v)$
$\mathcal{O}\{1\}$	$4+p$
$\mathcal{O}\{5\}$	$2+p$
$\mathcal{O}\{x+1\}$	$5+\frac{z}{p}$
$\mathcal{O}\{x+2p+1\}$	4

Similar to previous proofs, we can see that only vertices in $\mathcal{O}\{x+2p+1\}$ have degree 4, so $\mathcal{O}\{x+2p+1\}$ is invariant under $\mathrm{Aut}(\Gamma)$. The orbit $\mathcal{O}\{x+1\}$ is also invariant under $\mathrm{Aut}(\Gamma)$, because they are the only vertices adjacent to $\mathcal{O}\{x+2p+1\}$. The remaining two orbits have vertices of different degree, so they are also invariant under $\mathrm{Aut}(\Gamma)$. Therefore, $\mathrm{Aut}(\Gamma) = A(\Gamma, \mathrm{Dic}_n)$.

Define γ, g, and α as in the proof of Lemma 4.2, so that α fixes 1 and $x+2p+2$. From the two vertices that α is known to fix and the fact that $\mathrm{Aut}(\Gamma) = A(\Gamma, \mathrm{Dic}_n)$, we may sequentially deduce that:

- α leaves $\mathcal{O}_\sigma\{x+1\}$ invariant, because these vertices are the neighbors of 1 in $\mathcal{O}\{x+1\}$.
- Similarly, α leaves $\mathcal{O}_\sigma\{x+p+1\}$, the non-neighbors of 1 in $\mathcal{O}\{x+1\}$, invariant.
- α also leaves $\mathcal{O}_\sigma\{1\}$ and $\mathcal{O}_\sigma\{3\}$ invariant, as the former is the set of neighbors of $\mathcal{O}_\sigma\{x+1\}$ in $\mathcal{O}\{1\}$, and the latter is the set of non-neighbors.
- Similarly, α leaves $\mathcal{O}_\sigma\{5\}$ and $\mathcal{O}_\sigma\{7\}$ invariant, as the former is the set of neighbors of $\mathcal{O}_\sigma\{x+1\}$ in $\mathcal{O}\{5\}$, and the latter is the set of non-neighbors.
- α fixes $x+1$, the only common neighbor of fixed vertices 1 and $x+2p+2$.
- α fixes $x+p+3$, the only unfixed neighbor of $x+2p+2$ in $\mathcal{O}\{x+1\}$.
- α fixes $x+3$, the only common neighbor of fixed vertices 1 and $x+p+3$.
- α fixes 2, the only member of $\mathcal{O}_\sigma\{1\}$ other than 1.
- α fixes 5, the only neighbor of 1 in $\mathcal{O}_\sigma\{5\}$, and 6, the only neighbor of 2 in $\mathcal{O}_\sigma\{5\}$.
- α fixes 7 and 8 analogously.
- α fixes 3, the only neighbor of 6 in $\mathcal{O}_\sigma\{3\}$, and 4, the last unfixed vertex in $\mathcal{O}\{1\}$.

The argument given in the proof of Lemma 4.2 shows that the remaining vertices of $\mathcal{O}\{x+1\}$ and $\mathcal{O}\{x+2p+1\}$ are fixed by α, so α fixes all of the vertices of Γ. $\square$

Lemma 4.7. *Suppose $n > 1$, is a positive odd integer, and let p be a prime factor of n. Consider the embedding of $\phi : \mathrm{Dic}_n \to S_N$ where σ contains 4 2-cycles twisted by τ, 2 p-cycles pair-reversed together by τ, and a single n cycle self-reversed by τ. Then the graph Γ on $8+2p+n$ vertices with edges orbits generated by $\{\{1,2\},\{1,5\},\{1,7\},\{1,8+1\},\{5,8+1\},\{8+1,8+p+1\},\{8+1,8+2p+2\}\}$ under the action of $\phi(\mathrm{Dic}_n)$ has $\mathrm{Aut}(\Gamma) \cong \mathrm{Dic}_n$.*

Proof. We proceed as in the previous proof. Let $x = 2^{b+3}$. Calculation determines that vertices in the given orbits have the following degrees:

$\mathcal{O}\{v\}$	$\deg(v)$
$\mathcal{O}\{1\}$	$3+p$
$\mathcal{O}\{5\}$	$2+p$
$\mathcal{O}\{x+1\}$	$5+\frac{z}{p}$
$\mathcal{O}\{x+2p+1\}$	2

A similar argument to the one given in Lemma 4.6 shows that each vertex orbit is invariant under $\mathrm{Aut}(\Gamma)$.

Define γ, g, and ρ as in the previous proof.

A similar argument to the one given in the previous Lemma and in Lemma 4.3 suffices. $\square$

Lemma 4.8. *Suppose $n = p$ where $p > 7$ is a prime number. Let $\phi : \mathrm{Dic}_n \to S_N$ be the embedding where σ contains 4 2-cycles twisted by τ and 2 p-cycles pair reversed by τ. Then the graph Γ on $8+2p$ vertices with edges orbits generated by $\{\{1,2\},\{1,5\},\{1,7\},\{1,8+1\},\{5,8+1\},\{8+1,8+p+1\},\{8+1,8+p+3\},\{8+1,8+p+4\}\}$ under the action of $\phi(\mathrm{Dic}_n)$ has $\mathrm{Aut}(\Gamma) \cong \mathrm{Dic}_n$.*

Proof. Let $x = 2^{b+3}$. Calculation determines that vertices in the given orbits have the following degrees:

$\mathcal{O}\{v\}$	$\deg(v)$
$\mathcal{O}\{1\}$	$3+p$
$\mathcal{O}\{2^{b+2}+1\}$	$2+p$
$\mathcal{O}\{x+1\}$	7

Since $p > 5$, all vertex orbits in Γ have vertices of different degrees. The remainder of the proof is very similar to the proofs of Lemma 4.4 and Lemma 4.6, and is therefore omitted. $\square$

Lemma 4.9. *Suppose $n = p$ where $p \geq 7$ is a prime number. Let $\phi : \mathrm{Dic}_n \to S_N$ be the embedding where σ contains 4 2-cycles twisted by τ and 2 p-cycles pair reversed by τ. Then the graph Γ on $8+2p$ vertices with edges orbits generated by $\{\{1,2\},\{1,5\},\{1,7\},\{1,8+1\},\{5,8+1\},\{8+1,8+2\},\{8+1,8+p+1\},\{8+1,8+p+3\},\{8+1,8+p+4\}\}$ under the action of $\phi(\mathrm{Dic}_n)$ has $\mathrm{Aut}(\Gamma) \cong \mathrm{Dic}_n$.*

Proof. Let $x = 2^{b+3}$. Calculation determines that vertices in the given orbits have the following degrees:

$\mathcal{O}\{v\}$	$\deg(v)$
$\mathcal{O}\{1\}$	$3+p$
$\mathcal{O}\{5\}$	$2+p$
$\mathcal{O}\{x+1\}$	9

If $p > 7$, all vertex orbits in Γ have vertices of different degrees. On the other hand, in the case $p = 7$, a more careful analysis is required. In this case, vertices of $\mathcal{O}\{1\}$ have degree 11 while other vertices have degree 9, so $\mathcal{O}\{1\}$ is invariant under $\mathrm{Aut}(\Gamma)$. Vertices in $\mathcal{O}\{5\}$ have two neighbors in $\mathcal{O}\{1\}$, while vertices in $\mathcal{O}\{x+1\}$ have four neighbors in $\mathcal{O}\{1\}$, so $\mathcal{O}\{5\}$ and $\mathcal{O}\{x+1\}$ are also invariant under $\mathrm{Aut}(\Gamma)$.

The remainder of the proof proceeds as in Lemma 4.5 and Lemma 4.6. ☐

With these base cases in hand, we are finally able to exhibit our minimal graph.

Theorem 4.10. *Suppose n is not a power of 2. Let $\phi : \mathrm{Dic}_n \to S_N$ be the embedding where σ contains*

- *Four 2^{b+1} cycles twisted by τ.*
- *If 3 and 5 are both multiplicity 1 factors of n, two 3-cycles pair-reversed by τ and a 15-cycle self-reversed by τ.*
- *For each j from 1 to J, two q_j cycles pair-reversed by τ and a $q_j^{m_j}$ cycle self-reversed by τ, unless $q_j \in \{3,5\}$ and 3 and 5 are both multiplicity 1 factors of n.*
- *For each i from 1 to I, two p_i cycles pair-reversed by τ.*

Let

$$N' = 2^{b+3} + \sum_{i=1}^{I} 2p_i + \sum_{j=1}^{J} 2q_j + q_j^{m_j};$$

let $N = N' - 3$ if 3 and 5 are both multiplicity 1 factors of n, and $N = N'$ otherwise. Define

$$x_i = 2^{b+3} + \sum_{k=1}^{i-1} 2p_k.$$

Define y_j similarly, so that y_j is the first vertex in the pair-reversed vertex orbit of size $2q_j$. Explicitly: define $y_1 = x_{I+1}$. If 3 and 5 are both multiplicity 1 factors of n, define y_3 as $y_1 + 21$ and

$$y_j = y_2 + \sum k = 3^j 2q_j + q_j^{m_j}$$

for $j > 3$. Otherwise, define

$$y_j = y_1 + \sum k = 1^{j-1} 2q_j + q_j^{m_j}.$$

Let f be $\min 2^{b+1}, p_1, q_1$; define

$$s = \begin{cases} 0 & f = 2^{b+1} \\ x_1 & f = p_1 \\ y_1 & f = q_1 \end{cases}$$

Then the graph Γ on N vertices with edges generated by

- *If $b \geq 1$, the members of $\{\{1, 2^{b+1}+1\}, \{1, 2^{b+2}+1\}, \{1, 2^{b+2}+2\}, \{1, 3 \cdot 2^{b+1}+1\}\}$.*
- *If $b = 0$, the members of $\{\{1,5\},\{1,7\}\}$.*
- *If $b = 0$ and n is a prime power, $\{1,2\}$.*
- $\{1, s+1\}$

- *$\{s, x_i+1\}$ and $\{s+1, y_j+1\}$ for each i and j where the two vertices are distinct.*
- *If $b=0$, the edge $\{5, s+1\}$.*
- *For each i, the members of $\{\{x_i+1, x_i+p_i+1\}, \{x_i+1, x_i+p_i+3\}, \{x_i+1, x_i+p_i+4\}\}$.*
- *If $p_i = 7$, the edge $\{x_i+1, x_i+2\}$.*
- *For each j, the members of $\{\{y_j+1, y_j+p_j+1\}, \{y_j+1, y_j+2q_j+2\}\}$.*
- *For each j where $|\mathcal{O}\{y_j+2q_j+1\}|$ is not prime, $\{y_j+2q_j+1, y_j+2q_j+2\}$.*

under the action of $\phi(\mathrm{Dic}_n)$ has $\mathrm{Aut}(\Gamma) \cong \mathrm{Dic}_n$.

Proof. As in the proofs of the lemmas, we begin by showing that $\mathrm{Aut}(\Gamma) = A(\Gamma, \mathrm{Dic}_n)$. We begin by considering the vertices in $\mathcal{O}\{y_j+1\}$ and $\mathcal{O}\{y_j+2q_j+1\}$ for each j. Calculation shows that the vertices of $\mathcal{O}\{y_j+2q_j+1\}$ for each j are the only vertices of degree 2 or 4, so the union of these orbits must be left invariant under α. Notice that the only edges between vertices of degree 2 or 4 are generated by the edges $\{\{y_j+2q_j+1, y_j+2q_j+2\} : 1 \le j \le J\}$, so the edges between vertices of degree 4 always form a $q_j^{m_j}$-cycle or a 15-cycle for each orbit. All of these cycles are of different length, so they cannot be mapped onto each other by members of $\mathrm{Aut}(\Gamma)$, so $\mathcal{O}\{y_j+2q_j+1\}$ is invariant under $\mathrm{Aut}(\Gamma)$ for each j. As in previous proofs, this means that $\mathcal{O}\{y_j+1\}$, which contains all the other neighbors of vertices in $\mathcal{O}\{y_j+2q_j+1\}$, must be invariant for each j.

Now we intend to show that the vertices of $\mathcal{O}\{1\}$ are invariant. If $b>0$, then because vertices in $\mathcal{O}\{5\}$ have different degree than those in $\mathcal{O}\{1\}$ and no neighbors outside of $\mathcal{O}\{5\}$ and $\mathcal{O}\{1\}$, we know that $\mathcal{O}\{1\}$ must be invariant. If $b=0$, the fact that 2 is the smallest prime number ensures that $s=0$. If we also have $J>0$, we observe that the vertices of $\mathcal{O}\{1\}$ are the only external neighbors of $\mathcal{O}\{y_1+1\} \bigcup \mathcal{O}\{y_1+2q_1+1\}$, and so they must be invariant. If $b=J=0$, calculation reveals that vertices of $\mathcal{O}\{1\}$ have greater degree than all other vertices, and so they must be invariant.

Next, we show that $\mathcal{O}\{5\}$ is invariant. If $b>0$, then vertices in $\mathcal{O}\{5\}$ have degree 3, while those in $\mathcal{O}\{x_i+1\}$ have larger degree, so they may not be exchanged. If $b=0$ and q_1 is the least prime factor of n, then the only external neighbors of vertices in $\mathcal{O}\{5\}$ are in invariant orbits $\mathcal{O}\{1\}$ and $\mathcal{O}\{y_1+1\}$, so $\mathcal{O}\{5\}$ is invariant. On the other hand, if $b=0$ and p_1 is the least prime factor of n, things are slightly more complicated. Under these circumstances, the degree of a vertex in $\mathcal{O}\{5\}$ is $2+p_1$, while the degree of a vertex in $\mathcal{O}\{x_1+1\}$ is 7 if $p_1 \ne 7$ and 9 if $p_1 = 7$, and the degree of a vertex in $\mathcal{O}\{x_i+1\}$ is 5 for each $i \ge 2$. This means that the vertices in $\mathcal{O}\{5\}$ have different degree from those in every other orbit that is not already known to be invariant unless $p_1 = 7$, in which case vertices in $\mathcal{O}\{5\}$ and $\mathcal{O}\{x_1+1\}$ both have degree 9. In this special case, we will have to use other means to show that vertices of $\mathcal{O}\{5\}$ and $\mathcal{O}\{x_1+1\}$ may not be interchanged by α. If α interchanged a vertex v in $\mathcal{O}\{5\}$ with a vertex w in $\mathcal{O}\{x_1+1\}$, it would have to map the neighborhood of v onto the neighborhood of w. Calculation shows that w has a neighbor adjacent to no other neighbors of w (for $w = x_1+1$, the neighbor is x_1+p_1+1), and that no neighbor of v has this property, so α cannot map the neighborhood of v onto the neighborhood of w.

It is clear that $\mathcal{O}\{s+1\}$ is invariant: $\mathcal{O}\{1\}$ is invariant, and either $s+1=1$ or $\mathcal{O}\{s+1\}$ is the only orbit of vertices adjacent to $\mathcal{O}\{1\}$ that has not yet been shown to be invariant.

Finally, we show that the remaining orbits, $\{\mathcal{O}\{x_i+1\} : 1 \le i \le I, x_i \ne s\}$, are invariant. For each i except 1 if $x_1 = s$, the subgraph of Γ induced by the vertices of $\mathcal{O}\{x_i+1\}$ is connected, but there are no edges between vertices in $\mathcal{O}\{x_a+1\}$ and $\mathcal{O}\{x_b+1\}$ whenever $a \ne b$. Since $\mathrm{Aut}(\Gamma)$ leaves all orbits not of the form $\mathcal{O}\{x_i+1\}$ invariant, this means that α must map each orbit $\mathcal{O}\{x_a+1\}$ onto a unique $\mathcal{O}\{x_b+1\}$. But whenever $a \ne b$, we have $|\mathcal{O}\{x_a+1\}| \ne |\mathcal{O}\{x_b+1\}|$, so we know that every orbit $\mathcal{O}\{x_a+1\}$ is mapped onto itself. Hence $\mathrm{Aut}(\Gamma) = A(\Gamma, \mathrm{Dic}_n)$.

The remainder of the proof is essentially applying the relevant combination of Lemmas 4.2 through 4.9 to subgraph of Γ induced by the vertices of $\mathcal{O}\{1\} \bigcup \mathcal{O}\{2^{b+1}+1\} \bigcup \mathcal{O}\{x_i+1\}$ for each i and to $\mathcal{O}\{1\} \bigcup \mathcal{O}\{2^{b+1}+1\} \bigcup \mathcal{O}\{y_j+1\}$ for each j simultaneously. To see that for any γ in $\mathrm{Aut}(\Gamma)$ there exists a $g \in \mathrm{Dic}_n$ such that γ agrees with g on 1 and y_j+2q_j+2 or x_i+1 (as appropriate), simply observe that the g given by the six lemmas must exist and agree by the values of the stabilizers of vertex orbits given in Lemma 4.1 and the Chinese Remainder Theorem. Therefore, we may define $\alpha = \gamma g^{-1}$ as before.

Applying the foregoing lemmas to the various induced subgraphs is straightforward, because these induced subgraphs match the graphs in the hypotheses of the lemmas, except in three details. First, the edge orbits $\mathcal{O}\{1, x_i\}$ and $\mathcal{O}\{1, y_j\}$ are replaced by $\mathcal{O}\{s, x_i\}$, $\mathcal{O}\{s, y_j\}$, and $\mathcal{O}\{1, s\}$. The subgraph involving the vertex $s+1$ has the same form as is considered in one of the above lemmas, so we may show that its vertices are fixed by α. As for the remaining subgraphs, notice that the only use of the connection between 1 and the vertices of non-twisted cycles was to point out that, since $\mathcal{O}_\sigma\{1\}$ and $\mathcal{O}_\sigma\{2^{b+1}\}$ are invariant, so are the neighbors of $\mathcal{O}_\sigma\{1\}$; since we already know that α fixes the vertices of $\mathcal{O}\{s+1\}$, this is equally true of $\mathcal{O}_\sigma\{s+1\}$ and $\mathcal{O}_\sigma\{s+f+1\}$, so a similar argument is possible. Second, when $b = 0$ and p_i or q_j is not the minimal prime factor of n, the induced subgraph of Γ leaves out the edges of $\mathcal{O}\{5, p_i + 1\}$ or $\mathcal{O}\{5, q_j + 1\}$. Examination of the relevant arguments shows that these edges were only needed to demonstrate that $\mathcal{O}_\sigma\{5\}$ and $\mathcal{O}_\sigma\{7\}$ are invariant under the action of α, which is still true in Γ since either $\mathcal{O}\{5, x_1 + 1\}$ or $\mathcal{O}\{5, y_1 + 1\}$ was included. Third, when $b = 0$ and n has multiple prime factors, the induced subgraph of Γ leaves out $\mathcal{O}\{1, 2\}$. Note that the only reason these edges were needed for the proofs of the lemmas was to show that vertices of $\mathcal{O}\{1\}$ have different degree from those in $\mathcal{O}\{5\}$; if n has multiple multiple prime factors, this is true without the edges of $\mathcal{O}\{1, 2\}$. $\square$

5 Showing Our Embedding Is Vertex Minimal

Now we have a realized embedding $\mathrm{Dic}_n \to S_k$, which we assert is vertex minimal. To show it is vertex minimal, we will first prove several lemmas showing certain embeddings are not realizable. Then we will exhibit an enumeration of all embeddings not ruled out be these lemmas, which we will use to show that all except finitely many (for a given n) have more vertices than our previously exhibited embedding. We will prove these few are not realizable more directly.

The next two lemmas show that two more classes of embeddings of Dic_n are not realizable. These classes of embeddings were not dealt with previously because they only use fewer vertices than the minimal realizable embedding under specific conditions.

Lemma 5.1. *Suppose $\phi : \mathrm{Dic}_n \to S_k$ is an embedding such that $\phi(\sigma)$ contains a self-reversed cycle of length $2l$ where l is coprime to the length of every other cycle in $\phi(\sigma)$. Then ϕ is not realizable.*

Proof. Write the $2k$-cycle in question $(v+1, v+2, \dots v+2k)$, and suppose an arbitrary graph Γ with $\phi(\mathrm{Dic}_n) \subseteq \mathrm{Aut}(\Gamma)$. If $\mathrm{Aut}(\Gamma) \neq A_{\mathrm{Dic}_n}(\Gamma)$, it follows directly that $\mathrm{Aut}(\Gamma) \neq \mathrm{Dic}_n$. Otherwise, we claim the extra automorphism

$$\psi = \begin{cases} \sigma\tau(x) & x \in \mathcal{O}\{v+1\} \\ x & x \notin \mathcal{O}\{v+1\} \end{cases}$$

is an involution in $\mathrm{Aut}(\Gamma)$ other than $\phi(\tau^2)$, showing that $\mathrm{Aut}(\Gamma) \neq \phi(\mathrm{Dic}_n)$.

Because ψ agrees with $\sigma\tau$ on $\mathcal{O}\{v + 1\}$, we know that ψ preserves edges in $\Gamma[\mathcal{O}\{v + 1\}]$. Similarly, ψ fixes $V(\Gamma) \setminus \mathcal{O}\{v + 1\}$, and so preserves edges there as well. This leaves only the question of edges between a vertex in $\mathcal{O}\{v + 1\}$ and a vertex not in $\mathcal{O}\{v + 1\}$.

Suppose that the self-reverse of the $2k$-cycle has odd parity. From the definition of self-reverse, we may calculate that

$$\begin{aligned} \sigma\tau(v + 1 + r) &= \sigma(v + 2k - r) \\ \sigma\tau(v + 1 + r) &= v + 2k - r + 1 \\ \sigma\tau(v + 1 + r) &= v + r + 1 + 2(k - r) \\ \sigma\tau(v + 1 + r) &= \sigma^{2(k-r)}(v + r + 1) \end{aligned}$$

Consequently, all we need to do is show that there is an element of Dic_n restricting to σ^2 on $\mathcal{O}\{v+1\}$ and stabilizing $V(\Gamma) \setminus \mathcal{O}\{v+1\}$.

Let $\pi : \mathrm{Dic}_n \to S_{\mathcal{O}\{v+1\}}$ be the restriction homomorphism. By assumption, $|\pi(\sigma)| = 2k$. Hence $|\pi(\sigma^{2n/k})| = |\pi(\sigma^{\gcd(2k,2n/k)})| = k$, and we have $\pi(\sigma^2) = \pi(\sigma^{2n/k \cdot d})$ for some $d \in \mathbb{Z}^+$. But $\sigma^{2n/k}$ must fix $V(\Gamma) \setminus \mathcal{O}\{v+1\}$, because (by hypothesis) no cycle in σ other than $\mathcal{O}\{v + 1\}$ has length with a common factor with k.

Therefore, ψ leaves orbits of edges between $\mathcal{O}\{v+1\}$ and $V(\Gamma)\setminus\mathcal{O}\{v+1\}$ invariant. Explicitly, for any $w \notin \mathcal{O}\{v+1\}$ we have $\psi(\{w, v + 1 + r\}) = \sigma^{2drn/k}(\{w, v + 1 + r\}$. This means that $\psi \in \mathrm{Aut}(\Gamma)$, as desired.

In the case where the self-reverse has even parity, we instead restrict $\sigma^2\tau$. The calculations that result are the same. $\square$

Note that only the parity of the power of σ is important; in the even case, we could have restricted defined ψ as a combination of the identity and τ, with slightly different calculations later on.

Finally, we rule out embeddings of Dic_{3k} where k is coprime to 6 and the only vertex orbit with size a multiple of 3 is twisted.

Lemma 5.2. *Suppose k is coprime to 6. Then any embedding $\phi : \mathrm{Dic}_{3k} \to S_N$ in which σ has 2 2-cycles twisted by τ, 2 6-cycles twisted by τ, and some other non-twisted cycles with length a divisor of k is not realizable.*

Proof. For the reader who is willing to read ahead, a later result, Lemma 6.12, is a more enlightening and general variant on this one. We will give a direct proof here regardless.

First, consider the special case $k = 1$. Without loss of generality, assume $\sigma = (1,2)(3,4)(5,6,7,8,9,10)(11,12,13,14,15,16)$. Direct calculation shows that the permutation $\gamma = (1,2)(3,4)(5,10)(6,9)(7,8)(11,16)(12,15)(13,14)$ is a member of $\mathrm{Aut}(\Gamma)$ but not $\phi(\mathrm{Dic}_{3k})$. We claim that γ is still an extra automorphism regardless of the value of k. Let Γ_1 denote the subgraph induced by the vertices of the twisted cycles of $\phi(\mathrm{Dic}_{3k})$, and let Γ_2 denote the subgraph induced by all other vertices. Note that $\gamma \in A(\Gamma, \mathrm{Dic}_{3k})$. We already know that γ preserves adjacency within Γ_1, and γ clearly preserves adjacency within Γ_2, because γ leaves all vertices of Γ_2 fixed. This leaves only edges between vertices of Γ_1 and vertices of Γ_2, so let w be a vertex of Γ_2. If w is in a self-reversed cycle with length coprime to 3, there is only one orbit of edges between $\mathcal{O}\{w\}$ and $\mathcal{O}\{1\}$, so those adjacencies must be preserved; the same goes for edges between $\mathcal{O}\{w\}$ and $\mathcal{O}\{5\}$. On the other hand, if w is in a pair-reversed pair of cycles with length coprime to 3, then there are two orbits of edges between $\mathcal{O}\{w\}$ and $\mathcal{O}\{1\}$, namely the orbit of all edges between $\mathcal{O}_\sigma\{1\}$ and $\mathcal{O}\{w\}$ and those between $\mathcal{O}_\sigma\{3\}$ and $\mathcal{O}\{w\}$. Since $\gamma(v) \in \mathcal{O}_\sigma\{v\}$ for all vertices v, these orbits are also preserved; the same goes for the two orbits of edges between $\mathcal{O}\{5\}$ and $\mathcal{O}\{w\}$. □

We have now discovered enough restrictions on realizable embeddings of Dic_n to rule out all embeddings on fewer vertices than the graph exhibited in Theorem 4.10. However, the structure of the set of all possible embeddings of Dic_n is complicated, making the proof somewhat technical.

Theorem 5.3. *Suppose $n > 1$ is an integer, and let*

$$N' = 2^{b+3} + \sum_{i=0}^{s} 2p_i + \sum_{j=0}^{t} 2q_j + q_j^{m_j}.$$

If 3 and 5 are multiplicity 1 divisors of n, then let $N = N' - 3$; otherwise, $N = N'$. Then $\alpha(\mathrm{Dic}_n) = N$. Moreover, the embedding shown to be realizable in Theorem 4.10 is the unique embedding with N vertices.

Proof. In Theorem 4.10, we exhibited a graph with automorphism group Dic_n corresponding to an embedding on N vertices. Therefore, we only need to show that no other embedding on at most N vertices exists.

So far, we have discovered several constraints that realizable embeddings of Dic_n must satisfy. In particular, if $\langle\sigma, \tau\rangle \cong \mathrm{Dic}_n$ is realizable, we know that:

- σ must contain at least two twisted vertex orbits (Lemma 3.1). These twisted vertex orbits must have length a multiple of 2^{b+1}, according to Lemma 2.5.
- σ must have some cycle with length a multiple of p_i for each i, and some cycle with length a multiple of $q_j^{m_j}$ for each j, since p_i and $q_j^{m_j}$ divide $|\sigma| = 2n$.
- For each prime power divisor d of n (i.e., let d equal either p_i or $q_j^{m_j}$), there must be some sequence $l_0, l_1 \dots l_m$ of lengths of cycles in σ such that d divides l_0, some cycle of length l_m is not self-reversed, and $\gcd(l_k, l_{k+1}) > 1$ whenever $0 \le k < m$ (Lemma 3.4).

Call any embedding of Dic_n satisfying these three constraints potentially realizable. As the name suggests, we know that all realizable embeddings are potentially realizable. We will show that all potentially realizable embeddings of Dic_n other than the embedding given in Theorem 4.10 that are not ruled out by Lemmas 5.1 and 5.2 use at least $N + 1$ vertices.

In order to demonstrate that all potentially realizable embeddings have the stated number of vertices, we will first specify a non-deterministic procedure for generating realizable embeddings. (By non-deterministic, we mean that the procedure will generate a different embedding depending on which of several options are arbitrarily chosen at certain steps.) Then, we will prove that all potentially realizable embeddings may be generated by this procedure. Finally, we will prove that all embeddings generated by this procedure, except those excluded above, have at least $N + 1$ vertices.

1. Arbitrarily totally order the prime divisors of n other than 2, calling them $f_1, f_2 \dots f_r$ with respective multiplicities $m_1, m_2, \dots m_r$.
2. Define the embedding e_0 of Dic_{2^b} so that σ and τ have $t \ge 2$ pairs of twisted cycles of length 2^{b+1}, and for each k such that $1 \le k \le b$, σ and τ have any chosen number of 2^k cycles, self reversed or pair reversed as desired. Note that all three constraints above are satisfied, so e_0 is a potentially realizable embedding.

3. Produce e_k from e_{k-1} by any combination of the following changes, such that the listed constraints are satisfied and no cycle has length a multiple of $f_k^{m_k+1}$:

   - Multiply the lengths of existing cycles in σ by some power of f_k, preserving the behavior of τ (whether twist, pair-reverse, or self-reverse).
   - Add new cycles to σ with length a power of f_k, pair-reversed or self-reversed by τ as desired.

   Since the listed constraints are still satisfied and no cycle has length divided by too high a power of f_k, e_k is a potentially realizable embedding of Dic_{n_k} where $n_k = 2^b \prod_{i=0}^{k} f_i^{m_i}$.

4. Repeat step 3 until e_r, a potentially realizable embedding of Dic_n, has been obtained.

5. Produce e_{r+1} by adding a chosen number of 1-cycles to σ, pair-reversed or self-reversed as desired. Since 1-cycles have nothing to do with our three constraints, e_{r+1} is a potentially realizable embedding of Dic_n.

First, we must show that any potentially realizable embedding of Dic_n can appear as e_{r+1}, depending on what choices were made during the procedure. Since step 5 allowed us to add any configuration of 1-cycles, we may more simply show that any potentially realizable embedding of Dic_n without 1-cycles can appear as e_r. By Theorem 2.6 and Lemma 3.1, all potentially realizable embeddings of Dic_{2^b} may occur as e_0 whenever $b > 0$.

The special case $b = 0$ demands some attention. The group Dic_1 is, for our purposes, a degenerate case. Applying the presentation given in Definition 0.5, we may see that $\mathrm{Dic}_1 \cong \mathbb{Z}_2 \times \mathbb{Z}_2$, the Klein 4-group. In the case $b = 0$, the embedding e_0 is not actually an embedding of Dic_1, but the constraints are satisfied nonetheless. The theorems referenced in defining potentially realizable only discuss the case $n > 1$, but we can still verify the rest of their hypotheses. Since we are not considering Dic_1 in this paper, we will adopt the lexical convention that embeddings satisfying our above constraints for $n = 1$ are potentially realizable embeddings of Dic_1. Therefore, our procedure does in fact generate all potentially realizable embeddings for n a power of 2, and in particular (bearing in mind our convention regarding potentially realizable embeddings of Dic_1) when $n = 1$. This means we can proceed by induction.

Suppose that n is not a power of 2, and that for all $m < n$, our procedure generates all realizable embeddings of Dic_m, and let e_r denote some embedding of Dic_n without 1-cycles; we will now attempt to construct an e_{r-1} from which we may produce e_r. Picking any total ordering of the odd prime factors of n, we might work backwards, constructing e_{r-1} removing all factors of f_r from the lengths of cycles in σ in our embedding e_r, and removing any resulting 1-cycles entirely. Since e_r is assumed to be a working embedding, step 3 of our procedure will allow us to restore the missing factors and cycles. The only trouble is in showing that the resulting e_{r-1} is a potentially realizable embedding.

Since $f_r \neq 2$ and e_r has at least two pairs of twisted cycles, so does e_{r-1}. Since no factors besides those of f_r were removed, σ still has a cycle with length a multiple of $f_k^{m_k}$ for all $k < r$, regardless of the choice of f_r. However, it is possible that in some sequence demanded by applying Lemma 3.4 to e_r, we have $\gcd(l_k, l_{k+1})$ equal to a power of f_r, meaning that e_{r-1} no longer satisfies Lemma 3.4. We will show that there is always some choice of f_r, and hence some ordering of the prime factors of n which may be chosen in step 1, which avoids this difficulty.

Using the embedding e_r, define the graph X to have vertices $\{2\} \bigcup \{f_k\}$, and edges

$$\{(2, f_k) : f_k \text{ divides the length of a non-self-reversed cycle in } \sigma\} \bigcup \{(p, q) : pq \text{ divides the length of some cycle in } \sigma\}.$$

Every path in X from some f_k to 2 corresponds to a sequence of cycle lengths satisfying Lemma 3.4, and vice versa. If we are given the path $f_k, g_1, g_2, \ldots 2$ in X, each non-final edge (g_k, g_{k+1}) implies the existence of a cycle with length a multiple of $g_k g_{k+1}$, so the three consecutive vertices g_{k-1}, g_k, g_{k+1} imply the existence of cycles of lengths $g_{k-1}g_k$ and $g_{k+1}g_k$, with common factor g_k. Let X' be the corresponding graph for e_{r-1}. The graph X' can be obtained from X by simply removing the vertex f_r and all adjacent edges, so if f_r is not a cut-vertex, e_{r-1} will satisfy the restriction imposed by Lemma 3.4. Since X is a connected graph with at least two vertices, it has at least two non-cut vertices (take two leaves in an arbitrary spanning tree of X), and one of them is not 2; we choose this vertex as f_r.

Now, it remains to show that every embedding produced by our procedure besides the one given in Theorem 4.10 is either non-realizable or has at least $N + 1$ vertices. In an embedding e_k, some cycle of σ must have length a multiple of $f_k^{m_k}$, but no cycle in e_{k-1} has length a multiple of f_k. Therefore, when producing e_k from e_{k-1} in step 3 of our procedure, we must either add a new cycle (or cycles) with length $f_k^{m_k}$ or multiply the lengths of some existing cycle (or cycles) by f_k. These cycles may be twisted, pair-reversed, or self-reversed. If some of these cycles are not self-reversed, or they are preexisting self-reversed cycles, then e_k will immediately satisfy the conclusion of Lemma

3.4, because the cycle will give rise to an edge in X between f_k and some other vertex. If all of these cycles are new self-reversed cycles, then the Lemma requires that a factor of f_k be added to some existing cycle or cycles.

In total, depending on whether new cycles are added or existing cycles are enlarged and where the factor of $f_k^{m_k}$ is added, there are eight minimal choices (in the sense that whatever we do, we must add at least the cycles specified by one of these choices, and may add more). This allows for a case-by-case analysis of the number of vertices added. For readability, let $p = f_k$ and $m = m_k$. Then the eight minimal possibilities are:

1. Change a pair of twisted cycles of length l to length lp^m, adding $2l(p^m - 1)$ vertices.
2. Change a pair of pair-reversed cycles of length l to length lp^m, adding $2l(p^m - 1)$ vertices.
3. Add a pair of pair-reversed cycles of length p^m, adding $2p^m$ vertices.
4. Change a self-reversed cycle of length l to length lp^m, adding $l(p^m - 1)$ vertices.
5. Add a self-reversed cycle of length p^m and change a pair of twisted cycles of length l to length lp, adding $p^m + 2l(p - 1)$ vertices.
6. Add a self-reversed cycle of length p^m and change a pair of pair-reversed cycles of length l to length lp, adding $p^m + 2l(p - 1)$ vertices.
7. Add a self-reversed cycle of length p^m and change a self-reversed cycle of length l to length lp, adding $p^m + l(p-1)$ vertices.
8. Add a self-reversed cycle of length p^m and a pair of pair-reversed cycles of length p, adding $p^m + 2p$ vertices.

The embedding we have realized on N vertices results from using two twists, ordering the prime divisors of n in ascending order, always making choice 8 when $m > 1$ or $p \in \{3, 5\}$ and choice 3 otherwise, and never adding any extra cycles, except that when $p = 5$ we apply choice 4 to change a self-reversed 3 cycle to a 15 cycle if possible. The calculations below show that for most e_{k-1}, choice 8 leads to the smallest possible corresponding e_k if $m > 1$, and choice 3 leads to the smallest possible corresponding e_k otherwise. For some e_{k-1} where this is not the case, we will dismiss the possible e_k that are at least as small as our claimed bound by showing that they are non-realizable. For the rest, e_{k-1} will contain 2 or 4 cycles other than twists (which may be removed without making the embedding unrealizable), meaning that the smallest e_k corresponding to that e_{k-1} will contain more vertices than some e'_k resulting from an alternative e'_{k-1}; creating e_k will add fewer vertices, but to a larger preexisting embedding. Because choices 3 and 8 add entirely new cycles, the number of additional vertices does not depend on what cycles are already present in e_{k-1}, meaning that the smallest possible e_k must arise from the smallest possible e_{k-1}. By induction, the smallest possible e_k will result from making choice 8 or 3 (as determined by m_k) at each step. The case where, $p = 5$, and $m = 1$ is the lone exception to this rule, since the use of choice 4 depends on the existence of a self-reversed 3-cycle in e_{k-1}. However, the fact that this case can only ever arise once during the algorithm allows us to handle it separately.

Recall that in all cases, $p \geq 3$ and $l \geq 2$ (because 1 cycles are not added until the final step).

1. In case $m > 1$, this choice adds more vertices than choice 8, as demonstrated by the inequalities

$$\begin{aligned}
4p &> 4 \\
6p - 4 &> 2p \\
3p^m - 4 &> 2p \\
4p^m - 4 &> p^m + 2p \\
4(p^m - 1) &> p^m + 2p \\
2l(p^m - 1) &> p^m + 2p
\end{aligned}$$

In the case $m = 1$, we examine three subcases:

- If $p \neq 3$, we have

$$\begin{aligned}
p &> 4 \\
3p - 3 &> 2p \\
3(p^m - 1) &> 2p^m \\
2l(p^m - 1) &> 2p^m
\end{aligned}$$

Therefore, this choice adds more vertices than choice 3.

- If $l > 2$, we have

$$\begin{aligned} p &> 2 \\ 2p &> 4 \\ 4p - 4 &> 2p \\ 2l(p-1) &> 2p \\ 2l(p^m - 1) &> 2p^m \end{aligned}$$

  Again, this choice adds more vertices than choice 3.
- If $p = 3$ and $l = 2$, we discover that this choice adds $2l(p^m - 1) = 8$ vertices, 1 less than choice 8 would. However, the resulting graph is not realizable, as shown in Lemma 5.2.

2. This case adds the same number of vertices as the previous case, so the same calculations hold. However, in the case where $m = 1$, $p = 3$, and $l = 2$, we must rule out the resulting embedding in a different way. Since $l = 2$, we know that e_{r-1} contains a pair of pair-reversed 2-cycles; let e'_{r-1} be the embedding that results from removing these 2-cycles, and apply choice 3 to form e'_r. Since this choice only adds one fewer vertex than choice 3, while e'_{r-1} has 4 fewer vertices than e_{r-1}, we know that e'_r has fewer vertices than e_r, so this choice does not really produce the smallest vertex minimal embedding.

3. If $m = 1$, this choice is minimal (but if $p \in \{3, 5\}$, Lemma 3.7 says it is not realizable). If $m > 1$, calculation shows that

$$\begin{aligned} p^m &\geq p^2 \\ p^m &> 2p \\ 2p^m &> p^m + 2p \end{aligned}$$

   Therefore, choice 8 adds fewer vertices than choice 3 when $m > 1$.

4. This choice always adds more vertices than choice 8. In the case $m > 1$, we apply the inequalities

$$\begin{aligned} p &> 2 \\ (p-2)p &> 2 \\ p^2 &> 2p + 2 \\ p^m &> 2p + 2 \\ p^m - 1 &> 2p + 1 \\ (l-1)(p^m - 1) &> 2p + 1 \\ l(p^m - 1) &> p^m + 2p \end{aligned}$$

   In the special case $m = 1$, we must examine a few subcases.

   - If $l = 2$, the resulting embedding adds two fewer vertices than choice 3:

$$\begin{aligned} l(p^m - 1) &= 2(p^m - 1) \\ l(p^m - 1) &= [2p^m] - 2 \end{aligned}$$

     Similar to the reasoning seen in our treatment of choice 2, the fact that $l = 2$ implies the presence of a 2 cycle, so the e_k resulting from this choice can have, at best, exactly as many vertices as our existing vertex minimal embedding. By Lemma 5.1, this e_k is not realizable.
   - If $p \geq 7$ and $l > 2$, choice 3 uses fewer vertices:

$$\begin{aligned} (l-2)(p^m - 1) &> 2 \\ l(p^m - 1) &> 2p^m \end{aligned}$$

- If $p = 3$ and $l > 2$, we additionally know that $l \geq 4$; since we are currently adding factors of 3 to cycle lengths, no 3-cycles already exist. If $l = 4$, we observe that this choice adds $l(p^m - 1) = 8$ vertices compared to 9 from choice 8. Since $l = 4$ implies a preexisting self-reversed 4-cycle, we may dismiss this case along with the previous one.
- On the other hand, if $l > 4$, choice 8 uses fewer vertices, because

$$\begin{aligned} 10 &> 9 \\ 2l &> 3p \\ l(p^m - 1) &> p^m + 2p \end{aligned}$$

- In the case where $p = 5$ and $l \geq 4$, we find that choice 8 uses fewer vertices, because

$$\begin{aligned} p &> 4 \\ 4p - 4 &> 3p \\ l(p-1) &> 3p \\ l(p^m - 1) &> p^m + 2p \end{aligned}$$

- The only remaining case is where $p = 5$ and $l = 3$. In this case, we discover that this choice adds $l(p^m - 1) = 12$ vertices, while choice 8 adds 15 vertices. Assuming that a self-reversed 3-cycle exists and is not superfluous (such as when it was produced due to taking choice 8 when adding factors of 3), choice 4 indeed produces the minimal embedding. Because this is the only scenario in which the minimal embedding depends on the order primes are added, the minimal embedding over all orderings of the prime factors of n occurs when factors of 3 are added before factors of 5, allowing this case to arise.

5. The following inequalities demonstrate that this choice always adds more vertices than choice 8:

$$\begin{aligned} p &> 2 \\ 2p &> 4 \\ 4p - 4 &> 2p \\ 2l(p-1) &> 2p \\ p^m + 2l(p-1) &> p^m + 2p \end{aligned}$$

6. This argument for this choice is the same as in the previous case.
7. In case $m = 1$, this choice adds more vertices than choice 3, as demonstrated by the inequalities

$$\begin{aligned} (l-1)(p-1) &> 1 \\ (l+1)(p-1) + 1 &> 2(p-1) + 2 \\ p + l(p-1) &> 2p \\ p^m + l(p-1) &> 2p \end{aligned}$$

In the case $m > 1$, we must examine several subcases:

- If $l = 2$, we find that this choice adds 2 fewer vertices than choice 8, but the fact that $l = 2$ implies the existence of a self-reversed 2-cycle in e_{k-1}; we dismiss this case the same as we did when considering choice 4.
- If $l > 2$ and $p \neq 3$, we find that this choice uses more vertices than choice 8:

$$\begin{aligned} p &> 3 \\ p - 1 &> 2 \\ (l-2)(p-1) &> 2 \\ l(p-1) &> 2p \\ p^m + l(p-1) &> p^m + 2p \end{aligned}$$

- Finally, if $l > 2$ and $p = 3$, we again exploit the fact that p does not divide l, so $l \geq 4$:

$$\begin{aligned} p - 1 &= 2 \\ (l-2)(p-1) &> 2 \\ l(p-1) &> 2p \\ p^m + l(p-1) &> p^m + 2p \end{aligned}$$

8. If $m > 1$, this choice is generally minimal. If $m = 1$, calculation shows that

$$\begin{aligned} p^m &= p \\ 2p^m &= 2p \\ 2p^m &< p^m + 2p \end{aligned}$$

Therefore, choice 3 adds fewer vertices than choice 8 when $m = 1$.

$\square$

6 Edge Minimization

Our final task is to show that the vertex-minimal graph Γ given in Theorem 4.10 is also edge-minimal, and thereby determine $e(\mathrm{Dic}_n, \alpha(\mathrm{Dic}_n))$. For ease of notation, let $G_n \cong \mathrm{Dic}_n$ be the automorphism group of a graph on $\alpha(\mathrm{Dic}_n)$ vertices. By the uniqueness result in Theorem 5.3, we know the precise permutations that make up G, up to exchanging the names of vertices.

Orbits of edges included in the previously exhibited graph fall into two broad categories: those that are necessary to restrict the motion of the vertices of twisted vertex orbits, and those that restrict the rest of the graph. We begin by showing that the orbits in the latter category included in the graph exhibited by Theorem 4.10 are necessary. This is by far the easier of the two tasks, because many pairs of vertex orbits have only one or two orbits of edges between them.

Unlike in previous (and future) proofs, the extra automorphism exhibited in proving the following lemma will not have order 2; this departure allows for a simpler and more intuitive argument.

Lemma 6.1. *If Γ is a graph with $G_n \leq \mathrm{Aut}(\Gamma)$ and no edge orbit of the form $\mathcal{O}\{y_j + 2q_j + 1, y_j + 2q_j + 1 + k\}$ (where $0 < k < q_j$) for some $j : m_j > 1$, then $\mathrm{Aut}(\Gamma) > G_n$.*

Proof. Consider the map $\phi(x) = \begin{cases} \sigma^{q_j}(x) \text{ if } x \in \{y_j + 2q_j + 1 + kq_j : 0 \leq k < q_j^{m_j - 1}\} \\ x \text{ otherwise} \end{cases}$ We claim that $\phi \in \mathrm{Aut}(\Gamma)$ but $\phi \notin \langle \sigma, \tau \rangle$. For suppose the contrary, that $\phi \in \mathrm{Dic}_n$. Since ϕ agrees with σ^{q_j} on some vertices, it must be equal to some element of the coset $\sigma^{q_j}\,\mathrm{stab}(y_j + 2q_j + 1 + q_j k)$ for each $0 \leq k < q_j^{m_j - 1}$. Lemma 4.1 tells us that $\mathrm{stab}(y_j + 2q_j + 1 + r) = \langle \sigma^{q_j^{m_j}}, \sigma^{2r+1}\tau \rangle$, and the intersection of the stabilizers for $k = 0$ and $k = 1$ is just $\langle \sigma^{q_j^{m_j}} \rangle$. Therefore $\phi \in \sigma^{q_j} \langle \sigma^{q_j^{m_j}} \rangle$, and $\phi = \sigma^{q_j^{m_j}}$ on $\mathcal{O}\{y_j + 2q_j + 1\}$. But $\sigma^{q_j^{m_j}}$ does not fix any element of $\mathcal{O}\{y_j + 2q_j + 1\}$, a contradiction. $\square$

The next lemma is stated for the case where $m_j > 1$, because in the special case where $m_j = 1$ (and hence $q_j \in \{3, 5\}$), we in fact do not need the given orbit.

Lemma 6.2. *If Γ is a graph with $G_n \leq \mathrm{Aut}(\Gamma)$ and no edge orbit of the form $\mathcal{O}\{y_j + 1, y_j + 2q_j + 1 + k\}$ (where $0 \leq k < q_j$) for some $j \in [1, J]$, then $\mathrm{Aut}(\Gamma) > G_n$.*

Proof. By Lemma 3.2, we may assume there are no edges between $\mathcal{O}\{y_j + 2q_j + 1\}$ and the rest of Γ. Therefore, the restriction of τ to $\mathcal{O}\{y_j + 2q_j + 1\}$ is an involution in $\mathrm{Aut}(\Gamma)$ not equal to τ^2, and therefore not in the subgroup $\langle \sigma, \tau \rangle < \mathrm{Aut}(\Gamma)$. $\square$

The following lemma gives two options: connect the two pair-reversed cycles of length q_j or add a second orbit of edges between them and the self-reversed cycle of length $q_j^{m_j}$; the latter always uses at least as many edges as the former, more when $m_j > 1$, allowing us to ignore it. In fact, it is possible to avoid connecting the two pair-reversed cycles of length q_j, but only by including three edge orbits between $\mathcal{O}\{y_j + 1\}$ and $\mathcal{O}\{y_j + 2q_j + 1\}$, in a configuration similar to the one introduced in Lemma 4.4.

Lemma 6.3. *If Γ is a graph with $G_n \leq \mathrm{Aut}(\Gamma)$ and no edge orbit of the form $\mathcal{O}\{y_j+1, y_j+q_j+1+r\}$ and only one orbit of the form $\mathcal{O}\{y_j+1, y_j+2q_j+1+k\}$, then $\mathrm{Aut}(\Gamma) > G_n$.*

Proof. Define τ_r to be the permutation reversing the σ-cycles of $\mathcal{O}\{q_j+1\}$ onto themselves, as in Lemma 3.6. Let γ be given by

$$\gamma(v) = \begin{cases} v & v \notin \mathcal{O}\{y_j+1\} \bigcup \mathcal{O}\{y_j+2q_j+1\} \\ \sigma^{-2k}\tau_r(v) & v \in \mathcal{O}_\sigma\{y_j+1\} \\ \sigma^{2k}\tau_r(v) & v \in \mathcal{O}_\sigma\{y_j+q_j+1\} \\ \tau(v) & v \in \mathcal{O}\{y_j+2q_j+1\} \end{cases}$$

Clearly, $|\gamma| = 2$, yet $\gamma \neq \tau^2$. We will show that $\gamma \in \mathrm{Aut}(\Gamma)$.

Since γ preserves orbits under σ, all edges between the vertices in $\mathcal{O}\{y_j+1\} \bigcup \mathcal{O}\{y_j+2q_j+1\}$ and the rest of the graph are preserved by γ; since γ is the identity outside these two orbits, all edges outside these two orbits are preserved. Since $\gamma = \tau$ on $\mathcal{O}\{y_j+2q_j+1\}$, edges between vertices in that orbit are preserved. Therefore, the only remaining orbits of edges are those between vertices in $\mathcal{O}\{y_j+1\}$ not of the form excluded by hypothesis and $\mathcal{O}\{y_j+1, y_j+2q_j+1+k\}$. The chart below shows that these edges are also preserved.

$\mathcal{O}\{v_1, v_2\}$	$e \in \mathcal{O}\{v_1, v_2\}$	$\gamma(e)$	$\pi \in \langle\sigma, \tau\rangle : \pi(e*) = \gamma(e)$
$\mathcal{O}\{y_j+1, y_j+1+r\}$	$\{y_j+1+x, y_j+1+r+x\}$ $\{y_j+2q_j-x, y_j+2q_j-(r+x)\}$	$\sigma^{-2k}\{y_j+q_j-x, y_j+q_j-(r+x)\}$ $\sigma^{2k}\{y_j+q_j+1+x, y_j+q_j+1+r+x\}$	$\sigma^{-2k-r-x}\{y_j+1, y_j+1+r\}$ $\sigma^{2k+r+1+x}\tau\{y_j+1, y_j+1+r\}$
$\mathcal{O}\{y_j+1, y_j+2q_j+1+k\}$	$\{y_j+1+x, y_j+2q_j+1+k+x\}$ $\{y_j+2q_j-x, y_{j+1}-(k+x)\}$	$\{y_j+q_j-2k-x, y_{j+1}-(k+x)\}$ $\{y_j+q_j+1+2k+x, y_j+2q_j+k+x\}$	$\sigma^{-2k-x}\{y_j+1, y_j+2q_j+1+k\}$ $\sigma^{2k+x+1}\tau\{y_j+1, y_j+2q_j+1+k\}$

As noted above, γ may be thought of as reversing the $\{\mathcal{O}\{y_j+1, y_j+2q_j+1\} : k \in \mathbb{Z}_{q_j}\}$ onto itself. ☐

The next two lemmas essentially demand that the graph have only one component; because the automorphism group of a graph with several components is the direct product of the automorphism groups of the subgraphs induced by each component, this result is not very surprising. Notice that all self-reversed cycles have length coprime to the size of every other vertex orbit except for one pair-reversed vertex orbit, and hence there is only one orbit of edges between a self-reversed vertex orbit and any vertex orbit other than the corresponding pair-reversed vertex orbit. Lemma 3.2 shows that these single edge orbits will not be of any use when it comes to connecting the different pieces of the graph, so instead, edge orbits connecting pair-reversed and twisted vertex orbits to each other must be included.

Lemma 6.4. *If Γ is a graph with $G_n \leq \mathrm{Aut}(\Gamma)$ and no edge orbit of the form $\mathcal{O}\{1, y_j+1+k\}$ or $\mathcal{O}\{2^{b+2}+1, y_j+1+k\}$ (where $0 \leq k < q_j$) for any j, then $\mathrm{Aut}(\Gamma) > \mathrm{Dic}_n$.*

Proof. By Lemma 3.2, we may additionally assume there are no edges between $\mathcal{O}\{1\}$ and $\mathcal{O}\{y_j+2q_j+1\}$ or $\mathcal{O}\{2^{b+2}+1\}$ and $\mathcal{O}\{y_j+2q_j+1\}$. Therefore, the restriction of τ to the union of all the $\mathcal{O}\{y_j+1\}$ and $\mathcal{O}\{y_j+2q_j+1\}$ is an extra involution. ☐

In other words, there must be a connection between the twisted cycles and the chunks for each of the high-multiplicity primes q_j. For any given chunk, however, there are two possibilities: a direct connection from the twists to $\mathcal{O}\{y_j+1\}$ as in the above proof, or a connection between $\mathcal{O}\{y_i+1\}$ and $\mathcal{O}\{y_j+1\}$ where $i \neq j$. Both possibilities suffice, so we must compare the sizes of the needed edge orbits. An orbit of edges between $\mathcal{O}\{1\}$ and $\mathcal{O}\{y_j+1\}$ has $2^{b+2}q_j$ edges, while an orbit between $\mathcal{O}\{y_i+1\}$ and $\mathcal{O}\{y_j+1\}$ has $2q_iq_j$ edges. A similar argument explains why edges connecting vertices in $\mathcal{O}\{1\}$, $\mathcal{O}\{x_i+1\}$, and $\mathcal{O}\{y_j+1\}$ with the vertices of a particular non self-reversed orbit are included.

Lemma 6.5. *If Γ is a graph with $G_n \leq \mathrm{Aut}(\Gamma)$ and no edge orbit between some orbit of the form $\mathcal{O}\{1\}$, $\mathcal{O}\{x_i+1\}$ or $\mathcal{O}\{y_j+1\}$ and any other non-self-reversed orbit of vertices, then $\mathrm{Aut}(\Gamma) = \mathrm{Dic}_n$.*

Proof. The proof is similar to that of 6.4. ☐

The next lemma prepares us to verify that the edge orbits included by the structure first shown in Lemma 4.4 are needed.

Lemma 6.6. *If Γ is a graph with $G_n = \mathrm{Aut}(\Gamma)$, then for each j, Γ contains either have 3 orbits of the form $\mathcal{O}\{y_j+1, y_j+2q_j+1+r\}$ and $q_j \notin \{3,5\}$ or an orbit of the form $\mathcal{O}\{y_j+1, y_j+q_j+1+k\}$.*

Proof. In the vertex minimal embedding of Dic_n, the only vertex orbits with size a multiple of q_j are $\mathcal{O}\{y_j+1\}$ and $\mathcal{O}\{y_j+2q_j+1\}$. Orbits of edges between these two vertex are of the form $\mathcal{O}\{y_j+1, y_j+2q_j+1+k\}$; if no such orbits are included, then more than two orbits of the form $\mathcal{O}\{y_j+1, y_j+q_j+1+r\}$ must be included by Lemma 3.6. ☐

Recall that when exhibiting the vertex minimal graph of Theorem 4.10, we included an edge orbit of the form $\mathcal{O}\{x_i + 1, x_i + 1 + a\}$ precisely when $p_i = 7$. This special case was introduced in Lemma 4.5. The following lemma shows that this edge orbit was necessary.

Lemma 6.7. *If Γ is a graph with $G_n \leq \mathrm{Aut}(\Gamma)$ where 7 is a multiplicity 1 divisor of n, so that (without loss of generality) $p_1 = 7$, and Γ includes no edge orbit of the form $\mathcal{O}\{x_1 + 1, x_1 + 1 + a\}$, then $\mathrm{Aut}(\Gamma) > G_n$.*

Proof. For ease of reading, let $x = x_1$. Assume the contrary, that Γ includes no orbit of the form $\mathcal{O}\{x+1, x+1+a\}$ and yet $\mathrm{Aut}(\Gamma) = \mathrm{Dic}_n$. By Lemma 3.6, we know that Γ must include the edge orbits in some $\{\mathcal{O}\{x+1, x+7+1+a\} : a \in A\}$ such that the members of A form an asymmetric 2-coloring of the 7-gon. A simple exhaustive search reveals that, up to rotation of the names of vertices in $\mathcal{O}\{x+1\}$ and exchanging switching edges and non-edges in accordance with Lemma 3.2, there are only two choices for what set of orbits to take: either $A = \{0, 2, 3\}$ or $A = \{0, 1, 3\}$. In either case, we claim the same extra automorphism: $\gamma = (x+1, x+2)(x+5, x+7)(x+7+1, x+7+2)(x+7+3, x+7+5)$. Note that this automorphism preserves orbits under the action of σ. Whenever there are two orbits between chunks, the difference between them is which pairs of σ-orbits are connected, and either all edges or none between a pair of σ-orbits are included (because the lengths of the cycles in each chunk are coprime). Therefore, all orbits of edges between $\mathcal{O}\{x+1\}$ and the rest of the graph are preserved. Orbits of edges with neither vertex in $\mathcal{O}\{x+1\}$ are also preserved, since $\gamma = e$ on those vertices. Since we have assumed the absence of orbits of edges within each σ-orbit in $\mathcal{O}\{x+1\}$, the only edges we have to worry about are those between $\mathcal{O}_\sigma\{x+1\}$ and $\mathcal{O}_\sigma\{x+p+1\}$. We know precisely which 21 edges of this form have been included: the members of $\bigcup_{a \in A} \mathcal{O}\{x+1, x+8+a\}$. Calculation reveals that γ preserves these edges as well, regardless of the identity of A. $\square$

At this point, we have dealt with all the edge orbits that are required to restrict the motion of vertices in self-reversed and pair-reversed vertex orbits. Therefore, we turn to the question of the edge orbits that restrict the motion of the twisted vertex orbits. The minimal configuration when n is even is largely similar to that of the edge-minimal graph when n is a power of 2 (exhibited in [4]), and many of our results amount to showing that lemmas from that paper generalize. On the other hand, when n is odd and the length of a cycle in a twisted vertex orbit is only 2, a number of special cases arise. As a result, the special case $b = 0$ must often be handled separately.

Let $T = \{1, 2, \ldots 2^{b+3}\}$ denote the set of vertices of the twisted cycles of σ.

This following observation is made now in order to avoid giving repetitive arguments to the same effect.

Lemma 6.8. *Suppose Γ is a graph with with $G_n \leq \mathrm{Aut}(\Gamma)$ and let Γ' be the subgraph of Γ induced by the vertices in T. If $\phi \in \mathrm{Aut}(\Gamma') \setminus G_{2^b}$ and $\phi(v) \in \mathcal{O}_\sigma\{v\}$ for all $v \in T$, then $\mathrm{Aut}(\Gamma) > G_n$.*

Proof. Define

$$\phi' = \begin{cases} \phi(v) & v \in T \\ v & v \notin T \end{cases}$$

Clearly, ϕ' preserves adjacency between members of T and between members of $V(\Gamma) \setminus T$. Let $\{v, w\}$ be an edge (equivalently, non-edge) of Γ with $v \in T$ and $w \in V(\Gamma) \setminus T$. Then $|\mathcal{O}_\sigma\{v\}| = 2^{b+1}$ is coprime to $|\mathcal{O}_\sigma\{w\}|$, so no matter what $\phi(v)$ is, there is some k such that $\sigma^k\{v, w\} = \phi\{v, w\}$. $\square$

For even n, the only edge orbits of the graph given in Theorem 4.10 that we have not shown to be necessary are between vertices of the twisted vertex orbits. Helpfully, edge orbits between the twisted cycles only take on two sizes: $\mathcal{O}\{1, \tau^2(1)\}$ and $\mathcal{O}\{2^{b+2} + 1, \tau^2(2^{b+2} + 1)\}$ have 2^{b+1} edges, and all other orbits have 2^{b+2} edges. The extra automorphisms we will find here do not depend on edges between the twisted vertex orbits and the rest of the graph, so comparing the numbers of edges in different configurations of edge orbits will be simple. We will first show that two edge orbits between the two twisted vertex orbits are required, and then show that no single additional edge orbit will suffice, and neither will a pair of edge orbits where one or both contains 2^{b+1} edges. Hence, our selection of four edge orbits of size 2^{b+2} will be minimal.

Lemma 6.9. *A graph Γ with $\mathrm{Aut}(\Gamma) = G_n$ must have at least one edge-orbit of the form $\mathcal{O}\{1, 3 \cdot 2^{b+1} + 1 + a\}$ and one edge-orbit of the form $\mathcal{O}\{1, 2^{b+2} + 1 + a\}$.*

Proof. If either of the required orbits is not included, an extra automorphism outside of Dic_n permuting the vertices of the twisted cycles is exhibited in Lemma 20 from [4]. In each case, this automorphism is the restriction of τ^2 at some vertices and e at others. Since all vertices of non-twisted cycles are fixed by τ^2, this automorphism preserves those edges as well, as pointed out in Lemma 6.8. Note that the numbering may not be revision safe.

Suppose that no edge orbit of the form $\mathcal{O}\{1, 3\cdot 2^{b+1}+1+a\}$ is included in Γ. Then the extra automorphism is given by

$$\gamma(v) = \begin{cases} \tau^2(v) & v \in \mathcal{O}_\sigma\{1\} \bigcup \mathcal{O}_\sigma\{2^{b+2}+1\} \\ v & \text{otherwise} \end{cases}$$

As in previous cases, to see that γ is an automorphism, consult the following chart.

$\mathcal{O}\{v_1, v_2\}$	$e \in \mathcal{O}\{v_1, v_2\}$	$\gamma(e)$	$\pi \in \langle \sigma, \tau \rangle : \pi(e*) = \gamma(e)$
$\mathcal{O}\{1, 1+k\}$	$\{1+x, 1+k+x\}$ $\{2^{b+2}-x, 2^{b+2}-(k+x)\}$	$\tau^2\{1+x, 1+k+x\}$ $\{2^{b+2}-x, 2^{b+2}-(k+x)\}$	$\tau^2\{1+x, 1+k+x\}$ $\sigma^{-x}\tau\{1+x, 1+k+x\}$
$\mathcal{O}\{1, 1+k\}$	$\{2^{b+2}+1+x, 2^{b+2}+1+k+x\}$ $\{2^{b+3}-x, 2^{b+3}-(k+x)\}$	$\tau^2\{2^{b+2}+1+x, 2^{b+2}+1+k+x\}$ $\{2^{b+3}-x, 2^{b+3}-(k+x)\}$	$\tau^2\{2^{b+2}+1+x, 2^{b+2}+1+k+x\}$ $\sigma^{-x}\tau\{2^{b+2}+1+x, 2^{b+2}+1+k+x\}$
$\mathcal{O}\{1, 2^{b+1}+1+k\}$	$\{1+x, 2^{b+1}+1+k+x\}$ $\{2^{b+2}-x, 2^{b+1}-(k+x)\}$	$\{\tau^2(1+x), 2^{b+1}+1+k+x\}$ $\{2^{b+2}-x, \tau^2(2^{b+1}-(k+x))\}$	$\sigma^{k+1+x}\tau\{1, 2^{b+1}+k\}$ $\sigma^{-x}\tau\{1, 2^{b+1}+k\}$
$\mathcal{O}\{2^{b+2}+1, 3\cdot 2^{b+1}+1+k\}$	$\{2^{b+2}+1+x, 3\cdot 2^{b+1}+1+k+x\}$ $\{2^{b+3}-x, 3\cdot 2^{b+1}-(k+x)\}$	$\{\tau^2(2^{b+2}+1+x), 3\cdot 2^{b+1}+1+k+x\}$ $\{2^{b+3}-x, \tau^2(3\cdot 2^{b+1}-(k+x)\}$	$\sigma^{k+1+x}\tau\{2^{b+2}+1, 3\cdot 2^{b+1}+k\}$ $\sigma^{-x}\tau\{2^{b+2}+1, 3\cdot 2^{b+1}+k\}$
$\mathcal{O}\{1, 2^{b+2}+1+k\}$	$\{1+x, 2^{b+2}+1+k+x\}$ $\{2^{b+2}-x, 2^{b+3}-(k+x)\}$	$\tau^2\{1+x, 2^{b+2}+1+k+x\}$ $\{2^{b+2}-x, 2^{b+3}-(k+x)\}$	$\tau^2\sigma^x\{1, 2^{b+2}+1+k\}$ $\tau\sigma^x\{1, 2^{b+2}+1+k\}$

Orbits of the form $\mathcal{O}\{1, 1+k\}$ may be written as $\mathcal{O}\{1, \sigma^k(1)\}$; similarly, those of the form $\mathcal{O}\{2^{b+2}+1, 2^{b+2}+1+k\}$ may be written as $\mathcal{O}\{2^{b+2}+1, \sigma^k(2^{b+2}+1)\}$. Considering these forms for these edge orbits may make the effect of γ on their members more clear. $\square$

We have just shown that at least two edge orbits between twisted cycles are present; over the next several lemmas, we will see that adding any single additional orbit leaves some extra automorphism.

Lemma 6.10. *Let Γ be a graph with $G_n \leq \mathrm{Aut}(\Gamma)$ (on $\alpha(\mathrm{Dic}_n)$ vertices) such that the only orbits of edges between the first 2^{b+3} vertices are $\mathcal{O}\{1, 3\cdot 2^{b+1}+1+r\}$, $\mathcal{O}\{1, 2^{b+2}+1+s\}$, and any number of orbits of the form $\mathcal{O}\{1, 1+k\}$ or $\mathcal{O}\{2^{b+2}+1, 2^{b+2}+1+k\}$. Then $\mathrm{Aut}(\Gamma) > G_n$.*

Proof. Let τ^* be the permutation reversing the four 2^{b+1} cycles of T, as if they were self-reversed cycles. (That is, $\tau^* = (1, 2^{b+1})(2, 2^{b+1}-1)\cdots(2^{b+1}+1, 2^{b+2})(2^{b+1}+2, 2^{b+2}-1)\cdots(2^{b+2}+1, 3\cdot 2^{b+1})(2^{b+2}+2, 3\cdot 2^{b+1}-1)\cdots(3\cdot 2^{b+1}+1, 2^{b+3})(3\cdot 2^{b+1}+2, 2^{b+3}-1)\cdots$.)

We assert that an member of $\mathrm{Aut}(\Gamma) \setminus G_n$ is given by

$$\gamma(v) = \begin{cases} \tau^*(v) & 1 \leq v \leq 2^{b+1} \\ \sigma^{2(r+s)}\tau^*(v) & 2^{b+1}+1 \leq v \leq 2^{b+2} \\ \sigma^{2r}\tau^*(v) & 2^{b+2}+1 \leq v \leq 3\cdot 2^{b+1} \\ \sigma^{2s}\tau^*(v) & 3\cdot 2^{b+1}+1 \leq v \leq 2^{b+3} \end{cases}$$

By Lemma 6.8, we know that if γ preserves adjacency within the twisted vertex-orbits, then γ also preserves all other edges. That γ preserves adjacency between the vertices of the twisted-vertex orbits is shown below, as in previous proofs:

$\mathcal{O}\{v_1, v_2\}$	$e \in \mathcal{O}\{v_1, v_2\}$	$\gamma(e)$	$\pi \in \langle \sigma, \tau \rangle : \pi(e*) = \gamma(e)$
$\mathcal{O}\{1, 1+k\}$	$\{1+x, 1+k+x\}$ $\{2^{b+2}-x, 2^{b+2}-(k+x)\}$	$\{2^{b+1}-x, 2^{b+1}-(k+x)\}$ $\{2^{b+1}+x+2(r+s), 2^{b+1}+k+x+2(r+s)\}$	$\sigma^{-(k+2x)}\{1+x, 1+k+x\})$ $\sigma^{k+2x+2(r+s)}\{2^{b+2}-x, 2^{b+2}-(k+x)\}$
$\mathcal{O}\{2^{b+2}+1, 2^{b+2}+1+k\}$	$\{2^{b+2}+1+x, 2^{b+2}+1+k+x\}$ $\{2^{b+3}-x, 2^{b+3}-(k+x)\}$	$\{3\cdot 2^{b+1}-x+2r, 3\cdot 2^{b+1}-(k+x)+2r\}$ $\{3\cdot 2^{b+1}+1+x+2s, 3\cdot 2^{b+1}+1+k+x+2s\}$	$\sigma^{-(k+2x)+2r}\{2^{b+2}+1+x, 2^{b+2}+1+k+x\}$ $\sigma^{k+2x+2s}\{2^{b+3}-x, 2^{b+3}-(k+x)\}$
$\mathcal{O}\{1, 2^{b+2}+1+r\}$	$\{1+x, 2^{b+2}+1+r+x\}$ $\{2^{b+2}-x, 2^{b+3}-(r+x)\}$	$\{2^{b+1}-x, 3\cdot 2^{b+1}-x+r\}$ $\{2^{b+1}+1+x+2(r+s), 3\cdot 2^{b+1}+1+r+x+2s\}$	$\sigma^{-x}\{1, 2^{b+2}+1+r\}$ $\sigma^{x+2(r+s)}\{2^{b+2}, 2^{b+3}-r\}$
$\mathcal{O}\{1, 3\cdot 2^{b+1}+1+s\}$	$\{1+x, 3\cdot 2^{b+1}+1+s+x\}$ $\{2^{b+2}-x, 3\cdot 2^{b+1}-(s+x)\}$	$\{2^{b+1}-x, 2^{b+3}-x+s\}$ $\{2^{b+1}+1+x+2(r+s), 2^{b+2}+s+x+2r\}$	$\sigma^{-x}\{1, 3\cdot 2^{b+1}+1+s\}$ $\sigma^{x+2(r+s)}\{2^{b+2}, 3\cdot 2^{b+1}-s\}$

$\square$

Lemma 6.11. *Let Γ be a graph with $G_n \leq \mathrm{Aut}(\Gamma)$ (on $\alpha(\mathrm{Dic}_n)$ vertices) such that the only orbits between the first 2^{b+3} vertices are $\mathcal{O}\{1, 3\cdot 2^{b+1}+1+r\}$, $\mathcal{O}\{1, 2^{b+2}+1+s\}$, any orbits of the forms $\mathcal{O}\{1, 2^{b+1}+1+x\}$ or $\mathcal{O}\{2^{b+2}+1, 3\cdot 2^{b+1}+1+x\}$, and any combination of orbits in $\{\mathcal{O}\{1, \tau^2(1)\}, \mathcal{O}\{2^{b+2}+1, \tau^2(2^{b+2}+1)\}\}$. Then $\mathrm{Aut}(\Gamma) > G_n$.*

Proof. Define

$$\gamma = (1, \tau^2(1))(2^{b+1}+1+(r+s), \tau^2(2^{b+1}+1+(r+s)))(2^{b+2}+1+r, \tau^2(2^{b+2}+1+r))(3\cdot 2^{b+1}+1+s, \tau^2(3\cdot 2^{b+1}+1+s)).$$

We will show that γ is a member of $\mathrm{Aut}(\Gamma) \setminus G_n$. Again, we apply Lemma 6.8 to show that we need only worry about edge orbits between the first 2^{b+3} vertices.

Because γ is the restriction of τ^2, an involution, γ leaves the edge orbits $\mathcal{O}\{1, \tau^2(1)\}$ and $\mathcal{O}\{2^{b+2}+1, \tau^2(2^{b+2}+1)\}$ invariant, switching the endpoints of two edges in each and fixing all other vertices. Note that the orbit $\mathcal{O}\{1, 2^{b+1}+1+x\} = \mathcal{O}\{1, 2^{b+2}-(-x)\}$ may be written $\mathcal{O}\{1, \sigma^{-x}\tau(1)\}$. But since $(\sigma^k\tau)^2 = \tau^2$ for all k, we have $\sigma^{-1}\tau\{1, \sigma^{-x}\tau(1)\} = \{\tau^2(1), \sigma^{-x}\tau(1)\}$, and γ again leaves orbits of the form $\mathcal{O}\{1, 2^{b+1}+1+x\}$ invariant, switching the endpoints of two edges and leaving other vertices fixed. A symmetric argument holds for orbits of the form $\mathcal{O}\{2^{b+2}+1, 3\cdot 2^{b+1}+1+x\}$. For the remaining orbits, first note that since $\langle\sigma,\tau\rangle$ acts the same on each twisted vertex orbit, each vertex is incident to one edge in $\mathcal{O}\{1, 3\cdot 2^{b+1}+1+r\}$ and one edge in $\mathcal{O}\{1, 2^{b+2}+1+s\}$. Also note that $\tau^2 \in Z(\mathrm{Dic}_n)$. A simple calculation shows that the eight vertices permuted by γ form a component of the graph induced by the edges in these two orbits, and that γ is the restriction of τ^2 to the vertices of this component. □

Lemma 6.12. *Let Γ be a graph with $G_n \le \mathrm{Aut}(\Gamma)$ (on $\alpha(\mathrm{Dic}_n)$ vertices) such that the only orbits between the first 2^{b+3} vertices are three members of $\{\mathcal{O}\{1, 2^{b+2}+1+r : r \in \mathbb{Z}\}\} \bigcup \{\mathcal{O}\{1, 3\cdot 2^{b+1}+1+s : s \in \mathbb{Z}\}\}$ and any combination of orbits in $\{\mathcal{O}\{1, \tau^2(1)\}, \mathcal{O}\{2^{b+2}+1, \tau^2(2^{b+2}+1)\}\}$. Then $\mathrm{Aut}(\Gamma) > G_n$.*

Proof. If all the neighbors of 1 in the orbits of size 2^{b+1} are in the same cycle in σ, Lemma 6.9 applies. Otherwise, an extra automorphism of the restriction of Γ to T is given in Lemma 23 or Lemma 24 of [4], which remains in $\mathrm{Aut}(\Gamma)$ by Lemma 6.8.

□

Now, we turn to the structure of edge orbits when $b = 0$; that is, when n is odd. In this case, Lemma 6.9 still holds, so we already know two edge orbits that are required. The other two edge orbits whose inclusion are $\mathcal{O}\{5, x_1\}$ or $\mathcal{O}\{5, y_1\}$, whichever was included, and the $\mathcal{O}\{1, 2\}$ in the case that n was an odd prime power.

We begin by showing that the orbit of edges between $\mathcal{O}\{5\}$ and a pair-reversed vertex orbit is required.

Lemma 6.13. *Let Γ be a graph with $G_n \le \mathrm{Aut}(\Gamma)$ (on $\alpha(\mathrm{Dic}_n)$ vertices) such that there is no orbit of the form $\mathcal{O}\{5, x_i+1\}$ or $\mathcal{O}\{5, y_j+1\}$. Then $\mathrm{Aut}(\Gamma) > G_n$.*

Proof. Consider the permutation given by

$$\gamma(v) = \begin{cases} \tau^2(v) & v \in \{1,2\} \\ v & v \in \{3,4\} \\ \tau(v) & v \in \{5,6\} \\ \tau^{-1}(v) & v \in \{7,8\} \\ v & v \notin \mathcal{O}\{1\} \bigcup \mathcal{O}\{5\} \end{cases}$$

Calculation shows that all edge orbits between vertices of the twisted vertex orbits are preserved. Since the only members of twisted vertex orbits adjacent to members of non-twisted vertex orbits are in $\mathcal{O}\{1\}$ and $\gamma(v) \in \mathcal{O}_\sigma\{v\}$ for every $v \in \mathcal{O}\{1\}$, we may apply Lemma 6.8 to show that γ preserves the remaining edges. □

Since $\mathcal{O}\{1, 2\}$ is of the minimum possible size, for an additional orbit, we need only show that some additional orbit is necessary in the case where n is a prime power.

Lemma 6.14. *Suppose $n = p^m$ is an odd prime power. Let Γ be a graph with $G_n \le \mathrm{Aut}(\Gamma)$ (on $\alpha(\mathrm{Dic}_n)$ vertices) such that the only orbits of edges involving vertices of twisted vertex orbits are $\mathcal{O}\{1, 5+r\}$, $\mathcal{O}\{1, 7+s\}$, $\mathcal{O}\{1, 9\}$ or $\mathcal{O}\{1, 9+p\}$, and $\mathcal{O}\{5, 9\}$ or $\mathcal{O}\{1, 9+p\}$. Then $\mathrm{Aut}(\Gamma) > G_n$.*

Proof. Notice that under the above hypotheses, the relationship of $\mathcal{O}\{1\}$ to $\mathcal{O}\{5\}$ is symmetric. In fact, we can give an extra automorphism which switches $\mathcal{O}\{1\}$ and $\mathcal{O}\{5\}$. If 1 and 5 have a common neighbor in $\mathcal{O}\{9\}$, then let $\gamma = (1,5)(2,6)(3,8)(4,7)$. Clearly, γ is an order two permutation different from $\tau^2 = \sigma$. Since $\mathcal{O}_\sigma\{1\}$ is switched with $\mathcal{O}_\sigma\{5\}$, edges between the twisted cycles and $\mathcal{O}\{9\}$ are preserved. Calculation shows that, regardless of the values of r and s, the edges between the twisted vertex orbits are also preserved.

If instead 1 and 7 share a common neighbor in $\mathcal{O}\{9\}$, we let $\gamma = (1,7)(2,8)(3,5)(4,6)$; the proof is symmetric. □

Theorem 6.15. *The graph Γ with $\mathrm{Aut}(\Gamma) \cong \mathrm{Dic}_n$ exhibited in Theorem 4.10 has $e(\mathrm{Dic}_n, \alpha(\mathrm{Dic}_n))$ edges.*

Proof. This proof is largely a matter of combining the lemmas we have previously shown in this section.

First, note that the orbits of edges between vertices of $\mathcal{O}\{y_j + 2q_j + 1\}$ when $m_j > 1$ must be included by Lemma 6.1. The orbits between $\mathcal{O}\{y_j+1\}$ and $\mathcal{O}\{y_j + 2q_j + 1\}$ must be included by Lemma 6.2. The orbits between vertices of $\mathcal{O}\{y_j+1\}$ are required by Lemmas 6.6, 6.7, and 6.3. By Lemma 6.6, all orbits of edges between $\mathcal{O}_\sigma\{x_i+1\}$ and $\mathcal{O}_\sigma\{x_i + p_i + 1\}$ are necessary. A simple calculation shows that the orbits of edges included between the smallest of

$\mathcal{O}\{1\}$, $\mathcal{O}\{x_1+1\}$, or $\mathcal{O}\{y_1+1\}$, whichever is the smallest, and all other orbits of the forms $\mathcal{O}\{1\}$, $\mathcal{O}\{x_i+1\}$ and $\mathcal{O}\{y_j+1\}$ are the edge-minimal fulfillment of the requirements of Lemmas 6.4 and 6.5.

If n is even, the orbits of edges between $\mathcal{O}\{1\}$ and $\mathcal{O}\{2^{b+2}+1\}$ included in Γ are required by Lemma 6.9. By Lemmas 6.11, 6.10, and 6.12, no third edge orbit of edges amongst the vertices of $\mathcal{O}\{1\}$ and $\mathcal{O}\{2^{b+2}+1\}$ will suffice, nor will any trio of edge orbits amongst those vertices combined with an edge orbit of size 2^{b+1}. Hence Γ, with 4 edge orbits of size 2^{b+2} amongst the vertices of $\mathcal{O}\{1\}$ and $\mathcal{O}\{2^{b+2}+1\}$, has the minimum possible number of edges. If n is odd, a similar argument, aided by Lemmas 6.13 and 6.14 suffices. $\square$

Bibliography

[1] Laszlo Babai. On the minimum order of graphs with a given group. *Cannad. Math Bull*, 17(4):467–470, 1974.

[2] R. Frucht. Graphs of degree three with a given abstract group. *Canadian Journal of Mathematics*, 1:365–378, 1949.

[3] R. Frucht. Herstellung von graphen mit vorgegebener abstrakter gruppe. *Compositio math*, 6:239–250, 1939.

[4] Christina Graves, Stephen J. Graves, and L.-K. Lauderdale. Smallest graphs with given generalized quaternion automorphism group. Under review, 2014.

[5] Gert Sabidussi. On the minimum order of graphs with given automorphism group. *Monatsh. Math.*, 63:124–127, 1959.

[6] William C. Arlinghaus. The classification of minimal graphs with given abelian automorphism group. *Mem. Amer. Math. Soc.*, 57(330):viii+86, 1985.