FINDING DISJOINT PATHS IN THE NOVA GRAPH

JEFFE BOATS, LAZAROS KIKAS, JOHN OLEKSIK

ABSTRACT. For the purpose of large scale computing, we are interested in linking computers into large interconnection networks. In order for these networks to be useful, the underlying graph must possess desirable properties such as a large number of vertices, high connectivity and small diameter. In this paper, we are interested in the Nova graph as an interconnection network, and the k-Disjoint Path Problem. In 2007, Boats, Kikas and Oleksik showed that the Nova graph A_4^+ has the 3-Disjoint Path Property. In this paper, we extend the result to the general Nova graph A_n^+, and show that this class of interconnection networks has the $(n-1)$-Disjoint Path Property for $n \geq 4$. We discuss the significance of this result in comparison to other interconnection networks and close with remarks on possible future research directions.

Keywords: Interconnection networks, graphs, vertex disjoint paths

1. Introduction

The study of large interconnection networks has attracted a great deal of research over the last few years. In order for these networks to be useful, the underlying graph structure should have properties such as vertex symmetry, high connectivity, low diameter, and a large number of vertices (see [1,2]).

In this paper we study the following problem: Given k pairs of distinct nodes $(s_1, t_1), (s_2, t_2), \dots, (s_k, t_k)$ in a graph, can these pairs be connected with k disjoint paths? 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**.

Much has been studied about this problem. It has been shown for $k \geq 3$ the problem of finding k disjoint paths is NP-hard. Watkins in [9] showed that if a graph G has the k-Disjoint Path Property, then it must be $(2k-1)$-connected. Past work done by Cheng and Lipman shows that the split star graph, S_n^2, has the $(n-1)$-Disjoint Path Property [3]. In 2005, Cheng, Kikas, and Kruk showed that the alternating group graph AG_n has the $(n-2)$-Disjoint Path Property [4,8].

Both of the above proofs, for the split star and the alternating group, were existence proofs, not providing an algorithm for the construction

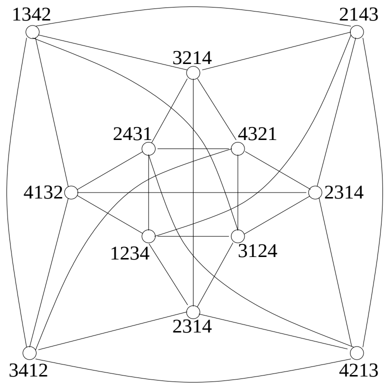

FIGURE 1. The Nova Graph A_4^+

of these paths. In 2006, Boats, Kikas, and Oleksik developed an algebraic/geometric algorithm that took advantage of the hierarchal structure of AG_5 to demonstrate the construction of the 3-disjoint paths [5].

In 2007, Boats, Kikas and Oleksik developed a generalized algebraic approach for finding disjoint paths in the alternating group graph AG_n [7]. In a separate paper in 2007, Boats, Kikas and Oleksik introduced the Nova graph as a new interconnection network, and showed that the Nova graph A_4^+ has the 3-Disjoint Path Property [6]. In this paper, we generalize our result to A_n^+ and show that this class of interconnection networks has the $(n-1)$-Disjoint Path Property for $n \geq 4$.

2. Composition and Connectivity of the Nova Graph

The Nova graph is formed from the alternating group graph, AG_n, $n \geq 4$, whose vertices are the even permutations of the first n natural numbers. Two vertices are connected by an edge if and only if one of the vertices' corresponding permutations can be obtained from the other by means of a permutation of the form $(12j)$, with $j \in \{3, 4, \ldots, n\}$. This makes AG_n a Cayley graph, with these $n-2$ edge permutations as generators. The degree of each vertex is $2n-4$, since there will be an adjacent vertex for each of the $n-2$ generators, and also each of the generators inverses, $(1j2)$.

The Nova Graph, A_n^+, is formed from AG_n by adding edges defined by the permutation $J = (12)(34)$. It was shown in [6] that this decreased the diameter of A_4^+ to two, enabling it to have the 3-disjoint path property. The Nova Graph is also a Cayley graph, with $n-1$ generators. The degree of each vertex is $2n-3$; that includes the $2n-4$ adjacent vertices from before, and one more adjacent vertex connected by a "J" edge. Figure 1 illustrates A_4^+.

Like AG_n, A_n^+ has a hierarchical structure, consisting of n "**substars**," denoted $H_1, H_2, \ldots, H_n$, each with identical structure to A_{n-1}^+. The index j in H_j indicates the last digit in the even permutation associated with the vertex. For this reason, every vertex in the graph is adjacent to $2n-5$ vertices within the same substar, and one more vertex in each of 2 other substars. Those two vertices are adjacent through $(12n)$ and $(1n2)$ edges, and they are also adjacent to each other, thus forming a 3-cycle. These $(12n)$ and $(1n2)$ edges are called "**conduits**."

There will be times when we wish to begin a path from s_i to t_i by immediately leaving s_i's substar via a conduit edge; we will refer to this as "**quickgating**." The vertex v_1 can "quickgate" to vertices $v_2 \in H_y$ and $v_3 \in H_z$, where y and z are the first two integers in the even permutation corresponding to v_1. The only 3-cycles which involve more than one substar necessarily involve three different substars, such as the $\{v_1, v_2, v_3\}$ example above, so it follows that two vertices in the same substar cannot be adjacent to the same vertex in another substar. This is useful to know, as it means whenever we choose to quickgate multiple vertices from the same substar, all of them will necessarily go to different places.

Between any 2 substars of A_n^+ are found $K = (n-2)!$ conduits. This is an important observation, since we will want to use conduits whenever routing a path for any (s_i, t_i) pair not contained within a single substar. If neither $s_i \in H_y$ nor $t_i \in H_z$ lies on a conduit connecting their substars, we may attempt to route s_i to a vertex $\sigma_i \in H_y$, and t_i to $\tau_i \in H_z$, where σ_i and τ_i occupy opposite ends of a conduit. This is called "**surrogate routing**," and σ_i and τ_i are the "**surrogates**" of s_i and t_i respectively.

3. Preliminary Considerations

Suppose M_{ij} is the number of mated pairs which must be routed between H_i and H_j. Before routing them, we must consider not only the value of K, but also B_{ij}, the number of conduit vertices occupied by elements of $S \cup T$ not to be so routed. The B_{ij} vertices are "blocks,"– they render their conduits useless since the M_{ij} paths cannot be routed through them.

We must also consider how the vertices of $S \cup T$ are distributed among the substars. The following definition effectively categorizes the possible arrangements.

Definition For the vertices $S \cup T \subset A_n^+$, their phi-count is the ordered n-tuple $(\phi_1, \phi_2, \ldots, \phi_n)$, where $\phi_\alpha = |(S \cup T) \cap H_\alpha|$. We denote by ϕ_{max} the largest number in a phi-count:

$$\phi_{max} = \max_\alpha |(S \cup T) \cap H_\alpha|.$$

There are two lemmas which we will need during the proof of the main result. The first establishes a condition on ϕ_{max} which will guarantee the existence of sufficient disjoint paths, while the second will help prove a difficult special case.

Lemma 1 *In A^+_{n+1}, with $n \geq 4$, if $\phi_{max} \leq n-1$, then the vertices of $S \cup T$ can be connected with n disjoint paths.*
Proof: Let K be the number of conduits between substars in A^+_{n+1}; for $n \geq 4$, we have $K = (n-1)! \geq 2(n-1)$.

Let M_{ij} be the number of mated pairs to be routed between substars H_i and H_j, and B_{ij} be the number of blocked conduits between those substars. Every vertex of $S \cup T$ could potentially block a conduit, but for every mated pair to be routed between the substars, at least one conduit (if not two) are still usable. Thus: $B_{ij} \leq \phi_i + \phi_j - M_{ij}$.

For routing the M_{ij} pairs via conduits to be impossible, it would take:

$$\begin{array}{rcl} M_{i,j} & > & K - B_{i,j} \\ & \geq & K - (\phi_i + \phi_j - M_{i,j}) \\ & = & M_{i,j} + K - \phi_i - \phi j \end{array}$$

$$\Rightarrow \quad \phi_i + \phi_j > K \geq 2(n-1)$$

Thus, $\phi_{max} \leq n-1$ is sufficient to guarantee there will be enough unblocked conduits for routing.

Once suitable open conduits are found for all inter-substar $S \cup T$ pairs, we can always use these conduits for surrogate routing. Each substar of A^+_{n+1} has the $(n-1)$-disjoint path property, and since $\phi_{max} \leq n-1$, it will always be possible to route vertices of $S \cup T$ to surrogates. $\square$

This lemma lends itself to a general strategy for proving A^+_{n+1} has the n-disjoint path property. If the vertices of $S \cup T$ are distributed such that $\phi_{max} \leq n-1$, or if we can route some or all of the $S \cup T$ vertices to new positions in such a way that the new $\phi^*_{max} \leq n-1$, then the result follows. Indeed, explaining how this can be done for various values of ϕ_{max} will constitute the bulk of the induction step of the proof.

Lemma 2: *For $n \geq 4$, if A_{n+1} has substars H_y and H_z such that $S \subset H_y$ and $T \subset H_z$, then it is always possible to find two mated (s_i, t_i) pairs for which both vertices can quickgate without being blocked.*

Proof: Since each vertex of $S \cup T$ can quickgate to a vertex in two different substars, one of the two substars must be empty. It follows that any given mated pair, s_i and t_i, can each quickgate without being blocked. This is because s_i has at least one open vertex to which it may quickgate, and the only way t_1 could have both its quickgates occupied is if s_1 was routed to one of them, but t_1 cannot be blocked by its mate.

Quickgate s_1 and t_1; we name their destinations w_1 and w_2 respectively. It is impossible for s_1 to block a quickgate destination for another vertex in S, because two vertices in the same substar can't quickgate to the same destination. The vertex s_1 could block one of the t_i's, but no more than one, since w_1 is adjacent to only two vertices from other substars, and s_1 is already one of them. Similarly, t_1 is also capable of blocking only one of the s_i's.

Thus the mated pair (s_1, t_1) is capable of blocking at most two other mated pairs. Since $n \geq 4$, there are at least three other mated pairs, and therefore a second pair will always be able to quickgate. $\square$

4. Proof of the Proposition

Proposition *For $n \geq 4$, A_n^+ has the $(n-1)$-disjoint path property.*

This will be a proof by induction, using $n = 3$ (i.e. A_4^+) as the base case. It has already been proven by Boats, Kikas, and Oleksik in [6] that A_4^+ has the 3-disjoint path property, so only the induction step remains; its proof will take up the remainder of this section.

For the induction step, we assume $n \geq 4$ and that A_n has the $(n-1)$-disjoint path property. It follows that the substars H_i, $i \in \{1, 2, \ldots, n+1\}$, which comprise A_{n+1}^+ each have the $(n-1)$-disjoint path property. The object will be to show that A_{n+1}^+ has the n-disjoint path property.

Let the $S \cup T$ configuration be given, and consider the value of ϕ_{max}.

<u>The $\phi_{max} \leq n-1$ cases:</u> By Lemma 1, n disjoint paths can be found.

<u>The $\phi_{max} = 2n$ case:</u> We can quickgate every $S \cup T$ vertex to one of two neighbors with no fear of blocking, and if the new $\phi^*_{max} \leq n-1$, then the n disjoint paths can be completed by Lemma 1.

Suppose $\phi^*_{max} \leq n-1$ isn't possible. Since any two $S \cup T$ vertices can be quickgated to different substars, it can only be that all or all but one of the $S \cup T$ vertices quickgate to the same two substars. In either case, the quickgate destinations can be selected so that $n-1$ mated (s_i, t_i) pairs are sent to some substar H^*, and the other two $S \cup T$ vertices are sent anywhere else. Since H^* has the $(n-1)$-disjoint path property, those $n-1$ disjoint paths can be completed; it is then a trivial matter to complete the last path through surrogate routing.

<u>The $\phi_{max} = 2n-1$ case:</u> Without loss of generality, say $s_1 \in H_2$, which means $\{(S \cup T) \setminus s_1\} \subset H_1$. We then quickgate some $S \cup T$ vertex in H_1 to a substar other than H_2, say H_3; if s_1 blocks a quickgate destination for some vertex in H_1 other than t_1, we necessarily choose this vertex to quickgate to H_3.

If $\phi_{max} \leq n-1$ is possible we can apply Lemma 1, so assume not. It follows that all or all but one of the remaining $S \cup T$ vertices in H_2 quickgate to the same two substars, and that at least one of those substars must be H_1 or H_3. As in the previous case, we can choose quickgate destinations so that $n-1$ mated pairs end up in the same substar (H_1 or H_3), with the other two vertices ending up elsewhere, and the result follows by the same argument as in the previous case.

<u>The $n \leq \phi_{max} \leq 2n-2$ cases:</u> We categorize these cases as $\phi_{max} = 2n-x$, where x is the number of vertices not in H_1. We will wish to connect some mated pairs within H_1, and route unmated $S \cup T$ vertices to conduits in order to send them to their mates' substars.

Since H_1 has the $(n-1)$-disjoint path property, we can route all we want provided we use no more than $2(n-1)$ total surrogate and $S \cup T$ vertices; hence the number of unmated vertices we can route to any conduits we wish is:

$$\begin{aligned} 2n-2-\phi_{max} &= 2n-2-(2n-x) \\ &= x-2 \end{aligned}$$

Define ψ to be the number of mated pairs in H_1; clearly $\psi \geq n-x$. Since H_1 has the $(n-1)$-disjoint path property, we can easily connect all ψ mated pairs with disjoint paths. After routing these mates together, the number of unmated vertices left is $\phi_{max} - 2\psi$.

Consider the subcase where not every $S \cup T$ pair has a representative in H_1; i.e. there exists a mated (s_1, t_1) pair with neither s_1 nor t_1 in H_1. It follows that $\psi \geq n-x+1$. This means the number of unmated vertices in H_1 is:

$$\begin{aligned} \phi_{max} - 2\psi &= 2n-x-2\psi \\ &\leq 2n-x-2(n-x+1) \\ &= x-2\ . \end{aligned}$$

Thus every $S \cup T$ vertex can either be routed to its mate in H_1, or routed to an open conduit from H_1 to the substar hosting its mate; the n disjoint paths are then easily completed.

All remaining subcases involve $S \cup T$ configurations for which every mated pair has a representative in H_1, and since switching the s_i and t_i designations doesn't change the problem, we can categorize these subcases, without loss of generality, as $S \subset H_1$. The number of unmated vertices in H_1 is:

$$\begin{aligned} \phi_{max} - 2\psi &= 2n-x-2\psi \\ &= 2n-x-2(n-x) \\ &= x\ , \end{aligned}$$

which means we can connect mated pairs in H_1 and direct all but two of the unmated $S \cup T$ vertices to appropriate conduits.

Therefore, the existence of n disjoint paths will be proven if we can demonstrate that it is always possible to quickgate two unmated $S \cup T$ vertices, avoiding blocks, and that doing so will create a new configuration in $\{A^+_{n+1} \setminus H_1\}$ for which $\phi^*_{max} \leq n-1$. This will be possible for most eligible values of x, and we will deal with the exception as a special case.

There are x unmated s_i's eligible for quickgating, with the x unmated t_i's as potential blocks. Two s_i's cannot quickgate to the same vertex, which implies that they cannot block each other, and also that it takes two t_i blocks to prevent an s_i from quickgating. Thus Q, the number of s_i's which can quickgate without blockage, is:

$$Q \geq x - \lfloor \frac{x}{2} \rfloor .$$

<u>The $x = 2$ subcase</u>: There are only two unmated s_i's in H_1, without loss of generality say s_1 and s_2. The vertex s_1 cannot be blocked from quickgating, as this would require two blocks, and the only potential block is t_2 since t_1 is s_1's mate and thus its destination, not a block. Similarly, t_2 cannot be blocked. If s_1 is sent to a different substar than t_1, and s_2 is sent to a different substar than t_2, then $\phi^*_{max} = 2 \leq n-1$ for $\{A^*_{n+1} \setminus H_1\}$, and the result follows from Lemma 1.

<u>The $3 \leq x \leq n-1$ subcases</u>: $x \geq 3$ implies that $Q = x - \lfloor \frac{x}{2} \rfloor \geq 2$, so it is possible to quickgate two or more unmated vertices from $S \cup T$.

Since $x \leq n-1$, we are guaranteed that the ϕ-value is no more than $n-1$ for any substar other than H_1. Suppose $\phi^*_{max} \leq n-1$ isn't possible after two quickgates. This implies one or more blocks outside of H_1 which force the quickgating s_i's to a substar H^* so that its ϕ-count becomes n or more. This would require the number of blocks plus the number of t_i's in H^* to be at least n, which contradicts $x \leq n-1$.

Hence $\{A^*_{n+1} \setminus H_1\}$ has $\phi^*_{max} \leq n-1$; the result follows from Lemma 1.

<u>The $x = n$ subcase</u>: This implies $S \subset H_1$ and $T \subset \{A^+_{n+1} \setminus H_1\}$.

If the ϕ-count is no more than $n-1$ for any substar other than H_1, then $\phi^*_{max} \leq n-1$ is possible after two quickgates, by the same argument as in the previous subcase.

This leaves only the special subcase where T is contained entirely within one substar; without loss of generality say $S \subset H_1$ and $T \subset H_2$. By Lemma 2, we can quickgate two mated pairs, say (s_1, t_1) and (s_2, t_2), without blockage. We can then use surrogate routing to send every other vertex in $S \cup T$ to any of the other $n-1$ substars we choose.

If $\{s_1, s_2, t_1, t_2\}$ all quickgate to the same substar, say H_3, then we can use surrogate routing to send every other vertex in S to H_4, and every other vertex in T to H_5 (for $n \geq 4$, there are at least 5 substars). The (s_1, t_1) and (s_2, t_2) paths can be easily completed within H_3, and the remaining paths can be completed by Lemma 1, since the new configuration in $H_4 \cup H_5$ has $\phi^*_{max} = n-2 \leq n-1$.

If no more than 3 of $\{s_1, s_2, t_1, t_2\}$ go to the same substar, then we choose surrogate routes for the rest of $S \cup T$ which spread the vertices out as much as possible over the substars of $\{A_{n+1}^{+} \setminus (H1 \cup H_2)\}$. Spreading $2n$ vertices over these $n-1$ substars guarantees that the resulting new ϕ-count has:

$$\phi_{max}^{*} = \min\left\{3, \lceil \frac{2n}{n-1} \rceil\right\} ,$$

which for $n \geq 4$ is guaranteed to be no more than 3. Thus $\phi_{max}^{*} \leq n-1$, and the n disjoint paths can then be completed by Lemma 1.

As we have covered all possible ϕ-counts, the proposition is proven. □

5. Performance Advantages of the Nova Graph

With the proposition proven, it is pertinent to consider the practicality of the result. Akers [2] first postulated numerous advantages to the use of symmetry groups for graph structures in interconnection networks. The alternating group graph AG_n, the split-star graph S_n^2, and now the Nova graph A_n^+ are the most recently-studied graphs of this type. Consider the following table, summarizing the disjoint path performance and several characteristics of these graphs.

Graph	AG_n	S_n^2	A_n^+
Disjoint Paths	$n-2$	$n-1$	$n-1$
Vertices	$\frac{n!}{2}$	$n!$	$\frac{n!}{2}$
Edges	$(n-2)n!$	$(\frac{2n-3}{2})n!$	$(\frac{2n-3}{4})n!$
Diameter	$\lfloor \frac{3(n-2)}{2} \rfloor$	$\lfloor \frac{3(n-2)}{2} \rfloor + 1$	$2 \quad , \quad n=4$ $\lfloor \frac{3(n-2)}{2} \rfloor \quad , \quad n \geq 5$
Vertex symmetry	yes	yes	yes
Edge Symmetry	yes	yes	no

The "J" edges used in constructing the A_n^+ from AG_n prevent the Nova graph from being edge-symmetric, but this is a small sacrifice which yields big dividends. In guaranteeing as many disjoint paths as the split-star, the Nova graph uses half as many vertices and half as many edges. It does so with as small a diameter as the alternating group graph, and in the case of A_4^+ the diameter is even smaller.

6. Conclusion

In this paper we have proven that the Nova graph, A_n^+, has the $(n-1)$-disjoint path property. Our proof is purely an existence proof, that is, we have not provided an algorithm which generates the paths. However, we believe that the generalized algebraic algorithm for finding disjoint paths in AG_n [7] can be adopted to find the disjoint paths in any Cayley graph, including A_n^+. This, combined with the performance advantages of the Nova graph, suggest it is desirable for use as an interconnection network.

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

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

[6] J.Boats, L.Kikas, and J.Oleksik. The nova graph: More disjoint paths with minimal graph augmentation. *Congressus Numerantium*, 184:71–83, 2007.

[7] J.Boats, L.Kikas, and J.Oleksik. An algebraic approach for finding disjoint paths in the alternating group graph. *Journal of Combinatorial Mathematics and Combinatorial Computing*, 64:109–119, 2008.

[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.

Jeffe Boats, Corresponding Author, Department of Mathematics and Computer Science, University of Detroit Mercy, Detroit, MI, 48221-3038, USA
E-mail address: boatsjj@udmercy.edu

Lazaros Kikas, Department of Mathematics and Computer Science, University of Detroit Mercy, Detroit, MI, 48221-3038, USA
E-mail address: kikasld@udmercy.edu

John Oleksik, Department of Mathematics and Computer Science, University of Detroit Mercy, Detroit, MI, 48221-3038, USA
E-mail address: oleksijj@udmercy.edu