Received: 15 March 2021 | Revised: 30 July 2023 | Accepted: 31 July 2023

DOI: 10.1002/jgt.23027

ARTICLE

WILEY

Acyclic graphs with at least $2\ell + 1$ vertices are ℓ-recognizable

Alexandr V. Kostochka[1] | **Mina Nahvi**[1] | **Douglas B. West**[1,2] | **Dara Zirlin**[3]

[1]University of Illinois, Urbana, Illinois, USA

[2]Zhejiang Normal University, Jinhua, China

[3]University of Illinois at Urbana–Champaign, Urbana, Illinois, USA

Correspondence
Douglas B. West, Zhejiang Normal University, Jinhua, China.
Email: dwest@illinois.edu

Funding information
National Natural Science Foundation of China; Beckman Campus Research Board; National Science Foundation, Grant/Award Numbers: DMS-1600592, DMS-2153507; Arnold O. Beckman Campus Research Board Award, Grant/Award Number: RB20003; NSF RTG, Grant/Award Number: DMS-1937241

Abstract

The $(n - \ell)$-*deck* of an n-vertex graph is the multiset of subgraphs obtained from it by deleting ℓ vertices. A family of n-vertex graphs is ℓ-*recognizable* if every graph having the same $(n - \ell)$-deck as a graph in the family is also in the family. We prove that the family of n-vertex graphs with no cycles is ℓ-recognizable when $n \geq 2\ell + 1$ (except for $(n, \ell) = (5, 2)$). As a consequence, the family of n-vertex trees is ℓ-recognizable when $n \geq 2\ell + 1$ and $\ell \neq 2$. It is known that this fails when $n = 2\ell$.

KEYWORDS

acyclic graph, deck, graph reconstruction, tree

1 | INTRODUCTION

The j-*deck* of a graph is the multiset of its j-vertex induced subgraphs. We write this as the $(n - \ell)$-deck when the graph has n vertices and the focus is on deleting ℓ vertices. An n-vertex graph is ℓ-*reconstructible* if it is determined by its $(n - \ell)$-deck. Since every member of the $(j - 1)$-deck arises $n - j + 1$ times by deleting a vertex from a member of the j-deck, the j-deck of a graph determines its $(j - 1)$-deck. Therefore, a natural reconstruction problem is to find for each graph the maximum ℓ such that it is ℓ-reconstructible. For this problem, Manvel [10, 11] extended the classical Reconstruction Conjecture of Kelly [5] and Ulam [16].

Conjecture 1.1 (Manvel [10, 11]). *For $\ell \in \mathbb{N}$, there exists a threshold M_ℓ such that every graph with at least M_ℓ vertices is ℓ-reconstructible.*

Manvel named this "Kelly's Conjecture" in honor of the final sentence in Kelly [6], which suggested that one can study reconstruction from the $(n-2)$-deck. Manvel noted that Kelly may have expected the statement to be false.

Many reconstruction arguments have two parts. First, one proves that the deck determines that the graph is in a particular class or has a particular property. When the $(n-\ell)$-deck determines this, the property is *ℓ-recognizable*. Separately, using the knowledge that the deck determines whether the graph has that property, one proves that only one graph with that property has that deck. This makes the family *weakly ℓ-reconstructible*, meaning that no two graphs in the family have the same deck. Bondy and Hemminger [1] articulated the distinction between these two steps for the case $\ell = 1$.

Here, toward ℓ-reconstructibility of trees, we consider ℓ-recognizability of acyclic graphs. We prove the following theorem.

Theorem 1.2. *For $n \geq 2\ell + 1$, except when $(n, \ell) = (5, 2)$, the family of n-vertex acyclic graphs is ℓ-recognizable.*

We forbid $(n, \ell) = (5, 2)$ due to the two graphs in Figure 1, which have the same 3-deck. Indeed, this possibility must be excluded from many of the claims we prove.

Since the $(n-\ell)$-deck determines the 2-deck when $n - \ell \geq 2$, in this setting we also know the number of edges. This yields the following corollary.

Corollary 1.3. *For $n \geq 2\ell + 1$, except when $(n, \ell) = (5, 2)$, the family of n-vertex trees is ℓ-recognizable.*

Spinoza and West [14] determined, for every graph G with maximum degree at most 2, the maximum ℓ such that G is ℓ-reconstructible. Their full result is complicated to state, but a special case is that for $n \geq 2\ell + 1$ (except $(n, \ell) = (5, 2)$), every n-vertex graph with maximum degree at most 2 is ℓ-reconstructible. They also show that a path with 2ℓ vertices has the same ℓ-deck as the disjoint union of an $(\ell + 1)$-cycle and a path with $\ell - 1$ vertices, so the threshold $n \geq 2\ell + 1$ in both [14] and Theorem 1.2 is sharp.

Nýdl [13] conjectured that trees with at least $2\ell + 1$ vertices are weakly ℓ-reconstructible. This conjecture would be sharp, since Nýdl [12] found two trees with 2ℓ vertices having the same ℓ-deck. The two trees are obtained from a path with $2\ell - 1$ vertices by adding one leaf, adjacent either to the central vertex of the path or to one of its neighbors. Kostochka and West [9] used the

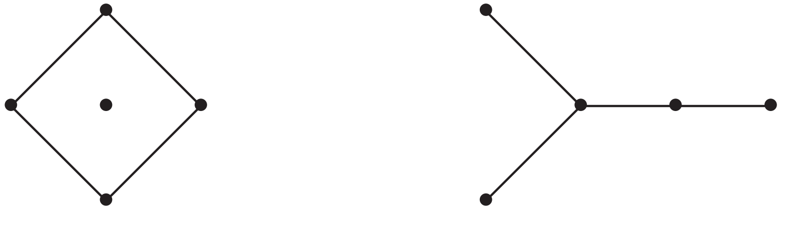

FIGURE 1 Five-vertex graphs with the same 3-deck.

results of [14] to give a short proof of this result of Nýdl. However, one counterexample to Nýdl's conjecture is known: Groenland, Johnston, Scott, and Tan [4] obtained two 13-vertex trees having the same 7-deck. Excluding this example and incorporating the ℓ-recognizability of trees leads to a modification of Nýdl's conjecture.

Conjecture 1.4. *For $n \geq 2\ell + 1$, except when $(n, \ell) \in \{(5, 2), (13, 6)\}$, every n-vertex tree is ℓ-reconstructible. The threshold on n is known to be sharp.*

For $\ell = 2$, the threshold on n must be at least 6 due to the graphs in Figure 1. Giles [2] proved that trees with at least six vertices are 2-reconstructible (using only the connected members of the deck). For general ℓ, Groenland et al. [4] proved that $n \geq 9\ell + 24\sqrt{2\ell} + o(\sqrt{\ell})$ suffices for ℓ-reconstructibility of n-vertex trees. In [8], the present authors prove that $n \geq 6\ell + 11$ suffices.

Besides acyclicity, another fundamental property of trees is connectedness. Spinoza and West [14] proved that connectedness is ℓ-recognizable for n-vertex graphs when $n > 2\ell^{(\ell+1)^2}$. Later, Groenland et al. [4] reduced the general threshold to $n \geq 10\ell$. Manvel [11] proved that connectedness is 2-recognizable for graphs with at least six vertices, and the present authors [7] proved that connectedness is 3-recognizable for graphs with at least seven vertices. Spinoza and West [14] suggested that (except for $(n, \ell) = (5, 2)$), connectedness is recognizable for n-vertex graphs when $n \geq 2\ell + 1$.

The $(n - \ell)$-deck of a graph is *acyclic* if each card in the deck is acyclic. As a step toward the threshold on n for ℓ-recognizability of connectedness, one can consider the special case of n-vertex graphs whose $(n - \ell)$-deck is acyclic. Our result in this paper settles the question for graphs with $n - 1$ edges, where connectedness and acyclicity are equivalent (the number of edges is known from the 2-deck). This suggests other detailed questions.

Problem 1.5. For $\ell \geq 1$ and $c \geq 0$, determine the smallest thresholds $N_{\ell,c}$ and $N'_{\ell,c}$ such that for all n-vertex graphs with $n + c$ edges whose $(n - \ell)$-deck $\mathcal{D}$ is acyclic,

(a) if $n \geq N_{\ell,c}$, then $\mathcal{D}$ determines whether the graph is connected, and
(b) if $n \geq N'_{\ell,c}$, then the graph is connected.

The thresholds when the deck is not required to be acyclic are also unknown.

We note that $N'_{\ell,1} = 2\ell$. For the upper bound, consider a disconnected n-vertex graph with an acyclic $(n - \ell)$-deck, where $n \geq 2\ell$. A smallest component must be acyclic, since a cycle would have length at most $n/2$ and be seen in the deck. Hence some other component H must have at least $|V(H)| + 2$ edges. However, a p-vertex graph with at least $p + 2$ edges has girth at most $\lfloor (p + 2)/2 \rfloor$ (see, e.g., Exercise 5.4.36 of [17]), yielding a cycle in a card. For the lower bound, we seek a disconnected graph with $2\ell - 1$ vertices whose $(\ell - 1)$-deck is acyclic. When ℓ is even, the graph consists of an isolated vertex plus four paths of length $\ell/2$ with common endpoints. When ℓ is odd, the nontrivial component consists of a cycle of length $2\ell - 2$ plus two diametric chords creating cycles of length ℓ (this example was contributed by a referee). It is possible that the threshold $N_{\ell,1}$ for determining connectedness from the $(n - \ell)$-deck is smaller.

For $c = 0$, we conjecture $N_{\ell,0} = 2\ell - 1$. The lower bound holds because a $(2\ell - 2)$-cycle and the disjoint union of two $(\ell - 1)$-cycles have the same $(\ell - 2)$-deck. Zirlin [18] proved $N_{\ell,0} \leq 2\ell + 1$ for $\ell \geq 3$, and she proved $N_{\ell,0} = 2\ell - 1$ for $\ell \geq 45$.

In Section 2 we develop tools that are useful for reconstructing information from acyclic decks. In Section 3 we prove that $n \geq 2\ell + 2$ suffices for ℓ-recognizability of n-vertex acyclic graphs. In Section 4 we obtain the sharp threshold, $2\ell + 1$.

2 | VINES, DIAMETER, AND MARKING

Let $\mathcal{D}$ be the $(n - \ell)$-deck of an n-vertex graph G (we henceforth just call it the "deck"). We will assume $n > 2\ell$. The members of $\mathcal{D}$ are the "cards" in the deck.

Definition 2.1. The *eccentricity* $\varepsilon_G(v)$ of a vertex v in a graph G is the maximum of the distances from v to other vertices. The *radius* is $\min_{v\in V(G)} \varepsilon_G(v)$, and the *diameter* is $\max_{v\in V(G)} \varepsilon_G(v)$. A *center* of G is a vertex of minimum eccentricity.

In G, the *j-ball* at a vertex v is the subgraph induced by all vertices within distance j of v in G. The *j-eball* at an edge e is the subgraph induced by all vertices within distance j of either endpoint of e. A *j-vine* or *j-evine* is a tree having diameter $2j$ or $2j + 1$, respectively. A *j-center* is a vertex that is the center of a j-vine; a *j-central edge* is the central edge of a j-evine (joining the two centers).

The term "j-vine" follows the botanical theme in terminology about trees; a vine grows from its main path. Note that if the j-ball at a vertex v in a graph G is a tree but does not contain a path with $2j + 1$ vertices, then v is not a j-center. When v *is* a j-center, the maximal j-vine with center v is just the j-ball at v.

We will be interested in j-vines and j-evines that are induced subgraphs of every reconstruction from the given deck. Our aim is to consider an acyclic and a nonacyclic graph having the same deck, show that they have the same number of j-centers for an appropriate value j, and obtain a contradiction by showing that they cannot have the same number of j-centers. The key property that will permit counting the j-centers in a reconstruction is in Lemma 2.2; it implies (under the girth condition) that maximal j-vines correspond bijectively to centers of j-vines (i.e., j-centers), and similarly for j-evines.

Lemma 2.2. *In a graph G with girth at least $2j + 2$, every j-vine lies in a unique maximal j-vine. If G has girth at least $2j + 3$, then every j-evine lies in a unique maximal j-evine.*

Proof. When H is a j-vine or a j-evine, let H' be the j-ball or j-eball in G at the center(s) of H, respectively. All vertices in any j-vine or j-evine containing H lie in H'. Thus H' is the desired unique maximal object unless it contains a cycle.

Let Q be a shortest cycle in H'. Because the vertex or vertices on Q that are farthest from the center of H' have distance at most j from the center, Q has at most $2j + 1$ vertices if H' has a unique center and at most $2j + 2$ vertices if H' has a central edge, contradicting the hypothesis on the girth of G. □

Example 2.3. To see that the girth condition in Lemma 2.2 is sharp, let G be a graph consisting of a $(2j + 1)$-cycle Q plus two paths of length j grown from a single vertex v on Q. Deleting from G the two vertices of Q that are farthest from v yields a j-vine H. Replacing either one of those vertices yields a maximal j-vine in G containing H, so the maximal j-vine containing H is not unique. An analogous example for j-evines consists of a $(2j + 2)$-cycle plus paths of length j grown from two consecutive vertices.

Definition 2.4. Given a family $\mathcal{F}$ of graphs, an *$\mathcal{F}$-subgraph* of a graph G is an induced subgraph of G belonging to $\mathcal{F}$. Let $s(F, G)$ denote the number of occurrences of F as an induced subgraph of G. Let $m(F, G)$ be the number of occurrences of F as a maximal $\mathcal{F}$-subgraph in G (with respect to induced subgraphs).

The special case $\ell = 1$ of Lemma 2.5 is due to Greenwell and Hemminger [3]. Similar statements for general ℓ appear, for example, in [4]. We include a proof for completeness; it is slightly simpler than proofs in the literature involving inclusion chains of subgraphs.

Lemma 2.5. *Fix an n-vertex graph G, and let $\mathcal{F}$ be a family of graphs such that every $\mathcal{F}$-subgraph of G lies in a unique maximal $\mathcal{F}$-subgraph of G. If the value of $m(F, G)$ is known for every $F \in \mathcal{F}$ with at least $n - \ell$ vertices, then for all $F \in \mathcal{F}$ the $(n - \ell)$-deck of G determines $m(F, G)$.*

Proof. Let $t = |V(G)| - |V(F)|$; we use induction on t. When $t \leq \ell$, the value $m(F, G)$ is given. When $t > \ell$, group the induced subgraphs of G isomorphic to F according to the unique maximal $\mathcal{F}$-subgraph of G containing them (as an induced subgraph). Counting all copies of F then yields

$$s(F, G) = \sum_{H \in \mathcal{F}} s(F, H) m(H, G).$$

Since $|V(F)| < n - \ell$, we know $s(F, G)$ from the deck, and we know $s(F, H)$ when F and H are known. By the induction hypothesis, we know all values of the form $m(H, G)$ when F is an induced subgraph of H except $m(F, G)$. Therefore, we can solve for $m(F, G)$. □

Before continuing with preparation for ℓ-recognizability of acyclicity, we note one application of Lemma 2.5 that was stated incorrectly in the paper by Kostochka and West [9]; it also illustrates the technique we use with j-vines. The special case for $\ell = 1$ was observed by Kelly [6] using different methods. Let P_n and C_n, respectively, denote a path and a cycle with n vertices, and let $G + H$ denote the disjoint union of graphs G and H.

Corollary 2.6. *If $n > 2\ell$, then n-vertex graphs having no component with more than $n - \ell$ vertices are ℓ-reconstructible, and this threshold on n is sharp. All n-vertex graphs having no component with at least $n - \ell$ vertices are ℓ-reconstructible, with no restriction on n.*

Proof. Let $\mathcal{F}$ be the family of connected graphs; $\mathcal{F}$ satisfies the property stated in the first sentence of Lemma 2.5 for any G.

Now consider $m(F, G)$ for $F \in \mathcal{F}$. When $n > 2\ell$, an n-vertex graph has at most one component with at least $n - \ell$ vertices, and it has no component with more vertices if and only if it has at most one connected $(n - \ell)$-card. Hence the hypothesized condition here is ℓ-recognizable. If some component has exactly $n - \ell$ vertices, then it is seen as a card. Hence $m(F, G)$ is known for $F \in \mathcal{F}$ with at least $n - \ell$ vertices, and by using Lemma 2.5 we obtain all the components of G under either hypothesis in the statement.

The result is sharp, since $P_\ell + P_\ell$ and $P_{\ell+1} + P_{\ell-1}$ have the same ℓ-deck. This follows from the result of Spinoza and West [14] that any two graphs with the same number of vertices and edges whose components are all cycles with at least $j + 1$ vertices or paths with at least $j - 1$ vertices have the same j-deck. □

When we speak of j-vines and j-evines in an n-vertex graph G, we always consider only induced subgraphs. For a given graph G, a particular value of j determined by the $(n - \ell)$-deck of G will be of interest. Recall that we require $n \geq 2\ell + 1$ and $\ell \geq 1$, so $n - \ell \geq 2$.

Definition 2.7. For a given graph G, let k denote the largest integer such that G contains a k-evine and, for $0 \leq j \leq k$, every j-evine and every j-vine in G has fewer than $n - \ell$ vertices. Since every edge is a 0-evine, k is well-defined. *This fixes k in terms of G for the remainder of the paper.*

We consider n-vertex reconstructions from an acyclic $(n - \ell)$-deck $\mathcal{D}$.

Lemma 2.8. *The value k is determined by the deck $\mathcal{D}$ of G. That is, all reconstructions from $\mathcal{D}$ have the same value of k.*

Proof. All subgraphs of G having at most $n - \ell$ vertices are visible in $\mathcal{D}$. The deck thus gives a candidate for k. Let k' be the largest integer such that some induced subgraph of a card is a k'-evine and, for $0 \leq j \leq k'$, every j-evine and every j-vine that appears in a card has fewer than $n - \ell$ vertices. Since this condition on k' is satisfied by k, we have $k' \geq k$.

The value k' is strictly greater than k if and only if G contains a j-evine or j-vine R with at least $n - \ell$ vertices for some j at most k'. By the definition of k', this R is not contained in a card and has strictly more than $n - \ell$ vertices. Without modifying a fixed longest path P, we can trim R to $n - \ell$ vertices by iteratively deleting leaves outside P unless we still have more than $n - \ell$ vertices when only P remains. Therefore, since G has no j-evine or j-vine with exactly $n - \ell$ vertices (by the definition of k'), we have $2j + 1 > n - \ell$.

On the other hand, the definition of k' gives us a k'-evine contained in a card and having fewer than $n - \ell$ vertices. A longest path in this subgraph has $2k' + 2$ vertices, so $2k' + 2 < n - \ell$. We thus have

$$2k' + 3 \leq n - \ell \leq 2j \leq 2k'.$$

This contradiction implies that k' must equal k. □

Lemma 2.9. *Every reconstruction from $\mathcal{D}$ has girth at least $2k + 4$.*

Proof. The claim holds for a reconstruction G having no cycle, so suppose that G has a cycle. Since $\mathcal{D}$ is acyclic, the girth of G is at least $n - \ell + 1$. Deleting some consecutive vertices from a shortest cycle yields an induced path R with $n - \ell$ vertices. Let $t = \lfloor (n - \ell - 1)/2 \rfloor$, so $n - \ell + 1 \geq 2t + 2$. The path R is a t-vine (if $n - \ell$ is odd) or a t-evine (if $n - \ell$ is even) with $n - \ell$ vertices. The definition of k thus requires $k < t$. We compute $n - \ell + 1 \geq 2t + 2 \geq 2k + 4$. Thus G has girth at least $2k + 4$. □

Corollary 2.10. *For $j \leq k$, the deck $\mathcal{D}$ determines the maximal j-evines and maximal j-vines, with multiplicity. Also, all reconstructions from $\mathcal{D}$ have the same numbers of j-centers and j-central edges.*

Proof. Fix j with $j \leq k$, and let $\mathcal{F}$ be the family of j-vines or the family of j-evines. By Lemma 2.8, all reconstructions have the same value of k. By the definition of k, we obtain $m(T, G) = 0$ whenever G is a reconstruction from $\mathcal{D}$ and T is a member of $\mathcal{F}$ having at least $n - \ell$ vertices. Since G has girth at least $2k + 4$ (by Lemma 2.9), every member of $\mathcal{F}$ lies in a unique maximal member of $\mathcal{F}$ (by Lemma 2.2). With these properties, Lemma 2.5 applies to compute $m(T, G)$ for all $T \in \mathcal{F}$.

With girth at least $2k + 4$, there is also a one-to-one correspondence between the maximal j-vines and the j-centers, and similarly for the maximal j-evines and j-central edges. Thus we obtain the total number of j-centers and the total number of j-central edges. □

Setting $j = 1$ in Corollary 2.10 almost provides the degree list. Groenland et al. [4] proved the strong result that the degree list is ℓ-reconstructible for all n-vertex graphs whenever $n - \ell > \sqrt{2n\log(2n)}$. Taylor [15] had shown that $n > f(\ell)$ suffices, where f is a particular function such that $f(\ell)$ is asymptotic to $e\ell$. For the context of acyclic decks we obtain a simpler intermediate threshold. For a vertex v in a graph, let $N[v]$ denote the closed neighborhood of v (the set of vertices equal or adjacent to v).

Corollary 2.11. *For $n \geq 2\ell + 1$ with $(n, \ell) \neq (5, 2)$, the degree list of any n-vertex graph with an acyclic $(n - \ell)$-deck is determined by its deck.*

Proof. The case $\ell = 1$ is well known: subtract the number of edges in each card from the total number of edges. Since $(n, \ell) \neq (5, 2)$ and $n \geq 2\ell + 1$, we may thus assume $n - \ell \geq 4$.

Let G be an n-vertex graph with an acyclic $(n - \ell)$-deck $\mathcal{D}$. Since $\mathcal{D}$ is acyclic and $n - \ell \geq 4$, all stars are induced subgraphs. Those with at least three vertices are the 1-vines.

Call a vertex $v \in V(G)$ *big* if $d_G(v) \geq n - \ell - 1$. A vertex with degree at least 2 in G is the center of a maximal 1-vine. For $t \geq 2$ the number of vertices with degree t is the number of maximal 1-vines with $t + 1$ vertices. Lemma 2.5 provides these values for $t \geq 2$ if we know the number of big vertices with each degree. There are no big vertices in G if and only if no card is a star.

Since $n - \ell \geq 4$ and $\mathcal{D}$ is acyclic,

$$G \text{ has no 3-cycles or 4-cycles.} \tag{1}$$

If G has exactly one big vertex and its degree is d, then exactly $\binom{d}{n-\ell-1}$ cards are stars.

Suppose that x and y are distinct big vertices in G. By (1), x and y have at most one common neighbor. When $xy \notin E(G)$,

$$n \geq |N[x] \cup N[y]| \geq 2n - 2\ell - 1 = n + (n - 2\ell - 1) \geq n. \tag{2}$$

It follows that $n = 2\ell + 1$ and $d(x) = d(y) = n - \ell - 1$ and $N[x] \cup N[y] = V(G)$. If there is a third big vertex, then since $n - \ell - 1 \geq 3$ it has at least two neighbors in $N[x]$ or in $N[y]$, contradicting (1). We conclude that in this case $\mathcal{D}$ has exactly two star cards.

If three big vertices are pairwise adjacent, then they induce a 3-cycle, contradicting (1). Hence three big vertices must include a nonadjacent pair, reducing to the previous case.

There remains only the case of exactly two big vertices x and y, adjacent. Now

$$n \geq |N[x] \cup N[y]| \geq 2n - 2\ell - 2 = n + (n - 2\ell - 2) \geq n - 1. \tag{3}$$

If $|N[x] \cup N[y]| = n - 1$, then x and y both have degree $n - \ell - 1$ and again there are exactly two star cards. If $|N[x] \cup N[y]| = n$, then either $n = 2\ell + 1$, with x and y having degrees $n - \ell - 1$ and $n - \ell$, or $n = 2\ell + 2$ with x and y both having degree $n - \ell - 1$. In the former case, there are $\ell + 2$ star cards (since $n - \ell = \ell + 1$) and in the latter case there are two.

We have shown that G can only have one or two big vertices, and if it has two, then $\mathcal{D}$ has either two or $\ell + 2$ star cards, and in the last case $n = 2\ell + 1$. Since $n - \ell \geq 4$, the degree d of a big vertex is at least 3. Hence $\binom{d}{n-\ell-1}$ cannot equal 2 and cannot equal $\ell + 2$ when $n = 2\ell + 1$. This makes all cases distinguishable from the others.

We now know $m(T, G)$ for any reconstruction G and every star T with at least $n - \ell$ vertices, and Lemma 2.5 applies to yield the number of vertices with degree t for each t at least 2. It remains to count vertices with degree at most 1. Since we know the number m of edges from the 2-deck, the number of vertices with degree 1 is given by subtracting the other known degrees from $2m$, and then the remaining vertices have degree 0. □

Lemma 2.12. *Let $\mathcal{D}$ be a deck having a connected card. Every connected card has diameter at least $2k + 2$, and some connected card has diameter at most $2k + 3$.*

Proof. Since the deck is acyclic, every connected card is a tree. A connected card with diameter at most $2k + 1$ would be a j-vine or j-evine with $j \leq k$ having $n - \ell$ vertices, contradicting the definition of k.

For the second claim, let C be a connected card. If C has diameter at least $2k + 3$, then C contains a path with $2k + 4$ vertices, which is a $(k + 1)$-evine. By the definition of k, some $(k + 1)$-evine or $(k + 1)$-vine R has at least $n - \ell$ vertices. Since $n - \ell \geq 2k + 4$, we can iteratively delete leaves outside a fixed longest path in R to trim it to $n - \ell$ vertices. We thus obtain a card that is a $(k + 1)$-evine or $(k + 1)$-vine, which have diameter $2k + 3$ or $2k + 2$, respectively. Hence some card has diameter at most $2k + 3$. □

When $n \geq 2\ell + 2$, we will show that $\mathcal{D}$ cannot have both an acyclic reconstruction F and a nonacyclic reconstruction H by showing that H would have more k-centers or $(k + 1)$-centers than F. We next introduce a tool for bounding the number of j-centers in a forest F.

Definition 2.13. *The marking process.* Let z be a central vertex of a connected $(n - \ell)$-card C with radius $j + 1$ in a forest F. Let Y be the set of neighbors of z that lie on paths of length $j + 1$ in C beginning at z, and let $d_C = |Y|$. In the component of F containing C, every j-center x that is not in Y *marks* one vertex x' at distance j from x along a path that extends the z, x-path in F (such a vertex exists, since x is a j-center).

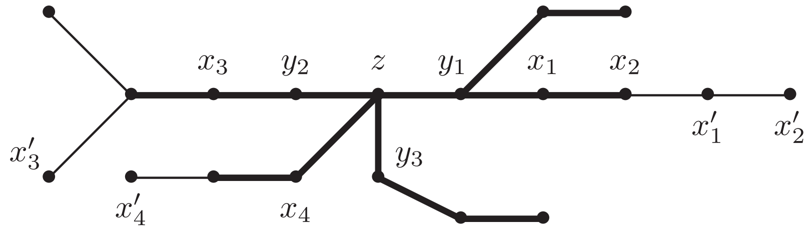

FIGURE 2 The marking process.

Note that d_C is the maximum number of edge-disjoint paths of length $j + 1$ in C with endpoint z. In particular, $d_C = 1$ when C has diameter $2j + 1$ and $d_C \geq 2$ when C has diameter $2j + 2$. Figure 2 illustrates the marking process for a card C (in bold) within a tree F. Here C has radius 3 with center z, we have $j = 2$ and $d_C = 3$, with $Y = \{y_1, y_2, y_3\}$, and x_i marks x'_i. Vertices of the form x_i, y_i, and z are j-centers in F.

Lemma 2.14. *If $j \geq 1$ and C is a connected card with radius $j + 1$ in the $(n - \ell)$-deck $\mathcal{D}$ of an n-vertex forest F, then the number of j-centers in F is at most $1 + d_C + \ell$. If equality holds, then in the marking process each vertex of F outside C is marked and F is a tree.*

Proof. Let F' be the component of F containing C, and let ℓ' be the number of vertices of F' outside C. All j-centers that are neighbors of z on paths of length $j + 1$ from z (and z itself) do not mark vertices. All other j-centers in F' mark a vertex that is outside C. Since F' has no cycles, every vertex of F' is marked by at most one j-center. Thus F has at most $\ell' + d_C + 1$ j-centers in F'.

There are $\ell - \ell'$ vertices of F outside F', and any that have degree at most 1 cannot be j-centers (since $j \geq 1$). Hence F has at most $1 + d_C + \ell$ j-centers, with equality only if F is a tree and all vertices outside C are marked. □

Note that the conclusion is false when $j = 0$, since every vertex is a 0-center.

3 | RESTRICTING TO $n = 2\ell + 1$

Given an acyclic $(n - \ell)$-deck $\mathcal{D}$ for $n \geq 2\ell + 1$, with k defined as in Section 2, we have proved that all n-vertex reconstructions from $\mathcal{D}$ have the same number of k-centers and have the same number of k-central edges. Our next aim is to prove this also for $(k + 1)$-centers when $\mathcal{D}$ has no card with diameter $2k + 2$. We will need connected cards, which are guaranteed when $\mathcal{D}$ has reconstructions both with and without cycles.

Definition 3.1. We say that a deck $\mathcal{D}$ is *ambiguous* if it is the $(n - \ell)$-deck of both an n-vertex acyclic graph F and an n-vertex nonacyclic graph H.

Remark 3.2. An ambiguous deck is acyclic, since F contains no cycle. Hence when $\mathcal{D}$ is ambiguous the graph H has girth at least $n - \ell + 1$, and thus $\mathcal{D}$ has connected cards (in particular, paths).

Lemma 3.3. *If $n \geq 2\ell + 1$ and an ambiguous deck $\mathcal{D}$ has no card with diameter $2k + 2$, then all reconstructions have the same number of $(k + 1)$-centers.*

Proof. Since G has a k-evine and all k-evines have fewer than $n - \ell$ vertices, we have $2k + 2 < n - \ell$. By Remark 3.2, $\mathcal{D}$ has a connected card. Thus if no card has diameter $2k + 2$, then some card has diameter $2k + 3$, by Lemma 2.12. A card with diameter $2k + 3$ is a $(k + 1)$-evine. In this card are two $(k + 1)$-vines, centered at the vertices of the central edge, so G has $(k + 1)$-vines.

If any $(k + 1)$-vine has at least $n - \ell$ vertices, then we obtain a $(k + 1)$-vine with $n - \ell$ vertices by iteratively deleting leaves outside a fixed longest path unless $2k + 3 > n - \ell$, but this contradicts the inequality $2k + 2 < n - \ell$. However, a $(k + 1)$-vine with $n - \ell$ vertices is a card with diameter $2k + 2$, which by hypothesis does not exist. Hence all $(k + 1)$-vines have fewer than $n - \ell$ vertices.

By Lemma 2.9, G has girth at least $2k + 4$. Hence every $(k + 1)$-vine lies in a unique maximal $(k + 1)$-vine, by Lemma 2.2. Therefore, the hypotheses of Lemma 2.5 hold for the family of $(k + 1)$-vines in G, and the deck determines all the maximal $(k + 1)$-vines in G, with multiplicity (as in Corollary 2.10 for k-vines). The maximal $(k + 1)$-vines correspond to the $(k + 1)$-centers, so we obtain the number of $(k + 1)$-centers. □

Lemma 3.4. *Let $\mathcal{D}$ be ambiguous, with $n \geq 2\ell + 1$ and $(n, \ell) \neq (5, 2)$. If C is a connected card in $\mathcal{D}$ and Q is a shortest cycle in H, then C and Q share at least four vertices and lie in a component of H with at most $n - 2$ vertices. Furthermore, C cannot be a star.*

Proof. Since $\mathcal{D}$ is acyclic, Q has at least $n - \ell + 1$ vertices. Since C has $n - \ell$ vertices and $2n - 2\ell + 1 > n$, subgraphs C and Q of H intersect. Since F has at most $n - 1$ edges and the component H' of H containing C and Q has at least as many edges as vertices, H' cannot be all of H. If H' has $n - 1$ vertices, then H has an isolated vertex and F has at least $n - 1$ edges and is a tree, with no isolated vertices. Since F and H have the same degree list (by Corollary 2.11), H' therefore has at most $n - 2$ vertices. With $t = |V(C) \cap V(Q)|$, we have $n - \ell + (n - \ell + 1) - t \leq n - 2$, so $t \geq 4$.

Since Q is a shortest cycle, three vertices of Q cannot have a common neighbor in or outside Q. Since $t \geq 4$, we conclude that C cannot be a star. □

Henceforth let $\mathcal{D}$ be the ambiguous $(n - \ell)$-deck of reconstructions F and H as in Definition 3.1, and let k be as in Definition 2.7. In the remainder of this section we restrict the possibility of ambiguous decks to the case $n = 2\ell + 1$, which completes the ℓ-recognizability proof when $n \geq 2\ell + 2$. We leave the boundary case $n = 2\ell + 1$ to Section 4.

When we want to use the marking process to compare the numbers of k-centers, we need to exclude the possibility $k = 0$, since the conclusion of Lemma 2.14 is false when $j = 0$.

Lemma 3.5. *Let $\mathcal{D}$ be ambiguous with $n \geq 2\ell + 1 \geq 5$ and $(n, \ell) \neq (5, 2)$. If $k = 0$, then $n = 2\ell + 1$ and a card with smallest diameter is a 1-evine that in a nonacyclic reconstruction intersects a cycle with $\ell + 2$ vertices exactly in four consecutive vertices.*

Proof. Since $\mathcal{D}$ is acyclic, H has girth more than $n - \ell$. Hence some card is a path; it has at least four vertices, since $n - \ell \geq 4$. Thus any reconstruction contains a 1-evine, so $k = 0$ requires a card that is a 1-evine or a 1-vine. By Lemma 3.4, no card is a 1-vine (a star).

Hence a card C with smallest diameter is a 1-evine. That is, C is a tree with two nonleaf vertices. By Lemma 3.4, C and a shortest cycle Q share at least four vertices in H. Since Q has at least five vertices and is a shortest cycle in H, in H there is no chord of Q and no vertex outside Q with two neighbors in Q. Hence each central vertex of C has only one neighbor that is in Q, so there are at most four vertices of C in Q, with equality only if they are two leaves and the two central vertices of C and occur consecutively along Q.

With $|V(C \cap Q)| = 4$, we have $|V(C \cup Q)| \geq 2n - 2\ell + 1 - 4$. Lemma 3.4 then implies $2n - 2\ell - 3 \leq n - 2$, which simplifies to $n \leq 2\ell + 1$. With $n = 2\ell + 1$, the cycle Q must have exactly $\ell + 2$ vertices. □

Theorem 3.6. *Given $n \geq 2\ell + 1$ with $\ell \geq 2$, let $\mathcal{D}$ be an ambiguous deck with acyclic reconstruction F and nonacyclic reconstruction H for which $k \geq 1$, and let C be a card with minimum diameter in $\mathcal{D}$. These conditions require $n = 2\ell + 1$, that H has girth $\ell + 2$, and that F is a tree with exactly $1 + d_C + \ell$ j-centers, where j is the radius of C.*

Proof. By Lemma 2.12, no card has diameter less than $2k + 2$, but some card has diameter at most $2k + 3$. Let Q be a shortest cycle in H, with length q. By Lemma 2.9, $q \geq 2k + 4$.

Case 1: *C has diameter* $2k + 2$. By Corollary 2.10, F and H have the same number s of k-central edges. Note that C has radius $k + 1$ and has d_C k-central edges incident to its unique center z. These edges are also k-central in F. An edge of F in the component containing z is a k-central edge if and only if its endpoint farther from z is a k-center. In other components, the number of k-central edges is less than the number of k-centers. Using $j = k$, Lemma 2.14 implies $s \leq d_C + \ell$.

Among the d_C k-central edges in C incident to z, only two can lie in Q. Since $q \geq 2k + 4$, every edge of Q is a k-central edge in H. Thus $s \geq q + d_C - 2$, and the bounds on s yield $q \leq \ell + 2$. With $q \geq n - \ell + 1$, we obtain $n \leq 2\ell + 1$. Since $n \geq 2\ell + 1$, we thus have $n = 2\ell + 1$ and $q = \ell + 2$. The bounds on s now yield $s = d_C + \ell$, so F has exactly $1 + d_C + \ell$ k-centers, which by Lemma 2.14 requires that F is a tree.

Case 2: *C has diameter* $2k + 3$. By Lemma 3.3, F and H have the same number s' of $(k + 1)$-centers. With diameter $2k + 3$, C has radius $k + 2$ and two centers. Let z be a center in C. By Lemma 2.14 with $j = k + 1$, we have $s' \leq 2 + \ell$, since $d_C = 1$. Since $q \geq 2k + 4$, every vertex of Q is a $(k + 1)$-center. Hence $n - \ell + 1 \leq q \leq s' \leq 2 + \ell$, which simplifies to $n \leq 2\ell + 1$. Since $n \geq 2\ell + 1$, we have equality throughout, so $n = 2\ell + 1$, and $q = 2 + \ell$, and $s' = 2 + \ell$, which by Lemma 2.14 requires that F is a tree. □

Corollary 3.7. *For $n \geq 2\ell + 2$, the family of n-vertex acyclic graphs is ℓ-recognizable.*

Proof. By Theorem 3.6, an ambiguous deck can exist only when $n = 2\ell + 1$. □

4 | THE EXTREME CASE $n = 2\ell + 1$

The arguments of Section 3 leave open the possibility of an ambiguous deck when $n = 2\ell + 1$, and the graphs in Figure 1 yield an ambiguous deck when $(n, \ell) = (5, 2)$. In this section we will prohibit ambiguous decks when $n = 2\ell + 1$ and $\ell \geq 3$, yielding the sharp threshold on n for ℓ-recognizability of acyclicity.

Comparing the numbers of k-centers and $(k+1)$-centers in an acyclic and a nonacyclic reconstruction only restricted us to $n \leq 2\ell + 1$. Now that we restrict to $n = 2\ell + 1$ and $\ell \geq 3$, we will distinguish these possibilities by counting the cards that are paths (with $n - \ell$ vertices). We will obtain a bound on this number for a special class of trees; this is a result that may be of independent interest. We will eventually use the marking process to restrict the acyclic reconstruction to this class when we have an ambiguous deck.

Definition 4.1. Fix the parameter ℓ. A *full path* in an n-vertex graph is a path with $n - \ell$ vertices (a card in the $(n - \ell)$-deck). A *branch vertex* in a tree is a vertex with degree at least 3. A *leg* of a nonpath tree is a path in the tree whose endpoints are a leaf and the branch vertex closest to it. A *spider* is a tree with at most one branch vertex (trees with no branch vertex are paths, in which we may designate any vertex as the "root" serving the role of a branch vertex). We denote a spider with legs of lengths $m_1, \ldots, m_d$ as $S_{m_1,\ldots,m_d}$; it has $1 + \sum_{i=1}^{d} m_i$ vertices (an n-vertex path can be described as $S_{m,n-1-m}$ for any m with $1 \leq m \leq n - 2$.) An n-vertex tree is *ℓ-spiderly* if it contains a spider such that all vertices not in the spider are within distance $(n - \ell - 2)/2$ of the branch vertex of the spider.

Note that all spiders are ℓ-spiderly. The upper bound that we can prove on the number of full paths in an n-vertex spider holds more generally for all ℓ-spiderly trees with n vertices.

Lemma 4.2. *For $n \geq 2\ell + 1 \geq 3$, every ℓ-spiderly n-vertex tree contains at most $\ell + 3$ full paths, except for the spider $S_{1,1,1,1}$ when $\ell = 2$.*

Proof. An ℓ-spiderly tree may allow many choices of the set U inducing the specified spider. We thus view an instance as a pair (T, U) and let $\overline{U} = V(T) - U$. We consider a counterexample (T, U) with the smallest n (over all ℓ), and with the smallest $|\overline{U}|$ among those minimizing n. When $\ell = 1$, an n-vertex tree has at most two full paths, except that $S_{1,1,1}$ has three full paths when $n = 4$. This is no problem, since $3 < 4 = \ell + 3$. Hence we may assume $\ell \geq 2$.

Let z be the branch vertex of $T[U]$ (we may designate any vertex of $T[U]$ as z when z is a path, as long as the distance condition is satisfied for vertices outside U). By the minimality of $|\overline{U}|$, neighbors of z are in U, and leaves of $T[U]$ are leaves of T.

Let v be a leaf of T. Since $(n-1) - (\ell - 1) = n - \ell$, the distance bound $(n - \ell - 2)/2$ satisfied by vertices in T also holds to make $T - v$ an $(\ell - 1)$-spiderly tree on $n - 1$ vertices. If $T - v = S_{1,1,1,1}$ with the forbidden parameters, then T has parameters $(n, \ell) = (6, 3)$, which violate $n \geq 2\ell + 1$. Therefore, we can apply the minimality of n to conclude that in $T - v$ there are at most $\ell + 2$ full paths. Full paths in T not containing v are full paths in $T - v$, so if v appears in at most one full path we have the desired bound. Hence we may assume that every leaf in T appears in at least two full paths.

First consider the case $\overline{U} = \varnothing$, where T is a spider. Since any leaf lies in at least two full paths, T has a branch vertex z. Let d be the degree of z in T, so $d \geq 3$. Let v be the leaf in a shortest leg of T, with length a. If $d \geq 4$, then two legs not containing v must each have length at least $n - \ell - 1 - a$, and some fourth leg with leaf w has length at least a. Summing the lengths of these four legs yields $2n - 2\ell - 2 \leq n - 1$, or $n \leq 2\ell + 1$. By the restriction to $n \geq 2\ell + 1$, equality holds, requiring $T = S_{a,a,\ell-a,\ell-a}$ and $n - \ell = \ell + 1$. If $a < \ell - a$, then exactly four full paths use v or w and a leg of length $\ell - a$, and $\ell - 2a + 1$ full paths use the two legs not containing v or w. The total is $\ell - 2a + 5$, which is at most $\ell + 3$ since $a \geq 1$. If $a = \ell - a$, then there is also one full path from v to w, but now exceeding $\ell + 3$ requires $a = 1$ and $\ell = 2$, which occurs precisely for the exceptional case $S_{1,1,1,1}$.

If $d = 3$, then $T = S_{a,b,c}$. To have each leaf in two full paths, the lengths of any two legs sum to at least $n - \ell - 1$. The union of two legs together having t vertices contains $t - (n - \ell - 1)$ full paths. Hence the number of full paths is $2(a + b + c) + 3 - 3(n - \ell - 1)$, which simplifies to $3\ell - n + 4$. Since $n \geq 2\ell + 1$, the value is at most $\ell + 3$.

Now we may assume $\overline{U} \neq \emptyset$. Let v be a leaf of T in $\overline{U}$, and let L be the leg of $T[U]$ closest to v (since neighbors of z are in U, L is well-defined). Let P be the path from v to z. The leg L must be at least as long as P, since otherwise we can enlarge U to obtain an earlier counterexample by replacing L with P in the spider without changing the tree or its number of full paths.

Since vertices of $\overline{U}$ are within distance $(n - \ell - 2)/2$ of z and full paths have length $n - \ell - 1$, the other end of any full path starting from v lies in U. The bound on the length of P implies that only one of these paths can end on L, so they end on distinct legs in $T[U]$. Let T^* be the tree obtained from $T - v$ by adding one leaf v^* to extend L. We obtain a new ℓ-spiderly instance (T^*, U^*), where $U^* = U \cup \{v^*\}$. Since T is a minimal counterexample, T^* has at most $\ell + 3$ full paths. Suppose that L has length at most $n - \ell - 3$. For any full path starting at v and ending on another leg L', we instead have a full path starting at v^* that ends on L' (since L is at least as long as P). Also $P \cup L$ may contain a full path starting at v, but it must end on L after turning away from z, and shifting the path to start closer to z on P and end closer to v^* replaces this path with one in T^* that we have not yet counted. Thus T, like, T^*, has at most $\ell + 3$ full paths.

We conclude that for any leaf $v \in \overline{U}$, the leg L in $T[U]$ closest to v has at least $n - \ell - 2$ edges. Any full path starting at v must use an edge of L. Therefore, if some full path P' in T shares no edges with L, then counting the edges in T at v, in L, and in P' yields

$$n - 1 \geq 1 + (n - \ell - 2) + (n - \ell - 1) = 2n - 2\ell - 2 \geq n - 1.$$

where the last inequality uses $n \geq 2\ell + 1$. Equality must hold throughout, so $n = 2\ell + 1$, L has $n - \ell - 2$ edges, v has distance 1 from L, and the only edge outside $P' \cup L$ is the edge e at v. Now L is not long enough to complete a full path starting at v. Since the only edge outside $P' \cup L$ is e, the path P' contains some vertex y of L. If $y \neq z$, then $V(P') - \{y\} \subseteq \overline{U}$, but one end of P' now has distance at least $(n - \ell + 1)/2$ from z. Hence P' contains z and lies in $T[U]$, but now one end of P' is within distance $(n - \ell - 1)/2$ of z, too close to v to finish a full path on that leg. Thus at most one full path starts at v, contradicting our earlier restriction. Thus every full path must share an edge with L.

Only one full path starting from v can end in L. To have two full paths from v, the spider $T[U]$ needs another leg L' where a full path from v ends after passing through z. Since v has distance at most $(n - \ell - 2)/2$ from z, the leg L' has length at least $(n - \ell)/2$. Since T has no full path edge-disjoint from L, any other leg L^* in $T[U]$ has length at most $(n - \ell - 3)/2$. Let w be the leaf at the end of L^*. Any full path from w shares an edge with L and hence must travel to z and then along L to reach length $n - \ell - 1$. Hence there can be only one such path, which contradicts the need for w to start two full paths.

We conclude that T has no such leg L^* in addition to L'. This means that a second full path from v, besides the one ending on L', must end on L. Thus $P \cup L$ has at least $n - \ell$ edges, since the edge of L incident to z is not in this full path.

Since all neighbors of z lie in U, we are now restricted to degree 2 at z. If L' contains a branch vertex, leading to a leaf v' outside U, then the argument we gave for v and L also yields at least $n - \ell$ edges in the union of L' and the path from v' to z. Now T has at least $2n - 2\ell$ edges, but with $n \geq 2\ell + 1$ this exceeds n. Hence L' contains no branch vertex. We can now shift z to the first branch vertex along L to obtain an earlier instance (T, U') with U' augmented by a neighbor of that branch vertex.

We have shown that there is no minimal counterexample. □

We have seen that when a card with the smallest diameter in an ambiguous deck has radius j, the reconstructions F and H may have the same number of j-centers when $n = 2\ell + 1$, but that forces F to be a tree. In this setting, we will forbid ambiguous decks by showing that F and H cannot have the same number of full paths. The marking process again forces F to be a tree, but to apply marking with $j = k$, again we must exclude the possibility $k = 0$.

Lemma 4.3. *If $\mathcal{D}$ is ambiguous and $n = 2\ell + 1$ with $\ell \geq 3$, then $k \geq 1$.*

Proof. Suppose $k = 0$. By Lemma 3.5, a card C with the smallest diameter is a 1-evine (double-star) that in a nonacyclic reconstruction H intersects the shortest cycle Q exactly in four consecutive vertices, and Q has length $\ell + 2$.

By Lemma 3.4, the component H' of H containing $C \cup Q$ has at most $n - 2$ vertices. With C having $n - \ell$ vertices and Q having $\ell + 2$ vertices, H' has at least $n - 2$ vertices, so equality holds. Thus H' consists of an $(\ell + 2)$-cycle plus $\ell - 3$ pendant edges at two consecutive vertices. The cards of H that are paths (with $\ell + 1$ vertices) lie along Q or start at a leaf of C not in Q. There are $\ell + 2$ such paths along Q and $2(\ell - 3)$ that start at leaves of C not in Q, for a total of $3\ell - 4$ cards that are paths.

In H' we have $n - 2$ edges. An n-vertex forest with $n - 2$ edges cannot have two isolated vertices. Since H and F have the same degree list (by Corollary 2.11), H has at most one isolated vertex. Hence the two vertices of H outside H' must be adjacent, giving both graphs $n - 1$ edges. Hence F is a tree.

The vertices of H not in C all have degree 2, except for the two leaves outside H'. Since F and H have the same degree list, the tree F grows from C only by appending edges at leaves to extend paths. That is, the central vertices of C cannot receive more incident edges, and no additional branch vertices can be created.

The 1-evine C may have only one branch vertex and be obtained from a star by appending one edge. In this case C and F are both spiders and F has at most $\ell + 3$ full paths. Otherwise, C and F have exactly two branch vertices, and they are adjacent. If at either of the two branch vertices of F there is at most one leg with length at least $(n - \ell - 3)/2$, then F is ℓ-spiderly, since outside the largest spider of degree 3 with legs emanating from the other branch vertex z are paths that reach distance at most $(n - \ell - 2)/2$ from z.

Otherwise, from the two branch vertices of F there are four legs that each have at least $\lceil(n-\ell-3)/2\rceil$ edges. Note that $n-\ell$ has opposite parity from ℓ, so these legs have at least $(n-\ell-2)/2$ edges when ℓ is odd, $(n-\ell-3)/2$ when ℓ is even.

When ℓ is odd, these four legs plus the central edge occupy $2n-2\ell-3$ edges. Since $2n-2\ell-3=n-2$, there remains only one edge to add. The subtree before adding that edge has four full paths, each consisting of a leg from each branch vertex plus the central edge. Adding one more edge creates at most three more full paths, achieved by extending one leg. Hence F has fewer than $\ell+3$ full paths when $\ell \geq 5$.

When ℓ is even, the four legs of length $(n-\ell-3)/2$ plus central edge occupy $n-4$ edges but do not create any full path. With only three edges to add, only $2+3+3$ full paths can be created, so F has fewer than $\ell+3$ full paths when $\ell \geq 6$. When $\ell=4$, in fact only seven full paths can be created, because when one of the legs is extended by two edges, the leaf will be too far away to start a full path that reaches past the opposite branch vertex.

We have shown that F has at most $\ell+3$ full paths when $\ell \geq 4$, but we found $3\ell-4$ full paths in H. Since $3\ell-4>\ell+3$ when $\ell \geq 4$, there is no ambiguous $\mathcal{D}$ with $k=0$ unless $\ell=3$. In that case $\ell-3=0$ and $n-\ell=4$, and C is just a 4-vertex path. We find that F is P_7 and contains four copies of P_4, while H is the disjoint union of a 5-cycle and an edge, containing five copies of P_4. □

Lemma 4.4. *Given $n=2\ell+1$, let $\mathcal{D}$ be an ambiguous deck with acyclic reconstruction F and nonacyclic reconstruction H. The graph F is a tree that has at most $\ell+3$ full paths.*

Proof. Let C be a card with minimum diameter; the diameter of C is $2k+2$ or $2k+3$, by Lemma 2.12, with radius $k+1$ or $k+2$, respectively. Let C have radius $j+1$, so $j \in \{k, k+1\}$. By Theorem 3.6, F has $1+d_C+\ell$ j-centers and is a tree. That is, the marking process from a center z of C marks all ℓ vertices of F outside C.

We claim that F cannot have vertices v and v' outside C at equal distance from z whose paths to z share an edge. Consider such a pair closest to z. Their paths to z cannot meet after traveling at most j steps toward z, because then there would be only one j-center that can mark them both, and it marks only one vertex. If the paths meet after traveling more than j steps toward z, then since z and its neighbors mark no vertices, v and v' have distance at least $j+3$ from z. Now the vertices next to v and v' on the paths to z are outside C and form such a pair closer to z.

We conclude that the subtree of F induced by all the vertices outside C and the vertices on their paths to z is a spider. That is, F grows from C only by extending edge-disjoint paths from z. If C has at most one branch vertex, then F is a spider, and by Lemma 4.2 F has at most $\ell+3$ full paths.

If C has at least two branch vertices, then any longest path in C has at most $n-\ell-2$ vertices. The longest path in C has $2k+3$ or $2k+4$ vertices, depending on the diameter. In either case, with radius $j+1$ for C, we obtain $j+1 \leq (n-\ell-2)/2$. All vertices of F not in the spider we have constructed are in C and hence are within distance $(n-\ell-2)/2$ of z. This makes F an ℓ-spiderly tree, and by Lemma 4.2 it has at most $\ell+3$ full paths. □

Theorem 4.5. *For $n \geq 2\ell+1 \geq 3$, the family of n-vertex acyclic graphs is ℓ-recognizable, except when $(n,\ell)=(5,2)$.*

Proof. Suppose that there is an ambiguous deck $\mathcal{D}$ with reconstructions F and H as we have been discussing. By Theorem 3.6, $n = 2\ell + 1$ and F is a tree and H has girth $\ell + 2$. By Lemma 4.4, F has at most $\ell + 3$ full paths. Hence at most $\ell + 3$ cards in the deck are paths; this is the key point that will yield a contradiction.

With $(n, \ell) = (5, 2)$ excluded and the result already known for $\ell = 1$ by the original work of Kelly [6], we may assume $n - \ell = \ell + 1 \geq 4$. Full paths have length ℓ.

By Theorem 3.6, the shortest cycle Q in H has length $\ell + 2$. There are $\ell + 2$ full paths in Q. Let C be a card with the smallest diameter. By Lemma 3.4, C and Q lie in the same component of H. Since this component has at most $n - 2$ vertices, C has at most $\ell - 3$ vertices outside Q, and hence each vertex of C is within distance $\ell - 3$ of $V(Q)$. Therefore, from each vertex of C outside $V(Q)$, one can travel to $V(Q)$ in C and then complete a full path in either direction along Q. Hence H contains at least $\ell + 2 + 2t$ full paths, where t is the number of vertices of C outside Q.

If $t \geq 1$, then this contradicts the previous conclusion that the deck has at most $\ell + 3$ cards that are paths. Since Q has no chords, we can only have $t = 0$ if C is a path contained in Q. If the smallest diameter card in $\mathcal{D}$ is a path, then F is a path. In that case F contains only $\ell + 1$ full paths, which again is fewer than the $\ell + 2$ full paths in H. □

ACKNOWLEDGMENTS

We thank the referees for their careful reading and many helpful comments. Alexandr V. Kostochka was supported by NSF grants DMS-1600592 and DMS-2153507. Douglas B. West was supported by National Natural Science Foundation of China grants NSFC 11871439, 11971439, and U20A2068. Dara Zirlin was supported by Arnold O. Beckman Campus Research Board Award RB20003 of the University of Illinois and by NSF RTG grant DMS-1937241.

REFERENCES

1. J. A. Bondy and R. L. Hemminger, *Graph reconstruction—A survey*, J. Graph Theory. **1** (1977), 227–268.
2. W. B. Giles, *Reconstructing trees from two-point deleted subtrees*, Discrete Math. **15** (1976), 325–332.
3. D. L. Greenwell and R. L. Hemminger, *Reconstructing the n-connected components of a graph*. Aequationes Math. **9** (1973), 19–22.
4. C. Groenland, T. Johnston, A. Scott, and J. Tan, *Reconstructing trees from small cards*, arXiv 2103.13359.
5. P. J. Kelly, *On isometric transformations*, Ph.D. thesis, University of Wisconsin-Madison, 1942.
6. P. J. Kelly, *A congruence theorem for trees*, Pacific J. Math. **7** (1957), 961–968.
7. A. V. Kostochka, M. Nahvi, D. B. West, and D. Zirlin, *Degree lists and connectedness are 3-reconstructible for graphs with at least seven vertices*, Graphs Combin. **36** (2020), 491–501.
8. A. V. Kostochka, M. Nahvi, D. B. West, and D. Zirlin, *Trees with at least* $6\ell + 11$ *vertices are* ℓ*-reconstructible*, arXiv 2307.10035.
9. A. V. Kostochka and D. B. West, *On reconstruction of graphs from the multiset of subgraphs obtained by deleting* ℓ *vertices*, IEEE Trans. Inform. Theory. **67** (2021), 3278–3286.
10. B. Manvel, *On reconstruction of graphs*, G. Chartrand, and S.F. Kapoor (eds.), The Many Facets of Graph Theory (Proc. Conf. Western Mich. Univ., Kalamazoo, MI, 1968), Springer, Berlin-New York, 1969, pp. 207–214.
11. B. Manvel, *Some basic observations on Kelly's conjecture for graphs*, Discrete Math. **8** (1974), 181–185.
12. V. Nýdl, *Finite graphs and digraphs which are not reconstructible from their large cardinality restricted subraphs*, Comment. Math. Univ. Carolin. **22** (1981), 281–287.
13. V. Nýdl, *A note on reconstructing of finite trees from small subtrees*, Acta Univ. Carol. Math. Phys. **31** (1990), 71–74.

14. H. Spinoza and D. B. West, *Reconstruction from the deck of k-vertex induced subgraphs*, J. Graph Theory. **90** (2019), 497–522.
15. R. Taylor, *Reconstructing degree sequences from k-vertex-deleted subgraphs*, Discrete Math. **79** (1990), 207–213.
16. S. M. Ulam, *A collection of mathematical problems*, Interscience Tracts in Pure and Applied Mathematics, vol. **8**, Interscience Publishers, New York-London, 1960.
17. D. B. West, Combinatorial mathematics, Cambridge University Press, Cambridge, 2021, xx+969pp.
18. D. Zirlin, *Cycle Structure of Graphs and Hypergraphs: Extremal Problems and Reconstruction*, Ph.D. thesis, University of Illinois, 2022.

How to cite this article: A. V. Kostochka, M. Nahvi, D. B. West, and D. Zirlin, *Acyclic graphs with at least $2\ell + 1$ vertices are ℓ-recognizable*, J. Graph Theory. **109** (2025), 184–200. https://doi.org/10.1002/jgt.23027