Strong local diagnosability of (n, k)-star graphs and Cayley graphs generated by 2-trees with missing edges

Eddie Cheng, László Lipták and Daniel E. Steffy

Department of Mathematics and Statistics
Oakland University, Rochester, MI 48309, USA
echeng@oakland.edu, liptak@oakland.edu, steffy@oakland.edu

Abstract

In this paper, we study the local diagnosability and strong local diagnosability properties for (n, k)-star graphs and Cayley graphs generated by 2-trees. Moreover, we also consider the corresponding problem with missing edges.

Keywords: (n, k)-star graphs, Cayley graphs, 2-trees, local diagnosability

1. Introduction

Large scale multiprocessor computing systems are ubiquitous, and many models have been proposed for the underlying network structure. One important characteristic of a network is that its structure facilitates the identification, or diagnosis, of faulty processors in the system. Ideally a precise diagnosis can be made, identifying exactly which processors have developed faults; this property is generally referred to as *diagnosability*. A variety of models have been proposed for measuring diagnosability and differ in the mechanism by which processors are tested for faults. Our study considers diagnosability under the comparison model of Maeng and Malek [1] and the corresponding notion of local diagnosability introduced by Hsu and Tan [2]. The classes of (n, k)-star graphs and Cayley graphs generated by 2-trees are analyzed and it is proven that they both have the strong local diagnosability property. Additionally, these networks are shown to still satisfy the strong local diagnosability property after a number of communication links are removed, as long as the number of links removed does not exceed some specified bounds.

2. Preliminaries and local diagnosability

A network is modeled as an undirected graph $G = (V, E)$, following standard terminology, with each processor represented as a node and each communication link represented as an undirected edge. A number of models have been proposed for diagnosing faulty processors in a network, and we follow the comparison model of Maeng and Malek [1]. Under this model, diagnosis is performed by carrying out the following operations. Each node w can send identical test signals to any pair u and v of its neighbors. It then compares their responses and returns failure if the results are different and success if they are the same. If w itself is faulty, its result is unreliable and may return either failure or success, regardless of the status of u, v. The following assumptions are made: If w returns success it is assumed that, unless w is faulty itself, both u and v are fully functional; if the responses are different it can be concluded that at least one of w, u or v is faulty.

After a set of tests has been performed, the result is referred to as a *syndrome*. Each syndrome can be represented by a function $\sigma : C \to \{0, 1\}$ where C is the set of all comparisons made, each of which can be indexed as $(u, v)_w$, denoting a comparison of u and v performed from w as described above. The value $\sigma((u, v)_w)$ is defined to be 1 if the comparison of u and v by w returned failure, and it is defined to be 0 if w returned success.

For a set of faulty processors $F \subseteq V$, we say that a syndrome σ is *consistent* with F if σ is a possible result of a test performed on the network with these faulty processors. For $F \subseteq V$ we define σ_F to be the set of all syndromes consistent with F. Note that syndromes may not correspond uniquely to sets of faulty nodes, we might have $\sigma \in \sigma_{F_1} \cap \sigma_{F_2}$ where $F_1 \neq F_2$, in such cases this test result would not be enough to confidently diagnose the faulty nodes. A graph is said to be *t-diagnosable* if for every $F, F' \subseteq V$ with $|F|, |F'| \leq t$, we have $\sigma_F \cap \sigma_{F'} = \emptyset$ whenever $F' \neq F$.

The notion of local diagnosability, introduced by Hsu and Tan [2], considers diagnosability locally on a node-by-node basis; this is especially relevant for networks with varying structure, or for studying a network with removed edges. A node v is *locally t-diagnosable* if for every $F, F' \subseteq V$ with $|F|, |F'| \leq t$ such that $v \in F$, the condition $\sigma_F \cap \sigma_{F'} \neq \emptyset$ implies that $v \in F'$. Node v has the *strong local diagnosability* property if v is locally $\deg_G(v)$-diagnosable, where $\deg_G(v)$ denotes the degree of node v in graph G. A graph G has the *strong local diagnosability* property if each node v in G is locally $\deg_G(v)$-diagnosable. It is shown in [2] that this notion generalizes the global notion of t-diagnosability in the following sense: a system $G = (V, E)$ is t-diagnosable if and only if G is locally t-diagnosable at every node.

The following proposition of [3] gives a sufficient condition for a node being n-diagnosable. Before we state this, we need a definition. Given a star $K_{1,r}$ centered at node v, replace each edge by a path of length 4. This is called an *extended star of order* r *at node* v, and it is denoted by $ES(v; r)$ as in [4]. The following proposition is a crucial part of our study.

Proposition 2.1 ([3])**.** *Let* v *be a node in* $G = (V, E)$*. If there exists an extended star of order* $\deg_G(v)$ *at node* v *in* G*, then* v *is locally* $\deg_G(v)$*-diagnosable.*

3. Diagnosability of (n, k)-star graphs

The (n, k)-star graph [5], denoted by $S_{n,k}$, is defined for positive integers n and k such that $n > k \geq 2$. The node set is all of the permutations on k elements of the set $\{1, 2, \ldots, n\}$. Two nodes corresponding to the permutations $[a_1, a_2, \ldots, a_k]$ and $[b_1, b_2, \ldots, b_k]$ are adjacent if and only if either:

(1) There exists an integer $2 \leq s \leq k$ such that $a_1 = b_s$ and $b_1 = a_s$ and for any $i \neq s$, $1 < i \leq k$, we have $a_i = b_i$. (That is, $[b_1, b_2, \ldots, b_k]$ is obtained from $[a_1, a_2, \ldots, a_k]$ by swapping a_1 and a_s.)

(2) For all i such that $2 \leq i \leq k$, we have $a_i = b_i$ and $a_1 \neq b_1$. (That is, $[b_1, b_2, \ldots, b_k]$ is obtained from $[a_1, a_2, \ldots, a_k]$ by replacing a_1 by an element in $\{1, 2, \ldots, n\} - \{a_1, a_2, \ldots, a_k\}$.)

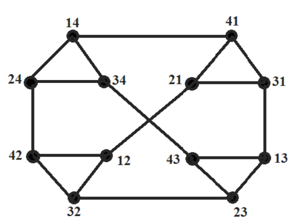

Figure 1: $S_{4,2}$

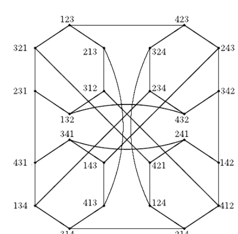

Figure 2: $S_{4,3}$

See Figures 1 and 2 for $S_{4,2}$ and $S_{4,3}$, respectively. (For simplicity, we use 123 instead of $[1, 2, 3]$ to represent the corresponding permutation.)

Let the set of nodes representing permutations whose kth element is i be H_i for $2 \leq i \leq n$. It can easily be seen that $S_{n,k}$ is $(n-1)$-regular as each node has $k-1$ neighbors by adjacency rule (1) and $n-k$ neighbors by adjacency rule (2). Let us first note some other preliminary facts about $S_{n,k}$: H_i is isomorphic to $S_{n-1,k-1}$ when $n > k > 2$; and for each pair H_i and H_j, there are exactly $\frac{(n-2)!}{(n-k)!}$ independent edges between them.

In order to determine the local diagnosability, the key is to apply Proposition 2.1. Two lemmas are needed, one for local diagnosability and one for local diagnosability with missing edges. However, one is simply a special case of the next result when no edges are missing. The ideal statement is: *Let $n \geq 4$ and $2 \leq k \leq n-1$. If F_e is an arbitrary set of missing edges with $|F_e| \leq n-3$, then for every node v in $S_{n,k}$, there exists an extended star of order $\deg_{S_{n,k}-F_e}(v)$ at v in $S_{n,k} - F_e$.* But this statement is not true, in fact it is easy to see that if G is r-regular with a p-clique then the maximum number of faulty edges we can allow is $r-p$. (See Figure 3. After the two red edges are deleted, there is no extended star of order 6 at the square vertex.)

Lemma 3.1. *Let $n \geq 4$ and $2 \leq k \leq n-1$ such that $(n,k) \neq (4,2)$. If F_e is an arbitrary set of missing edges in $S_{n,k}$ with $|F_e| \leq k-2$, then for every node v in $S_{n,k}$, there exists an extended star $ES(v; \deg_{S_{n,k}-F_e}(v))$ of order $\deg_{S_{n,k}-F_e}(v)$ at v in $S_{n,k} - F_e$.*

Proof. $S_{4,2}$ is too small to have the required extended star. We verify the result for $S_{5,2}$ and $S_{4,3}$. $S_{5,2}$ is easy as the set F_e must be empty and checking this property for any single node v is sufficient since $S_{5,2}$ is node transitive. For $S_{4,3}$ we note that, although it can also easily be verified by hand, $S_{4,3} \cong S_4$ and Lemma 7 of [4] establishes precisely the same result for S_n with $n \geq 4$ as long as the set of faulty edges has size at most $n-3$.

Next consider the case $k = 3$, when $|F_e| \leq 1$. We will use a less well-known decomposition on $S_{n,k}$ applied to $S_{n,3}$. The node set of $S_{n,3}$ can be partitioned into four sets as follows. The first set, denoted by W_{n-1}, is the set of nodes in $S_{n,3}$ with n in the first position. Clearly W_{n-1} forms an independent set, and it has $(n-1)(n-2)$ nodes. For $i = 2, 3$, let Y_i be the set of nodes with symbol n in position i. These are the second and third sets, and each induces a subgraph isomorphic to $S_{n-1,2}$. Finally, the fourth set is the set of nodes without the symbol n. This last set induces a subgraph isomorphic to $S_{n-1,3}$, and each of its nodes has exactly one neighbour in W_{n-1}. Moreover, if each set of nodes in the partition is identified to a single node, the resulting graph is a $K_{1,3}$ with the center obtained from W_{n-1}. We now apply induction on n. We have already checked the result for $S_{4,3}$, so consider $S_{n,3}$ with $n \geq 5$. Since $S_{n,3}$ is node-transitive, we may assume that v is in the subgraph H isomorphic to $S_{n-1,3}$. Since this subgraph contains at most $k-2=1$ element of F_e, we can apply the induction hypothesis to obtain an extended star A in H. Let y be the unique neighbour of v in W_{n-1}. If $(v,y) \in F_e$, then we are done, so we may assume that $(v,y) \notin F_e$. Now y has exactly two neighbours, one in Y_1 and the other in Y_2. Since $|F_e| \leq 1$, it is clear A can be extended to a desired extended star.

The case $k=2$ can be done similarly, and in fact it is easier since $|F_e|=0$, so there are no faults.

We proceed with induction on k, so assume $k \geq 4$. Let H_i be the subgraph of $S_{n,k}$ with i in the kth position for $1 \leq i \leq n$. Then each H_i is isomorphic to $S_{n-1,k-1}$, and every node in H_i has exactly one neighbour not in H_i. For notational convenience, we may assume that v is in H_n. Let e be the edge between v and the node y obtained by swapping the symbols in the first position and the kth position of v. So y is in H_j for some $j \neq n$. We consider two cases.

Case 1: H_n contains at most $k-3$ elements of F_e.
We apply the induction hypothesis to obtain an extended star A of order $\deg_{H_n-F_e}(v)$ at v in $H_n - F_e$. If $e \in F_e$, then A is our desired extended star. Suppose $e \notin F$. Then this extended star is incomplete as it does not contain the node w, which is in H_j. Since $|F_e| \leq k-2$, the subgraph $H_j - F_e$ is connected as H_j has edge-connectivity $n-2$. So we can find a 3-path in H_j starting at y. Attach this 3-path to A via e and we obtain the desired extended star.

Case 2: H_n contains all the elements of F_e.
Let f be an arbitrary element of F_e and let $F'_e = F_e - \{f\}$. Apply the induction hypothesis to obtain an extended star A of order $\deg_{H_n-F'_e}(v)$ at v in $H_n - F'_e$. If A does not contain f or if f is incident with v (so $\deg_{H_n-F_e}(v) = \deg_{H_n-F'_e}(v) - 1$), then we can complete the proof as in Case 1 to get the desired extended star. So we may assume otherwise. Now A contains $f = (w, w')$, where w is closer to v than w' is to v in A. Let p be the distance between w and v in A. So $1 \leq p \leq 3$. Node w is adjacent to a node z in $S_{n,k}$ where z is not in H_n. Since H_n contains all the elements of F_e, we have $(w,z) \notin F_e$. Consider two subcases. The first is when z is not in H_j, say z is in $H_i = H_i - F_e$. Find a $(4-p-1)$-path starting at z in $H_i = H_i - F_e$ together with a 3-path starting at y in $H_j = H_j - F_e$ as in Case 1. The second case is when z is in H_j. Note that $y \neq z$ as it is not possible for both v and w to be adjacent to the same node in H_j (since every node can have only one "outside" neighbour.) We are done if we can find two node-disjoint paths in H_j, one being a 3-path starting at y and one being a $(4-p-1)$-path

starting at w. One can easily check that this claim is true for $S_{n,3}$ since H_j contains $S_{n,3}$ as a subgraph, so we are done. (Note that $4 - p - 1 \leq 2$, so these two paths have at most eight nodes.) $\square$

We are now ready to obtain our main result for $S_{n,k}$ using Proposition 2.1 and Lemma 3.1. (Note that taking $F_e = \emptyset$ in the following theorem gives the corresponding result with no missing edges.)

Theorem 3.2. *Let $n \geq 4$ and $2 \leq k \leq n-1$ with $n + k \geq 7$. If F_e is an arbitrary set of missing edges with $|F_e| \leq k - 2$, then for each node v in $S_{n,k}$ with missing edges F_e, node v has the strong local diagnosability property in $S_{n,k} - F_e$, hence $S_{n,k} - F_e$ has the strong local diagnosability property.*

4. Diagnosability of Cayley graphs generated by 2-trees

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 node 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 node-transitive. 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 arcs by an edge.

Here, we choose the finite group to be the alternating group Γ_n, the set of even permutations on $\{1, 2, \ldots, n\}$, and the generating set Δ to be a set of 3-cycles. In this paper, a permutation may be recorded and referenced either via its cycle decomposition or as a rearrangement of symbols. For example, $(12)(34)(567)(8)$ is the cycle decomposition of the permutation whose rearrangement representation is $[2, 1, 4, 3, 6, 7, 5, 8]$. Thus the nodes of the corresponding Cayley graph $\Gamma_n(\Delta)$ 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 graph H with node set $\{1, 2, \ldots, n\}$, where a triangle K_3 on nodes a, b, and c corresponds to each pair of a 3-cycle (abc) and its inverse in Δ.

It is easy to see that the Cayley graph generated by the 3-cycles in Δ is connected if its corresponding graph is connected. Since an interconnection network needs to be connected, we require H graph to be connected. In

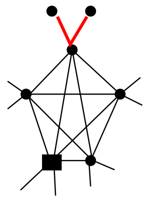

Figure 3: An example

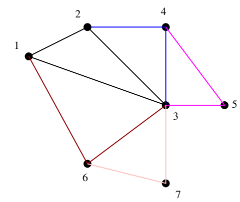

Figure 4: A 2-tree

general, this graph may have superfluous K_3s formed by nodes that do not correspond to a 3-cycle in Δ. We will avoid this possibility by considering a simpler case when H has a tree-like structure. Such a graph is built by the following procedure. We start from K_3, then repeatedly add a new node, joining it to exactly two adjacent nodes of the previous graph. Any graph obtained by this procedure is called a 2-*tree*. If v is a node of a 2-tree H with the property that H can be generated in such a way that v is the last node added, then v is called a *leaf* of the 2-tree. Figure 4 provides an example of a 2-tree (with leaves 5 and 7).

Now let H be a 2-tree, in which the K_3s correspond to the 3-cycles of Δ. The graph H will be called the 3-*cycle generating graph* of $\Gamma_n(\Delta)$ or simply its *generating graph* if it is clear from the context. We call $\Gamma_n(\Delta)$ the *Cayley graph generated by* H. Graph H is simply a pictorial representation of the elements in Δ.

We need the following lemma. Since in the lemma graph G contains triangles and its nodes each have degree $2n-4$, we can allow at most $2n-7$ missing edges in G.

Lemma 4.1. *Let G be a Cayley graph generated by a 2-tree on $\{1, 2, \ldots, n\}$ with $n \geq 5$, and let F_e be an arbitrary set of missing edges with $|F_e| \leq 2n-7$. For every node v in G, there exists an extended star of order $\deg_{G-F_e}(v)$ at v in $G - F_e$.*

Proof. For $n = 5$, there are only two possibilities, each generates a graph with 60 nodes where $|F_e| \leq 3$. They were checked by an ad hoc argument as well as a computer check. The computer check was done using the NetworkX python package (available at `http://networkx.github.com/`). For each of the two graphs it was verified that fixing any single node and deleting any set of edges F_e with size at most 3, the graphs still contain an extended star $ES(v; \deg_{G-F_e}(v))$. Although there were a large number of cases to check, the entire check could be completed in only a few minutes because it was often very easy to grow the extended star from v using a randomized greedy construction algorithm.

Now assume the result is true for all Cayley graphs obtained from 2-trees on $\{1, 2, \ldots, n-1\}$ with $n \geq 6$ and let G be a Cayley graph obtained from a 2-tree T on $\{1, 2, \ldots, n\}$. We may assume that n is a leaf in T. Let T' be obtained from T by deleting n, and let G' be the Cayley graph generated by T'. Let H_i be the subgraph of G induced by nodes with i in the last position for each i, $1 \leq i \leq n$. Then each H_i is isomorphic to G' and every node in H_i has exactly two neighbours not in H_i, and they are in different H_js. For notational convenience, we may assume that v is in H_n. Node v is adjacent to two nodes not in H_n, let these be node y in H_j and node z in H_q, where n, j, q are distinct. We consider three cases.

Case 1: H_n contains at most $2n-9$ elements of F_e.
Apply the induction hypothesis for $H_n - F_e$ to obtain an extended star A of order $\deg_{H_n - F_e}(v)$ at v in $H_n - F_e$. If $(v, y) \in F_e$, then we do not need to extend A via (v, y). Suppose that $(v, y) \notin F_e$. Now y is in H_j. Since $|F_e| \leq 2n-5$, graph $H_j - F_e$ is connected as H_j has edge-connectivity $2n-4$. So we can find a 3-path in H_j starting at y and attach it to A via (v, y). Repeating the argument for (v, z) yields the desired extended star in $G - F_e$.

Case 2: H_n contains exactly $2n-8$ elements of F_e.
Let f be an arbitrary element of F_e and let $F'_e = F_e - \{f\}$. Apply the induction hypothesis to obtain an extended star A of order $\deg_{H_n - F'_e}(v)$ at v in $H_n - F'_e$. If A does not contain f, or f is incident with v, then we can complete the proof as in Case 1 to get the desired extended star. So we may assume that A contains edge $f = (w, w')$, where node w is closer to v along A than w'. Let p be the distance between v and w in A, so $1 \leq p \leq 3$. Since H_n contains all but one elements of F_e, and w is adjacent to two nodes not in H_n, there is a node x adjacent to w such that $(w, x) \notin F_e$ and x is not in H_n. We consider two subcases. The first is if x is in neither H_j nor H_q. So, say, x is in $H_i - F_e$. Find a $(4-p-1)$-path starting at x in $H_i - F_e$ together with a 3-path (if necessary) starting at y in $H_j - F_e$ and a 3-path (if necessary) starting at z in $H_q - F_e$ to complete the extended star as in Case 1. We can easily do these since there is only one edge of F_e outside H_n.

The second subcase is when x has to be in either H_j or H_q. Since F_e has only one element outside H_n, we may assume that x is in H_j and F_e has no elements in H_j. Note that $y \neq x$ as it is not possible for both v and w to be adjacent to the same node in H_j. We are done if we can find two node-disjoint paths in $H_j = H_j - F_e$, a 3-path starting at y and a $(4-p-1)$-path starting at w. One can easily check that this claim is true for $n = 5$, and since H_j contains one of these as a subgraph, we are done.

Case 3: H_n contains all the elements of F_e.
Let f, f' be arbitrary elements of F_e and let $F'_e = F_e - \{f, f'\}$. Apply the induction hypothesis to obtain an extended star A of order $\deg_{H_n - F'_e}(v)$ at v in $H_n - F'_e$. We may assume that A contains both f and f' and neither f nor f' is incident with v, otherwise we can complete the proof as in Case 1 to get the desired extended star. Moreover, we may assume that f and f' belong to different paths in A, otherwise we can apply the same proof as in Case 2. Let $f = (w, w')$ where w is closer to v along A than w', and let $f' = (t, t')$ where t is closer to v than t'. Each of w and t has two neighbours outside H_n in distinct H_ls. Choose one neighbour for each in different H_l's. The argument in Case 2 applies and the proof is complete. $\square$

Proposition 2.1 and Lemma 4.1 yields our main result for Cayley graphs generated by 2-trees. (Again, letting $F_e = \emptyset$, we have the corresponding result with no missing edges.)

Theorem 4.2. *Let G be a Cayley graph generated by a 2-tree on $\{1, 2, \ldots, n\}$ with $n \geq 5$, and let F_e be an arbitrary set of missing edges with $|F_e| \leq 2n-7$. For each node v in G, node v has the strong local diagnosability property in $G - F_e$, so $G - F_e$ has the strong local diagnosability property.*

References

[1] J. Maeng, M. Malek, A comparison connection assignment for self-diagnosis of multiprocessor systems, in: Proceedings of the 11th International Symposium on Fault-Tolerant Computing, 1981, pp. 173–175.

[2] G.-H. Hsu, J.-M. Tan, A local diagnosability measure for multiprocessor systems, IEEE Trans. on Parallel Distrib. Syst. 18 (5) (2007) 598–607.

[3] C.-F. Chiang, J.-M. Tan, Using node diagnosability to determine t-diagnosability under the comparison diagnosis model, IEEE Trans. Comput. 58 (2009) 251–259.

[4] C.-F. Chiang, G.-H. Hsu, L.-M. Shih, J.-M. Tan, Diagnosability of star graphs with missing edges, Information Sciences 188 (2012) 253–259.

[5] W.-K. Chiang, R.-J. Chen, The (n, k)-star graph: a generalized star graph, Information Processing Letters 56 (1995) 259–264.