

Linearly Many Faults in 2-Tree-Generated Networks

Eddie Cheng and László Lipták

Department of Mathematics and Statistics, Oakland University, Rochester, Michigan 48309

Fred Sala

University of Michigan, Ann Arbor, Michigan 48109

In this article we consider a class of Cayley graphs that are generated by certain 3-cycles on the alternating group A_n . These graphs are generalizations of the alternating group graph AG_n . We look at the case when the 3-cycles form a “tree-like structure,” and analyze its fault resiliency. We present a number of structural theorems and prove that even with linearly many vertices deleted, the remaining graph has a large connected component containing almost all vertices. © 2009 Wiley Periodicals, Inc. NETWORKS, Vol. 55(2), 90–98 2010

Keywords: interconnection network; Cayley graph; alternating group graph; fault resiliency

1. INTRODUCTION

A static interconnection network has point-to-point communication links among the processors. In the design of such a network for parallel computing, one has to decide how to link these processors. Several classes of graphs have been suggested for this purpose. The n -cubes formed the first major class of interconnection networks. The star graph, proposed by [1], has many advantages over the n -cube such as lower degree and smaller diameter. Since its introduction, a number of additional classes have been introduced; the most popular ones are the alternating group graphs [17], their companion graphs, the split-stars [7], and further generalizations such as the (n, k) -star graphs [13] and the arrangement graphs [14]. These interconnection networks have gathered considerable attention with many papers written on them in a wide range of areas such as topological properties, broadcasting issues, fault-tolerant routing, strong connectivity, and vulnerability issues as well as fault-tolerant Hamiltonian properties.

Fault-tolerance can be especially important for interconnection networks, since computers may fail, creating faults in the network. To be reliable, the rest of the network should

stay connected. Obviously, this can only be guaranteed if the number of faults is smaller than the smallest degree in the network. When the number of faults is the smallest degree, a new concept of connectivity is necessary, since the graph may become disconnected. Thus several types of superconnectivity were introduced and examined for various classes of graphs in [4, 7, 9, 11, 12, 16, 20], including superconnectedness and tightly superconnectedness, where only singletons can appear in the remaining network, and restricted connectivity, where each remaining component must have a certain minimum size.

As we increase the number of faults in the graph, it is desirable that most of the network stays connected, with at most a few processors separated from the rest, since then the network will continue to be able to function. Many interconnection networks have been examined in this aspect, when the number of faults is about twice the smallest degree; see [4, 5, 8, 10, 21]. One can even go further and ask what happens when more, even linearly many vertices are deleted. This was examined for the hypercube in [22] and for certain Cayley graphs generated by transpositions in [10], and it was shown that the resulting network will have a large component containing almost all vertices.

We consider the same questions for a class of graphs that is a generalization of the alternating group graph. These graphs are Cayley graphs generated by 3-cycles, introduced in [6]. We show that they have similar properties, i.e., they are superconnected and tightly superconnected; moreover, when linearly many vertices are deleted, the remaining graph stays connected apart from a few vertices. In Section 2 we introduce the necessary terminology, in Section 3 we discuss connectivity and superconnectivity, in Section 4 we consider deleting roughly twice or three times as many vertices as the connectivity, and finally in Section 5 we examine the case of linearly many faults.

2. PRELIMINARIES

We will follow the usual graph terminology, which can be found in [19]. Let $G = (V, E)$ be a graph with vertex set $V = V(G)$ and edge set $E = E(G)$. The complete graph K_n

Received May 2008; accepted December 2008

Correspondence to: L. Lipták; e-mail: liptak@oakland.edu

DOI 10.1002/net.20319

Published online 25 August 2009 in Wiley InterScience (www.interscience.wiley.com).

© 2009 Wiley Periodicals, Inc.




FIG. 1. A connected hypergraph.

is the graph on n vertices in which there is an edge between any two different vertices. A set of edges is called *independent* if no two of them have a common endpoint. Graph G is *k-regular* if the degree of every vertex is k . Deleting a set of vertices $H \subset V$ means removing these vertices and all edges incident to them. A noncomplete graph G is *k-connected* if deleting any $k - 1$ vertices results in a connected graph. A *k-regular* graph is *maximally connected* if it is *k-connected*. A maximally connected *k-regular* graph is *tightly superconnected* if deleting any k vertices results in either a connected graph or a graph with exactly two components, one of which is a single vertex (or *singleton*). Note that in the latter case the deleted vertices are exactly the neighbors of the singleton.

Let Γ be a finite group, and let Δ be a set of elements of Γ such that the identity of the group does not belong to Δ . The *Cayley graph* $\Gamma(\Delta)$ is the directed graph with vertex set Γ with an arc directed from u to v if and only if there is an $s \in \Delta$ such that $u = vs$. The Cayley graph $\Gamma(\Delta)$ is strongly connected if and only if Δ generates Γ . A Cayley graph is always vertex-transitive, so it is maximally arc-connected if it is connected (see [18]); however, it may not be maximally connected (see [15]). If whenever $u \in \Delta$, we also have its inverse $u^{-1} \in \Delta$, then for every arc, the reverse arc is also in the graph. So we can treat this Cayley graph as an undirected graph by replacing each pair of such arcs by an edge.

In this article, we choose the finite group to be the alternating group A_n , the set of even permutations on $[n] := \{1, 2, \dots, n\}$, and the generating set Δ to be a set of 3-cycles. Thus the vertices of the corresponding Cayley graph $A_n(\Delta)$ are the even permutations. To get an undirected Cayley graph, we will assume that whenever a 3-cycle (abc) is in Δ , so is its inverse, (acb) . Since (abc) , (bca) , and (cab) represent the same permutation, the set $\{a, b, c\}$ uniquely represents this 3-cycle and its inverse. So we can depict Δ via a hypergraph with vertex set $[n]$, where a hyperedge of size 3 corresponds to each pair of a 3-cycle and its inverse in Δ .

It is easy to see that the Cayley graph generated by the 3-cycles in Δ is connected if and only if its corresponding hypergraph is connected. Since an interconnection network needs to be connected, we require this hypergraph to be connected; an example is shown in Figure 1. We would like to represent such a hypergraph by a graph on $[n]$ whose edges are obtained by forming a K_3 (triangle) among

vertices a, b, c for every hyperedge $\{a, b, c\}$. Notice that in general this procedure may form extra K_3 's, e.g., for the hypergraph in Figure 1 the corresponding graph becomes the graph shown in Figure 2, in which the K_3 formed by vertices 1, 2, and 4 does not correspond to a hyperedge in the original hypergraph. We will avoid this possibility by considering a simpler case when this graph has a tree-like structure, such as the graph shown in Figure 3, which consists of K_3 's on $\{1, 2, 3\}, \{1, 2, 4\}, \dots, \{1, 2, n\}$. The graph formed by K_3 's corresponding to the 3-cycles of Δ in this manner will be called the *3-cycle generating graph* of $A_n(\Delta)$ or simply its *generating graph* if it is clear from the context.

The alternating group graph AG_n can be viewed as the Cayley graph generated by the graph in Figure 3 having a tree-like (in fact, star-like) structure of triangles. Such a generalization of trees has been studied extensively, first in [2, 3]. A *k-tree* $T_{k,n}$ with n vertices is defined recursively as follows: A set of k pairwise adjacent vertices constitutes a *k-tree* $T_{k,k}$ and a *k-tree* $T_{k,n+1}$ is any graph obtained by joining a new vertex to k pairwise adjacent vertices of a *k-tree* $T_{k,n}$. So the graph in Figure 3 is a 2-tree. (Note that the star graph is generated by a 1-tree—i.e., a tree—that is a star; see [1] for details.) We call a vertex in $T_{k,n}$ with $n \geq k + 1$ the *tail* if it is the last vertex added in the recursion (this may not be uniquely determined just by the picture of $T_{k,n}$). Notice that in a *k-tree*, no $k + 1$ vertices can be pairwise adjacent unless one of those vertices were joined to the other k in one step of the recursion. Hence when the generating graph is a *k-tree*, it uniquely determines Δ , so with a slight abuse of terminology we will identify Δ with the generating graph.

In this article we will consider those Cayley graphs whose generating graphs are 2-trees, so let Δ be a 2-tree on n vertices. Clearly, the number of 3-cycles in Δ is exactly $2n - 4$, so $A_n(\Delta)$ is $(2n - 4)$ -regular on $n!/2$ vertices.

It is easy to prove that if two 2-trees are isomorphic, then the corresponding Cayley graphs will be isomorphic as well; hence without loss of generality we may assume that vertex n is the tail of the 2-tree. Then for $n \geq 4$ the vertices corresponding to even permutations ending with i induce a subgraph H_i that is also a Cayley graph generated by a 2-tree Δ' , which is obtained by deleting the edges corresponding to the two 3-cycles in Δ containing n . Thus we obtain the following easy result concerning the recursive structure of $A_n(\Delta)$:




FIG. 2. A graph that is not a 2-tree.




FIG. 3. The generating graph for AG_n .

Proposition 1. Let $A_n(\Delta)$ be a Cayley graph generated by the 2-tree Δ , and let $\Delta' = \Delta - \{n\}$, $n \geq 4$. Then

- I. $A_n(\Delta)$ consists of n vertex-disjoint subgraphs, H_1, H_2, \dots, H_n , each isomorphic to $A_{n-1}(\Delta')$.
- II. H_i has $(n-1)!/2$ vertices, and it is $(2n-6)$ -regular for all i .
- III. There are exactly $(n-2)!$ independent edges between H_i and H_j for all $i \neq j$.
- IV. Each vertex in H_i has exactly two neighbors outside H_i (which are called its outside neighbors); these two outside neighbors are in different H_k 's, and there is an edge between them. Thus every vertex forms a triangle with its two outside neighbors.

For the rest of this article we assume that the Cayley graph $A_n(\Delta)$ has this structure. Notice that Properties III and IV mean that there are a lot of edges between different H_i 's, and every vertex has neighbors in two other H_i 's. We can use these properties when the faults in the graph concentrate in up to two H_i 's:

Lemma 2. Let $G = A_n(\Delta)$ be a Cayley graph generated by a 2-tree Δ . If we delete vertices in at most two H_i 's in G , then the resulting graph is connected.

Proof. Assume we are deleting vertices in only H_i and H_j , and let G' be the resulting graph. Clearly all the remaining H_k 's, $k \neq i, j$, belong to the same component of G' . Since every vertex in H_i has two neighbors not in H_i , one of which is not in H_j , every remaining vertex in H_i will be adjacent to a vertex in the remaining H_k 's. The same idea applies to H_j , so G' is indeed connected. ■

This lemma will be quite useful in proving our results in the following sections.

3. CONNECTIVITY AND SUPER CONNECTIVITY

We noted that the Cayley graph $A_n(\Delta)$ is connected and vertex-transitive, so it is maximally edge-connected. We will now prove a stronger connectivity property that is also an intermediate step to our next result.

Theorem 3. Let $G = A_n(\Delta)$ be a Cayley graph generated by a 2-tree Δ for $n \geq 4$. Then G is maximally connected.

Proof. The proof will be by induction on n . For $n = 4$, the graph G is the alternating group graph AG_4 , shown in Figure 4. By Proposition 1, AG_4 consists of four triangles H_1, H_2, H_3, H_4 and two independent edges between each pair of triangles. The degree of each vertex is 4, so we must show that removing at most three vertices will result in a connected graph. By Lemma 2, if the deleted vertices are in at most two H_i 's, then the resulting graph is connected. The remaining possibility is that we delete one vertex from each of three triangles, leaving three K_2 's and a triangle. Since originally there are two edges between any two triangles, there is an edge between the complete triangle and every other remaining K_2 , so the graph is connected, and for $n = 4$ the theorem holds.

Assume now that $n \geq 5$ and that the claim is true for $n-1$. Recall that H_i is the subgraph of G spanned by vertices with i in the last position for $1 \leq i \leq n$, and each H_i is $(2n-6)$ -regular and is isomorphic to a Cayley graph generated by a 2-tree on $n-1$ vertices. Let T be a set of vertices in G such that $|T| = 2n-5$. We must show that $G-T$ is connected. Let $T_i = T \cap V(H_i)$ and $t_i = |T_i|$, thus $\sum_{i=1}^n t_i = 2n-5$.

If there are at most two H_i 's containing all the deleted vertices, then $G-T$ is connected by Lemma 2. Otherwise we must have $t_i \leq 2n-7$ for all $1 \leq i \leq n$, hence $H_i - T_i$ is connected for all i by the induction hypothesis. There are $(n-2)!$ independent edges between H_i and H_j for any $i \neq j$, and $(n-2)! > 2n-5$ for $n \geq 5$, so there is at least one edge




FIG. 4. The alternating group graph AG_4 .

remaining between $H_i - T_i$ and $H_j - T_j$ for any $i \neq j$. Thus $G - T$ is connected, and our claim is proven. ■

Since G is $(2n - 4)$ -connected, our next step will be to look at all disconnecting sets in G of size $2n - 4$, and show that they must be neighborhoods of one vertex:

Theorem 4. *Let $G = A_n(\Delta)$ be a Cayley graph generated by a 2-tree Δ . Then G is tightly superconnected for $n \geq 5$. For $n = 4$, if we delete four vertices in AG_4 , then there are three possibilities for the remaining graph: (1) it is connected; (2) it has two components, one of which is a singleton; or (3) it has two components, each of which is a 4-cycle.*

Proof. For $n = 4$ the graph G is just AG_4 , which is not tightly superconnected, because if vertices 2431, 3124, 4213, and 1342 are deleted in Figure 4, the remaining graph consists of two 4-cycles. We show that this is the only exceptional case for $n = 4$. Recall that AG_4 contains four triangles, H_1, H_2, H_3, H_4 , and there are two independent edges between each pair of triangles. If the deleted vertices are in at most two H_i 's, then the remaining graph is connected by Lemma 2. If we delete two vertices from (without loss of generality) H_1 and one from each of H_2 and H_3 , then each of the remaining K_2 's still has an edge to the remaining triangle (since there are two independent edges between each pair of H_i 's). Thus the graph has at most two components, and if it does have two components, one of them is a singleton. Since every degree is 4, the neighbors of this singleton must be the deleted vertices. Finally, we may delete one vertex from each of the H_i 's. In this case, we have four K_2 's remaining, each of which must have an edge remaining to at least one of the other K_2 's, since each K_2 has two neighbors in one of the other H_i . Thus we either get a connected graph, or two components of size 4, which are easily seen to be 4-cycles.

We defer the proof for $n \geq 5$ until the next section, because the result will follow from Theorem 5. ■

In the next section we consider what happens when we delete more vertices.

4. BEYOND SUPERCONNECTEDNESS

When we delete more vertices than the connectivity of the graph, it is still important that we get few vertices disconnected from the rest of the graph. Theorem 4 says that if $2n - 4$ vertices are deleted from $G = A_n(\Delta)$, the resulting graph is either connected or it has two components, one of which is a singleton. It is natural to ask the following question:

What is the largest number K such that if at most K vertices are deleted from $G = A_n(\Delta)$, the resulting graph is either connected or it has two components, one of which is a singleton?

Our next structural theorem answers this question:

Theorem 5. *Let $G = A_n(\Delta)$ be a Cayley graph generated by a 2-tree Δ for $n \geq 4$, and let T be a set of vertices in G*

such that $|T| \leq 4n - 11$. If $n \geq 5$, then $G - T$ satisfies one of the following conditions:

1. $G - T$ is connected.
2. $G - T$ has two components, one of which is a singleton.
3. $G - T$ has two components, one of which is K_2 . Moreover, $|T| = 4n - 11$, and the set T is formed by the neighbors of the two vertices in the K_2 .

When $n = 4$, there are two additional possibilities. In both cases $G - T$ has two components, one of which is a 4-cycle. The other component is either a 4-cycle if $|T| = 4$ or a path of length 3 if $|T| = 5$.

Proof. We will prove the theorem by induction on n . We base our induction on the case $n = 4$, even though our theorem is only valid for $n \geq 5$, because we get an exceptional possibility just as in Theorem 4. Recall that when $n = 4$, the graph G is just AG_4 , shown in Figure 4, composed of four triangles H_1 through H_4 , with two independent edges between each pair of triangles. If we delete up to four vertices, then the claim follows from Theorems 3 and 4. Now assume that we delete five vertices altogether. If the deleted vertices lie in at most two H_i 's, then the remaining graph is connected by Lemma 2. If we delete three vertices from (without loss of generality) H_1 , then H_1 is completely deleted, and we delete one vertex each from (say) H_2 and H_3 . Then there will still be an edge remaining between H_4 and each of the K_2 's remaining from H_2 and H_3 (since originally there are two independent edges between each), so $G - T$ is connected. If we delete two vertices from H_1 , two from H_2 , and one from H_3 , the remaining K_2 from H_3 and the triangle H_4 have an edge between them. By the usual argument, the K_2 and H_4 belong to the same component C in the remaining graph. Thus the two singletons remaining in H_1 and H_2 are in C , or at most one of them is not in C , or both have its outside neighbor in H_3 deleted, in which case they form a K_2 by Property IV. Finally, we have to look at the case where we delete two vertices from H_1 and one from each of H_2, H_3 , and H_4 . The remaining graph consists of a singleton and three K_2 's, with some edges in between them. If the three K_2 's belong to the same component in $G - T$, then $G - T$ is either connected or has two components, one of which is a singleton. Otherwise, just as in the proof of Theorem 4, two of the three remaining K_2 's must belong to the same component forming a 4-cycle. If the singleton from H_1 has an edge to the 4-cycle, then the graph is either connected or has two components, one of which is a K_2 . Finally, if the 4-cycle has no edges to the singleton or the K_2 , then it is easy to see that the singleton and the K_2 must belong to one component, so the graph has two components, a 4-cycle and a path of length 3.

Now let $n \geq 5$, and assume that the result is true for $n - 1$ with the above exceptions for $n - 1 = 4$. As before, let H_i be the subgraph of G spanned by vertices with i in the last position for $1 \leq i \leq n$. Then H_i is isomorphic to a Cayley graph generated by a 2-tree on $n - 1$ vertices. Let $T_i = V(H_i) \cap T$ and $t_i = |T_i|$, so $\sum_{i=1}^n t_i \leq 4n - 11$. We will break up the proof into several cases.

CASE 1. $t_i \leq 2n - 7$ for all $1 \leq i \leq n$.

Since H_i is $(2n - 6)$ -connected, $H_i - T_i$ is connected by Theorem 3 for all i . There are $(n - 2)!$ edges between H_i and H_j , and since $(n - 2)! > 4n - 14 = 2 \cdot (2n - 7)$ for $n > 5$, there is an edge remaining between $H_i - T_i$ and $H_j - T_j$ for all pairs $i \neq j$, so $G - T$ is connected. For the $n = 5$ case we note that if there are no edges left between $H_i - T_i$ and $H_j - T_j$ for some i, j , then $t_i = t_j = 3$, so we can only delete three more vertices in the other H_k 's. Thus the remaining $H_k - T_k$'s belong to the same component in $G - T$, and both $H_i - T_i$ and $H_j - T_j$ will have an edge to $H_k - T_k$ for at least two other $k \notin \{i, j\}$ (whenever $t_k < 3$), hence $G - T$ is connected.

CASE 2. $2n - 6 \leq t_i \leq 4n - 15$ for some i , and $t_j \leq 2n - 7$ for all $j \neq i$.

As in the last case, all the $H_j - T_j$, $j \neq i$, belong to a single large component of $G - T$, say C . From the induction hypothesis, $H_i - T_i$ has at most two components, where the smaller one has up to two vertices (except possibly for $n = 5$). Let B be the largest component of $H_i - T_i$. We will show that B belongs to C . We know that B has at least $(n - 1)!/2 - (4n - 15) - 2$ vertices, which have $(n - 1)! - 2(4n - 15) - 4$ different outside neighbors. Since there are at most $(4n - 11) - (2n - 6) = 2n - 5$ deleted vertices outside H_i , and $(n - 1)! - 2(4n - 15) - 4 > 2n - 5$ for $n \geq 5$, there must be an edge between B and C , so cases (1) through (3) of the statement in the theorem cover this case. When $n = 5$, the same argument applies unless it is the exceptional case. Since $4 \leq t_i \leq 5$, $H_i - T_i$ has two 4-cycles or a 4-cycle and a path of length 3. There are at most five deleted vertices outside H_i , and since every vertex in H_i has two outside neighbors (all distinct), the path of length 3 and the 4-cycle have six and eight different outside neighbors, respectively. Thus $G - T$ is connected.

CASE 3. $4n - 14 \leq t_i$ for some i .

Then $\sum_{j \neq i} t_j \leq 3$. By our previous argument, $H_j - T_j$ belong to a large component C of $G - T$ for all $j \neq i$ (even for $n = 5$). If a component B of $H_i - T_i$ has an edge, then its endpoints have exactly four distinct outside neighbors, so B belongs to C . Thus only singletons of $H_i - T_i$ may not be part of C . Since each singleton of $H_i - T_i$ has two outside neighbors, which are all different, there can be only one such singleton. Hence $G - T$ is either connected or has two components, one of which is a singleton.

So far, all our cases have had at most one i with $t_i \geq 2n - 6$. We cannot have more than two $t_i \geq 2n - 6$, since $3(2n - 6) > 4n - 11$ for $n \geq 5$. So we have one more case to deal with:

CASE 4. $t_i, t_j \geq 2n - 6$ for some $i \neq j$.

Then $\sum_{k \neq i, j} t_k \leq 1$, so all $H_k - T_k$ belong to a large connected component C of $G - T$ for all $k \neq i, j$ (even for $n = 5$). Every vertex of $H_i - T_i$ has at least one neighbor outside

$H_i \cup H_j$, so any nonsingleton component of $H_i - T_i$ is part of C . The same applies to $H_j - T_j$. By the same argument, at most one singleton of $H_i - T_i$ may not be part of C , and similarly for $H_j - T_j$. Thus $G - T$ satisfies one of the conditions in (1)–(3) since by Property IV in Proposition 1 the two singletons cannot both remain singletons in $G - T$.

This covers all possibilities. Moreover, it is obvious that in case (3), if the small component is K_2 , then T must be the neighbors of these two vertices, and thus the proof is complete. ■

Theorem 5 tells us that for $n \geq 5$ the maximum number of vertices we can delete and still always get at most one vertex disconnected from the rest of the graph is $4n - 12$. Moreover, Theorem 5 implies Theorem 4 for $n \geq 5$, since $4n - 12 > 2n - 4$ for $n \geq 5$.

The next question that we may ask is how many vertices we can delete to get at most two vertices disconnected. We have the following result:

Theorem 6. *Let $G = A_n(\Delta)$ be a Cayley graph generated by a 2-tree Δ for $n \geq 5$, and let T be a set of vertices in G such that $|T| \leq 6n - 20$. Then $G - T$ satisfies one of the following conditions:*

1. $G - T$ is connected.
2. $G - T$ has two components, one of which is a singleton.
3. $G - T$ has two components, one of which is K_2 .
4. $G - T$ has three components, two of which are singletons.

Proof. Again, our proof will be by induction on n . For $n = 4$ we delete at most 4 vertices, so by Theorem 4 either $G - T$ has two components, each of which is a 4-cycle, or 1 or 2 holds.

Now let $n \geq 5$ and assume that the result is true for $n - 1$ with the above exception for $n - 1 = 4$. Again, let H_i be the subgraph of G spanned by vertices with i in the last position for $1 \leq i \leq n$. Then H_i is isomorphic to a Cayley graph generated by a 2-tree on $n - 1$ vertices. Let $T_i = V(H_i) \cap T$ and $t_i = |T_i|$, so that $\sum_{i=1}^n t_i \leq 6n - 20$. We will break up the proof into several cases depending on the location of the faults.

CASE 1. $t_i \leq 2n - 7$ for all $1 \leq i \leq n$.

Since H_i is $(2n - 6)$ -connected, $H_i - T_i$ is connected by Theorem 3 for all i . There are $(n - 2)!$ edges between H_i and H_j , and since $(n - 2)! > 4n - 14 = 2 \cdot (2n - 7)$ for $n > 5$, there is an edge remaining between $H_i - T_i$ and $H_j - T_j$ for all pairs $i \neq j$, so $G - T$ is connected. For the $n = 5$ case we note that if there are no edges left between $H_i - T_i$ and $H_j - T_j$ for some i, j , then $t_i = t_j = 3$. Since we are deleting 10 vertices overall, there is an H_k containing at most two deleted vertices; thus this $H_k - T_k$ will have an edge to every other $H_m - T_m$, and hence $G - T$ is connected.

Next we consider those cases when exactly one t_i is larger than $2n - 7$.

CASE 2. $2n - 6 \leq t_i \leq 6n - 26$ for some i and $t_j \leq 2n - 7$ for all $j \neq i$.

Again, all the $H_j - T_j$'s are connected for $j \neq i$, and they belong to a single connected component C in $G - T$. By the induction hypothesis, there are at most three components in $H_i - T_i$, with two of them having at most two vertices in total (except for $n = 5$). Let D be the largest component in $H_i - T_i$. It has at least $\frac{(n-1)!}{2} - (6n - 26) - 2$ vertices, which have at least $(n - 1)! - 2(6n - 24)$ distinct outside neighbors. We have deleted at most $6n - 20 - (2n - 6) = 4n - 14$ of these neighbors, and the rest of them all belong to C , hence there will be at least one edge between the vertices in D and those in C outside H_i , since $(n - 1)! - 2(6n - 24) > 4n - 14$ for $n \geq 5$. Thus D belongs to C , and we are done. In the exceptional case $H_i - T_i$ consists of two 4-cycles, each of which has eight different outside neighbors. Since we delete at most six vertices outside H_i , both 4-cycles must belong to C as well, and we are done.

CASE 3. $6n - 25 \leq t_i$ for some i .

Then $\sum_{j \neq i} t_j \leq 5$. If there is a $j \neq i$ such that $t_j = 5$, then $G - T$ is connected by Lemma 2. Otherwise, by our previous argument either all of the $H_j - T_j$'s with $j \neq i$ will belong to one connected component C of $G - T$, or there is a $j \neq i$ such that $H_j - T_j$ has at least two components. Assume first that all $H_j - T_j$'s with $j \neq i$ belong to C . Since every vertex of $H_i - T_i$ has two outside neighbors, and these neighbors are all different, there can be at most two vertices of $H_i - T_i$ not belonging to C , and we are done. On the other hand, if there is a $j \neq i$ such that $H_j - T_j$ has at least two components, then we must have $t_j = 4$ and $n = 5$, and the remaining H_k 's still belong to a large connected component C of $G - T$. Every component of $H_i - T_i$ and $H_j - T_j$ of size at least two will also belong to C , since it will have at least two distinct neighbors in $G - H_i - H_j$, and we delete at most one vertex outside H_i and H_j . By the same reason, there can be at most one singleton each in $H_i - T_i$ and $H_j - T_j$ not belonging to C , and we are done.

Next we still have to consider the case when there are two t_i 's greater than $2n - 7$. Note that it is impossible for three t_i 's to be greater than $2n - 7$, since $6n - 18 > 6n - 20$.

CASE 4. $t_i \geq t_j \geq 2n - 6$ for some $i \neq j$.

Then clearly $t_j \leq t_i \leq 4n - 14$. If $t_i = 4n - 14$, then $t_j = 2n - 6$, and we are done by Lemma 2. Otherwise we have $t_i \leq 4n - 15$. First consider $t_i = 4n - 15$. If $t_j = 2n - 5$, then we are again done by Lemma 2, so we are left with the case $t_j = 2n - 6$ and $t_k = 1$ for some $k \notin \{i, j\}$. For all $w \notin \{i, j, k\}$ we have $t_w = 0$, and thus $H_k - T_k$ and every $H_w - T_w$ with $w \notin \{i, j, k\}$ belong to a large connected component C of $G - T$. Since every vertex in H_i and H_j has two outside neighbors in different H_z 's, and there is one deleted vertex outside $H_i \cup H_j$, in each of them at most one vertex can be disconnected from C in $G - T$, and we are done.

Finally, we consider those cases where $t_j \leq t_i \leq 4n - 16 = 4(n - 1) - 12$. Then by Theorem 5, $H_i - T_i$ and $H_j - T_j$ have at most one singleton and one large component with the exception for $n = 5$, since case (3) cannot occur. For all $k \notin \{i, j\}$ we have $t_k \leq 2n - 7$, so as before, all these $H_k - T_k$'s belong to a large component C of $G - T$. Assume first that $n = 5$. Then $t_i = t_j = 4$, and we have two deleted vertices outside $H_i \cup H_j$. Since every vertex in H_i or H_j has two outside neighbors in different H_k 's, each of $H_i - T_i$ and $H_j - T_j$ can have at most two vertices not belonging to C . Thus if either $H_i - T_i$ or $H_j - T_j$ has two 4-cycles as components, the 4-cycles must belong to C . Thus each of $H_i - T_i$ and $H_j - T_j$ can have at most one singleton disconnected from C , and we are done. When $n \geq 6$, the deletion of vertices can remove at most $2(6n - 20)$ edges of the $(n - 2)(n - 2)!$ edges that H_i or H_j has to vertices in the other H_k 's ($k \neq i, j$). Since two edges may be incident to the singleton in $H_i - T_i$ or $H_j - T_j$, and $(n - 2)(n - 2)! > 2(6n - 20) + 2$ for $n \geq 6$, the large components in $H_i - T_i$ and $H_j - T_j$ must be part of C . The two possible singletons in $H_i - T_i$ and $H_j - T_j$ may either remain singleton components in $G - T$, or they could belong to one component forming a K_2 , or they could also belong to C . In each case we are finished, and the proof of the theorem is complete. ■

Note that Theorem 6 is sharp, since we can take three vertices of a 4-cycle and delete all their neighbors, altogether $6n - 19$ vertices, creating a big connected component and a path of length 3. This shows that we can delete up to $6n - 20$ vertices without creating three vertices disconnected from the rest of the graph, and this is best possible.

So far we have proved Theorem 3 ($A_n(\Delta)$ is maximally connected), Theorem 5 (structural theorem when we have roughly $2(2n - 4)$ faults in $A_n(\Delta)$), and Theorem 6 (structural theorem when we have roughly $3(2n - 4)$ faults in $A_n(\Delta)$). All the proofs have a similar theme, using induction and case analysis based on the distribution of faults. Although the proofs are not particularly difficult, care must be taken in the case analysis. A fair question is to ask whether separate proofs are necessary. Could we prove Theorem 6 right from the start? The answer is no. In Case 3 of Theorem 6 we use Theorem 5, and in the proofs of both Theorem 5 and Theorem 6 we use Theorem 3.

A natural question to ask next is what happens if we delete roughly $4(2n - 4)$ faults. Indeed, such a theorem can be established. One can easily come up with the form of the statement; the difficult part is what "roughly" means, that is, what should p be if we allow $4(2n - 4) - p$ faults. Naturally, one would want the bound to be sharp. However, this may cause the result to be false for small n . Given that induction is the technique of choice here, one may have to overcome the difficulty of checking the base case for a relatively large n . Even if we establish such a theorem here, the next question is to replace 4 by 5. This leads us to consider deleting $k(2n - 4) - f(k)$ vertices for every k . Of course, compromises must be made. The first is that we will no longer describe the possibilities of the structures of the "small components" in the resulting

graphs when vertices are deleted. Instead we will focus on the total number of vertices in the small components. The second is that our bound will be good only asymptotically, that is, we will not be overly aggressive in choosing $f(k)$, to ensure that the statement is true for small n . This result will be presented in the next section.

5. LINEARLY MANY FAULTS

In this section we look at what happens when the number of faults is linear in n . Clearly, we can arbitrarily choose k vertices in the Cayley graph $A_n(\Delta)$, and isolate them by deleting all their neighbors. This gives us k singletons and (probably) a large connected component in the remaining graph by deleting at most $k(2n - 4) = 2kn - 4k$ vertices. Thus, we look for a result saying that if you delete $2kn - f(k)$ vertices, where $f(k)$ is a constant depending only on k , then the resulting graph has a large connected component and small components with up to $k - 1$ vertices in total. Finding the best value of $f(k)$ becomes more difficult as k gets bigger, and is not worth the effort, so we will only aim for an asymptotically sharp result.

We have already observed that the proof of Theorem 6 ($k = 3$) depends on Theorem 5 ($k = 2$) and Theorem 3 ($k = 1$). Moreover, in each of these proofs induction on n is used. So it is natural to use double induction to prove our next result, first on k (strong form), then on n . Note that the proof is similar to the proof of the main result in [10].

Theorem 7. *Let $G = A_n(\Delta)$ be a Cayley graph generated by a 2-tree Δ for $n \geq 4$, and let T be a set of vertices in G such that $|T| \leq k(2n - 4) - 2k(k - 1) - 1$, where $k \geq 1$. Then $G - T$ has at most k components: one large (connected) component and small components containing up to $k - 1$ vertices in total.*

Proof. We will use induction on k . For $k = 1$, we delete up to $2n - 5$ vertices, so the claim follows from the fact that G is $(2n - 4)$ -regular and maximally connected. If $k = 2$ or $k = 3$, we delete up to $4n - 13$ or $6n - 25$ vertices, respectively, so the claim holds by Theorem 5 and Theorem 6 (even for $n = 4$).

Let $k' \geq 4$, and assume that the claim is true for all $k < k'$ and all $n \geq 4$. We will prove the claim for k' . If $n \leq k' + 1$, then we are not deleting any vertices ($|T| < 0$), so the claim automatically holds. Thus, we can assume that $n \geq k' + 2$ and that T is a set of vertices in G with $|T| \leq k'(2n - 4) - 2k'(k' - 1) - 1$. We must show that $G - T$ has one large connected component and small components containing at most $k' - 1$ vertices in total. For all $1 \leq i \leq n$, let H_i be defined as before, and let $T_i = T \cap V(H_i)$ and $t_i = |T_i|$. We will look at cases depending on the distribution of faults in the H_i 's, beginning with the case when we remove many vertices from a single H_i :

CASE 1. $t_i \geq k'(2n - 6) - 2k'(k' - 1)$ for some $1 \leq i \leq n$.

Then $\sum_{j \neq i} t_j = |T| - t_i \leq 2k' - 1 \leq 2n - 5$. We know that H_j is maximally connected and $(2n - 6)$ -regular, so if we also

have $t_j \leq 2n - 7$ for all $j \neq i$, then $H_j - T_j$ stays connected for all $j \neq i$. Also, there are $(n - 2)!$ independent edges between each pair of H_j 's, and $(n - 2)! \geq (k')! > 2k' - 1$ for $k' \geq 4$, hence at least one of these edges remains in $G - T$ for every pair of H_j 's, so all the graphs $H_j - T_j$ ($j \neq i$) belong to a large component C of $G - T$. Now, each vertex in $H_i - T_i$ has two outside neighbors lying in different H_j 's. Two vertices in H_i cannot have a common neighbor outside of H_i , since each vertex has exactly two outside neighbors in different H_k 's. This implies that there are at most $k' - 1$ vertices not belonging to C .

On the other hand, if there is a j such that $t_j \geq 2n - 6$, then there is at most one deleted vertex outside $H_i \cup H_j$, so clearly all H_k 's, $k \neq i, j$, belong to the same component of $G - T$. Since every vertex in $H_i \cup H_j$ has at least one outside neighbor not in $H_i \cup H_j$, there can be at most two vertices not belonging to the big component and the claim is proven.

CASE 2. $t_i \leq k'(2n - 6) - 2k'(k' - 1) - 1$ for all $1 \leq i \leq n$.

Now we will need a second induction on n . We have $t_i \leq k'(2k' - 2) - 2k'(k' - 1) - 1 < 0$ for $n = k' + 2$, so this case is trivial. When n is $n' = k' + 3$, we have

$$\begin{aligned} t_i &\leq k'(2k') - 2k'(k' - 1) - 1 = 2k' - 1 \\ &< 4k' - 5 = (k' - 1)(2n' - 6) - 2(k' - 1)(k' - 2) - 1 \end{aligned}$$

for all i , where $n \geq 7$. Using the induction hypothesis on k with $k \leq k' - 1 = (n' - 2) - 1$ on the subgraph H_i , we get that if $t_i \leq k(2n - 6) - 2k(k - 1) - 1$, then $H_i - T_i$ has one large component and small components containing at most $k - 1$ vertices in total. If $k = k' - 1$, each $H_i - T_i$ has one large component and small components with at most $k' - 2$ vertices in total. Thus, if $H_i - T_i$ has a large component and small components containing at least k vertices in total, then $t_i \geq k(2n' - 6) - 2k(k - 1)$.

To prove the base case for n , we will have to show that for all $H_i - T_i$, their largest components stay all connected, belonging to a single large component, and that the number of total vertices in the small components does not exceed $k' - 1$. We will start with the first claim:

There are $(n' - 2)!$ independent edges between H_i and H_j for every pair (i, j) with $1 \leq i, j \leq n$. If this number of edges is larger than the number of deleted vertices plus the number of vertices in the small components, then the large components of H_i and H_j will be part of the same component. Thus, we require

$$(n' - 2)! > k'(2n' - 4) - 2k'(k' - 1) - 1 + 2(k' - 2).$$

For $n' = k' + 3$, this is equivalent to $(k' + 1)! > (6k' - 5)$, which is true for $k' \geq 4$. Thus, the largest components of the $H_i - T_i$'s are part of the same component in $G - T$. Furthermore, as n' gets larger than $k' + 3$, the left-hand side of the inequality grows much faster than the right-hand side, so the conclusion remains true for $n' > k' + 3$ as well.

Now we will prove that the number of vertices in the small components of all $H_i - T_i$ is no larger than $k' - 1$. We will prove

the claim by contradiction. Assume that the total number of vertices in the small components is at least k' . Then we can choose nonnegative integers $0 \leq k_1, k_2, \dots, k_{n'} \leq k' - 2$ such that $k' = \sum_{i=1}^{n'} k_i$ and for each i , $H_i - T_i$ has one large component and at least k_i vertices in total in the small components. As shown earlier, this implies that $t_i \geq k_i(2n' - 6) - 2k_i(k_i - 1)$ for all i . We know that at least two of the k_i 's are positive, since $k_i \leq k' - 2$ for all i . Consider H_i and H_j with $i \neq j$ and $k_i, k_j > 0$. Together, they have at least $k_i + k_j$ vertices in the small components, and

$$\begin{aligned} t_i + t_j &\geq (k_i + k_j)(2n' - 6) - 2k_i(k_i - 1) - 2k_j(k_j - 1) \\ &\geq (k_i + k_j)(2n' - 4) - 2(k_i + k_j)(k_i + k_j - 1) \end{aligned}$$

since the latter inequality simplifies to $(2k_i - 1)(2k_j - 1) \geq 1$, which is true. This number is at least $(k_i + k_j)(2n' - 6) - 2(k_i + k_j)(k_i + k_j - 1)$, which is of the form $(\square)(2n' - 6) - 2(\square)(\square - 1)$ (here \square represents an expression). Hence we can apply this repeatedly for each positive k_i to get

$$\begin{aligned} |T| &= \sum_{i=1}^{n'} t_i \geq \left(\sum_{i=1}^{n'} k_i \right) (2n' - 4) \\ &\quad - 2 \left(\sum_{i=1}^{n'} k_i \right) \left(\left(\sum_{i=1}^{n'} k_i \right) - 1 \right) \\ &= k'(2n' - 4) - 2k'(k' - 1), \end{aligned}$$

a contradiction. Thus, the base case $n' = k' + 3$ is proven.

Now we have to consider the case $n' > k' + 3$. Let $1 \leq i \leq n$. If $t_i \leq (k' - 1)(2n' - 6) - 2(k' - 1)(k' - 2) - 1$, we can apply the induction hypothesis with $k = k' - 1$ and all admissible n as we did before. If $(k' - 1)(2n' - 6) - 2(k' - 1)(k' - 2) - 1 < t_i \leq k'(2n' - 6) - 2k'(k' - 1) - 1$, then we can use the induction hypothesis with $k = k'$ and $n = n' - 1$. In both cases, $H_i - T_i$ has one large component and small components containing at most $k' - 1$ vertices in total. Then just as before, for each $k \leq k' - 1$, if $H_i - T_i$ has one large component and small components containing at least k vertices in total, then $t_i \geq k(2n' - 6) - 2k(k - 1)$. The remainder of the argument is the same, and our proof is complete. ■

6. CONCLUDING REMARKS

We have shown that the Cayley graph $A_n(\Delta)$ generated by the 2-tree Δ has good structural properties: for any $k \geq 1$, if we delete asymptotically kn vertices, the remaining graph has a large connected component and small components containing at most $k - 1$ vertices in total. This means that networks based on these graphs are extremely fault-resistant, since deleting roughly kn vertices creates at most $k - 1$ vertices disconnected from the rest of the graph.

Note that in our proofs we only used Properties I–IV in Proposition 1 about the structure of these Cayley graphs, so our results extend to every graph that is built up according to those properties starting from AG_4 (in fact, the H_i 's do not even need to be isomorphic). This also means that in

the recursive definition of 2-trees one could allow adding a 3-cycle (abc) and its inverse whenever exactly two of a , b , and c occur in previous 3-cycles, and the same result would follow. We chose to keep the more restrictive definition to keep the nice correspondence between the 3-cycles in Δ and the triangles in the generating graph. A natural question is then what happens if we delete more vertices. How many vertices do we need to delete to create a component of size n ? Clearly we can delete all neighbors of any n vertices, thus deleting $n^2 - O(n)$ vertices is always enough, but is this best possible asymptotically?

Acknowledgments

The authors thank the editor and the anonymous referees for their helpful comments.

REFERENCES

- [1] S.B. Akers, D. Harel, and B. Krishnamurthy, The star graph: An attractive alternative to the n -cube, Proceedings of the International Conference on Parallel Processing, University Park, PA, 1987, pp. 393–400.
- [2] L.W. Beineke and R.E. Pippert, The number of labeled k -dimensional trees, J Combinatorial Theory 6 (1969), 200–205.
- [3] L.W. Beineke and R.E. Pippert, Properties and characterizations of k -trees, Mathematika 18 (1971), 141–151.
- [4] E. Cheng and M.J. Lipman, Increasing the connectivity of split-stars, Congressus Numerantium 146 (2000), 97–111.
- [5] E. Cheng and M.J. Lipman, Increasing the connectivity of the star graphs, Networks 40 (2002), 165–169.
- [6] E. Cheng, M.J. Lipman, L. Lipták, and D. Stiebel, Hamiltonian connectivity of 2-tree-generated networks, Math Comput Model 48 (2008), 787–804.
- [7] E. Cheng, M.J. Lipman, and H.A. Park, Super connectivity of star graphs, alternating group graphs and split-stars, Ars Combinatoria 59 (2001), 107–116.
- [8] E. Cheng and L. Lipták, Structural properties of Cayley graphs generated by transposition trees, Congressus Numerantium 180 (2006), 81–96.
- [9] E. Cheng and L. Lipták, Fault resiliency of Cayley graphs generated by transpositions, Int J Found Comput Sci 18 (2007), 1005–1022.
- [10] E. Cheng and L. Lipták, Linearly many faults in Cayley graphs generated by transposition trees, Inform Sci 177 (2007), 4877–4882.
- [11] Y.-C. Chen and J.J.M. Tan, Restricted connectivity for three families of interconnection networks, Appl Math Comput 188 (2007), 1848–1855.
- [12] Y.-C. Chen, J.J.M. Tan, L.-H. Hsu, and S.-S. Kao, Super-connectivity and super-edge-connectivity for some interconnection networks, Appl Math Comput 140 (2003), 245–254.
- [13] W.K. Chiang and R.J. Chen, The (n, k) -star graph: A generalized star graph, Inform Proc Lett 56 (1995), 259–264.
- [14] K. Day and A. Tripathi, Arrangement graphs: A class of generalized star graphs, Inform Proc Lett 42 (1992), 235–241.

- [15] C.D. Godsil, GRR's for non-solvable groups, *Coll Math Soc János Bolyai* 25 (1981), 221–229.
- [16] A. Hellwig, D. Rautenbach, and L. Volkmann, Cuts leaving components of given minimum order, *Discrete Math* 292 (2005), 55–65.
- [17] J.S. Jwo, S. Lakshmivarahan, and S.K. Dhall, A new class of interconnection networks based on the alternating group, *Networks* 23 (1993), 315–326.
- [18] *Combinatorial problems and exercises*, Elsevier and Akadémiai Kiadó, Budapest, 1993.
- [19] D.B. West, *Introduction to graph theory*, Prentice Hall, Upper Saddle River, NJ, 1996.
- [20] X. Yang, D.J. Evans, B. Chen, G.M. Megson, and H. Lai, On the maximal connected component of hypercube with faulty vertices, *Int J Comput Math* 81 (2004), 515–525.
- [21] X. Yang, D.J. Evans, and G.M. Megson, On the maximal connected component of hypercube with faulty vertices II, *Int J Comput Math* 81 (2004), 1175–1185.
- [22] X. Yang, D.J. Evans, and G.M. Megson, On the maximal connected component of a hypercube with faulty vertices III, *Int J Comput Math* 83 (2006), 27–37.