Isomorphism of Graphs Derived from Voltage Graphs

Deniz Sarioz
sarioz@acm.org

January 7, 2003

Contents

1 Introduction

In this thesis, familiarity with some abstract algebra (terms such as group, ring, field) and elementary graph theory (terms such as graph, digraph, graph isomorphism) is assumed. Definition of *voltage graph* and related terminology will be provided. This section accomplishes that in addition to introducing our problem.

1.1 Definition of *Voltage Graphs*: cyclic voltages and regular voltages

The motivation behind voltage graphs is the ability to encode or capture symmetries and other features of a large graph by a much smaller one. Some graphs consisting of a million vertices and three million edges can be compactly represented by a voltage graph consisting of a single node and 3 edges.

Even though cyclic voltage graphs are the simplest to grasp, only proper definitions pertaining to regular-voltage graphs will be provided, since cyclic voltage graphs are special cases. These definitions come from my main reference text [1] and I have modified some to add clarity. Voltage graphs were first presented in [2] and I have also seen [3] referred to about voltage graphs. Permutation voltage graphs are out of the scope of this paper.

`DEFINITION`: Let $G = (V_G, E_G)$ be a digraph (V_G refers to the vertex set and E_G refers to the arc set), and let B be a group. A *regular-voltage assignment* on G in the group B is a function α that assigns to every arc $e \in E_G$ an element $\alpha(e) \in B$. The element $\alpha(e) \in B$ is called the *voltage* on e.

`DEFINITION`: A *regular-voltage graph* is a pair $\langle G, \alpha : E_G \to B \rangle$ such that $G = (V, E)$ is a digraph, B is a group and α is a B-voltage assignment on G.

`DEFINITION`: The *voltage group* of a voltage graph $\langle G, \alpha : E_G \to B \rangle$ is the group B in which the voltages are assigned.

`DEFINITION`: Let $\langle G, \alpha : E_G \to B \rangle$ be a regular-voltage graph. The *derived vertex-set* is $V^\alpha = \{v_b \mid v \in V_G \text{ and } b \in B\}$

`DEFINITION`: Let $\langle G, \alpha : E_G \to B \rangle$ be a regular-voltage graph. The *derived arc-set* is $E^\alpha = \{e_b \mid e \in E_G \text{ and } b \in B\}$

`DEFINITION`: Let $\langle G, \alpha : E_G \to B \rangle$ be a regular-voltage graph. The *derived digraph* is the graph $G^\alpha = (V^\alpha, E^\alpha)$. The *derived graph* is the underlying graph of the derived digraph.

After these tersely presented definitions, we can now give some examples to make them more intuitive. A cyclic voltage graph is one in which the voltage group is cyclic. Without loss of generality, our voltage group is (isomorphic to) Z_n where $n \in Z$. Informally, the derivation process is as simple as follows: We make n copies of each vertex in the voltage graph, with subscripts corresponding to the n elements of the voltage group. For each edge from a vertex u to a vertex v with a voltage assignment of m (and all edges in our voltage graph have a voltage, a tail and a head), we have an edge in the derived graph pointing from u_i to v_{i+m} for each i.

Example: C_2 with voltages (0,1) in Z_3 specifies C_6.

Example: The dumbbell graph with voltages in Z_5 (1 and 2 on the self-edges, 0 on the edge joining the two vertices) specifies Petersen's graph.

Notice that every graph can trivially be specified by a voltage graph (isomorphic to it) with voltages in the zero group ($\{0\}$ with $0+0=0$). But we cannot find a non-trivial way to specify every graph with a voltage graph.

A note on terminology: the statement *"A voltage graph G specifies the graph H"* is equivalent to the statement *"The graph H is derived from a voltage graph G"* where G is understood to be a digraph and that we are talking about a certain voltage assignment on G in a certain group. When we say *"derived graph H"* in isolation, it is understood that there is some voltage graph G that specifies H.

We leave the terms *'generate'* and *'generator'* to be used exclusively in the traditional group-theoretic sense. We will refrain from using the term *'obtained'* (synonym of *'derived'*) to alleviate further proliferation of redundant terms.

Despite the fact that derived graphs of voltage graphs are formally defined to be digraphs, for the scope of this paper, we consider solely the underlying graphs of the specified digraphs. e.g., graph H as in the above statements is to be considered an undirected graph.

1.2 Problem Statement: What we are aiming at

According to the definition above, every voltage graph specifies exactly one graph. We can ask the question, is this relation one-to-one? We can immediately see the answer to be a firm "no". If we consider a graph with a single self-edge (B_1) with voltages in Z_5, we can see that the graph with 5 copies of B_1 is derived if an only if the voltage on the single edge is 0, and that we get C_5 for any other voltage assignment.

Hence, we see that the graphs derived from the 5 possible voltage graphs—provided the digraph B_1 and the group Z_5—correspond to only 2 isomorphism classes!

It is clear that it would be computationally very expensive to calculate by brute force how many of these graphs are isomorphic. We know a priori that all the graphs derived from the same voltage graph $G(V,E)$ and group B will have exactly the same number of vertices ($|V||B|$) and the same number of edges ($|E||B|$). So, for all we know, every pair could be isomorphic, or non-isomorphic. (We use '$|.|$' to just mean the number of elements in these finite sets, treating all the arguments simply as sets.)

If we took the pure brute-force approach, we would need (in the best case and with a lot of luck, such as the example given above in which we have only two classes) $N-1$ comparisons to establish isomorphism classes, where N is the number of derived graphs. In the worst case scenario (think about every pair of graphs having to "shake hands"), this is $N(N-1)/2$ comparisons! While this polynomial does not seem that intimidating, we are talking about graph isomorphism testing (given two graphs with vertex and edge sets of the same size, as in all of our pairs), for which no polynomial-time algorithm is known yet. The number of possible bijections between two such graphs is factorial in the number of vertices [1, pg.66]! We can definitely try to do better than brute force approaches, since we see symmetry in all derived graphs with voltages in a non-trivial group.

Without further ado, we pose the question in its most general form:

"Given any digraph $G(V,E)$ and any finite group B, determine, for the graphs derived from all possible voltage assignments, the number of isomorphism classes and the number of voltage graphs corresponding to each class, making as few graph isomorphism tests as possible."

In this paper, we make strong statements for the graphs derived from B_1 (the *bouquet* graph consisting of one vertex with a single self-edge) with voltage assignments in Z_n (all cyclic groups) and B_2 (the *bouquet* graph consisting of one vertex with two self-edges) with voltage assignments in Z_p (all cyclic groups of prime order).

Any element a in the group B can be assigned to any edge e in E. Note that some voltage assignments can be thought of as equivalent due to symmetry in the digraph, however, for our discussion, we will consider every edge in E as labelled in advance (from e_1 to $e_{|E|}$). Hence, we will represent every voltage assignment as a unique $|E|$-tuple. We can immediately see that there are $|E||B|$ such tuples to consider, which is our naive upperbound on the number of isomorphism classes.

Everything beyond this point is my original work. The conclusion makes what I mean by this clearer.

2 Cyclic Voltages

2.1 Graphs encoded by B_1 with voltage in Z_p

We first consider the case where our voltage graph has a single vertex with a single self-edge with voltage in Z_p where p is prime. We can drop the mainscripts of the derived graph since they would all be the same, and casually refer to the vertices in our derived graph with the elements in Z_p. We will use negative numbers wherever convenient, understood to be equivalent to an official element of Z_p. Let us use i to denote the voltage.

If i is 0, this means that there is an edge leading from every vertex j to vertex $j+0$, which is another way of saying j. Hence, what we get in this case is p copies of B_1.

If i is 1, then there is an edge leading from every vertex j to vertex $j+1$. Since we are in a cyclic group, it is clear that we have a p-cycle in this case.

For every other value of i, we have an edge leading from vertex j to vertex $j+1$. We again have a p-cycle, but we will need to invoke some algebra to prove this.

Every cyclic group can also be thought of as a commutative ring with unity. This means that alongside addition, we can also use the familiar multiplication modulo n in Z_n. *Commutative ring* means that for all pairs of elements r and s in the ring, $rs = sr$. *With unity* implies that we have an element 1 (called the unity) in the ring which is the multiplicative identity. Now, notice that Z_p can further be thought of as a field. This means that every non-zero element will have a unique multiplicative inverse. For instance, in Z_7, the multiplicative inverse of 3 (denoted with 3^{-1}) is 5. This not only means $(3)(5) = (5)(3) = 15 = 1 \pmod 7$, it also means that there is no other element which we can multiply by 3 to obtain 1. Moreover, in a field, we can be sure that the product of all pairs of non-zero elements is non-zero.

Just like in the set of integers, multiplication in Z_n is the same as successive addition, and we make heavy use of this fact. We know that adding together m copies of i will yield the same result as $(m \ mod \ n)i$.

The easy way to see this is that we start at vertex 0 and follow the edges in the $+i$ direction. (Whenever we have a self-edge with voltage i on some vertex, we can consider the $-i$ direction as well as the $+i$ direction for that vertex.) So, we visit the vertices $\langle 0, i, 2i, \ldots \rangle$. We will reach 0 again after exactly p edges and after having visited every vertex at least once. For every non-zero voltage i in Z_p, and the mth vertex on this walk (0 counting as the 0th), the vertex is mi. Due to the fact that we are working in a field, mi for each m such that $1 <= m <= p-1$ will attain distinct values in Z_p. This means that our walk $\langle 0, i, 2i, \ldots \rangle$ with p edges is isomorphic to C_p.

The tricky way to see this hopefully gives us more insight. The strategy is to show that what we got is isomorphic to the case where the voltage was 1 and we easily showed to be isomorphic to C_n. Since isomorphism is an equivalence relation, it will suffice to show that the graph derived from every nonzero voltage i is isomorphic to the graph derived from voltage 1. Let us consider the graph derived from B_1 with any non-zero voltage i on the single edge. Now, what happens if we multiply every vertex label with i^{-1}? Applied to all vertices in the walk $\langle 0, i, 2i, \ldots \rangle$, we get $\langle 0, 1, 2, \ldots, -2, -1, 0 \rangle$. We can think of this as a permutation of the original vertex labels. In the following chapters too, we will use this idea of permuting vertex labels to find familiar graphs.

2.2 Graphs encoded by B_1 with voltage in Z_n

Now, what if we have any cyclic group as our voltage group? For a voltage of 0 and a voltage of 1, we get exactly what we got in the previous case. Consider Z_6 as the voltage group, and let the voltage be 4. We can see that we will have two components, one being $\langle 0, 4, 2 \rangle$ and the other being $\langle 1, 5, 3 \rangle$. (We can recognize the vertex labels in these 3-cycles as being members of additive cosets of Z_6). Notice that the lengths of these cycles is the smallest non-zero element of Z_6 which yields 0 when multiplied by 4. In Z, we can represent that as $LCM(6,4)/4$. Then, in general, when we have a voltage i in Z_n, the derived graph will have cycles of length $LCM(n,i)/i$ which is $n/GCD(n,i)$. These cycles are necessarily disjoint, so since we have n vertices, we know that we will have $ni/LCM(n,i)$ of these cycles, which is equal to $GCD(n,i)$. The number of different isomorphism classes (for all voltages) is the same as the unique number of factors of n because for every element i of Z_n, $GCD(n,i)$ will be the same as $GCD(n,j)$ for some factor of n.

2.3 Graphs encoded by B_2 with voltages in Z_p

2.3.1 Warming up: the complete explication of the cases Z_3, Z_5, Z_7

In Z_3, we have a class for the pair of voltages on B_2 (0,0) which is a graph with 3 vertices and two self-edges at every vertex. In other words, 3 x B_2. For all the other pairs (0,i) or (i,0), we have a 3-cycle and on top of that, we have one self-edge at every vertex. For (1,1) and (2,2), we clearly get doubly-linked 3-cycles. For (1,2) and (2,1), we can notice that these graphs are simply (1,1) with one of the voltages reversed, since $1 = -2$ in Z_3, so they are isomorphic to the one derived from (1,1). In total, this is one class derived by 1 pair, one class derived by 4 pairs, and another class derived by 4 pairs.

In Z_5, we again have a class with just the pair (0,0) that is 5 x B_2. We have another class for the other pairs (0,i) and (i,0). (1,1) and (-1,-1) (or (4,4)) will get us the 5-doubly-linked-cycle. In fact, whenever we have (i,i) where i is non-zero, we get the 5-doubly-linked-cycle. In addition, whenever we have (i,-i), we get the same 5-doubly-linked cycle due to voltage reversal not affecting derived graph topology. All other pairs derive a graph which can be illustrated with a star-polygon, in this case with 5 vertices, a pentagram inside a pentagon. We can see that (1,2) would be the most intuitive recipe for that structure. And just by taking additive inverses and switching the voltages, we also have the pairs (4,2), (4,3), (1,3), (2,1), (2,4), (3,4), (3,1) in the same class. And we have mentioned all the $5^2 = 25$ pairs in some way or other. If we are to count, here are the number of pairs in each equivalence class (by what graph is derived from them): 1 (5 x B_2), 8 (self-edges plus a cycle), 8 (doubly-linked cycle), 8 (Star-Polygon). (We can add these to make sure we get p^2 as a check.)

By now, we realize that these types are ubiquitous: that for any voltage group with a greater prime order, we will have 1 pair derive p x B_2, we will have 2p-1 pairs derive the p self-edges and a cycle, 2p-2 pairs derive the p-doubly-linked-cycle, and the rest will be Star-Polygon graphs.

In Z_7, the interesting case is (1,2). From now on, we adopt the convention of referring to the graphs in B_2 with voltages (i,j) with SP(i,j) (SP for Star-Polygon). We will do this regardless of whether the derived graph is topologically a Star-Polygon or not. In the same class with (1,2), we have (6,2), (6,5), (1,5), (2,1), (2,6), (5,6), and (5,1) just by taking additive inverses. And in the same class with (1,3), we have (6,3), (6,4), (1,4), (3,1), (3,6), (4,6), (4,1). Doing the vertex-permutation that we mentioned in the previous section (multiplying each vertex of the graph derived from (1,2) by 2), we can see that SP(1,2) is isomorphic to SP(2,4). Immediately follows that (5,4), (5,3), (2,3), (4,2), (4,5), (3,5), (3,2) are in the same class. And from (1,3), we can likewise find all the rest. But how do we know that SP(1,2) is not isomorphic to SP(1,3)? We don't, because they actually are isomorphic! You can convince yourself of this by realizing that both of these Star-Polygon graphs consist of 7 vertices and 14 edges. But the complete graph K_7 has $\binom{7}{2}$=21 edges. Hence, the edge complement of both of these Star-Polygon graphs consist of a $21 - 14 =$ 7-cycle (C_7), and we know that graphs whose edge complements are isomorphic are isomorphic. Then the sizes of classes for Z_7 are: 1 (7 x B_2), 12 (self-edges plus a 7-cycle), 12 (7-doubly-linked-cycle), 24 (Heptagon-Heptagram). Note that the heptagon-heptagram for this can look different depending on which fiber we have as the heptagon and which one we take as the heptagram. By fiber over an arc in the voltage graph, we essentially mean the edges in the derived graph that correspond to that voltage.

We should realize that there is a less klug way to see why SP(1,2) is isomorphic to SP(1,3) for Z_7. We can permute the vertex labels in SP(1,2) by multiplying each with 4. This shows that SP(1,2) is isomorphic to SP(4,8) or if we use symbols in the valid range, SP(4,1). Notice that by switching the voltages and reversing one (that is, 4), we get SP(1,3). In general, we realize that we can represent every pair (i,j) with a pair (1,ji^{-1}) and furthermore, every pair (1,i) with (1,i^{-1}). This will be more formally explained in the next section.

We could explicate the case of Z_{11}, however, we have the basic idea by now. All we can do is to note that SP(1,2) is NOT isomorphic to SP(1,3) in this case. We can see this by noting that the former has 3-cycles whereas the latter does not. Notice that for the general case of Z_n where $n > 2$ that SP(1,2) ought to contain a 3-cycle. We make use of this in the proof of our main theorem which we have not stated yet.

We can already see that for the general case of Z_p, Each of the p^2 derived graphs belongs to one of at most $p/2$ classes. We can see this right away due to the fact that they every SP(i,j) is isomorphic to some SP(1,k) for some k. Furthermore, we have the additive and the multiplicative inverses of k, as well as $-k^{-1}$. This means that actually, the number of classes is on the order of $p/4$.

Something that we should pay attention to is that there are some cyclic groups with some elements whose additive inverse is the same as their multiplicative inverse. Due to this, we can (slightly unexpectedly) have a few more classes, since we have only p-1 as opposed to 2p-2 elements in classes for the case of (1,i) where i's additive inverse is the same as its multiplicative inverse.

The following section will address the general case from scratch, and having looked at specific cases will have strengthened our intuition and should make it more palatable.

2.3.2 The General Case of Z_p

Let us consider the graph derived from B_2 with voltage assignments (r, s) in the cyclic group Z_p on the two self-edges, with prime $p > 2$, and $r \neq 0 \neq s$. In the derived graph H, we consider the vertex labels to be the replaced with their subscripts. Notice that the graph derived from the voltage assignments (r, -s) is isomorphic to this (note that we are looking only at the underlying graph of the derived digraph). The reason is simple: in the derived digraph for the pair (r,s), there is an edge pointing from every vertex i to vertex $i + s$. After either voltage is reversed, there is an edge pointing from vertex $i + s$ to vertex i.

Now, consider the voltage graph with the same topology and voltage group, this time, with voltages (kr, ks) with $k \neq 0$. Since p is prime, for every non-zero value of k, (kr, ks) is a unique ordered pair. Let us denote the derived graph from this voltage graph with Hk. Notice that there is an isomorphism ϕ that maps the vertices n of H to vertices of Hk such that $\phi(n) = kn$. This enables us, by using $k = r^{-1}$ (since we can talk about multiplicative inverses in a field), to identify every graph derived H derived from a pair of voltage assignments (r, s) neither of which are zero with the pair $(r^{-1}r, r^{-1}s) = (1, r^{-1}s)$. For the sake of simplicity, let us define $t = r^{-1}s$ so that we can talk about the pair (1, t) as being representative of any pair (r, s). We know that the pair (1, -t) is equivalent. In addition, we can do to (1, t) what we did to (r, s), and multiply each pair by a non-zero element of Z_p (again) to come up with an equivalent pair (in the sense that a graph isomorphic to it can be derived from it). More specifically, if we choose $k = t^{-1}$, now we have the pair $(t^{-1}1, t^{-1}t) = (t^{-1}, 1)$. Notice that we can always swap the pairs to come up with a bouquet with equivalent voltage assignments. We have (somewhat densely) demonstrated that for any $t \in Z_p^*$, there is an equivalence relation among the pairs:

$$(1, t) \sim (1, -t) \sim (1, t^{-1}) \sim (1, -t^{-1})$$

in the sense that if they are interpreted as voltage assignments on the bouquet B_2, the graphs derived from these four different voltage graphs will be isomorphic. (Notice that the additive inverse of the multiplicative inverse of an element is the same as the multiplicative inverse of its additive inverse.) If we consider that for any p, most elements in Z_p do not have their additive inverse as their multiplicative inverse, we have at most on the order of p/4 equivalence classes of pairs based on the isomorphism of their derived voltage graphs, out of a possible p^2. If these are indeed the entire equivalence classes (for $r, s \neq 0$), we have exactly $2p - 2$ unique equivalent pairs in the case that we have two representative ordered (1, t) pairs (the additive inverse different from the multiplicative inverse), and we have exactly $p - 1$ unique equivalent pairs otherwise (we have only two representative ordered (1, t) pairs). It is easy to see that (0, 0) is unique in that the graph derived from it is simply p copies of B_2, and from any other pair (r, s) for which $r = 0$ or $s = 0$, we obtain a polygon with self-edges at every edge (There are exactly $2p - 2$ pairs in this equivalence class). The pairs (r, r) and (r, -r) where $r \neq 0$ can be seen to form another equivalence class where the derived graph has a doubly-linked-polygon structure (with exactly $2p - 2$ in this class). We can speak so surely because for whenever $r, s \neq 0$ and $r \neq s$, we are bound to get a derived structure that looks like a Star-Polygon graph. It is those that are really the interesting case.

Now, what we did is to establish an upperbound on these classes. We still do not really know that the Star-Polygon deriving classes are disjoint. To illustrate what we mean by this, in Z_{37},

$$(1,4) \sim (1,33) \sim (1,28) \sim (1,9)$$

because $33 = -4$, $28 = 4^{-1}$ and $9 = -4^{-1}$. And similarly,

$$(1,7) \sim (1,30) \sim (1,16) \sim (1,21)$$

We are talking about star-polygon graphs with 37 vertices each. We know that the four pairs each encode the same graph, but can we say that we cannot reduce the number of classes further? What makes us so sure that the pair (1, 4) does not encode a graph isomorphic to the graph that the pair (1, 7) encodes for? This is the question that will be addressed in this section–and we are seeking a result purely for pairs that encode for star-polygon graphs.

2.3.3 Main Theorem establishing equivalence classes

I am using the notation SP(1,i) to represent the Star-Polygon graph derived from B_2 with voltage assignments 1 and i in the cyclic group Z_p on the two self-edges. The following restrictions on the values of i and p make sure that such a derived graph indeed has a Star-Polygon structure.

Let us assume that there is a graph isomorphism ϕ from SP(1,a) to SP(1,b) both with voltages in Z_p for the same prime p, such that $p > 11$; $a, b \in Z_p \setminus \{-1, 0, 1\}$. We have already looked at the cases for which $p <= 11$ and seen that there are no exotic isomorphisms. We need p to be so large so that it is meaningful to talk about SP(1,b) with some of the following restrictions. We have already discussed that the graphs derived from SP(1,a) with $a \in \{-1, 0, 1\}$ are not Star-Polygon graphs at all, but that they are either (a=0) polygons with self-edges or (otherwise) have a doubly-linked-polygon structure.

This would be a good point to stop and state the theorem: *given such a construction,* $a \in \{b, -b, b^{-1}, -b^{-1}\}$. Notice that it is meaningful to talk about the multiplicative inverse of b only because we are in a field, thanks to p for being prime, and because $\pm b$ is non-zero, hence a unit. These 4 values do not have to be distinct. For instance, the following argument still holds for the case where $b = 5$ and $p = 13$. Since 5 and 8 are additive as well as multiplicative inverses in Z_{13}, a can assume only those two values.

To prove this theorem, it would suffice to look only at:

$$a, b \in Z_p \setminus \{-1, 0, 1, -2, 2, 2^{-1}, -2^{-1}\}.$$

We can think we are losing generality by demanding a and b to be so restricted, but it is easily verified that SP(1,b) with $b \in \{-2, 2, 2^{-1}, -2^{-1}\}$ has 3-cycles (and we already know that the 4 graphs derived from these values are isomorphic), and no SP(1,a) with $a \in Z_p \setminus \{-2, 2, 2^{-1}, -2^{-1}, -1, 0, 1\}$ has a 3-cycle—which implies that such an SP(1,a) cannot be isomorphic to such an SP(1,b) anyway.

Realizing that this might not be so obvious, here is why 3-cycles can exist in Star-Polygon graphs SP(1,i) if and only if i is related to 2. A 3-cycle can exist in a Star-Polygon graph only if three consecutive directed voltages add up to 0. (By "consecutive directed voltages", we mean that if we are travelling an edge opposite to the edge direction that comes with the SP graph, we subtract that voltage.) This would mean that for SP(1,i), 3-cycles can exist only if $x+y+z = 0$ for $x, y, z \in \{-i, -1, 1, i\}$. This implies at least one of the following, based on all possible values for $x + y + z$: $\pm 3 = 0$, $\pm 3i = 0$, $\pm 1 = 0$, $\pm i = 0$, $\pm i \pm 2 = 0$, and $\pm 2i \pm 1 = 0$. The first 4 of these equality bundles are absurd in a cyclic group of prime order greater than 3, and the latter 2 equality bundles (or densely written 8 equalities, whichever way you look at it) are equivalent to one of the 4 equalities packed into: $i = \pm 2^{\pm 1}$. It is much easier to show the converse, i.e., that 3-cycles exist for $i = \pm 2^{\pm 1}$. We know that in SP(1,2), vertex 0 is connected to vertex 1, 1 is connected to 2, and 2 is connected to 0—so we have a 3-cycle. We also know from the previous section that SP(1,2) is isomorphic to all SP(1,i) with $i = \pm 2^{\pm 1}$, therefore we have our if and only if.

We can also make a mental note that not having a 3-cycle in a Star-Polygon graph has the consequence that the shortest cycle to be found is of length 4. Even though we do not make use of this fact directly, I consider it useful to note this for providing some cohesion into our thinking about these graphs.

Furthermore, we also make the seemingly unreasonable assumption that $b \in Z_p \setminus \{-3, 3, 3^{-1}, -3^{-1}\}$. We will be able to justify this assumption only after we have proven our theorem for this limited case, and show that our statement is valid for $b \in \{-3, 3, 3^{-1}, -3^{-1}\}$ as well.

We will continue the convention to drop the mainscripts (which would be the same as whatever the single vertex was labelled as in our respective B2 graphs were begin with) and using the subscripts to denote the vertices of the derived graphs. This has the meaning that the vertex labels are not arbitrary, but that we keep the subscripts produced by the original derivations of our respective Star-Polygon graphs as the vertex labels. This will enable us to talk about subtracting pairs of connected vertices to obtain integers in Z_p (which represent the "directed voltage" on that edge), and adding integers to vertices to obtain other vertices.

For SP(1,a), then, we can say that vertex n is connected to vertices n-1, n+1, n-a, and n+a. These 4 vertices are distinct because $p > 2$ imply that $-1 \neq 1$ and $a \neq -a$ (because $a = -a$ would imply $2a = p$ while all primes greater than 2 are odd), and we started out with further conditions such that $a \neq \pm 1$ which immediately implies $-a \neq \pm 1$. (And only these 6 conditions need to be checked, since $\binom{4}{2} = 6$.) The same argument holds for any vertex m of SP(1,b) and its four neighbors m-1, m+1, m-b and m+b being distinct. Now, we know that the major axes of SP(1,a) are hamiltonian cycles. Namely,

$$A1 = \langle 0, 1, 2, \ldots, -2, -1, 0 \rangle$$

is one of these axes and

$$A2 = \langle 0, a, 2a, \ldots, -2a, -a, 0 \rangle$$

is the other. I will refer to these cycles as A1 and A2, and consider them to be directed in the left-to-right order noted here. I am calling these axes "major" because even though other hamiltonian cycles could exist in our derived graph SP(1,a), the arcs in A1 and A2 correspond respectively to the mutually exclusive arc subsets derived from the voltages on the two edges of our voltage group, and their union results in the entire derived arc-set. It is easy to see that both of these cycles contain p distinct points: In the case of A1, we clearly have all the vertices since we have a vertex corresponding to every element of our voltage group. In the second, it might not be as clear, but we have all multiples of a, and since our voltage group happens to be a field, we can simply divide through by a to see that we have p distinct vertices in that construct.

Now, a graph isomorphism is a bijection between vertices and edges of two graphs. Of course, we can suspect that the vertex labels can be different. Let us represent the image of every vertex n of SP(1,a) under the isomorphism ϕ with $\phi(n)$. Note that it would be incorrect to say $\phi(n) = n$, even though it could well be the case for some value of n. However, making a statement like the latter will not even be necessary. Let us fix n and look at the respective neighborhoods of vertex n in SP(1,a) and vertex $\phi(n)$ in SP(1,b). The four neighbors of $\phi(n)$ are $\phi(n-1)$, $\phi(n+1)$, $\phi(n-a)$, and $\phi(n+a)$; which are necessarily distinct due to one-to-oneness.

Observe that vertex $n + a + 1$ in SP(1,a) has edges connecting it to the distinct vertices $n + 1$ and $n + a$. Hence, we have a structure isomorphic to C_4: $\langle n, n+1, n+a+1, n+a, n\rangle$. How do we know that these 4 vertices are indeed distinct? We have checked all conditions except one, can it be that $n = n + a + 1$? This would mean $0 = a + 1$, hence, $a = -1$, violating one of our initial conditions. After the previous section, this should nothing new, but it is a building block of this proof. The isomorphism takes the aforementioned 4-cycle to the 4-cycle in SP(1,b): $\langle \phi(n), \phi(n+1), \phi(n+a+1), \phi(n+a), \phi(n)\rangle$.

Let us make the following definitions:

$$\begin{aligned} w &= \phi(n+1) - \phi(n), \\ x &= \phi(n+a) - \phi(n), \\ y &= \phi(n+a+1) - \phi(n+1), \quad \textit{and} \\ z &= \phi(n+a+1) - \phi(n+a). \end{aligned}$$

Where by '$-$' we mean subtraction in the group Z_p. $\{-1, 1, -b, b\}$.

It is easy to verify using the definitions and simple algebra that $w+y = x+z$. Since all of these vertices are distinct, we have the following inequalities (that can be easily inferred from the definitions):

$$\begin{array}{lll} w \neq x & \text{because} & \phi(n+1) \neq \phi(n+a), \\ y \neq z & \text{because} & \phi(n+1) \neq \phi(n+a), \\ z \neq -x & \text{because} & \phi(n) \neq \phi(n+a+1), \quad \textit{and} \\ y \neq -w & \text{because} & \phi(n) \neq \phi(n+a+1). \end{array}$$

We can notice that the condition $y \neq -w$ requires all of our legitimate choices for $w + y$ and consequentially $x + z$ to be non-zero. The following table contains ALL possible values for x, z, w, and y:

x	z	w	y
1	b	b	1
1	-b	-b	1
-1	b	b	-1
-1	-b	-b	-1
b	1	1	b
b	-1	-1	b
-b	1	1	-b
-b	-1	-1	-b

Now we enter a long discussion to show that this table is complete. The idea here is that the values of x and z determine the values of w and y uniquely. We already know $x \neq -z$ and it is explained below why x and z cannot assume the same value. Beyond those two constraints, the table above certainly contains all possible combinations of x and z. The determination of the values of w and y is in the light of the fact that every sum x+z with $x, z \in \{-1, 1, -b, b\}$ and $x \neq -z$, i.e., $x + z \neq 0$ can be uniquely factored into pairs of summands over the set $\{-1, 1, -b, b\}$.

We have no control over the ordering of those factors, since we are using the addition operation of a commutative group. However, the requirement that $w \neq x$ forces the order of the factorization (into w and y) to be unique too. To convince ourselves, we can think of the following addition table (basically a part of the definition of binary addition in the group Z_p, which is commutative, which is why we can leave half of it blank):

+	-1	1	b	-b
-1	-2			
1	0	2		
b	b-1	b+1	2b	
-b	-b-1	-b+1	0	-2b

In other words (and ignoring the zeros): 2, -2, 2b, -2b, b+1, -b+1, b-1, and -b-1 have to be distinct, and they will be so based on our conditions. To really convince ourselves of this, we have $\binom{8}{2} = 28$ different inequalities to show. The easiest way to show this is by making a difference table for these 8 sums that I am claiming to be distinct, using the idea that $i = j$ if and only if $i - j = 0$ (in Z_p), and going on to show that only the main diagonal of our subtraction table has zeros. The main diagonal is absolutely unnecessary to show, but it is hoped to be aesthetically pleasing nonetheless.

−	2	-2	2b	-2b	b+1	-b+1	b-1	-b-1
2	0							
-2	-4	0						
2b	2b-2	2b+2	0					
-2b	-2b-2	-2b-2	-4b	0				
b+1	b-1	b+3	-b+1	3b+1	0			
-b+1	-b-1	-b+3	-3b+1	b+1	-2b	0		
b-1	b-3	b+1	-b-1	3b-1	-2	2b-2	0	
-b-1	-b-3	-b+1	-3b-1	b-1	-2b-2	-2	-2b	0

The upper-right half of this table has been left blank intentionally in order to avoid clutter because $i - j = 0$ if and only if $j - i = 0$, and we need to show inequality to zero only once for each pair of sums (i,j). It is clear that the 6 uppermost differences below the main diagonal can be obtained by successively multiplying (some of) -1, 2, b, b-1, and b+1. We know that none of these 5 values are zero in Z_p, since we start out with $b \in Z_p \setminus \{-1, 0, 1\}$ and $p > 2$.

Since p is prime, there are no divisors of zero and we can be sure that these 6 products are hence non-zero. Exactly the same idea applies to the six rightmost differences below the main diagonal as well. About the 16 remaining values, it is easy to see that in order for them to be zero, one of the below would have to be true: $b = \pm 1$, which is false, $b = \pm 3$, which we excluded from our construction, or $3b = \pm 1$, which would imply that $b = \pm 3^{-1}$, which we also excluded from our construction.

Due to unique factorization of all the possible sums $x + z$ into pairs of summands in the set $\{-1, 1, -b, b\}$ which we just finished demonstrating, $x = z$ would imply $w = y$ which would in turn imply $w = x$, which violates one of our inequalities immediately inferred from the definitions of w, x, y, and z. Therefore, w, x, y, and z can assume only those 8 possible assignments provided in the initial table.

You might be wondering where I am going with all these addition and subtraction tables. What I would like to point out is that for whatever legitimate values of x and z, $w = z$ and $x = y$. If we go one step back to the definitions of w, x, y, and z, I have just shown for any value of n, that:

$$\begin{aligned}\phi(n+1) - \phi(n) &= \phi(n+a+1) - \phi(n+a) and \\ \phi(n+a) - \phi(n) &= \phi(n+a+1) - \phi(n+1).\end{aligned}$$

The first equation is a statement concerning the voltages on the edges of the image of the directed hamiltonian cycle A1 under ϕ, and the second equation concerns the voltages on the edges of the image of the directed hamiltonian cycle A2 under ϕ. Why must this be so? We can simply use mathematical induction. We can simply derive:

$$w = \phi(n+a+1) - \phi(n+a) = \phi(n+2a+1) - \phi(n+2a)$$

from the first equation, and in general, that (where $k * a$ is meant to represent the sum of k a's in Z_p)

$$w = \phi(n+k*a+1) - \phi(n+k*a) = \phi(n+(k+1)*a+1) - \phi(n+(k+1)*a)$$

for any $k >= 0$, by induction. We can see that $n + k * a$ assumes all values in Z_p because a is not a divisor of p, since there are no divisors of p. This means, due to our initial isomorphism assumption, that we have managed to say that for all vertices in SP(1,b),

$$w = \phi(m+1) - \phi(m)$$

which is a statement about the voltages on ALL the edges in $\phi(A1)$.

Similarly, we can use induction on our second equation, since it implies

$$x = \phi(n + a + 1) - \phi(n + 1) = \phi(n + a + 2) - \phi(n + 2)$$

and in general, by induction and the same idea of multiplying by scalar k as above:

$$x = \phi(n + a + k * 1) - \phi(n + k * 1) = \phi(n + a + (k + 1) * 1) - \phi(n + (k + 1) * 1)$$

This is even easier to see than above that varying k, $n+k*1$ spans all elements in Z_p, hence, $\phi(n+k*1)$ spans all vertices of SP(1,b). This statement, equivalently reads, that $\phi(A2)$ where A2 is the directed hamiltonian cycle in SP(1,a) with voltages $+a$ on all the edges, has only a single voltage on the edges, in the same direction. We can neatly represent this result with:

$$x = \phi(m + a) - \phi(m)$$

for all m. Moreover, this value x is not equal to, nor is the additive inverse of the value on the edges of $\phi(A1)$, w.

Now, what can we say that relates a to b? Given the properties derived above, let us consider the telescoping series:

$$\sum_{j=0}^{a-1} \phi(m + j + 1) - \phi(m + j) = \phi(m + a) - \phi(m)$$

Notice that the right hand side is x, which we showed to be is an invariant of the directed cycle $\phi(A2)$, hence, same for all m. Likewise, notice that each term of the series on the left hand side equals w, and that there are a such terms. Therefore this equality can be written alternatively as:

$$aw = x$$

where $x, w \in \{-1, 1, +b, -b\}$ such that $x \neq \pm w$. If $w = \pm 1$, then $x = \pm b$. Which implies that $a(\pm 1) = \pm b$, in other symbols, $a = \pm b$. The other possibility is that $w = \pm b$, then $x = \pm 1$. This implies that $a(\pm b) = \pm 1$, equivalent to stating that $a = \pm b^{-1}$. We're done.

One thing that perhaps makes our proof appear flawed is that we excluded from our discussion the possibility that $b = \pm 3^{\pm 1}$. But our proof was such that a CAN assume those values, and it turns out that no isomorphism exists between SP(1,a) to SP(1,b) unless a and b are related under the composition of taking additive and multiplicative inverses. Therefore, if $b \neq \pm 3^{\pm 1}$, then $a \neq \pm 3^{\pm 1}$. Our proof would not have been complete if we had made a second exception of this sort. Note that requiring $a, b \neq \pm 2^{\pm 1}$ was not an exception of this sort because we already explained how those values derive the only Star-Polygon graphs with 3-cycles and are already known to be isomorphic to one another.

3 Conclusion

This was my first attempt at mathematical writing, and thank you for bearing with me if you read it thus far. Whoever you may be, feel free to contact me about it. As of yet, I really know little about how this voltage graphs relate to physics or Kirchoff's laws. You can read an introduction in [1]. I also have not taken any classes in topology (yet) so it is seriously out of the question for me to know how voltage graphs pertain to covering spaces.

3.1 About the results

As of January 2003, it is known (via Prof. Jonathan L. Gross, personal communication) that a paper by Alspach and Parsons [4] shows results that subsume those in this paper, albeit the starting point has a different specification (circulent graphs). This information was obtained after the submission and evaluation of this thesis. This relatively inelegant effort was completely independent from Alspach and Parsons' and any study of Star-Polygon graphs.

3.2 Further research suggestions

Logically, what is there next to look at? I have scraped the surface of the following two cases, but have not come up with mathematical results as solid as above. It appears that some of our findings are applicable to B_2 with voltages in Z_n. (This means ALL possible cyclic voltage groups.) I have started looking at these for simple composites $n = p^2$ and $n = pq$ where p and q are different primes greater than 2. So far, I do not have a presentable amount of research on this. Classes derived from B_3 and other bouquets (B_n) with voltages in Z_p can be approached in a similar way to B_2. For instance, in the case of B_3 with voltages in Z_p, we can continue insisting on a "standard" notation using 1 as at least one of the elements (where we have some non-zero voltage). It is clear (in the light of some results shown in this paper) that we will get equivalent triplets if we multiply the triplet through by the multiplicative inverse of any of the three elements of the triplet. (1,2,3) in Z_{43} is then equivalent (multiplying through by $2^{-1} = 22$) to (22, 1, 23) and (multiplying through by $3^{-1} = 29$) to (29, 15, 1). I have not had nearly enough time to prove general results to present an entire argument of the type provided for B_2. I am sure that my results will make it much easier.

Beyond these, one could look at classes of graphs derived from topologically more involved (hint: having more than one vertex) voltage graphs. It would probably be most interesting to look at classes of graphs derived from regular (non-cyclic) voltage graphs and permutation voltage graphs, and it seems plausible that getting a better grasp of the cyclic case would facilitate the process.

3.3 Acknowledgements

This is the hardest part to write, since without any of these people, this work would not have been possible, or at least turn out all right—and I could count the things I yet have been involved with that I would say that for with some fingers of one hand.

I am seriously indebted to:

Professor Jonathan L. Gross at the Computer Science department of Columbia University who agreed to advise me for this thesis for his time, guidance, enthusiasm, and his relentless motivation for me to push my limits. Working with him was extraordinary, and I consider it to be the raison d'être of my Columbia experience in retrospect. I was introduced to voltage graphs in the graph theory course that I took with him (which clearly had the most fun lectures at Columbia by far), which alone would have gotten him mentioned here.

Professor Bruce Peterson, for agreeing to be my advisor, and making it possible for this work to be recognized as a legitimate Middlebury College Mathematics Thesis. And somehow letting me know that I was somewhat on the right track in terms of what was required.

Professor Michael Olinick, for suggesting and making possible for me to receive a mathematics degree from Middlebury College after I was supposed to be done, i.e, by writing this thesis while I have been attending the School of Engineering and Applied Science at Columbia University, which was probably the single action that put coherence into my dual degree program. I believe it was the epitome of flexibility where getting a real degree is concerned (I am known to some as a troublemaker and even I would not have considered it appropriate enough to suggest it), and I appreciate it a lot.

Professor Peter Schumer, who ran the senior seminar this year, for communicating with me often, keeping me posted, suggesting great mathematical writing books (that I was able to make little use of), and making sure that I was doing work.

All those who have collectively contributed to the person that I am—first of all, all my current and former advisors and instructors in the Mathematics and Computer Science departments of all my alma maters, i.e., *Middlebury College*, Columbia University, and Robert College of Istanbul. And zeroth of all, my parents who are entirely responsible for my existence, and have dedicated an enormous proportion of their lives to my education for reasons I will never know.

References

[1] Jonathan L. Gross and Jay Yellen. *Graph Theory and its Applications.* CRC Press LLC, 2000 Corporate Blvd., N.W., Boca Raton, Florida 33431, 1999.

[2] Jonathan L. Gross and T. W. Tucker. Generating All Graph Coverings by Permutation Voltage Assignments. *Discrete Mathematics* 18 (1977), 273-283.

[3] Jonathan L. Gross and T. W. Tucker. *Topological Graph Theory.* Wiley Interscience Series in Discrete Mathematics and Optimisation. First Edition. 1987.

[4] Brian Alspach and T. D. Parsons. Isomorphism of circulant graphs and digraphs. *Discrete Mathematics* 25 (1979), 97-108.