CHAPTER 3

Digraph Roots

1. Matrices and Digraphs, Powers and Roots

Consider a directed graph D (or digraph, for short) on n vertices together with its adjacency matrix, i.e., the $n \times n$ matrix $A = (a_{ij})$ with $a_{ij} = 1$ iff there is an arc $j \to i$ in D and $a_{ij} = 0$ otherwise. We can use iterated multiplications of the adjacency matrix with itself to find paths in the digraph. Precisely, the (i, j)-entry of the kth power A^k of A is positive iff there is a walk of length *exactly* k from vertex i to vertex j in D. By *walk* of *length* k we mean a sequence $(v_0, v_1, \ldots, v_k)$ of $k + 1$ vertices with an arc from v_{i-1} to v_i for $1 \leq i \leq k$, where vertices may appear several times; in contrast to a *path*, which is a walk with all vertices distinct. We are only interested in the existence of such walks, not their number—which is counted by the respective entry of A^k—so we interpret A as a Boolean 0/1-matrix with the product $C = A \cdot B$ defined in the usual way:

$$c_{ij} = \bigvee_{h=1}^{n} a_{ih} \wedge b_{hj}.$$

Identifying a digraph with its adjacency matrix, we define the kth power, $k \in \mathbb{N}$, of a digraph D to be the digraph D^k on the same vertex set and with an arc from a to b if and only if there is a directed walk of length *exactly* k from a to b in D (possibly visiting some vertices several times). Figure 1 shows an example.

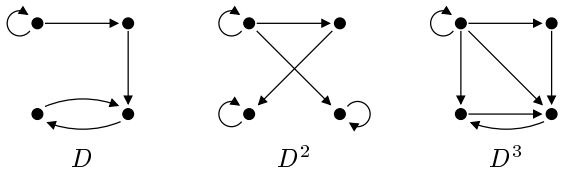

FIGURE 1. Powers of a digraph.

Note that the interpretation of digraphs as Boolean matrices implies that our digraphs may have loops but no multiple arcs. It is easy to see that the adjacency matrix of D^k is in fact the kth Boolean power of the adjacency matrix of D (see, for example, [**38**]).

Boolean matrix algebra serves as a fundamental tool in algorithmic graph theory. The correspondence between graphs and matrices lies at the heart of many fundamental algorithms for transitive-closure or shortest-path computations [**35, 1, 12**] (where usually powers of the matrix $A + I$, with I the identity matrix, are considered to account for all paths *up to* a certain length k).

We are interested in the inverse operation to exponentiation: root finding. The complexity of the following problem was open until now.

THE BOOLEAN-MATRIX-ROOT PROBLEM. Given a Boolean $n \times n$ matrix A and an integer $k \geq 2$, does there exist a kth *root* B of A, that is, an $n \times n$ matrix B with $B^k = A$.

Or equivalenty, stated in terms of digraphs:

THE DIGRAPH-ROOT PROBLEM. Given a digraph D and an integer $k \geq 2$, does there exist a kth *root* R of D, that is, a digraph R on the same vertex set, with $R^k = D$.

Twenty years ago, in the open-problems section of his book [**26**], Kim inquired for the special case $k = 2$, whether the Boolean-matrix-root problem might perhaps be NP-complete. We answer this question in the affirmative.

1. THEOREM. *Deciding whether a square Boolean matrix or, equivalently, a digraph has a* kth *root is* NP*-complete for each single parameter* $k \geq 2$.

With the right computational problem for the reduction, the proof of this result turns out surprisingly simple. This is quite remarkable since it thus relates digraph roots very closely to a well-known NP-complete problem. It allows to identify quite accurately "the reason" for the hardness of the problem. In an attempt to isolate and inhibit these computationally difficult aspects, we shall discover a close connection between digraph roots and graph isomorphism, which eventually leads to a further complexity result (Theorem 3). But let us postpone these issues till after the proof and discussion of Theorem 1.

Related work—related questions. Over the field of complex numbers or the reals, matrix roots are a well-studied and still up-to-date topic of linear algebra [**29, 24, 36**]. But results from that field of research do generally not apply to Boolean matrices. While it is known, for example, that every regular matrix over the complex numbers has a kth root for any $k \geq 2$ [**36**], this is not true for Boolean matrices, as the invertible matrix $\left(\begin{smallmatrix} 0 & 1 \\ 1 & 0 \end{smallmatrix}\right)$ shows. Further, complex or real matrices are amenable to numerical methods like Newton iteration [**22**], whereas such techniques clearly do not apply to Boolean matrices. When it comes to roots, Boolean matrices don't seem to have much in common with matrices over $\mathbb{C}$ since the former behave much more rigidly than the latter.

The situation is, however, different if we ask for powers of a matrix instead of roots. There are theoretical results on Boolean matrix powers [**13**] and in practice we can, of course, compute the kth power of a Boolean matrix A by treating it as a matrix over the reals. We calculate A^k over $\mathbb{R}$ and afterwards replace each positive entry with 1. This simple reformulation allows us, for example, to apply fast matrix multiplication methods such as Strassen's to path problems in graphs [**35, 1**]. But this simulation through matrices over the reals clearly only works because there cannot happen cancellation between positive and negative entries. For root finding, such simulation over $\mathbb{R}$ or $\mathbb{C}$ would lead into major problems.

FIGURE 2. A directed square root (left) of a symmetric digraph (right) which does not have a symmetric square root.

Alternative notions of graph powers. A problem similar to the one at hand has been discussed by Motwani and Sudan. In [**34**] they showed that computing square roots of undirected graphs is NP-hard. But their notion of graph powers differs from ours in two important points.

They consider undirected graphs only, which in our setting would correspond to symmetric digraphs, i.e., all edges are bidirectional. This not only restricts the set of possible inputs but also—and this is the decisive difference—the solutions. For example, the symmetric digraph on the right of Figure 2 has the digraph to its left as a square root, but it is not the square of any symmetric digraph. To see this, observe that any square root of an undirected graph with maximum degree strictly greater than 2 must also have a vertex of degree at least 3. Such a vertex would in turn induce a triangle in the square. The digraph in the figure has maximum degree 3 but it does not contain a triangle.

Further, Motwani and Sudan define squaring to maintain existing edges, which in our setting would correspond to attaching loops to all vertices. This monotonicity ensures that much information of the underlying graph can be read off from its square and the hardness proof of [**34**] makes essential use of this property. In contrast to this, squaring a digraph under the rules derived from Boolean matrix multiplication can almost completely destroy the neighborhood information and may even decompose the digraph. Actually, most of our arguments depend crucially on such vanishing edges. So apparently, the squares in [**34**] and our notion of powers are fundamentally different concepts.

Nomenclature. In the light of the preceding discussion, Boolean matrices form the right framework to ask questions about roots in the sense we defined them. They do not leave the ambiguities that the expression "graph root" obviously has and locate the problem correctly in the context of semigroups. However, for the actual work, the proof of Theorem 1, we will resort to the language of graph theory since our arguments will extensively use respective notions like paths, cycles, and vertex neighborhoods. Moreover, after the NP-completeness proof we shall emphasize the link to graph theory even more by relating our roots to graph-isomorphism.

So let us agree on the precise meanings of some common graph theoretic notions whose exact distinction will be crucial in certain situations. A *walk* is simply a sequence $(a_0, a_1, \ldots, a_r)$ of vertices with an arc $a_i \to a_{i+1}$ for $0 \leq i < r$, whereas a *path* is a walk of pairwise distinct vertices. The parameter r is the length of the walk respectively path. A *cycle* is a closed walk, that means, $a_0 = a_r$ and vertices may be traversed several times. By

isolated cycle we mean a strongly connected component of a digraph where each vertex has indegree and outdegree 1, i.e., a single non-self-touching cycle without further arcs.

For a digraph R on vertex set V we let

$$R(v) := \left\{ w \in V \;\middle|\; v \xrightarrow{R} w \right\}$$

denote the set of *outneighbors* of v in R. Defining $\bar{R}$ to be the digraph obtained from R by inverting all arcs, we write $\bar{R}(v)$ for the *inneighbors* of v. Note that our generalization

$$R(U) := \bigcup_{u \in U} R(u)$$

to subsets $U \subseteq V$ diverts from standard notation (as for example in [**3**, sec. 1.2]) as $R(U)$ need not be disjoint from U. In other words, we let a digraph act on vertex sets just like its adjacency matrix acts on the characteristic vectors of such sets.

These definitions help simplify our notation. For example, we write $x \in \bar{R}^j(Y)$ to state that there is a walk of length j from x to some vertex in $Y \subseteq V$ and expressions like $R^3\bar{R}^8R$ make perfect sense, encoding some kind of zig-zag walk through the digraph R.

2. NP-Completeness

This section comprises the proof of Theorem 1; but before turning to the details, presenting a suitable NP-complete problem which we can reduce to digraph roots, let us collect some motivating observations about digraph square roots.

Consider some set X of vertices of a digraph D and let Z denote all outneighbors of vertices in X. Assume for simplicity that X and Z are disjoint, so in particular, there are no loops or cycles on these vertices. In a square root of the digraph D, any of the arcs from X to Z must be realized as paths of length two. Hence, the root must provide a set Y of intermediate vertices through which all these paths can pass. If now—for whatever reason—there is only a small number of such intermediate vertices available, $|Y| \leq r$, say, with r a little smaller than $|X|$ and $|Z|$, these paths must intersect in order to ship all their information from X to Z. This situation is almost exactly captured by the following decision problem, which is already listed in Garey and Johnson's classic [**18**] (p. 222).

The Set-Basis Problem. Let $\mathfrak{C}$ be a collection of subsets of some finite set S. A *set basis* for $\mathfrak{C}$ is another collection $\mathfrak{B}$ of subsets of S such that each $C \in \mathfrak{C}$ can be written as a union of sets from $\mathfrak{B}$. Given a finite set S, a collection $\mathfrak{C}$ of subsets of S, and an integer $r \leq |S|$, the *set-basis problem* asks whether there exists a set basis $\mathfrak{B}$ for $\mathfrak{C}$ consisting of at most r sets. This problem is known to be NP-complete [**40**].

We claim that the local configuration of the above square-root problem is nothing but a set-basis instance. The sets X and Z correspond to the given collection $\mathfrak{C}$ and the ground set S, respectively, while the intermediate vertex set Y takes the place of the sought-after collection $\mathfrak{B}$.

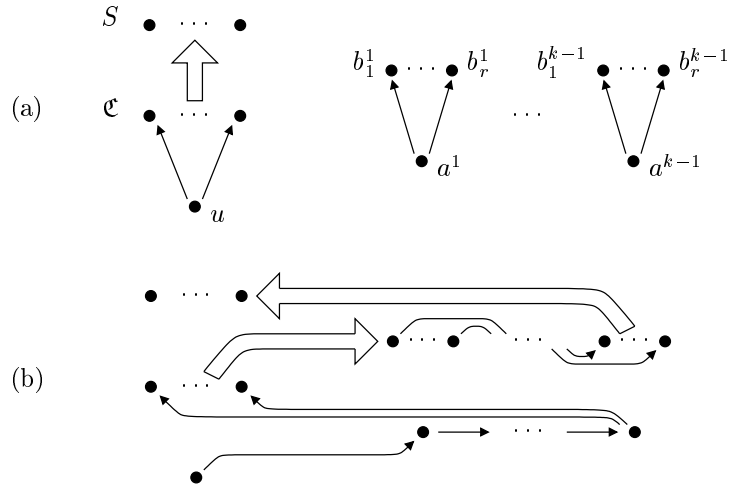

FIGURE 3. Reducing set basis to kth root (a) and encoding a set basis as a root (b). (Wide arrows represent collections of arcs that depend on the actual instance.)

Our precise proof of this claim, which also treats the general case of arbitrary kth roots, comes in the three customary parts: a reduction from a set-basis instance to a kth-digraph-root instance and the two complementary transformations between valid solutions.

The reduction. From a set-basis-problem instance $(\mathfrak{C}, S, r)$ we construct a digraph D such that D has a kth root iff $\mathfrak{C}$ has a set basis $\mathfrak{B}$ of size at most r. We may assume w.l.o.g. that neither the collection $\mathfrak{C}$ nor any $C \in \mathfrak{C}$ be empty, that all $C \in \mathfrak{C}$ be pairwise distinct, and further that $\bigcup \mathfrak{C} = S$, i.e., each $s \in S$ lie in some set $C \in \mathfrak{C}$.

As suggested by the above discussion, our construction essentially draws the containment graph of the set system $\mathfrak{C}$ on S and provides the right number of intermediate vertices. Surprisingly few framework arcs will have to be added in order to ensure that any root uses them as intended.

We start with the containment relations. The digraph D possesses the sets $C \in \mathfrak{C}$ and the elements $s \in S$ as vertices and additionally an "anchor vertex" u. Define the containment arcs

$$C \xrightarrow{D} s \qquad \text{for all pairs } (C, s) \in \mathfrak{C} \times S \text{ with } s \in C \tag{14}$$

and additionally the grounding arcs

$$u \xrightarrow{D} C \qquad \text{for each } C \in \mathfrak{C}.$$

Compare the left component of Figure 3(a).

The intermediate vertices come in $k-1$ isomorphic components which are simply stars. The μth component consists of the $r+1$ vertices $a^\mu, b_1^\mu, b_2^\mu, \ldots, b_r^\mu$ connected via

$$a^\mu \xrightarrow{D} b_i^\mu \qquad \text{for } i \in \{1, \ldots, r\},$$

as shown in the right half of Figure 3(a).

Constructing a root from a set basis. To show that our construction works, we describe how to obtain a kth root R of the digraph D from a set basis of size r for $\mathfrak{C}$. Therefore we first need a lot of framework arcs that are independent of the actual basis $\mathfrak{B}$: the horizontal paths

$$u \xrightarrow{R} a^1 \xrightarrow{R} a^2 \xrightarrow{R} \cdots \xrightarrow{R} a^{k-1}$$

and

$$b_i^1 \xrightarrow{R} b_i^2 \xrightarrow{R} \cdots \xrightarrow{R} b_i^{k-1} \qquad \text{for each } i \in \{1, \ldots, r\},$$

and also the back connections

$$a^{k-1} \xrightarrow{R} C \qquad \text{for each } C \in \mathfrak{C};$$

drawn as thin arcs in Figure 3(b).

The remaining arcs depend on the given set basis $\mathfrak{B} = \{B_1, \ldots, B_r\}$, which comes with a representation

$$C = \bigcup_{i \in I_C} B_i, \qquad I_C \subseteq \{1, \ldots, r\} \tag{15}$$

of each set $C \in \mathfrak{C}$.

Note that a basis with less than r sets can be extended to one of size r by adding singleton sets $\{s\} \subseteq S$ and it is also clear that we can pick the collection $\mathfrak{B}$ and the index sets I_C in such a way that each index $i \in \{1, \ldots, r\}$ appears in at least one I_C.

The set basis $\mathfrak{B}$ is now wired via

$$b_i^{k-1} \xrightarrow{R} s \qquad \text{for each pair } (i, s) \text{ with } s \in B_i,$$

while the corresponding representations are realized as

$$C \xrightarrow{R} b_i^1 \qquad \text{for each index } i \in I_C.$$

These connections appear bundled as wide arrows in Figure 3(b).

These definitions guarantee that there exists an R-walk of length k from a certain C to some $s \in S$ iff there exists any basis set B_i with $s \in B_i$ and $i \in I_C$. By the definition of a set basis, the latter condition is equivalent to $s \in C$, which, by construction of the digraph D, means just that there is a D-arc from C to s. Thus we have shown that R^k equals D on $\mathfrak{C} \times S$. The identity of these two digraphs on the remaining vertices is immediate.

Getting a set basis from a root. We turn to the other, slightly more intricate implication. Let D be the digraph constructed from a given set-basis instance $(\mathfrak{C}, S, r)$ and let R be any kth root of D. From this root we must obtain a set basis $\mathfrak{B}$ for $\mathfrak{C}$ with at most r sets. The basic idea is, of course, to show that the root R must look essentially as the one we constructed in the preceding paragraph.

First of all, observe that cycles in R would induce cycles in any positive power of R. Thus, R contains no cycles. Now consider an arbitrary vertex $C \in \mathfrak{C}$. Since $u \to C$ in D, there must be an R-walk of length k from u to C. We claim that all interior vertices of any such walk P are from the set $\{a^1, \ldots, a^{k-1}\}$. To see this, pick any interior vertex x on P. Clearly x must have positive outdegree in D because C has. So x can only be some a^μ or from the set $\mathfrak{C}$; the remaining alternative $x = u$ would yield a cycle. Assume for contradiction that $x \in \mathfrak{C}$. Then there is a path Q of length k in

R from u to x. Because x was assumed to be an inner vertex on the path P, a certain inner vertex y on Q is at distance $-k$ from C. This means $y = u$, which implies that the vertex u lies on an R-cycle—a contradiction.

So all interior vertices of R-walks from u to some $C \in \mathfrak{C}$ are from the set $\{a^1, \ldots, a^{k-1}\}$. Obviously, any such path must use each of these a^μ exactly once since otherwise there would be cycles. Furthermore, all such paths pass the a^μ in the same order, again because two different orders would yield cycles. We may assume by symmetry that the a^μ are traversed from a^1 through a^{k-1}. Thus we see that $R(a^{k-1}) = \mathfrak{C}$ and conclude

$$R^{k-1}(\mathfrak{C}) = R^{k-1}R(a^{k-1}) = D(a^{k-1}) = \{b_1^{k-1}, \ldots, b_r^{k-1}\}.$$

So all R-walks from $\mathfrak{C}$ to S pass through these b_i^{k-1}. We focus on the ultimate edges on any such walk and define

$$B_i := R\big(b_i^{k-1}\big) \qquad \text{for } 1 \le i \le r.$$

We claim that $\mathfrak{B} := \{B_1, \ldots, B_r\}$ is a set basis for $\mathfrak{C}$. This is easily verified. Reading the defining relation (14) as

$$C \xrightarrow{R^k} s \quad \iff \quad s \in C,$$

one sees that the index sets

$$I_C := \big\{ i \bigm| b_i^{k-1} \in R^{k-1}(C) \big\}$$

yield basis representations of the sets $C \in \mathfrak{C}$ as in Equation (15).

This concludes the proof of Theorem 1.

We emphasize that the given set-basis instance is completely maintained by our reduction. Its containment relations are encoded one-to-one by arcs of the digraph. Thus, on the large scale, an instance of the digraph-root problem can be seen as a collection of many interacting set-basis problems. One might well argue that finding digraph roots is actually a generalized set-basis problem.

As a corroboration for this point of view we mention that the set-basis problem already appeared before in connection with Boolean matrix algebra. Markowsky [**31**] used it in a very economic proof for the NP-completeness of Schein-rank computation.[1]

3. Roots and Isomorphism

Let us carry the concluding remarks of the preceding section a little further and have a closer look at Figure 3 from page 75 again. The construction there required only paths of length 2, which then induced a few long paths in the root. One could say that the computational complexity of root finding results from the described interaction of many very short paths. In some sense, our proof has exploited a *local* phenomenon. If we suppress the local interaction by some restriction on the digraph, maybe we can find some further properties of digraph roots that live on a *global* scale. Here is our approach.

[1] Analogous to the matrix rank over fields, the *Schein rank* of a Boolean matrix A is the minimal integer ρ such that A can be represented as a Boolean sum $A = \bigvee_{i=1}^{\rho} c_i r_i$, where the c_i are column and the r_i row vectors with zero-one entries [**26**, Sec. 1.4].

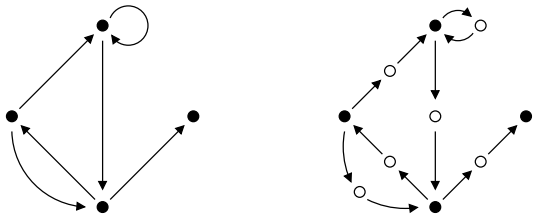

FIGURE 4. The complete subdivision (right) of a digraph (left).

2. DEFINITION. The *complete subdivision* of a digraph D is the digraph obtained from D by replacing each arc $a \to b$ of D by a new vertex x_{ab} and the two arcs $a \to x_{ab} \to b$. (See Figure 4.) We call a digraph a *subdivision digraph* if it is (isomorphic to) the complete subdivision of some digraph.

Subdivisions are a fundamental notion in graph theory but opposed to their common usage in relation with topological minors, we employ them here to equip our digraphs with a certain stiffness. The effect is the desired inhibition of the local interaction we exploited in the NP-completeness proof. However, the problem of root finding for such subdivision digraphs does not become trivial. Instead, the following surprising relation to graph isomorphism shows up.

3. THEOREM. *Deciding whether a subdivision digraph with positive minimal indegree and outdegree has a* kth *root, is graph-isomorphism complete for each single parameter* $k \geq 2$.

The *graph-isomorphism problem* asks whether two given (di)graphs are isomorphic or not, i.e., whether there exists an arc-preserving bijection between their vertex sets.[2] No polynomial-time algorithm for this problem is known, neither is it known to be NP-complete. On the contrary, it is a prime candidate for a problem strictly between P and NP-completeness (cf. [**27**] and [**30**]). Computational problems of the same complexity as the graph-isomorphism problem are called *graph-isomorphism complete*, or simply *isomorphism complete* because isomorphism problems for several algebraic or combinatorial structures fall into this class. For example, isomorphism of semigroups and finite automata [**9**], finitely represented algebras, or convex polytopes [**25**]. Other problems ask for properties of the automorphism group of a graph, for example, computing the order of this group or its orbits [**33**].[3] Finally, several restrictions of the graph-isomorphism problem are known to remain isomorphism complete, as for example isomorphism of regular graphs [**9**].

As the above list indicates, actually all problems known to be isomorphism complete are more or less obviously isomorphism problems of various combinatorial structures. Hence, the relation between digraph roots and

[2]One usually considers undirected graphs but it is well-known and easily seen that with respect to their computational complexity the undirected and directed version of the problem are equivalent.

[3]The latter two problems are known to be isomorphism complete only in the weaker sense of Turing reduction, as opposed to the concept of many-one reduction.

graph isomorphism we are going to establish in our proof of Theorem 3 may come quite as surprise.

From isomorphisms to roots. Theorem 3 rests on a structural result (Theorem 6) which states that any kth root of a subdivision digraph D establishes isomorphisms between the components of D. This is just the kind of global structure we wanted to find when we defined subdivision digraphs.

The starting point is the following connection between digraph roots and digraph isomorphism, which holds for arbitrary digraphs. Subdivisions will then be needed to obtain a converse of this result.

4. Proposition. *Let $D = D_1 \dot\cup D_2 \dot\cup \cdots \dot\cup D_k$ be the disjoint union of k isomorphic digraphs $D_1, \ldots, D_k$. Then D has a kth root.*

Proof. We construct a digraph R on the vertices of D with $R^k = D$. Pick isomorphisms $\varphi_i \colon D_1 \to D_i$, $1 \le i \le k$ (φ_1 being simply the identity). For each vertex a of D_1 we let R contain the path

$$\varphi_1(a) \xrightarrow{R} \varphi_2(a) \xrightarrow{R} \cdots \xrightarrow{R} \varphi_k(a) \tag{16}$$

and additionally the arcs

$$\varphi_k(a) \xrightarrow{R} \varphi_1(b) \quad \text{for all } b \in D_1(a). \tag{17}$$

Figure 5 shows a local picture of this construction.

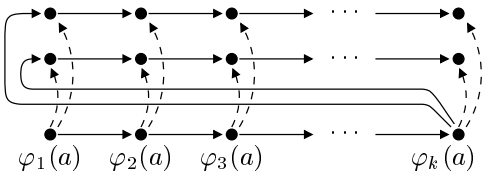

Figure 5. Constructing a kth root (continuous lines) for a disjoint union of k isomorphic digraphs (dashed lines).

We claim that $R^k = D$. To see this, pick any $v \in D_i$, $1 \le i \le k$, and compute

$$\begin{aligned} R^k(v) &= R^i \varphi_k \varphi_i^{-1}(v) && \text{by (16)} \\ &= R^{i-1} D_1 \varphi_i^{-1}(v) && \text{by (17)} \\ &= \varphi_i D_1 \varphi_i^{-1}(v) && \text{by (16)} \\ &= D_i(v) = D(v), \end{aligned}$$

treating digraphs and isomorphisms equally as mappings between subsets of the vertex set. □

From roots to isomorphisms. Note how the root arcs in the above construction encode the isomorphism between the components of the digraph D. Our goal is to show that for a subdivision digraph, *any* root establishes isomorphisms between the weakly connected components of this digraph in exactly the same way. Before we can embark on this venture, however, we have to take care of some degenerate cases that do not fit into this picture.

Usually in a subdivision digraph one can easily distinguish the original vertices, sometimes called *branching vertices*, from the newly inserted *subdivision vertices*. In fact, a subdivision digraph is obviously bipartite and as soon as every weakly connected component contains at least one vertex whose indegree or outdegree differs from 1, the two classes can be uniquely identified.

A problem arises with subdivision digraphs that contain isolated cycles. In such components, all vertices look like subdivision vertices and this absence of clearly identifiable branching vertices leads to untypical behavior with respect to root finding. Fortunately, isolated cycles are simple objects and we can completely describe their powers.

5. Lemma. *The* kth *power of an isolated cycle of length* r *is the disjoint union of* $\gcd(r,k)$ *isolated cycles of length* $r/\gcd(r,k)$.

Proof. For every vertex x on an isolated cycle C, the sets $C^k(x)$ and $\bar{C}^k(x)$ are singletons. So each vertex of C^k has in– and outdegree 1, that means, C^k is the disjoint union of isolated cycles and by symmetry, all these cycles are of the same length. To determine this common length, start at an arbitrary vertex a and walk around C until you first reach a again in a multiple ℓ of k steps. Clearly, ℓ is the least common multiple of r and k; so the length of a cycle in C^k is

$$\frac{\ell}{k} = \frac{\operatorname{lcm}(r,k)}{k} = \frac{r}{\gcd(r,k)}. \qquad \square$$

As a consequence of Lemma 5, isolated cycles cannot have the isomorphism property we are looking for. But this is no problem. We shall show later that any vertex on an isolated cycle of a subdivision digraph D must also lie on an isolated cycle in any root of D. Thus, with respect to roots, cycle vertices do not interact with vertices from the other components of a subdivision digraph and we may in the following restrict our attention to subdivision digraphs without isolated cycles.

Ignoring isolated cycles we can show that subdivision digraphs bear the desired isomorphism structure—under the unfortunately indispensable additional condition that each vertex has at least one inneighbor and one outneighbor. We shall prove the following theorem.

6. Theorem. *A subdivision digraph without isolated cycles and with positive minimal indegree and outdegree has a* kth *root if and only if it is the disjoint union of* k *isomorphic digraphs.*

The basic idea for the proof of Theorem 6 is to show that in any kth root of a subdivision digraph, subdivision vertices and branching vertices appear in blocks of length k. More precisely, we will show that any subdivision

vertex of D lies on an R-path of length k that consists only of subdivision vertices (of D) and likewise for branching vertices.

A direct proof of this statement, however, appears quite difficult since "subdivision vertex" is a semantic concept depending on the global structure of the digraph. Therefore we work with the simple local properties of subdivision vertices that can easily be dealt with.

7. Definition. We call a vertex of a digraph *thin* if its indegree and outdegree are 1; otherwise we call it *proper*.

The second step in our analysis will be to identify root arcs that are unique for their incident vertices, thus establishing unique correspondences that will be needed to identify the sought-after isomorphisms.

8. Definition. We call an arc ab of a digraph R *strong* if no further arcs leave a or enter b, i.e., $R(a) = \{b\}$ and $\bar{R}(b) = \{a\}$. More generally, a walk is called *strong* if all of its arcs are strong.

Most of the forthcoming proofs will be indirect, leading to contradictions to the following trivial observation about subdivision digraphs, which expresses the simple fact that digraphs, as we define them, cannot have parallel edges.

9. Observation. *No two vertices in a subdivision digraph have a common inneighbor and a common outneighbor.* □

A general remark to avoid confusion. As before, we shall deal with two different digraphs on the same vertex set. When we talk about subdivision and branching vertices or thin and proper vertices, these notions shall always refer to (the arcs of) the subdivision digraph D. On the other hand, the term "strong" will always refer to arcs of the root R.

For technical reasons we provide the lemmas about unique arcs first and construct the long paths afterwards, since the latter rely on the former. Here is our first criterion for strongness of root arcs:

10. Lemma. *In a root R of a subdivision digraph D, any R-arc between two D-thin vertices is strong.*

Proof. Consider any pair a, b of D-thin vertices with $a \to b$ in R. As a thin vertex, a must also have at least one outneighbor in R, so assume for contradiction that $\deg^+_R(a) > 1$, i.e., there exists some $c \neq b$ with $a \to c$ in R. By symmetry, the case $\deg^-_R(b) \neq 1$ reduces to this situation by reversing all arcs.

The unique vertex u in $\bar{R}^{k-1}(a)$ has at least two D-outneighbors, b and c. Hence, this u is proper and therefore c is thin. So b and c are both thin and the sets $R^{k-1}(b)$ and $R^{k-1}(c)$ must therefore be nonempty. From $R^{k-1}(b) \cup R^{k-1}(c) \subseteq R^k(a)$ we thus conclude that $R^{k-1}(b) = R^{k-1}(c) = \{v\}$, where v is the unique D-outneighbor of a. Altogether, we have found two vertices, b and c, with common in– and common outneighbors—a contradiction to Observation 9. □

One could actually relax the preconditions in Lemma 10 but its present form is sufficient for our purposes and it will fit quite naturally into its later applications.

There is an analog of Lemma 10 for proper vertices but it requires an explicit minimal-degree condition that was trivially met by thin vertices. Actually there can be non-strong arcs between pairs of proper vertices. So it is in the following lemma where the additional degree condition of Theorem 3 enters.

11. Lemma. *In a root R of a subdivision digraph D, any R-arc between two D-proper vertices that have each at least one in– and one outneighbor is strong.*

Proof. Consider any pair a, b of D-proper vertices with $a \to b$ in R. Assume for contradiction that there exists some $c \neq b$ with $a \to c$ in R. Again, the case $\deg^-_R(b) > 1$ reduces to this situation. Since a has a D-inneighbor, the set $\bar{R}^{k-1}(a)$ is nonempty. But any vertex from this set is an inneighbor of two vertices, one of which is proper. An impossible configuration in a subdivision digraph. □

The preceding two lemmas provide us with a simple procedure to identify R-walks of D-thin or D-proper vertices. Starting from a thin vertex a_0 of D, we check whether there is some D-thin outneighbor a_1 of a_0 in R. If such an a_1 exists it must be unique by Lemma 10. Next check for a D-thin outneighbor a_2 of a_1 and iterate this process until some ultimate a_t has no further D-thin outneighbors in R. Likewise we may search for inneighbors, altogether constructing a unique maximal R-walk of D-thin vertices containing a_0—provided we don't run into cycles. Analogously, we can find unique maximal walks of proper vertices.

We have now all necessary prerequisites to prove that thin vertices and proper vertices come in blocks.

12. Lemma. *Let R be a kth root of a subdivision digraph D and let $a_0 \to a_1 \to \cdots \to a_\ell$ be an R-walk of length $\ell \leq k$ between two D-thin vertices a_0 and a_ℓ. Then all intermediate a_i, $0 < i < \ell$, are also thin.*

Proof. We pick an arbitrary index j between 0 and ℓ and show that a_j is a thin vertex. Therefore first observe that the sets $R^k(a_j)$ and $\bar{R}^k(a_j)$ are nonempty because a_0 and a_ℓ are thin. We assume for contradiction that a_j is a proper vertex, so one of those two sets must contain at least two elements. By symmetry assume that $|R^k(a_j)| > 1$; so let x, y be two different elements from this set.

Denote the unique vertex in $R^k(a_0)$ by v. Since $R^{k-j}(a_j), R^{k-\ell}(a_\ell) \subseteq R^k(a_0)$, we get precisely

$$R^{k-j}(a_j) = \{v\} = R^{k-\ell}(a_\ell).$$

The first identity tells us that from a_j the two vertices $x, y \in R^k(a_j)$ are only reachable via v, i.e., $x, y \in R^j(v)$, and together with the second identity this implies

$$x, y \in R^{k-\ell+j}(a_\ell). \tag{18}$$

See Figure 6.

Since a_ℓ is thin, the set $R^k(a_\ell)$ contains exactly one vertex, w, say. Thus, by (18), we have $R^{\ell-j}(x) \cup R^{\ell-j}(y) \subseteq \{w\}$. As neighbors of the proper vertex

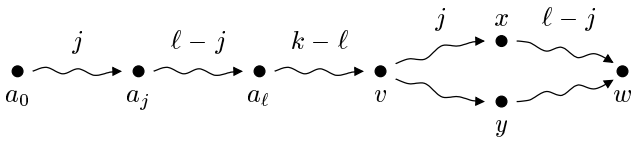

FIGURE 6. Path construction from the proof of Lemma 12.

a_ℓ the vertices x and y must be thin, so the sets $R^{\ell-j}(x)$ and $R^{\ell-j}(y)$ are nonempty and we actually get $R^{\ell-j}(x) = R^{\ell-j}(y) = \{w\}$, which implies $R^k(x) = R^k(y)$. Altogether, x and y have the common D-inneighbor a_ℓ and also a common outneighbor, in contradiction to Observation 9. □

13. LEMMA. *Let R be a kth root of a subdivision digraph D and let $a_0 \to a_1 \to \cdots \to a_\ell$ be an R-walk of length $\ell \leq k$ between two D-proper vertices a_0 and a_ℓ. Then all intermediate a_i, $0 < i < \ell$, are also proper.*

PROOF. Assume for contradiction that some a_j is a thin vertex. Then $R^k(a_j)$ is nonempty, so we may pick some $u \in R^{k-j}(a_j)$ together with some R-walk P of length $k - j$ from a_j to u. As a D-outneighbor of the proper vertex a_0 the vertex u is thin. Thus, by Lemma 12, all vertices on the walk P are in fact thin and Lemma 10 then implies that this walk is strong. Therefore the set $R^{\ell-i}(a_i)$ contains exactly one vertex, which can only be a_ℓ. But this vertex was assumed to be proper. □

The proofs of Lemmas 10 through 13 show very graphically how the local properties of subdivision digraphs are exploited on the way to Theorems 3 and 6. They all employ a kind of squeezing technique along R-paths, leading to the unique identification of certain vertices or a contradiction involving too many neighbors of a subdivision vertex.

Combining the homogeneous paths provided by Lemmas 12 and 13 with the uniqueness statements of Lemmas 10 and 11, we are now able to construct isomorphisms from roots.

PROOF OF THEOREM 6. We already know from Proposition 4 that the disjoint union of k isomorphic digraphs has a kth root. So it remains to decompose D into k isomorphic subgraphs $D_1, \ldots, D_k$ and to provide isomorphisms between them. We do this by partitioning the whole vertex set into blocks of size k, such that each block contains exactly one vertex from each D_i.

For each proper vertex a of D, determine the maximal R-walk P_a through a that consists entirely of D-proper vertices, as described in connection with Lemmas 10 and 11. Such a walk cannot extend indefinitely, precisely, it consists of at most k vertices because all vertices at distance k from a proper vertex are thin. On the other hand, P_a must have at least k vertices because otherwise its thin neighbors would, by Lemma 12, force all its vertices to be thin, too.

For a thin vertex b we proceed similarly. Determine the maximal R-walk Q_b through b that consists entirely of D-thin vertices. Again, such a walk is bounded by some proper vertices to its left and right because otherwise

we would get a cycle of thin vertices, which we excluded in the statement of the theorem. As in the case of proper vertices, the length of Q_b is at least $k-1$ (i.e., it contains at least k vertices) because by Lemma 13 the proper neighbors at the two ends must be at least $k+1$ steps apart. To determine its exact length, we turn back to the original concept of subdivision and branching vertices. Observe that by what we already know about proper vertices, Q_b is adjacent to a sequence of k branching vertices at each end. Hence, the first k and also the last k vertices of Q_b must be subdivision vertices of D. The next k vertices, on either end of Q_b, are then by definition branching vertices again, followed by another sequence of k subdivision vertices, etc. Clearly, this pattern only works out even if Q_b contains exactly $(2t+1)k$ vertices, for some nonnegative integer t.

We then subdivide all paths Q_b into paths of size k so that afterwards each vertex v of D lies on a unique strong path P_v of k thin respectively proper vertices and any two such paths P_b, P_c are either vertex disjoint or identical.

The obvious idea to identify isomorphic subgraphs now, is to put each vertex v of D into the subgraph D_i that corresponds to the position of v on the path Q_v, i.e., the ith vertex goes into D_i. The sought-after isomorphisms $\phi_{ij}\colon D_i \to D_j$ are also induced by the partition. Simply let ϕ_{ij} map a vertex $v \in D_i$ to the unique vertex of D_j that lies on the path Q_v. Clearly this mapping is well-defined. In order to check that it is also an isomorphism, we essentially only have to revisit the proof of Proposition 4, which constructed a root from isomorphisms. The crucial observation is again the strongness of our paths. Any walk of length k in R passes exactly once from one path P_a to a some path P_b with $a \to b$ in D, the remaining $k-1$ steps using only strong arcs. From this correspondence we see immediately that two vertices from the same path P_a have D-neighbors in the same set of adjacent paths. □

For computational purposes we note the following simple reformulation of Theorem 6.

14. Corollary. *Let D be a subdivision digraph without isolated cycles and with positive minimal indegree and outdegree. Let further $D_1, \dots, D_m$ be the different isomorphism classes of weakly connected components appearing in D and let d_i count the components in D isomorphic to D_i, $1 \le i \le m$. Then D has a kth root if and only if $k \mid d_i$ for all $i \in \{1, \dots, m\}$.* □

Counting cycles. We already discovered in Lemma 5 that powers of cycles are again cycles. To justify our hitherto ignorance towards cycles, we now also establish the converse: cycles have cycles as roots.

15. Lemma. *All vertices that lie on isolated cycles of a subdivision digraph D also lie on isolated cycles in any root of D.*

Proof. Let R be some kth root of D. We show that for any vertex c on a D-cycle, the sets $R^i(c)$ and $\bar{R}^i(c)$, $1 \le i < k$, are all singletons. This means that two D-adjacent vertices are connected through a strong walk in R, which proves the lemma.

So assume for contradiction that there exist two different vertices x, y in $R^j(c)$, $1 \le j < k$. (For $\bar{R}$ the statement is completely symmetric to this

case.) There exists some $u \in \bar{R}^k(x) \cap \bar{R}^k(y)$ because $\bar{R}^k(c)$ is nonempty. With two outneighbors in the subdivision digraph D, this u must be a branching vertex, hence, x and y are subdivision vertices. Therefore the sets $R^k(x)$ and $R^k(y)$ are nonempty and since $R^k(c)$ consists of exactly one vertex, we even have $R^{k-i}(x) = R^{k-i}(y)$, which now implies $R^k(x) = R^k(y) \neq \emptyset$. Hence, the two vertices x and y yield a contradiction to Observation 9. □

Lemma 5 told us that a single isolated root cycle yields only cycles of the same length in D. When we want to decide whether a collection of cycles in a given subdivision digraph D has a root, we may thus treat cycles of different lengths separately.

So assume that that D is the disjoint union of isolated cycles, all of a common length ℓ, and that R is a kth root of D. Let C be a cycle in R of some length r. We write

$$\ell = \prod p_i^{\ell_i}, \quad k = \prod p_i^{k_i}, \quad r = \prod p_i^{r_i}, \tag{19}$$

where $p_1, p_2, \ldots$ are the prime numbers. Lemma 5 tells us $r = \ell \cdot \gcd(r, k)$; expressed in terms of prime factorizations this reads $r_i = \ell_i + \min\{r_i, k_i\}$, which yields the implications

$$\ell_i > 0 \quad \Rightarrow \quad r_i = \ell_i + k_i, \tag{20}$$

$$\ell_i = 0 \quad \Rightarrow \quad 0 \leq r_i \leq k_i. \tag{21}$$

So the length r of the root cycle C is determined up to the order r_i at p_i for those indices i that satisfy $\ell_i = 0$ and $k_i > 0$.

We now argue that for root checking we may restrict our attention to root cycles with $r_i = 0$ in (21). Assume that some root cycle C of length r has $r_j > 0$ for some index j with $\ell_j = 0$. Replace C by $p_j^{r_j}$ many cycles of length

$$r' := \frac{r}{p_j^{r_j}} = \prod_{i \neq j} p_i^{r_i}$$

each. One easily checks $r'/\gcd(r', k) = r/\gcd(r, k)$ to see that the new cycles together have the same kth power as the old cycle C. Hence, the new digraph is also a root of D. By repeating this transformation until all root cycles satisfy $r_i = 0$ in (21) for all primes, we may assume that all cycles in R have the same (minimal) length

$$r = \prod_{l_i > 0} p_i^{\ell_i + k_i}.$$

How many D-cycles of length ℓ does one R-cycle of length r give? By Lemma 5 this number is exactly

$$\gcd(r, k) = \prod p_i^{\min\{r_i, k_i\}} = \prod_{\ell_i > 0} p_i^{\min\{\ell_i + k_i, k_i\}} = \prod_{\ell_i > 0} p_i^{k_i}.$$

This shows that a disjoint union of m cycles of length ℓ has a kth root if and only if

$$\prod_{\ell_i > 0} p_i^{k_i} \text{ divides } m, \tag{22}$$

where ℓ_i and k_i are the orders of ℓ resp. k at p_i as defined in (19).

16. Proposition. *Given a subdivision digraph D that consists of isolated cycles only and a parameter $k \geq 2$, we can check in polynomial time whether D has a kth root.*

Proof. We sum up the results of the preceding discussion in a simple algorithm. For each integer ℓ that appears as the length of a cycle in D, compute the prime factorization $\ell = \prod p_i^{\ell_i}$ and then the order k_i of k at each prime p_i with positive ℓ_i, i.e., the maximal k_i so that $p_i^{k_i} | k$. The digraph D has a kth root iff (22) is satisfied for each length ℓ (the integer m there counting the number of length-ℓ cycles).

The ℓ_i can be obtained in polynomial time since ℓ is bounded by the size of D and the relevant k_i are determined efficiently by simple division, even if k is exponential in the input size. □

Reducing isomorphism to subdivision roots. It remains to merge the results of the preceding sections into a proof of our isomorphism-completeness theorem. We now give the details of both polynomial-time reductions between digraph-isomorphism and subdivision-digraph roots.

Proof of Theorem 3. Let us first show that digraph roots are no easier to compute than digraph isomorphism, by giving a many-one reduction from the latter problem to the former.

For a given pair D_1, D_2 of digraphs, we construct a subdivision digraph D as follows.

(i) Make $k-2$ isomorphic copies $D_3, \ldots, D_k$ of D_2
(ii) Extend each D_i, $1 \leq i \leq k$, to a digraph D_i' by adding two new "super vertices" s_i, t_i, introducing the double connections $s_i \to a \to s_i$ for each $a \in D_i$, equipping t_i with a self-loop $t_i \to t_i$, and attaching it via $s_i \to t_i$.
(iii) Form the complete subdivision D_i'' of each extended D_i'.
(iv) Let $D := D_1'' \dot\cup D_2'' \dot\cup \cdots \dot\cup D_k''$ be the disjoint union of the D_i''.

Clearly D is a subdivision digraph and the vertices s_i guarantee that it has positive minimal in– and outdegree and consists of exactly k components, none of which is an isolated cycle. Hence, Theorem 6 tells us that D has a kth root iff all D_i'' are isomorphic or, equivalently, all D_i' are isomorphic. Since the t_i are distinguishable from all other vertices in the respective D_i' (because they are the only self-looped vertices with outdegree 1) this is the case iff all D_i are isomorphic or, by step (i), simply iff $D_1 \simeq D_2$.

We turn to the other reduction from subdivision-digraph roots to digraph isomorphism, which, by means of Proposition 16 and Theorem 6, is now very easy to formulate; but only as a Turing reduction, as opposed to the stronger notion of many-one reduction. That is, we describe a polynomial-time algorithm for the subdivision-digraph-root problem that may use a digraph-isomorphism oracle arbitrarily often.

Given a subdivision digraph D with positive minimal in– and outdegree, together with an integer k, we first use Proposition 16 to test in polynomial time whether the union of all isolated cycles of D has a kth root. Then we group the non-cycle components of D into isomorphism classes and apply Corollary 14. The independent treatment of isolated cycles and non-cycle components was justified by Lemmas 5 and 15. □

Outlook. While the original problem, the open complexity status of Boolean matrix root computation, is now settled, our search for further structure has lead to new questions. First of all, it would be desirable to get rid of the degree condition in Theorem 3. Let us indicate what can happen in a subdivision digraph that contains vertices without in– or outneighbors. Figure 7 shows such a digraph D together with a square root R. The two final root arcs can touch each other because the topmost vertex has no outneighbor and Lemma 11 about strong root arcs does not apply. Consequently, the minimal-degree condition is in fact indispensable for Theorem 6. But could it still be possible to remove it from the complexity result of Theorem 3? Observe that instead of being the disjoint union of two isomorphic subgraphs, the digraph D in Figure 7 can be decomposed into two parts, A and B (the former consisting of the two paths on the left, the latter containing the remaining five vertices), such that there exists a surjective *homomorphism* (i.e., an arc-preserving map) from A onto B. This homomorphism corresponds exactly to those arcs of R that go from A to B.

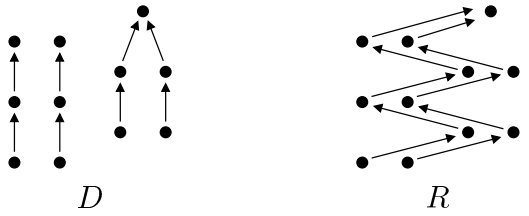

FIGURE 7. Dropping the degree condition in Theorem 6.

Though the general situation seems more difficult to analyze, this simple example indicates that when the degree condition is dropped, we have to deal with several interacting homomorphism problems. Thus, it is not at all clear whether the relaxed digraph root problem remains isomorphism complete since the general homomorphism problem for graphs is NP-complete [**21**]. (3-Colorability can be written as a homomorphism problem, for example).

More generally, we might ask for stronger versions of Theorem 3 showing isomorphism completeness of root finding for larger classes of digraphs. Although the structural result of Theorem 6 requires the special appearance of subdivision digraphs, their strict regularity should not ultimately be needed to deactivate the computationally hard aspects of the root problem established through Theorem 1. Yet, the concept of subdivisions and the techniques we employed throughout the proofs of Lemmas 10 to 13 might serve as a guideline for such generalizations.