THE NOVA GRAPH: MORE DISJOINT PATHS WITH MINIMAL GRAPH AUGMENTATION

JEFFERY J. BOATS, LAZAROS KIKAS, JOHN OLEKSIK

Abstract. For the purpose of large scale computing, we are interested in linking computers into large interconnection networks. For such networks to be useful, the underlying graph must possess desirable properties such as a large number of vertices, high connectivity, and small diameter. In these graphs, we consider the "k-Disjoint Path Problem": given a graph G, and any k pairs of distinct nodes $\{(s_1, t_1), (s_2, t_2),...,(s_k, t_k)\}$, can there be found k-disjoint paths, each one connecting a pair?

In this paper, we will demonstrate that a modification to the alternating group graph, which we have named the "Nova Graph," improves upon previous results regarding the k-Disjoint Path Problem. While maintaining the alternating group structure, which is useful in parallel-processing, the addition of a minimal number of edges increases the number of guaranteed disjoint paths. Furthermore, the Nova Graph gives the same number of disjoint paths as previously-observed graphs, while using far fewer vertices and edges.

Keywords: Interconnection networks, graphs, vertex disjoint paths

1. Introduction

For the purpose of large scale computing, we are interested in linking computers/processors into large symmetric interconnection networks. The papers [1,2] list some desirable properties that we wish the underlying graph of our interconnection networks to have: vertex symmetry, large number of vertices, high connectivity, and small diameter and degree.

Informally, a graph is vertex symmetric if each vertex in the graph can be viewed as identical to the other vertices in the graph. Recall that the connectivity of a non-complete graph refers to the minimum number of vertices that must be deleted in order to disconnect the graph. In terms of interconnection networks, the deletion of a vertex can be viewed as equivalent to processor failure. Hence, it is desirable to have networks that allow for a large number of processor failure but still be operational. In graph theoretic terms, we want to be able to delete a large number of vertices and still have a connected graph; that is, we want the underlying graph to have high connectivity.

Another measure of performance for interconnection networks deals with communication delay. When one sends a message from one processor to another in an interconnection network, the communication is never instantaneous, as it must be sent over a path in the graph. We want the graph to have a small diameter to minimize potential delay.

For a long time, the boolean n-cube served as a standard interconnection network model. In 1988, [1] introduced the star graph as a competitive alternative to the n-cube. In the 1990's, the split-star graph and the alternating group graph were introduced into the literature; see [4,7]. The vertex set of the split-star is, the symmetric group S_n, this time with vertex adjacency if and only if one can get from one permutation to the other by either the 2-exchange (12) or a 3-rotation of the form $(12k)$, where $k \in \{3, 4, .., n\}$. The alternating group graph consists of A_n, the even permutations of S_n, with vertex adjacency only given by the 3-rotations mentioned above.

2. Motivation

Suppose we have four processors (or computers), A, B, C, and D, within an interconnection network, and that we want A to communicate with B and C with D simultaneously. Communication between two processors is accomplished by sending the message across a path in the network. If the path of communication between A and B shares a computer with the path between C and D,then a resource must be shared during the simultaneous communication. The sharing is called **signal collision**, and can cause communication delay.

Signal collision is a major factor in the performance of parallel networks, for many signal collisions can slow down a network. The question is: Given an interconnection network, how many simultaneous signals can be routed through a particular network topology while avoiding signal collisions?

In graph theoretic terms we want to study the following problem: Given k pairs of distinct nodes $(s_1, t_1), (s_2, t_2), ..., (s_k, t_k)$, do there exist k-disjoint paths, one connection each pair? This problem is called the k-Disjoint Path Problem, and has generated much research. If for a graph G we can do this for any selection of k pairs of distinct nodes, then G is said to have the k-Disjoint Path Property.

The graphs mentioned above have all been studied with regard to this property, as they are favorable when it comes to parallel processing [1]. In 2003, [3] demonstrated that the split-star S_n^2 topology had the $(n-1)$-Disjoint Path Property for $n \geq 4$. Then [5], in 2005, showed that AG_n, the group graph of the alternating group A_n, has the $(n-2)$-Disjoint Path Property for $n \geq 5$. In 2006, [6] presented an algorithm for actually constructing the paths in AG_5, and their research continues on algorithms for path construction within Cayley graphs in general.

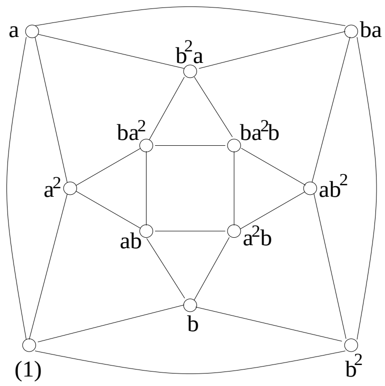

FIGURE 1. The Alternating Group Graph AG_4

In this paper, we create the "Nova Graph," or A_4^+, a modification on the alternating group graph which yields 3 disjoint paths using as few vertices and edges as possible, while still preserving the symmetry group-like structure favorable in parallel processing.

3. Construction and Initial Observations

As detailed in [7], the alternating group A_4 is generated by $a = (123)$ and $b = (124)$, and edge-connection in the group graph AG_4 can be defined as follows: there exists an edge connecting $v_1, v_2 \in G$ if and only if $v_1 = xv_2$, for some permutation $x \in \{a, b, a^2, b^2\}$. By this definition, the four vertices adjacent to a given vertex v are: $\{av, bv, a^2v, b^2v\}$.

The structure of the A_4 graph becomes apparent from exploring the algebra. The graph features eight 3-cycles created by repetition of a or b, since $a^3 = b^3 = (1)$, and six $abab$ 4-cycles since $abab = baba = (1)$. [8] showed that for $k \geq 5$, A_k has the $(k-2)$-disjoint path property, but that this result fails for $k = 4$ because a counterexample can be found on any $abab$ 4-cycle: for any $v \in G$, take $s_1 = v$, $s_2 = av$, $t_1 = bav$, and $t_2 = abav = b^2v$. To guarantee a greater number of disjoint paths, we will need to add more edges.

Watkins' Criterion [9] indicates that a vertex symmetric graph must have a minimum vertex degree of $(2k-1)$ for the graph to have the k-disjoint path property, so it follows that if the vertex degree is increased from four to five, the new graph could potentially have the 3-disjoint path property. It turns out that it is possible to do this for AG_4 with a minimum augmentation of six edges – we call the resulting graph the "Nova Graph."

Define $J = (12)(34)$, and consider the Cayley graph generated by $\{a, b, J\}$. It is a supergraph of A_4, augmented by six J edges, and defined by the following algebra:

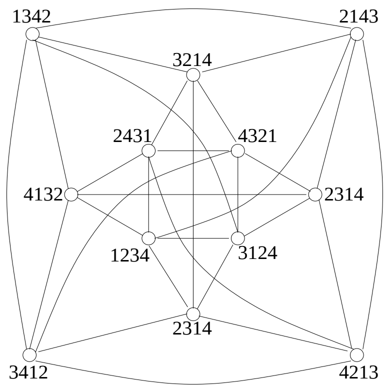

FIGURE 2. The Nova Graph A_4^+

$$\forall x, y \in \{a, b\}, \begin{cases} \mathbf{1}. \quad xyx = & \begin{cases} (1) & \text{if} \quad x = y; \\ y^2 & \text{if} \quad x \neq y; \end{cases} \\ \mathbf{2}. \quad \quad J = & ab^2a. \end{cases}$$

The first criterion describes the 3- and 4-cycles which comprise the A_4 group graph, while the second defines the generator J in terms of a and b. An algebraic analysis shows that $J = ab^2a = ba^2b = a^2ba^2 = b^2ab^2$. The augmentation of the graph by these J edges not only increases vertex degree to five, but it also reduces the graph diameter from three to two. This is an important observation, as it will simplify several parts of the proof that A_4^+ has the 3-Disjoint Path Property. The addition of J edges also greatly increases the number of 4-cycles, though it does not increase the number of 3-cycles, nor does it change the fact that the 3-cycles are pairwise edge-disjoint.

Let $S = \{s_1, s_2, s_3\}$ and $T = \{t_1, t_2, t_3\}$. If we let the six vertices of $S \cup T \subset A_4^+$ be given, any (s_i, t_i) pair can be connected with a path of at most two edges, unless of course they are a distance two apart and all possible intermediate vertices are "blocked" by other members of $S \cup T$, or by previously-drawn paths. The technique of constructing three disjoint paths involves analyzing the initial set-up of the $S \cup T$ vertices and deciding in which order to plan the routes. We'll need a few definitions to clarify such decisions.

Definition 1: An (s_i, t_i) pair is <u>initially connectible</u> if there exists a path from one to the other, disjoint from the other vertices in $S \cup T$, using only one or two edges.

Definition 2: A vertex is <u>block-surrounded</u> is every adjacent vertex is occupied by an element of $S \cup T$ or by a vertex used in a previously-constructed path.

Every time a vertex is used in constructing a connecting path for some (s_i, t_i) pair, it is a block to any not-yet-constructed paths, as are the other elements of $S \cup T$ whether their paths have been constructed yet or not. A block-surrounded vertex has been disconnected from the rest of the graph because every adjacent vertex has already been used. While attempting to construct k disjoint paths, we must be careful not to block-surround any elements of $S \cup T$ in the process.

4. Disjoint Paths in the Nova Graph

Theorem 1: *A_4^+ has the 3-Disjoint Path Property.*

The proof is lengthy, as many different contingencies can arise due to the possible configurations of $S \cup T$. We will attack it by first proving six lemmas, which we will then use to construct a coherent proof of the theorem.

Lemma 1 *Given (s_i, t_i), if neither s_i nor t_i is block-surrounded, then no less than 7 blocks are needed to disconnect them.*

<u>Proof</u>: Let (s_1, t_1) be given, and suppose that a number of other vertices in A_4^+ are declared blocks, causing the graph to be disconnected so that the path from s_1 to t_1 is not connectible. Since s_1 and t_1 are not block-surrounded, let x and y be the non-block vertices connected to s_1 and t_1 respectively. We must consider two cases: $x = Js_1$, and $x \neq Js_1$.

If $x \neq Js_1$, then $x = zs_1$, where $z \in \{a, a^2, b, b^2\}$. Without loss of generality, say $x = bs_1$; the other choices for z are handled similarly. Then the following seven vertices are adjacent to x or s_1 or both:

$$\{b^2s_1, as_1, a^2s_1, Js_1, ba^2s_1, a^2bs_1, abs_1\}.$$

Since t_1 and y cannot be any of the aforementioned vertices (as this would make (s_1, t_1) connectible), and since they are adjacent to each other, this forces them to be bas_1 and ab^2s_1 in some order. Each one of the seven vertices in the above list is adjacent to t_1 or y or both, so it follows that all seven must be blocks in order to disconnect s_1 from t_1.

If $x = Js_1$, then the list of vertices adjacent to x or s_1 or both has eight members: $\{as_1, bs_1, a^2s_1, b^2s_1, ab^2s_1, ba^2s_1, a^2bs_1, b^2as_1\}$. This leaves only abs_1 and bas_1 to be s_2 and y in some order. Again, each of the vertices in the above list is adjacent to either s_2 or y, so every vertex in the list must be a block. Hence, in either case, at least seven blocks are required to prevent the (s_1, t_1) path. $\square$

Lemma 2 *Two blocking paths each consisting of two edges or less cannot disconnect* A_4^+.

Proof: By contradiction. Suppose two paths of two edges or less can be found which disconnect A_4^+. Then there would be two vertices, $v_1, v_2 \in A_4^+$, which would be blocked from being connected by the at most six vertices in these paths. By Lemma 1, six vertices are not sufficient to disconnect v_1 and v_2 unless one or both of them are block-surrounded.

Without loss of generality, say v_1 is block-surrounded by the blocking paths. This implies that the blocking paths contain the five vertices adjacent to v_1: $\{av_1, bv_1, a^2v_1, b^2v_1, Jv_1\}$. One of the paths must therefore contain three of the vertices in the list, which implies at least one vertex in the list is adjacent to two others in the list. An algebraic analysis shows this to be untrue. By contradiction, the desired result follows. □

Lemma 3 *If each of the three* (s_i, t_i) *pairs are initially connectable, then at least two of the pairs can be connected with disjoint paths of two edges or less.*

Proof: If $dist(s_i, t_i) = 1$ for more than one index i, then the result holds trivially.

If $dist(s_i, t_i) = 1$ for exactly one pair, without loss of generality say (s_1, t_1), then its path can be connected without blocking any other path. Since (s_2, t_2) is initially connectible, the path can still be constructed in two edges or less, and the result holds.

Suppose $dist(s_i, t_i) = 2$ for all three paths. Then the (s_1, t_1) path can be connected through some intermediate vertex v. If this prevents (s_2, t_2) from being connected in two steps or less, it can only be that (s_2, t_2) was only initially connectible through v. If the (s_3, t_3) path is also prevented, then it also needed v as a intermediate vertex. This makes v adjacent to all six members of $S \cup T$, which contradicts the fact that vertices in A_4^+ have degree five. Thus a second (s_i, t_i) path must still be connectible in two edges or less. □

Lemma 4 *If* (s_1, t_1) *are not initially connectible, and are not on opposite corners of an* abab 4*-cycle, then two of the* (s_i, t_i) *paths can be connected using a total of four edges or less, without disconnecting the graph.*

Proof: It must be that $t_1 = ws_1$, where $w \in \{ab^2, ba^2, a^2b, b^2a\}$. By symmetry, the proofs for each of these four cases are the same – we present the $t_1 = ab^2s_1$ case.

Since (s_1, t_1) is not initially connectible, the three vertices adjacent to both s_1 and t_1, $X = \{a^2s_1, ba^2bs_1, b^2s_1\}$, must be members of $S \cup T$. It

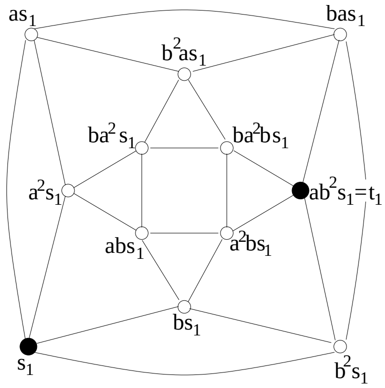

FIGURE 3. Illustration for Lemma 4

must be that two of them form an (s_i, t_i) pair – without loss of generality say (s_2, t_2).

The vertex $y = ba^2s_1$ is also adjacent to every member of X. If y is the last remaining element of $S \cup T$, then it can be connected to its mate in one edge, since its mate must be in X. Meanwhile, the (s_1, t_1) path can be made in three edges as follows: $s_1 \to bs_1 \to a^2bs_1 \to t_1$. These two paths set up only six blocks for the remaining (s_2, t_2) path. Furthermore, s_2 and t_2 are elements of X, and none of X's elements are block-surrounded, since a^2s_1 is adjacent to the unused as_1, vertex ba^2bs_1 is adjacent to b^2as_1, and b^2s_1 is adjacent to bas_1. By Lemma 1, the result follows.

If $y \notin S \cup T$, then the (s_2, t_2) path can be made through it, using two edges. The other member of X is either s_3 or t_3, and must be paired with a vertex in one of the sets $Z_1 = \{bs_1, abs_1, a^2bs_1\}$ and $Z_2 = \{as_1, bas_1, b^2as_1\}$, both of which are 3-cycles and thus have pairwise-adjacent vertices. By inspection, every element of X is adjacent to exactly one vertex in each of Z_1 and Z_2, so it follows that (s_3, t_3) can also be connected in two edges or less. By Lemma 2, the two paths do not disconnect A_4^+. □

Lemma 5 *If (s_1, t_1) occupy opposite corners of an* abab *4-cycle, and are not initially connectible due to blocks from different (s_i, t_i) pairs, then (s_2, t_2) and (s_3, t_3) can be connected with disjoint paths of two edges or less.*

Proof: By symmetry, and without loss of generality, say s_2 and s_3 are the blocks of the (s_1, t_1) paths, so that:

$$s_2 = as_1,\ t_1 = bas_1, \text{ and } s_3 = abas_1 = b^2s_1.$$

The remaining eight vertices are:

$$\{a^2s_1, ba^2s_1, b^2as_1, ba^2bs_1, abs_1, bs_1, a^2bs_1, ab^2s_1\}.$$

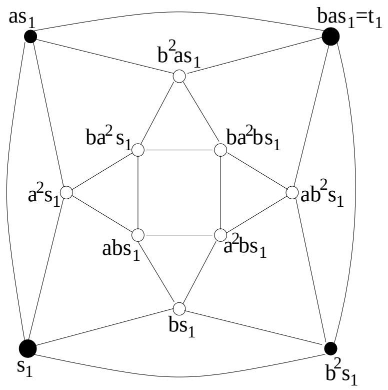

FIGURE 4. Illustration for Lemmas 5 and 6

Each can be reached in two edges or less from either s_2 or s_3. When not adjacent, they can be reached in two edge-disjoint ways, with the intermediate vertices of said paths being non-adjacent, as demonstrated in the following table.

$$\begin{array}{ccccccccc}
 & & a\mathbf{s_2} & = & a^2\mathbf{s_1} & = & b^2J\mathbf{s_3} & = & Ja\mathbf{s_3} \\
ba\mathbf{s_2} & = & a^2b^2\mathbf{s_2} & = & ba^2\mathbf{s_1} & = & J\mathbf{s_3} & & \\
 & & b^2\mathbf{s_2} & = & b^2a\mathbf{s_1} & = & Jb^2\mathbf{s_3} & = & aJ\mathbf{s_3} \\
bJ\mathbf{s_2} & = & ab^2\mathbf{s_2} & = & ba^2b\mathbf{s_1} & = & b^2a\mathbf{s_3} & = & a^2J\mathbf{s_3} \\
b^2a\mathbf{s_2} & = & a^2J\mathbf{s_2} & = & ab\mathbf{s_1} & = & bJ\mathbf{s_3} & = & ab^2\mathbf{s_3} \\
Jb^2\mathbf{s_2} & = & aJ\mathbf{s_2} & = & b\mathbf{s_1} & = & b^2\mathbf{s_3} & & \\
 & & J\mathbf{s_2} & = & a^2b\mathbf{s_1} & = & a^2b^2\mathbf{s_3} & = & ba\mathbf{s_3} \\
b^2J\mathbf{s_2} & = & Ja\mathbf{s_2} & = & ab^2\mathbf{s_1} & = & a\mathbf{s_3} & &
\end{array}$$

If s_2 and t_2 are adjacent, their path requires one edge and is unblockable. If they are non-adjacent, then there are two edge-disjoint paths between them, and t_3 can block no more than one of them. Either way, the (s_2, t_2) path can be made in two edges or less.

Suppose now that the (s_3, t_3) path is blocked from connection in two edges or less. It would follow that s_3 and t_3 are non-adjacent, and that the two paths between them, listed in the table above, are both blocked – one by t_2 and the other by the intermediate vertex of the (s_2, t_2) path (call it v). This makes the set $\{v, t_2, s_3\}$ a 3-cycle, since they are pairwise-adjacent. Similarly, $\{v, t_2, t_3\}$ must be a 3-cycle. The edge from v to t_2 appears in both 3-cycles, which contradicts the observation that 3-cycles are edge-disjoint in A_4^+. By contradiction, it follows that (s_3, t_3) is also connectible in two edges or less. □

Lemma 6 *If (s_1, t_1) and (s_2, t_2) occupy opposite corners of the same* abab *4-cycle, blocking each other from being initially connectible, then each* *of their paths can be connected using three edges, in a way which does not disconnect s_3 from t_3.*

Proof: By symmetry, and without loss of generality, we can restrict our consideration to the configuration $s_2 = as_1$, $t_1 = bas_1$, and $t_2 = abas_1 = b^2s_1$. We will need to consider a few cases, based upon the locations of s_3 and t_3.

Define $U = \{Js_1, Js_2, Jt_1, Jt_2\} = \{ba^2bs_1, a^2bs_1, abs_1, ba^2s_1\}$, an $abab$ 4-cycle, and $V = \{a^2s_1, b^2as_1, ab^2s_1, bs_1\}$. This makes $U \cup V$ the complement of $\{s_1, s_2, t_1, t_2\}$ in A_4^+.

The (s_1, t_1) path requires at least 3 edges, but there are six possible paths; similarly, six possible (s_2, t_2) paths exist. We number them (#1) through (#12) below, for easy reference. The proper choice of (s_1, t_1) and (s_2, t_2) paths will avoid disconnecting s_3 and t_3.

(#1)	$\mathbf{s_1}$	$\to$	a^2s_1	$\to$	abs_1	$\to$	$\mathbf{t_1}$
(#2)	$\mathbf{s_1}$	$\to$	a^2s_1	$\to$	ab^2s_1	$\to$	$\mathbf{t_1}$
(#3)	$\mathbf{s_1}$	$\to$	bs_1	$\to$	abs_1	$\to$	$\mathbf{t_1}$
(#4)	$\mathbf{s_1}$	$\to$	bs_1	$\to$	b^2as_1	$\to$	$\mathbf{t_1}$
(#5)	$\mathbf{s_1}$	$\to$	Js_1	$\to$	b^2as_1	$\to$	$\mathbf{t_1}$
(#6)	$\mathbf{s_1}$	$\to$	Js_1	$\to$	ab^2s_1	$\to$	$\mathbf{t_1}$
(#7)	$\mathbf{s_2}$	$\to$	a^2s_1	$\to$	ba^2s_1	$\to$	$\mathbf{t_2}$
(#8)	$\mathbf{s_2}$	$\to$	a^2s_1	$\to$	ab^2s_1	$\to$	$\mathbf{t_2}$
(#9)	$\mathbf{s_2}$	$\to$	b^2as_1	$\to$	ba^2s_1	$\to$	$\mathbf{t_2}$
(#10)	$\mathbf{s_2}$	$\to$	b^2as_1	$\to$	bs_1	$\to$	$\mathbf{t_2}$
(#11)	$\mathbf{s_2}$	$\to$	a^2bs_1	$\to$	bs_1	$\to$	$\mathbf{t_2}$
(#12)	$\mathbf{s_2}$	$\to$	a^2bs_1	$\to$	ab^2s_1	$\to$	$\mathbf{t_2}$

If $s_3, t_3 \in U$, paths (#2) and (#10) are used. These paths avoid U altogether, and since U is a connected subgraph, the (s_3, t_3) path can be constructed.

If $s_3, t_3 \in V$ and are non-adjacent (i.e. $s_3 \neq Jt_3$), then path (#3) or (#5) is used for (s_1, t_1), while path (#7) or (#12) is used for (s_2, t_2) – at least one of the four possible combinations is guaranteed to miss s_3 and t_3 and leave them connected.

Finally, if $s_3, t_3 \in V$ and are adjacent (i.e. $s_3 = Jt_3$), or if one is in U and the other in V, then one of the following four combinations will miss s_3 and t_3 and leave them connected: (#5) with (#11), (#3) with (#9), (#6) with (#7), or (#1) with (#12). □

The lemmas have explored all possible configurations of $S \cup T$, and described how to deal with every contingency. We are now ready to assemble them and prove that A_4^+ has the 3-Disjoint Path Property.

Proof of Theorem 1: Let $S \cup T$ be given. Consider whether or not all three (s_i, t_i) pairs which are initially connectible. If all three pairs are initially connectible, then by Lemma 3, at least two of the paths can be constructed disjointly. Furthermore, these paths will have two edges or less, so by Lemma 2 they do not disconnect the graph, thus allowing a third disjoint path connecting the third pair.

If one of the pairs is not initially connectible, without loss of generality say (s_1, t_1), then we must consider whether the pair occupies opposite vertices in an *abab* 4-cycle. If not, then Lemma 4 guarantees the other two pairs can be connected with paths which don't diconnect the graph, and again the third path can be constructed disjointly.

On the other hand, if the non-initially connectible pair occupies opposite vertices of an *abab* 4-cycle, then we must consider whether the blocks on this 4-cycle belong to different (s_i, t_i) pairs or the same pair. If they belong to different pairs, then Lemma 5 guarantees (s_2, t_2) and (s_3, t_3) can be connected in two edges or less, and we again apply Lemma 2. If the blocks belong to the same pair, say (s_2, t_2), then Lemma 6 assures that both pairs can be connected without disconnecting s_3 from t_3, thus allowing three disjoint paths. □

5. A_n^+ and its Minimality of Edges and Vertices

The split-star graph S_4^2 has the 3-Disjoint Path Property, and can be extended to S_n^2, which possesses the $(n-1)$-Disjoint Path Property for any $n \geq 4$. Similarly, the alternating group graph is extendable to AG_n, having the $(n-2)$-Disjoint Path Property for any $n \geq 5$. The ability to extend graphs to higher values of n is important in demonstrating the graphs' usefulness in larger-scale computations.

The Nova Graph A_4^+ can also be extended, to A_n^+. The construction is simple – take AG_n and add the J edges defined as $J = (12)(34)$. The hierarchical structure of AG_n is maintained, as the graph of A_n^+ consists of n "substars" which are each isomorphic to A_{n-1}^+.

The hierarchical structures of S_n^2 and AG_n facilitate inductive proofs that their guaranteed number of disjoint paths increases with n. A similar induction proof, using A_4^+ as the base case, demonstrates that A_n^+ has the $(n-1)$-Disjoint Path Property for any $n \geq 4$; this proof is left for a later paper.

A more important consideration involves efficiency of network resources. Given that A_4^+ has the 3-Disjoint Path Property, the question naturally arises: Is A_4^+ the smallest graph with this property? To answer this, we must first specify what is meant by "smallest." If it is advantageous to have as many guaranteed disjoint paths as possible, then it is advantageous to do so with as few edges and/or vertices a possible. How small, in terms of the number of edges and vertices, can a network be and still guarantee a given number of disjoint paths?

The article [1] expounds on the advantages of using symmetry group graphs in parallel processing. The graphs of this type which have been studied are S_n^2 and AG_n, and now A_n^+. Given the criterion that a graph must guarantee k disjoint paths, the avantage of the Nova Graph becomes clear when we count the number of vertices and edges necessary. For $k \geq 3$, the previous smallest number of each was S_{k+1}^2, giving k disjoint paths with $(k+1)!$ vertices and $\frac{2k-1}{2}(k+1)!$ edges. The graph A_{k+1}^+ guarantees k disjoint paths with half as many vertices, $\frac{1}{2}(k+1)!$, and half as many edges, $\frac{2k-1}{4}(k+1)!$.

The Nova Graph is the best graph in this respect so long as we restrict ourselves to the symmetry group graph family. We must mention, however, that it is possible to obtain k disjoint paths with still fewer vertices and edges, provided we are willing to be less particular about which types of graphs we use. For example, the complete graph K_6 obviously has the 3-disjoint path property. However, the purpose of studying the k-disjoint path property in graphs is to determine how best to create interconnection networks when we don't have the luxury of linking every processor to every other processor. We will restrict our further study to non-complete, vertex-symmetric graphs.

If a graph $G = (V, E)$ is vertex-symmetric, and contains $|V|$ vertices, then the number of edges is given by $|E| = \frac{1}{2}d|V|$, where d is the degree of each vertex. Since $|E|$ must be an integer, it follows that $|V|$ and d cannot both be odd. Since G must be non-complete, it must have $|V| > 6$ if it is to guarantee 3 disjoint paths. Even with these restrictions, it is possible to guarantee 3 disjoint paths with a slightly smaller graph than A_4^+, which requires 12 vertices and 30 edges.

Theorem 2: If a non-complete, vertex-symmetric graph $G = (V, E)$ has the 3-Disjoint Path Property, then $|V| \geq 9$ and $|E| \geq 25$.

Proof: Since G is non-complete, we already know $|V| > 6$. Since G has the 3-Disjoint Path Property, Watkins' Criterion assures that the degree d of each vertex must be at least 5. If $d = |V| - 1$, that would make G complete, so it must be that $5 \leq d < |V| - 1$. It follows that $|V| \neq 7$, since the inequality would force $d = 5$, yet d and $|V|$ can't both be odd.

For $|V| = 8$, consider the graph complement of G, which must also be vertex-symmetric. If we select s_1 and t_1 so that they are adjacent in the complement, then they are non-adjacent in G, and the (s_1, t_1) path would require at least two edges and three vertices. If (s_2, t_2) and (s_3, t_3) can be chosen similarly, then G cannot have the 3-Disjoint Path Property, since the three paths would require no less than nine vertices.

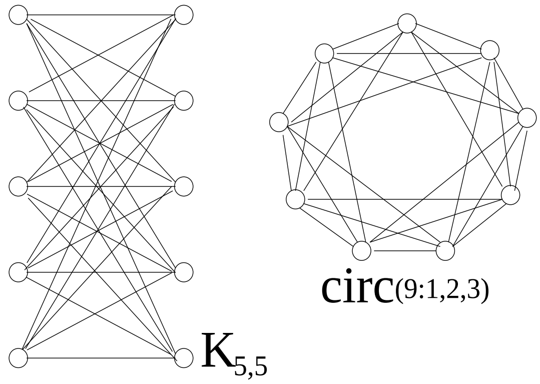

FIGURE 5. Minimal graphs with the 3-Disjoint Path Property

If $|V| = 8$, the inequality forces $d = 5$ or $d = 6$. If $d = 6$, then the degree of every vertex in the graph complement of G is 1, which means the complement of G is a set of four disjoint P_2 path graphs. If we select each (s_i, t_i) pair to occupy opposite ends of one of these P_2's, we see by counterexample that G does not have the 3-disjoint path property. If instead $d = 5$; then vertices in the complement have degree 2. The only vertex-symmetric graphs with 8 vertices of degree 2 are the cycle graph C_8 and the graph consisting of two disjoint C_4 cycles. In either case, it is simple to place $S \cup T$ so that s_i is adjacent to t_i for $i = 1, 2, 3$; thus G does not have the 3-Disjoint Path Property.

Therefore $|V| \geq 9$. If $|V| = 9$, then d must be even, so $d = 6$, which makes $|E| = 27$. If $|V| = 10$, d can be as low as 5, allowing $|E| = 25$. The formula $|E| = \frac{1}{2}d|V|$, combined with Watkins' Criterion $d \geq 5$, forces $|E| > 25$ whenever $|V| \geq 11$. Hence, $|V| = 9$ and $|E| = 25$ are minimal requirements. $\square$

These minimal requirements are in fact both necessary and sufficient, as examples are known which exhibit these minimal requirements and have the 3-Disjoint Path Property. The circulant graph $circ(9 : 1, 2, 3)$ is an example with $|V| = 9$ and $|E| = 27$. The complete bitartite graph $K_{5,5}$ is an example with $|V| = 10$ and $|E| = 25$.

These minimal graphs are shown in Figure 5; we leave as exercises the proofs that they have the 3-Disjoint Path Property.

It should be mentioned that we have paid a price in order to get away with fewer vertices and edges than A_4^+. Both $K_{5,5}$ and $circ(9 : 1, 2, 3)$ are more than half-complete; they have more edges than their graph complements. In other words, a randomly selected pair of vertices are more likely than not to be connected by an edge. One might ask: If we are to use graphs so heavily edge-laden as this, why not use complete graphs?

As the whole point of the exploration was to find useful non-complete graphs, another desirable characteristic of the Nova Graph becomes clear. If we wish our vertex-symmetric graph $G = (V, E)$ to have the 3-Disjoint Path Property, than d must be at least 5. If we wish it also to have fewer edges than its complement, then its complement's vertices must have degree at least 6; this forces $|V| \geq 12$. It follows that the Nova Graph, with its 12 vertices, is optimal among graphs which are less than half-complete.

6. Conclusion

The Nova Graph has been shown to guarantee 3 disjoint paths using the minimum number of edges and vertices while maintaining the graph properties desirable for parallel processing networks. A future paper will show it is extendable to A^+_{k+1}, which has the k-Disjoint Path Property.

References

[1] S. B. Akers, D. Harel, and B. Kirshnamurthy. The star graph: An attractive alternative to the n-cube. *Proceedings of the International Conference on Parallel Processing*, pages 393–400, 1987.

[2] S. B. Akers and B. Kirshnamurthy. A group theoretic model for symmetric interconnection networks. *IEEE Transactions on Computers*, 38(4):555–566, 1989.

[3] E. Cheng and M. Lipman. Disjoint paths in split-stars. *Congressus Numerantium*, 137:47–63, 1999.

[4] E. Cheng, M. Lipman, and H. A. Park. Super connectivity of star graphs, alternating group graphs and split-stars. *Ars Combinatoria*, 59:107–116, 2001.

[5] E.Cheng, L.D.Kikas, and S.Kruk. A disjoint path problem in the alternating group graph. *Congressus Numerantium*, 175:117–159, 2005.

[6] J.Boats, L.D.Kikas, and J.Oleksik. Algorithm for finding disjoint paths in the alternating group graph. *Congressus Numerantium*, 181:97–109, 2006.

[7] J. S. Jwo, S. Lakshimivarahan, and S. K. Dhall. A new class of interconnection network based on the alternating group. *Networks*, 23:315–325, 1993.

[8] Lazaros D. Kikas. *Interconnection Networks and the k-Disjoint Path Problem.* PhD thesis, Oakland University, 2004.

[9] M. E. Watkins. On the existence of certain disjoint arcs in graphs. *Duke Mathematics Journal*, 35:231–246, 1968.

Dr. Jeffery J. Boats, Corresponding Author, Department of Mathematics and Computer Science, University of Detroit Mercy, Detroit, MI, 48221, USA
E-mail address: boatsjj@udmercy.edu

Lazaros Kikas, Department of Mathematics and Computer Science, University of Detroit Mercy, Detroit, MI, 48221, USA
E-mail address: kikasld@udmercy.edu

John Oleksik, Department of Mathematics and Computer Science, University of Detroit Mercy, Detroit, MI, 48221, USA
E-mail address: oleksijj@udmercy.edu