
This is an electronic reprint of the original article.
This reprint may differ from the original in pagination and typographic detail.

Björklund, Andreas; Kaski, Petteri; Koutis, Ioannis

Directed hamiltonicity and out-branchings via generalized laplacians

Published in:

44th International Colloquium on Automata, Languages, and Programming, ICALP 2017

DOI:

[10.4230/LIPIcs.ICALP.2017.91](https://doi.org/10.4230/LIPIcs.ICALP.2017.91)

Published: 01/07/2017

Document Version

Publisher's PDF, also known as Version of record

Published under the following license:

CC BY

Please cite the original version:

Björklund, A., Kaski, P., & Koutis, I. (2017). Directed hamiltonicity and out-branchings via generalized laplacians. In *44th International Colloquium on Automata, Languages, and Programming, ICALP 2017* (pp. 1-14). Article 91 (Leibniz International Proceedings in Informatics ; Vol. 80). Schloss Dagstuhl - Leibniz-Zentrum für Informatik. <https://doi.org/10.4230/LIPIcs.ICALP.2017.91>

This material is protected by copyright and other intellectual property rights, and duplication or sale of all or part of any of the repository collections is not permitted, except that material may be duplicated by you for your research use or educational purposes in electronic or print form. You must obtain permission for any other use. Electronic or print copies may not be offered, whether for sale or otherwise to anyone who is not an authorised user.

Directed Hamiltonicity and Out-Branchings via Generalized Laplacians^{*†}

Andreas Björklund¹, Petteri Kaski², and Ioannis Koutis³

1 Department of Computer Science, Lund University, Lund, Sweden
andreas.bjorklund@yahoo.se

2 Department of Computer Science, Aalto University, Helsinki, Finland
petteri.kaski@aalto.fi

3 Department of Computer Science, University of Puerto Rico – Rio Piedras,
San Juan, Puerto Rico
ioannis.koutis@upr.edu

Abstract

We are motivated by a tantalizing open question in exact algorithms: can we detect whether an n -vertex directed graph G has a Hamiltonian cycle in time significantly less than 2^n ?

We present new randomized algorithms that improve upon several previous works:

1. We show that for any constant $0 < \lambda < 1$ and prime p we can **count** the Hamiltonian cycles modulo $p^{\lfloor (1-\lambda)\frac{n}{3p} \rfloor}$ in expected time less than c^n for a constant $c < 2$ that depends only on p and λ . Such an algorithm was previously known only for the case of counting modulo two [Björklund and Husfeldt, FOCS 2013].
2. We show that we can detect a Hamiltonian cycle in $O^*(3^{n-\alpha(G)})$ time and **polynomial space**, where $\alpha(G)$ is the size of the maximum independent set in G . In particular, this yields an $O^*(3^{n/2})$ time algorithm for bipartite directed graphs, which is faster than the exponential-space algorithm in [Cygan *et al.*, STOC 2013].

Our algorithms are based on the algebraic combinatorics of “incidence assignments” that we can capture through evaluation of determinants of Laplacian-like matrices, inspired by the Matrix–Tree Theorem for directed graphs. In addition to the novel algorithms for directed Hamiltonicity, we use the Matrix–Tree Theorem to derive simple algebraic algorithms for detecting out-branchings. Specifically, we give an $O^*(2^k)$ -time randomized algorithm for detecting out-branchings with at least k internal vertices, improving upon the algorithms of [Zehavi, ESA 2015] and [Björklund *et al.*, ICALP 2015]. We also present an algebraic algorithm for the directed k -Leaf problem, based on a non-standard monomial detection problem.

1998 ACM Subject Classification F.2.2 Nonnumerical Algorithms and Problems, G.2.1 Combinatorics, G.2.2 Graph Theory

Keywords and phrases counting, directed Hamiltonicity, graph Laplacian, independent set, k -internal out-branching

Digital Object Identifier 10.4230/LIPIcs.ICALP.2017.91

* A full version of the paper is available at <http://arxiv.org/abs/1607.04002>.

† The research leading to these results has received funding from the Swedish Research Council grants VR 2012-4730 “Exact Exponential-Time Algorithms” and VR 2016-03855 “Algebraic Graph Algorithms” (A.B.), the European Research Council under the European Union’s Seventh Framework Programme (FP/2007-2013) / ERC Grant Agreement 338077 “Theory and Practice of Advanced Search and Enumeration” (P.K.), and grant NSF CAREER CCF-1149048 (I.K.). Work done in part while the authors were at Dagstuhl Seminar 17041 in January 2017 and at the Simons Institute for the Theory of Computing in December 2015.



© Andreas Björklund, Petteri Kaski, and Ioannis Koutis;
licensed under Creative Commons License CC-BY

44th International Colloquium on Automata, Languages, and Programming (ICALP 2017).

Editors: Ioannis Chatzigiannakis, Piotr Indyk, Fabian Kuhn, and Anca Muscholl;

Article No. 91; pp. 91:1–91:14



Leibniz International Proceedings in Informatics

Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany



1 Introduction

The Hamiltonian cycle problem has played a prominent role in development of techniques for the design of *exact algorithms* for hard problems. The early $O^*(2^n)$ algorithms based on dynamic programming and inclusion-exclusion [1, 20, 19], remained un-challenged for several decades. In 2010, Björklund [3], gave a randomized algorithm running in $O(1.657^n)$ time for the case of undirected graphs. The algorithm taps into the power of algebraic combinatorics, and in particular determinants that enumerate cycle covers.

Despite this progress in the undirected Hamiltonian cycle problem, a substantial improvement in the more general directed version of the problem remains an open problem and a key challenge in the area of exact algorithms. The currently best known general algorithm runs in $O^*(2^{n-\Theta(\sqrt{n/\log n})})$ time [4], and there are no known connections with the theory of SETH-hardness [18] that would – at least partly – dash the hope for a faster algorithm.

A number of recent works have attempted to crack directed Hamiltonicity, revealing that the problem is indeed easier in certain restricted settings. Cygan and Pilipczuk [13] showed that the problem admits an $O^*(2^{(1-\epsilon_d)n})$ time algorithm for graphs with average degree bounded by d , where ϵ_d is a constant with a doubly exponential dependence on d . Cygan *et al.* [12] showed that the problem admits an $O^*(1.888^n)$ time randomized algorithm for bipartite graphs and that the parity of directed Hamiltonian cycles can also be computed within the same time bound. Björklund and Husfeldt [6] showed that the parity of Hamiltonian cycles can be computed in $O^*(1.619^n)$ randomized time in general directed graphs. Finally, Björklund *et al.* [5] showed that the problem can be solved in $O^*((2-\Theta(1))^n)$ time when the graph contains less than 1.038^n Hamiltonian cycles, via a reduction to the parity problem. In this paper we improve or generalize all of these works.

Our results. As one would expect, all recent “below-2ⁿ” algorithm designs for the Hamiltonicity problem rely on algebraic combinatorics and involve formulas that enumerate Hamiltonian cycles. But somewhat surprisingly, none of these approaches employs the directed version of the Matrix–Tree Theorem (see e.g. Gessel and Stanley [17, §11]), one of the most striking and beautiful results in algebraic graph theory. The theorem enables the enumeration of spanning out-branchings, that is, rooted spanning trees with all arcs oriented away from the root, via a determinant polynomial. Our results in this paper derive from a detailed combinatorial understanding and generalization of this classical setup.

The combinatorial protagonist of this paper is the following notion that enables a “two-way” possibility to view each arc in a directed graph:

► **Definition 1** (Incidence assignment). Let G be a directed graph with vertex set V and arc set E . For a subset $W \subseteq V$ we say that a mapping $\mu : W \rightarrow E$ is an *incidence assignment* if for all $u \in W$ it holds that $\mu(u)$ is incident with u .

In particular, looking at a single arc $uv \in E$, an incidence assignment μ can assign uv in two¹ possible ways: as an out-arc $\mu(u) = uv$ at u , or as an in-arc $\mu(v) = uv$ at v .

From an enumeration perspective the serendipity of this “two-way” possibility to assign an arc becomes apparent when one considers how an incidence assignment μ can realize a

¹ Strictly speaking we are here assuming that both $u \in W$ and $v \in W$. To break symmetry in our applications we do allow also situations where uv has only one possible assignment due to either $u \notin W$ or $v \notin W$.

directed cycle in its image $\mu(W)$. Indeed, let

$$u_1u_2, u_2u_3, \dots, u_{\ell-1}u_\ell, u_\ell u_1 \in E$$

be the arcs of a directed cycle C of length $\ell \geq 2$ in G with $V(C) \subseteq W$. It is immediate that there are now exactly two² ways to realize C in the image $\mu(W)$. Namely, we can realize C either (i) using only in-arcs with

$$\mu(u_1) = u_\ell u_1, \quad \mu(u_2) = u_1 u_2, \quad \mu(u_3) = u_2 u_3, \quad \dots, \quad \mu(u_\ell) = u_{\ell-1} u_\ell, \quad (1)$$

or (ii) using only out-arcs with

$$\mu(u_1) = u_1 u_2, \quad \mu(u_2) = u_2 u_3, \quad \mu(u_3) = u_3 u_4, \quad \dots, \quad \mu(u_\ell) = u_\ell u_1. \quad (2)$$

Incidence assignments thus enable two distinct ways to realize a *directed* cycle. Furthermore, it is possible to *switch* between (1) and (2) so that only the images of u_1, u_2, \dots, u_ℓ under μ are affected. The algebraization of this combinatorial observation is at the heart of the directed Matrix–Tree Theorem (which we will review for convenience of exposition in Sect. 2) and all of our results in this paper.

Our warmup result involves a generalization of the directed Hamiltonian path problem, namely the *k-Internal Out-Branching* problem, where the goal is to detect whether a given directed graph contains a spanning out-branching that has at least k internal vertices. This is a well-studied problem on its own, with several successive improvements the latest of which is an $O^*(3.617^k)$ algorithm by Zehavi [24] and an $O^*(3.455^k)$ algorithm by Björklund *et al.* [9] for the undirected version of the problem.

Using a combination of the directed Matrix–Tree Theorem and a monomial-sieving idea due to Floderus *et al.* [15], in Sect. 3 we show the following:

► **Theorem 2** (Detecting a *k-Internal Out-Branching*). *There exists a randomized algorithm that solves the k -internal out-branching problem in time $O^*(2^k)$ and with negligible probability of reporting a false negative.*

In the full version of the paper [10] we give a further application for the *k-Leaf* problem, that is, detecting a spanning out-branching with at least k leaves. We note that Gabizon *et al.* [16] have recently given another application of the directed Matrix–Tree Theorem for the problem of detecting out-branchings of bounded degree.

Proceeding to our two main results, in Sect. 4 we observe that the directed Matrix–Tree Theorem leads to a formula for computing the number of Hamiltonian paths in arbitrary characteristic by using a standard inclusion–exclusion approach, which leads to a formula that involves the summation of 2^n determinants. To obtain a below- 2^n design, we present a way to randomize the underlying Laplacian matrix so that the number of Hamiltonian paths does not change but in expectation most of the summands vanish modulo a prime power. Furthermore, to efficiently list the non-vanishing terms, we use a variation of an algorithm of Björklund *et al.* [8] that was used for a related problem, computing the permanent modulo a prime power. This leads to our first main result:

► **Theorem 3** (Counting directed Hamiltonian cycles modulo a prime power). *For all $0 < \lambda < 1$ there exists a randomized algorithm that, given an n -vertex directed graph and a prime p as input, counts the number of Hamiltonian cycles modulo $p^{\lfloor (1-\lambda)n/(3p) \rfloor}$ in expected time $O^*(2^{n(1-\lambda^2/(19p \log_2 p))})$. The algorithm uses exponential space.*

² Again strictly speaking it will be serendipitous to break symmetry so that certain cycles will have only one realization instead of two.

A corollary of Theorem 3 is that if G has at most d^n Hamiltonian cycles, we can detect one in time $O(c_d^n)$, where d is any fixed constant and $c_d < 2$ is a constant that only depends on d . As a further corollary we obtain a randomized algorithm for counting Hamiltonian cycles in graphs of bounded average (out-)degree d in $O(2^{(1-\epsilon_d)n})$ time. The constant ϵ_d has a polynomial dependency in d . Previous algorithms had a constant ϵ_d with an exponential dependency on d [7, 13]. (The proofs of these results are relegated to the full version of the paper [10].)

Returning to undirected Hamiltonicity, a key to the algorithm in [3] was the observation that determinants enumerate all non-trivial cycle covers an even number of times. This is due to the fact that each undirected cycle can be traversed in both directions. By picking a special vertex, one can break symmetry and force this to happen only for non-Hamiltonian cycle covers, so that the corresponding monomials cancel in characteristic 2. In Sect. 5 we present a “quasi-Laplacian” matrix whose determinant enables a similar approach for the directed case via algebraic combinatorics of incidence assignments, and furthermore enables one to accommodate a speedup assuming the existence of a good-sized independent set. We specifically prove the following as our second main result:

► **Theorem 4** (Detecting a directed Hamiltonian cycle). *There exists a randomized algorithm that solves the directed Hamiltonian cycle problem on a given directed graph G with a maximum independent set of size $\alpha(G)$, in $O^*(3^{n(G)-\alpha(G)})$ time, polynomial space and with negligible probability of reporting a false negative.*

Theorem 4 improves and generalizes the exponential-space algorithm of Cygan *et al.* [12].

Terminology and conventions. All graphs in this paper are directed and without loops and parallel arcs unless indicated otherwise. For an arc e starting from vertex u and ending at vertex v we say that u is the *tail* of e and v is the *head* of e . The vertices u and v are the *ends* of e . A directed graph is *connected* if the undirected graph obtained by removing orientation from the arcs is connected. A subgraph of a graph is *spanning* if the subgraph has the same set of vertices as the graph. A connected directed graph is an *out-branching* if every vertex has in-degree 1 except for the *root* vertex that has in-degree 0. We say that a vertex is *internal* to an out-branching if it has out-degree at least 1; otherwise the vertex is a *leaf* of the out-branching. The (*directed*) *Hamiltonian cycle* problem asks, given a directed graph G as input, whether G has a spanning directed cycle as a subgraph. The notation $O^*(\)$ suppresses a multiplicative factor polynomial in the input size. We say that an event parameterized by n has *negligible* probability if the probability of the event tends to zero as n grows without bound.

2 The symbolic Laplacian of a directed graph

This section develops the relevant preliminaries on directed graph Laplacians.

Permutations and the determinant. A bijection $\sigma : U \rightarrow U$ of a finite set U is called a *permutation* of U . A permutation σ *moves* an element $u \in U$ if $\sigma(u) \neq u$; otherwise σ *fixes* u . The *identity* permutation fixes every element of U . A permutation σ of U is a *cycle* of length $k \geq 2$ if there exist distinct $u_1, u_2, \dots, u_k \in U$ with $\sigma(u_1) = u_2, \sigma(u_2) = u_3, \dots, \sigma(u_{k-1}) = u_k, \sigma(u_k) = u_1$ and σ fixes all other elements of U . Two cycles are *disjoint* if every point moved by one is fixed by the other. The set of all permutations of U forms the *symmetric group* $\text{Sym}(U)$ with the composition of mappings as the product operation of the group.

Every nonidentity permutation factors into a unique product of pairwise disjoint cycles. The *sign* of a permutation σ that factors into c disjoint cycles of lengths k_1, k_2, \dots, k_c is $\text{sgn}(\sigma) = (-1)^{\sum_{j=1}^c (k_j - 1)}$. The sign of the identity permutation is 1.

The *determinant* of a square matrix A with rows and columns indexed by U is the multivariate polynomial

$$\det A = \sum_{\sigma \in \text{Sym}(U)} \text{sgn}(\sigma) \prod_{u \in U} a_{u, \sigma(u)}.$$

The punctured Laplacian determinant via incidence assignments. Let G be a directed graph with n vertices. Associate with each arc $uv \in E = E(G)$ an indeterminate x_{uv} . The *symbolic Laplacian* $L = L(G)$ of G is the $n \times n$ matrix with rows and columns indexed by the vertices $u, v \in V = V(G)$ and the (u, v) -entry defined³ by

$$\ell_{uv} = \begin{cases} \sum_{w \in V: wu \in E} x_{wu} & \text{if } u = v; \\ -x_{uv} & \text{if } uv \in E; \\ 0 & \text{if } u \neq v \text{ and } uv \notin E. \end{cases} \quad (3)$$

Observe that for each $v \in V$ we have that column v of L sums to zero because the diagonal entries cancel the negative off-diagonal entries. Furthermore, for each $u \in V$ we have that the monomials on row u of L correspond to the arcs incident to u . Indeed, each monomial at the diagonal corresponds to an in-arc to u , and each monomial at an off-diagonal entry corresponds to an out-arc from u . Thus, selecting one monomial from each row corresponds to selecting an incidence assignment.

To break symmetry, select an $r \in V$. The symbolic Laplacian of G *punctured at r* is obtained from L by deleting both row r and column r . We write $L_r = L_r(G)$ for the symbolic Laplacian of G punctured at r . Let us write $\mathcal{B}_r = \mathcal{B}_r(G)$ for the set of all spanning out-branchings of G with root $r \in V$. The following theorem is well-known (see e.g. Gessel and Stanley [17, §11]) and is presented here for purposes of displaying a proof that presents the cancellation argument using incidence assignments.

► **Theorem 5** (Directed Matrix–Tree Theorem). $\det L_r = \sum_{H \in \mathcal{B}_r} \prod_{uv \in E(H)} x_{uv}$.

Proof. Let us abbreviate $V_r = V(G) \setminus \{r\}$ and study the determinant

$$\det L_r = \sum_{\sigma \in \text{Sym}(V_r)} \text{sgn}(\sigma) \prod_{u \in V_r} \ell_{u, \sigma(u)}. \quad (4)$$

In particular, let us fix an arbitrary permutation $\sigma \in \text{Sym}(V_r)$ and study the monomials of the polynomial $\prod_{u \in V_r} \ell_{u, \sigma(u)}$ with the assumption that this polynomial is nonzero. From (3) it is immediate for each $u \in V_r$ that $\ell_{u, \sigma(u)}$ expands either (i) to the diagonal sum $\sum_{w \in V: wu \in E} x_{wu}$, which happens precisely when σ fixes u with $\sigma(u) = u$, or (ii) to the off-diagonal $-x_{uv}$, which happens precisely when σ moves u with $\sigma(u) = v$.

Let us write $M(\sigma)$ for the set of all incidence assignments $\mu : V_r \rightarrow E$ with the properties that (i) each $u \in V_r$ fixed by σ is assigned to an in-arc $\mu(u) = wu \in E$ for some $w \in V$, and (ii) each $u \in V_r$ moved by σ is assigned to the unique out-arc $\mu(u) = uv \in E$ with $\sigma(u) = v$.

³ Recall that we assume that G is loopless so the entries with $u = v$ are well-defined.

Let us write $f = f(\sigma)$ for the number of elements in V_r fixed by σ . It is immediate by (i) and (ii) that we have

$$\prod_{u \in V_r} \ell_{u, \sigma(u)} = \sum_{\mu \in M(\sigma)} (-1)^{n-1-f(\sigma)} \prod_{u \in V_r} x_{\mu(u)}. \quad (5)$$

Next observe that from μ we can reconstruct $\sigma = \sigma(\mu)$ by (i) setting $\sigma(u) = u$ for each u assigned to an in-arc in μ , and (ii) setting $\sigma(u) = v$ for each u assigned to an out-arc uv in μ . Thus the union $M = \bigcup_{\sigma \in \text{Sym}(V_r)} M(\sigma)$ is disjoint. Let us call the elements of M *proper incidence assignments*. By (4) and (5) we have

$$\det L_r = \sum_{\mu \in M} (-1)^{n-1-f(\sigma)} \text{sgn}(\sigma(\mu)) \prod_{u \in V_r} x_{\mu(u)}. \quad (6)$$

We claim that an incidence assignment μ is proper if and only if for every $u \in V_r$ there is exactly one $u' \in V_r$ such that $\mu(u')$ is an in-arc to u . For the “only if” direction, let σ be the permutation underlying a proper μ , and observe that vertices moved by σ partition to cycles so a σ never moves a vertex to a fixed vertex. Thus, we have $u' = u$ for the points fixed by σ , and $u' = \sigma^{-1}(u)$ is the vertex preceding u along a cycle of σ for points moved by σ . For the “if” direction, define $\sigma(u) = u$ if $u = u'$ and $\sigma(u') = u$ if $u' \neq u$. In the latter case we have $\mu(u') = u'u$, which means that $u'' \neq u'$ and thus $\sigma(u'') = u'$; by uniqueness of u' eventually a cycle must close so σ is a well-defined permutation underlying μ and thus μ is proper.

Let us write \mathcal{H}_r for the set of all spanning subgraphs of G with the property that every vertex in V_r has in-degree 1 and the root r has in-degree 0. From the previous claim it follows that we can view the set $\mu(V_r) = \{\mu(u) : u \in V_r\}$ for a proper μ as an element of \mathcal{H}_r . Furthermore, $\mu(V_r)$ is connected (and hence a spanning out-branching with root r) if and only if $\mu(V_r)$ is acyclic.

Consider an arbitrary $H \in \mathcal{H}_r$. If H has a cycle, let C be the least cycle in H according to some fixed but arbitrary ordering of the vertices of G . (Observe that any two cycles in H must be vertex-disjoint and cannot traverse r because r has in-degree 0.) Now consider an arbitrary proper μ that realizes H by $\mu(V_r) = H$. The cycle C is realized in μ by either (1) (in which case $\sigma(\mu)$ fixes all vertices in C), or (2) (in which case $\sigma(\mu)$ traces the cycle C). Furthermore, we may switch between realizations (1) and (2) so that the number of fixed points in the underlying permutation changes by $|V(C)|$ and the sign of the underlying permutation gets multiplied by $(-1)^{|V(C)|-1}$. It follows that the realizations (1) and (2) have different signs and thus cancel each other in (6). If H does not have a cycle, that is, $H \in \mathcal{B}_r$, it follows that there is a unique proper μ that realizes H . Indeed, first observe that H can be realized only by assigning in-arcs since any assignment of an out-arc in μ implies a cycle in $H = \mu(V_r)$, a contradiction. Second, the in-arcs are unique since each $u \in V_r$ has in-degree 1 in H . Finally, since μ assigns only in-arcs the underlying permutation $\sigma(\mu)$ is the identity permutation which has $\text{sgn}(\sigma(\mu)) = 1$ and $(-1)^{n-1-f(\sigma(\mu))} = 1$. Thus, each acyclic H contributes to (6) through a single $\mu \in M$ with coefficient 1. The theorem follows. ◀

3 Corollary for k -internal out-branchings

This section proves Theorem 2. We rely on a substitution idea of Floderus *et al.* [15, Theorem 1] to detect monomials with at least k distinct variables.

Let G be an n -vertex directed graph given as input together with a nonnegative integer k . Without loss of generality we may assume that $k \leq n - 1$. Iterate over all choices for a root vertex $r \in V$. Introduce an indeterminate y_u for each vertex $u \in V$ and an

indeterminate z_{uv} for each arc $uv \in E$. Introduce one further indeterminate t . Construct the symbolic Laplacian L of G given by (3) and with the assignment $x_{uv} = (1 + ty_u)z_{uv}$ to the indeterminate x_{uv} for each $uv \in E$. Puncture L at r to obtain L_r . Using, for example, Berkowitz’s determinant circuit design [2] for an arbitrary commutative ring with unity, in time $O^*(1)$ build an arithmetic circuit \mathcal{C} of size $O^*(1)$ for $\det L_r$. Viewing $\det L_r$ as a multivariate polynomial over the polynomial ring $R[t, y_u, z_{uv} : u \in V, uv \in E]$ where R is an abstract ring with unity, from Theorem 5 it follows that G has a spanning out-branching rooted at r with at least k internal vertices if and only if the coefficient of t^k in $\det L_r$ (which is a polynomial that is either identically zero or both (i) homogeneous of degree k in the indeterminates y_u and (ii) homogeneous of degree $n - 1$ in the indeterminates z_{uv}) has a monomial that is multilinear of degree k in the indeterminates y_u . Indeed, observe that the substitution $x_{uv} = (1 + ty_u)z_{uv}$ tracks in the degree of the indeterminate y_u whether u occurs as an internal vertex or not; the indeterminates z_{uv} make sure that distinct spanning out-branchings will not cancel each other.

To detect a multilinear monomial in \mathcal{C} restricted to the coefficient of t^k we can invoke [11, Lemma 1] or [21, Lemma 2.8]. This results in a randomized algorithm that runs in time $O^*(2^k)$ and has a negligible probability of reporting a false negative. This completes the proof of Theorem 2. ◀

4 Modular counting of Hamiltonian cycles

This section proves Theorem 3. Fix an arbitrary constant $0 < \lambda < 1$. Let $0 < \beta < 1/2$ be a constant whose precise value is fixed later. Let p be a prime and let G be an n -vertex directed graph with vertex set V and arc set E given as input. Without loss of generality (by splitting any vertex u into two vertices, s and t , with s receiving the out-arcs from u , and t receiving the in-arcs to u) we may count the spanning paths starting from s and ending at t instead of spanning cycles. Similarly, without loss of generality we may assume that $2 \leq p < n$. (Indeed, for $p \geq n$ the counting outcome from Theorem 3 is trivial.)

Sieving for Hamiltonian paths among out-branchings. Let $s, t \in V$ be distinct vertices. Let us write $\text{hp}(G, s, t)$ for the set of spanning directed paths that start at s and end at t in G . Recall that we write $V_t = V \setminus \{t\}$ for the t -punctured version of the vertex set V . Let us also write $V_{st} = V \setminus \{s, t\}$. For $O \subseteq V_t$, let L_s^O be the matrix obtained from the Laplacian (3) by first puncturing at s and then substituting $x_{uv} = 0$ for all arcs $uv \in E$ with $u \in V_t \setminus O$. Since a path $P \in \text{hp}(G, s, t)$ is precisely a spanning out-branching rooted at s such that every vertex $u \in V_t$ has out-degree 1, we have, by Theorem 5 and the principle of inclusion and exclusion,

$$\sum_{P \in \text{hp}(G, s, t)} \prod_{uv \in E(P)} x_{uv} = \sum_{O \subseteq V_t} (-1)^{|V_t \setminus O|} \det L_s^O. \tag{7}$$

In particular observe that (7) holds in any characteristic.

Cancellation modulo a power of p . With foresight, select $k = \lfloor (1 - \lambda)(1/2 - \beta)n/p \rfloor$. Our objective is next to show that by carefully injecting entropy into the underlying Laplacian we can, in expectation and working modulo p^k , cancel all but an exponentially negligible fraction of the summands on the right-hand side of (7). Furthermore, we can algorithmically narrow down to the nonzero terms, leading to an exponential improvement to 2^n .

Let us assign $x_{uv} = 1$ for all $uv \in E$ with $u \neq t$. Since no spanning path that ends at t may contain an arc $tu \in E$ for any $u \in V_t$, we may without loss of generality assume that G contains

all such arcs, and assign, independently and uniformly at random $x_{tu} \in \{0, 1, \dots, p-1\}$. Thus, the summands $\det L_s^O$ for $O \subseteq V_t$ are now integer-valued random variables and (7) evaluates to $|\text{hp}(G, s, t)|$ with probability 1.

Let us next study a fixed $O \subseteq V_t$. Let F_O be the event that L_s^O has no more than k rows where each entry is divisible by p . In particular, $\det L_s^O \not\equiv 0 \pmod{p^k}$ implies F_O . To bound the probability of F_O from above, observe that L_s^O is identically zero at each row $u \in V_{st} \setminus O$ except possibly at the diagonal entries. Furthermore, because of the random assignment to the indeterminates x_{tu} , each diagonal entry at these rows is divisible by p with probability $1/p$. Let us take this intuition and turn it into a listing algorithm for (a superset of the) sets $O \subseteq V_t$ that satisfy F_O .

Bipartitioning. For listing we will employ a meet-in-the-middle approach based on building each set $O \subseteq V_t$ from two parts using the following bipartitioning. Let $V_t^{(1)} \cup V_t^{(2)} = V_t$ be a bipartition with $|V_t^{(1)}| = \lceil n/3 \rceil$ and $|V_t^{(2)}| = n - 1 - \lceil n/3 \rceil$. Associate with each $O_1 \subseteq V_t^{(1)}$ a vector $z^{O_1} \in \{0, 1, \dots, p-1, \infty\}^{V_{st}}$ with the entry at $u \in V_{st}$ defined by

$$z_u^{O_1} = \begin{cases} \infty & \text{if } u \in O_1; \\ (x_{tu} + \sum_{w \in O_1: wu \in E} x_{wu}) \bmod p & \text{otherwise.} \end{cases} \quad (8)$$

Similarly, associate with each $O_2 \subseteq V_t^{(2)}$ a vector $z^{O_2} \in \{0, 1, \dots, p-1, \infty\}^{V_{st}}$ with the entry at $u \in V_{st}$ defined by

$$z_u^{O_2} = \begin{cases} \infty & \text{if } u \in O_2; \\ (-\sum_{w \in O_2: wu \in E} x_{wu}) \bmod p & \text{otherwise.} \end{cases} \quad (9)$$

Suppose now that we have $O_1 \subseteq V_t^{(1)}$ and $O_2 \subseteq V_t^{(2)}$ with $O = O_1 \cup O_2$. We claim that F_O holds only if the vectors z^{O_1} and z^{O_2} agree in at most k entries. Indeed, observe that $z_u^{O_1} = z_u^{O_2}$ holds only if both $u \in V_{st} \setminus O$ and the (u, u) -entry of L_s^O is divisible by p . That is, $z_u^{O_1} = z_u^{O_2}$ implies the entire row u of L_s^O consists only of elements divisible by p . Thus it suffices to list all pairs (O_1, O_2) such that z^{O_1} and z^{O_2} have at most k agreements.

Balanced and unbalanced sets. To set up the listing procedure, let us now partition the index domain V_{st} of our vectors into $b = \lceil 3 \log_2 p \rceil$ pairwise disjoint sets S_1, S_2, \dots, S_b such that we have $\lfloor (n-2)/b \rfloor \leq |S_i| \leq \lceil (n-2)/b \rceil$.

Let us split the sets $O \subseteq V_t$ into two types. Let us say that O is *balanced* if $(1/2 - \beta)n/b \leq |(V_{st} \setminus O) \cap S_i| \leq (1/2 + \beta)n/b$ holds for all $i = 1, 2, \dots, b$; otherwise O is *unbalanced*. Recalling that $\sum_{j=0}^{\ell} \binom{n}{j} \leq 2^{nH(\ell/n)}$ holds for all integers $1 \leq \ell \leq n/2$, where $H(\rho) = -\rho \log_2 \rho - (1-\rho) \log_2 (1-\rho)$ is the binary entropy function, observe that there are in total at most

$$\begin{aligned} 2^{n+1-\min_i |S_i|} b^{\lceil (1/2-\beta)n/b \rceil} \sum_{j=0}^{\lfloor n/b+2 \rfloor} \binom{\lfloor n/b+2 \rfloor}{j} &\leq 2^{n-(n-2)/b+3} 2^{(n/b+2)H(1/2-\beta)} b \\ &\leq 2^{n(1-(1-H(1/2-\beta))/b)+7} b \end{aligned} \quad (10)$$

sets O that are unbalanced.

Precomputation and listing. Suppose that $O_1 \subseteq V_t^{(1)}$ and $O_2 \subseteq V_t^{(2)}$ are *compatible* in the sense that z^{O_1} and z^{O_2} agree in at most k entries. For $S \subseteq V_{st}$ and a vector z whose entries

are indexed by V_{st} , let us write z_S for the restriction of z to S . If O_1 and O_2 are compatible, then by an averaging argument there must exist an $i = 1, 2, \dots, b$ such that $z_{S_i}^{O_1}$ and $z_{S_i}^{O_2}$ agree in at most k/b entries. In particular, this enables us to iterate over O_2 and list all compatible O_1 by focusing only on each restriction to S_i for $i = 1, 2, \dots, b$. Furthermore, the search inside S_i can be precomputed to look-up tables. Indeed, for each $i = 1, 2, \dots, b$ and each key $g \in \{0, 1, \dots, p-1, \infty\}^{S_i}$, let us build a complete list of all subsets $O_1 \subseteq V_t^{(1)}$ such that $z_{S_i}^{O_1}$ and g agree in at most k/b entries. These b look-up tables can be built by processing in total at most

$$\sum_{i=1}^b 2^{V_t^{(1)}} (p+1)^{|S_i|} \leq 2^{n/3+7} 2^{(n/(3 \log_2 p)+2) \log_2(p+1)} \log_2 p = O(2^{0.87n})$$

pairs (O_1, g) . This takes time $O^*(2^{0.87n})$ in total.

The main listing procedure now considers each $O_2 \subseteq V_t^{(2)}$ in turn, and for each $i = 1, 2, \dots, b$ consults the look-up table for direct access to all O_1 such that $z_{S_i}^{O_1}$ and $z_{S_i}^{O_2}$ agree in at most k/b entries. In particular this will list all compatible pairs (O_1, O_2) and hence all sets $O = O_1 \cup O_2$ such that F_O holds.

Expected running time. Let us now analyze the expected running time of the algorithm. We start by deriving an upper bound for the expected number of pairs (O_1, O_2) considered by the main listing procedure. First, observe that the total number of pairs (O_1, O_2) considered by the procedure with $O = O_1 \cup O_2$ unbalanced is bounded from above by our upper bound (10) for the total number of unbalanced O . Indeed, $O_1 = O \cap V_t^{(1)}$ and $O_2 = O \cap V_t^{(2)}$ are uniquely determined by O .

Next, for a pair (O_1, O_2) with balanced $O = O_1 \cup O_2$ and $i = 1, 2, \dots, b$, let $G_{O_1, O_2, i}$ be the event that $z_{S_i}^{O_1}$ and $z_{S_i}^{O_2}$ agree in at most k/b entries. We seek an upper bound for the probability of $G_{O_1, O_2, i}$ to obtain an upper bound for the expected number of pairs with balanced $O = O_1 \cup O_2$ considered by the main listing procedure. Let $A_{O_1, O_2, i}$ be the number of entries in which $z_{S_i}^{O_1}$ and $z_{S_i}^{O_2}$ agree. We observe that $A_{O_1, O_2, i}$ is binomially distributed with expectation $|(V_t \setminus O) \cap S_i|/p$. Since O is balanced, we have $(1/2 - \beta)n/b \leq |(V_t \setminus O) \cap S_i| \leq (1/2 + \beta)n/b$. We also recall that $k = \lfloor (1 - \lambda)(1/2 - \beta)n/p \rfloor$. A standard Chernoff bound now gives

$$\begin{aligned} \Pr(G_{O_1, O_2, i}) &\leq \Pr(A_{O_1, O_2, i} \leq k/b) \\ &\leq \Pr(A_{O_1, O_2, i} \leq (1 - \lambda)|(V_t \setminus O) \cap S_i|/p) \\ &\leq \exp(-\lambda^2 |(V_t \setminus O) \cap S_i|/(2p)) \\ &\leq \exp(-\lambda^2 (1/2 - \beta)n/(2pb)). \end{aligned}$$

Recalling that $b = \lfloor 3 \log_2 p \rfloor$, the main listing procedure thus considers in expectation at most $2^n \exp(-\lambda^2 (1/2 - \beta)n/(2p(3 \log_2 p)))$ pairs (O_1, O_2) with $O = O_1 \cup O_2$ balanced. Recalling our upper bound for the total number of unbalanced sets (10), we thus have that the main listing procedure runs in $O^*(2^n \exp(-\lambda^2 (1/2 - \beta)n/(2p(3 \log_2 p)))) + 2^{n(1 - (1 - H(1/2 - \beta))/(3 \log_2 p))}$ expected time. Recalling that precomputation runs in $O^*(2^{0.87n})$ time, we thus have for $\beta = 1/6$ that the entire algorithm runs in $O^*(2^{n(1 - \lambda^2/(19p \log_2 p))})$ expected time and computes $|\text{hp}(G, s, t)|$ modulo $p^{\lfloor (1 - \lambda)n/(3p) \rfloor}$. This completes the proof of Theorem 3. \blacktriangleleft

5 Directed Hamiltonicity via quasi-Laplacian determinants

This section proves Theorem 4. Let G be a directed n -vertex graph given as input.

Finding a maximum independent set. Let $B \cup Y = V(G)$ be a partition of the vertex set into two disjoint sets B (“blue”) and Y (“yellow”) such that no arc has both of its ends in Y . That is, Y is an independent set.

We can find an Y of the maximum possible size as follows. First, in time polynomial in n compute the maximum-size matching in the undirected graph obtained from G by disregarding the orientation of the arcs. This maximum-size matching must consist of at least $\lfloor n/2 \rfloor$ edges or G does not admit a Hamiltonian cycle. (Indeed, from a Hamiltonian cycle we can obtain a matching with $\lfloor n/2 \rfloor$ edges by taking every other arc in the cycle.) Since for each matching edge it holds that both ends of the edge cannot be in an independent set, we can in time $O^*(3^{n/2})$ find a maximum-size independent set Y of G . Furthermore, $\alpha(G) = |Y| \leq \lfloor n/2 \rfloor + 1$, so we are within our budget of $O^*(3^{n(G)-\alpha(G)})$ in terms of running time. In fact, $|Y| \leq \lfloor n/2 \rfloor$ or otherwise G trivially does not admit a Hamiltonian cycle.

The symbolic quasi-Laplacian. We will first define the quasi-Laplacian and then give intuition for its design. Let us work over a field of characteristic 2. For each $y \in Y$ introduce a copy y_{in} and let Y_{in} be the set of all such copies. Similarly, for each $y \in Y$ introduce a copy y_{out} and let Y_{out} be the set of all such copies. We assume that Y_{in} and Y_{out} are disjoint. For each $j \in Y_{\text{in}} \cup Y_{\text{out}}$ let us write $\underline{j} \in Y$ for the underlying element of Y of which j is a copy. Let B_* be a set of $n - 2|Y|$ elements that is disjoint from both Y_{in} and Y_{out} . For each $uv \in E$ and each $j \in B_* \cup Y_{\text{in}} \cup Y_{\text{out}}$, introduce an indeterminate $x_{uv}^{(j)}$.

Select an arbitrary vertex $s \in B$ for purposes of breaking symmetry and let $I, O \subseteq B$. The *quasi-Laplacian* $Q^{I,O,s} = Q^{I,O,s}(G)$ of G with skew at s be the $n \times n$ matrix whose rows are indexed by $u \in B \cup Y$ and whose columns are indexed by $j \in B_* \cup Y_{\text{in}} \cup Y_{\text{out}}$ such that the (u, j) -entry is defined by

$$q_{uj}^{I,O,s} = \begin{cases} \sum_{w \in O: wu \in E, u \in I} x_{wu}^{(j)} & \text{(a) if } u \in B \text{ and } j \in B_*; \\ + \sum_{w \in I: uw \in E, u \in O \setminus \{s\}} x_{uw}^{(j)} & \\ x_{uj}^{(j)} & \text{(b) if } u \in O \setminus \{s\} \text{ and } j \in Y_{\text{in}} \text{ with } u\underline{j} \in E; \\ x_{\underline{j}u}^{(j)} & \text{(c) if } u \in I \text{ and } j \in Y_{\text{out}} \text{ with } \underline{j}u \in E; \\ \sum_{w \in O: wu \in E} x_{wu}^{(j)} & \text{(d) if } u \in Y \text{ and } j \in Y_{\text{in}} \text{ with } u = \underline{j}; \\ \sum_{w \in I: uw \in E} x_{uw}^{(j)} & \text{(e) if } u \in Y \text{ and } j \in Y_{\text{out}} \text{ with } u = \underline{j}; \\ 0 & \text{otherwise.} \end{cases} \quad (11)$$

Let us next give some intuition for (11) before proceeding with the proof.

Analogously to the Laplacian (3), the quasi-Laplacian (11) has been designed so that the monomials of each row $u \in B \cup Y$ of $Q^{I,O,s}$ control the assignment of either an in-arc or an out-arc to u in an incidence assignment, and the skew at s is used to break symmetry so that s is always assigned an in-arc to s . In particular, without the skew at row s and with $I = O$, each column of $Q^{I,O,s}$ would sum to zero, in analogy with the (non-punctured) Laplacian.

Let us now give intuition for the columns $j \in Y_{\text{in}}$ and $j \in Y_{\text{out}}$. First, observe by (b) and (d) in (11) that selecting a monomial from column $j \in Y_{\text{in}}$ corresponds to making sure that the in-degree of \underline{j} is 1. Such a selection may be either a “quasi-diagonal” assignment of the in-arc $w\underline{j}$ to $u = \underline{j} \in Y$ via (d) for some $w \in B$; or an “off-diagonal” assignment of the out-arc $u\underline{j}$ to $u \in B$ via (b). Second, observe by (c) and (e) in (11) that selecting a monomial from column $j \in Y_{\text{out}}$ corresponds to making sure that the out-degree of \underline{j} is 1. Thus, the columns $j \in Y_{\text{in}} \cup Y_{\text{out}}$ enable us to make sure that an incidence assignment has both in-degree 1 and out-degree 1 at each $u \in Y$ without the use of sieving. This gives us the speed-up from $O^*(3^n)$ to $O^*(3^{n-|Y|})$ running time. Observe also that the structure for the quasi-Laplacian

$Q^{I,O,s}$ is enabled precisely because Y is an independent set and thus no arc contributes to both in-degree and out-degree in Y .

The quasi-Laplacian determinant sieve. Recalling that we are working over a field of characteristic 2, let us study the sum

$$\sum_{\substack{I,O \subseteq B \\ I \cup O = B \\ s \in I}} \det Q^{I,O,s} = \sum_{I \subseteq B} \sum_{O \subseteq B} \det Q^{I,O,s} = \sum_{\substack{\sigma: B \cup Y \rightarrow B_* \cup Y_{\text{in}} \cup Y_{\text{out}} \\ \sigma \text{ bijective}}} \sum_{I \subseteq B} \sum_{O \subseteq B} \prod_{u \in B \cup Y} q_{u,\sigma(u)}^{I,O,s}. \quad (12)$$

Observe that the first equality in (12) holds because $Q^{I,O,s}$ has by (11) an identically zero row unless $I \cup O = B$ and $s \in I$; the second equality holds by definition of the determinant in characteristic 2 and changing the order of summation. From (11) and the right-hand side of (12) it is immediate that (12) is either identically zero or a homogeneous polynomial of degree n in the $n|E|$ indeterminates $x_{uv}^{(j)}$ for $j \in B_* \cup Y_{\text{in}} \cup Y_{\text{out}}$ and $uv \in E$.⁴ We claim that (12) is not identically zero if and only if G admits at least one spanning cycle. Furthermore, each spanning cycle in G defines precisely $|B_*|!$ distinct monomials in (12).

To establish the claim, fix a bijection $\sigma : B \cup Y \rightarrow B_* \cup Y_{\text{in}} \cup Y_{\text{out}}$. Let us write $M(\sigma)$ for the set of all incidence assignments $\mu : B \cup Y \rightarrow E$ that are *proper* in the sense that all of the following six requirements hold (cf. (11)):

- (s): $\mu(s)$ is an in-arc to s ;
- (a): for all $u \in B$ with $\sigma(u) \in B_*$ it holds that $\mu(u)$ has both of its vertices in B ;
- (b): for all $u \in B$ with $\sigma(u) \in Y_{\text{in}}$ it holds that $\mu(u)$ is an in-arc to $\sigma(u)$;
- (c): for all $u \in B$ with $\sigma(u) \in Y_{\text{out}}$ it holds that $\mu(u)$ is an out-arc from $\sigma(u)$;
- (d): for all $u \in Y$ with $\sigma(u) \in Y_{\text{in}}$ it holds that $\mu(u)$ is an in-arc to u and $u = \sigma(u)$; and
- (e): for all $u \in Y$ with $\sigma(u) \in Y_{\text{out}}$ it holds that $\mu(u)$ is an out-arc from u and $u = \sigma(u)$.

Observe that each $\mu \in M(\sigma)$ defines a collection of n arcs $\mu(B \cup Y) = \{\mu(u) : u \in B \cup Y\}$. Let us write Z_{in}^μ (respectively, Z_{out}^μ) for the set of vertices in $B \cup Y$ with zero in-degree (out-degree) with respect to the arcs in $\mu(B \cup Y)$. Since σ is a bijection and thus has a preimage for each $j \in Y_{\text{in}} \cup Y_{\text{out}}$, from (b,c,d,e) above it follows that $Z_{\text{in}}^\mu \subseteq B$ and $Z_{\text{out}}^\mu \subseteq B$. Furthermore, from (a,b,c,d,e) it follows that for the arcs in $\mu(B \cup Y)$ the sum of the in-degrees (and the sum of the out-degrees) of the vertices in B is $|B| = |B_*| + |Y|$. Thus, we have that Z_{in}^μ and Z_{out}^μ are both empty if and only if for the arcs $\mu(B \cup Y)$ both the in-degree and the out-degree of every vertex $u \in B \cup Y$ is 1. (Note that the claim is immediate for $u \in Y$ by (b,c,d,e) and bijectivity of σ .)

Let us now study the right-hand side of (12) for a fixed σ . Using (a,b,c,d,e) and (11) to rearrange in terms of incidence assignments, we have

$$\sum_{I \subseteq V} \sum_{O \subseteq V} \prod_{u \in B \cup Y} q_{u,\sigma(u)}^{I,O,s} = \sum_{\mu \in M(\sigma)} \prod_{u \in B \cup Y} x_{\mu(u)}^{(\sigma(u))} \sum_{I \subseteq Z_{\text{in}}^\mu} \sum_{O \subseteq Z_{\text{out}}^\mu} 1. \quad (13)$$

Since we are working in characteristic 2, all other $\mu \in M(\sigma)$ except those for which $\mu(B \cup Y)$ is a cycle cover will cancel in the right-hand side of (13).

Take the sum of (13) over all bijections σ . Consider an arbitrary cycle cover of $B \cup Y$. Let C be a cycle in this cycle cover. Assuming that C does not contain s , we can realize C in an incidence assignment $\mu : B \cup Y \rightarrow E$ either using (1) or (2). If C contains s and μ is proper,

⁴ Our algorithm for deciding Hamiltonicity will naturally not work with a symbolic representation of (12) but rather in a homomorphic image under a random evaluation homomorphism.

only the realization (1) is possible by (s). To see that realization with a proper $\mu \in M(\sigma)$ for some σ is possible, consider an arbitrary $u \in B \cup Y$ and observe that each image $\mu(u)$ by (b,c,d,e) uniquely determines the image $\sigma(u) \in Y_{\text{in}} \cup Y_{\text{out}}$ when $\mu(u)$ has one vertex in Y ; when $\mu(u)$ has both vertices in B , an unused $\sigma(u) \in B_*$ may be chosen arbitrarily so that $\mu \in M(\sigma)$. It follows that any cycle cover with c cycles is realized as exactly $|B_*|!$ distinct monomials $\prod_{u \in B \cup Y} x_{\mu(u)}^{(\sigma(u))}$ in (12), each with coefficient 2^{c-1} . This coefficient is nonzero if and only if $c = 1$.

Completing the algorithm. To detect whether the given n -vertex directed graph G admits a Hamiltonian cycle, first decompose the vertex set into disjoint $V = B \cup Y$ with Y an independent set of size $|Y| = \alpha(G)$ using the algorithm described earlier. Next, in time $O^*(1)$ construct an irreducible polynomial of degree $2 \lceil \log_2 n \rceil$ over \mathbb{F}_2 (see e.g. von zur Gathen and Gerhard [23, §14.9]) to enable arithmetic in the finite field of order $q = 2^{2 \lceil \log_2 n \rceil} \geq n^2$ in time $O^*(1)$ for each arithmetic operation. Next, assign an independent uniform random value from \mathbb{F}_q to each indeterminate $x_{uv}^{(j)}$ with $j \in B_* \cup Y_{\text{in}} \cup Y_{\text{out}}$ and $uv \in E$. Finally, using the assigned values for the indeterminates, compute the left-hand side of (12) using, for example, Gaussian elimination to compute each determinant $\det Q^{I,O,s}$ in $O^*(1)$ operations in \mathbb{F}_q . Let us write $r \in \mathbb{F}_q$ for the result of this computation. In particular, we can compute r from a given G in total $O^*(3^{|B|}) = O^*(3^{n(G)-\alpha(G)})$ operations in \mathbb{F}_q , and consequently in total $O^*(3^{n(G)-\alpha(G)})$ time. If (12) is identically zero, then clearly $r = 0$ with probability 1. If (12) is not identically zero (and hence a homogeneous polynomial of degree $d = n$ in the indeterminates) then by the DeMillo–Lipton–Schwartz–Zippel lemma [14, 22, 25] we have $r \neq 0$ with probability at least $1 - d/q \geq 1 - n/n^2 \geq 1 - 1/n = 1 - o(1)$. Thus we can decide whether G is Hamiltonian based on whether $r \neq 0$. In particular this gives probability $o(1)$ of reporting a false negative. This completes the proof of Theorem 4. ◀

References

- 1 Richard Bellman. Dynamic programming treatment of the travelling salesman problem. *J. ACM*, 9(1):61–63, January 1962. doi:10.1145/321105.321111.
- 2 Stuart J. Berkowitz. On computing the determinant in small parallel time using a small number of processors. *Information Processing Letters*, 18:147–150, 1984. doi:10.1016/0020-0190(84)90018-8.
- 3 Andreas Björklund. Determinant sums for undirected hamiltonicity. *SIAM Journal on Computing*, 43(1):280–299, 2014. doi:10.1137/110839229.
- 4 Andreas Björklund. Below All Subsets for Some Permutational Counting Problems . In Rasmus Pagh, editor, *15th Scandinavian Symposium and Workshops on Algorithm Theory (SWAT 2016)*, volume 53 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 17:1–17:11, Dagstuhl, Germany, 2016. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik. doi:10.4230/LIPIcs.SWAT.2016.17.
- 5 Andreas Björklund, Holger Dell, and Thore Husfeldt. The parity of set systems under random restrictions with applications to exponential time problems. In *Automata, Languages, and Programming: 42nd International Colloquium, ICALP 2015, Kyoto, Japan, July 6-10, 2015, Proceedings, Part I*, pages 231–242, Berlin, Heidelberg, 2015. Springer Berlin Heidelberg. doi:10.1007/978-3-662-47672-7_19.
- 6 Andreas Björklund and Thore Husfeldt. The parity of directed hamiltonian cycles. In *2013 IEEE 54th Annual Symposium on Foundations of Computer Science*, pages 724–735, Oct 2013. doi:10.1109/FOCS.2013.83.

- 7 Andreas Björklund, Thore Husfeldt, Petteri Kaski, and Mikko Koivisto. The traveling salesman problem in bounded degree graphs. *ACM Trans. Algorithms*, 8(2):18:1–18:13, April 2012. doi:10.1145/2151171.2151181.
- 8 Andreas Björklund, Thore Husfeldt, and Isak Lyckberg. Computing the permanent modulo a prime power. Unpublished manuscript, 2015.
- 9 Andreas Björklund, Vikram Kamat, Lukasz Kowalik, and Meirav Zehavi. Spotting trees with few leaves. In Magnús M. Halldórsson, Kazuo Iwama, Naoki Kobayashi, and Bettina Speckmann, editors, *Automata, Languages, and Programming – 42nd International Colloquium, ICALP 2015, Kyoto, Japan, July 6-10, 2015, Proceedings, Part I*, volume 9134 of *Lecture Notes in Computer Science*, pages 243–255. Springer, 2015. doi:10.1007/978-3-662-47672-7_20.
- 10 Andreas Björklund, Petteri Kaski, and Ioannis Koutis. Directed Hamiltonicity and Out-Branchings via Generalized Laplacians. *CoRR*, abs/1607.04002, 2017. URL: <http://arxiv.org/abs/1607.04002>.
- 11 Andreas Björklund, Petteri Kaski, and Łukasz Kowalik. Constrained multilinear detection and generalized graph motifs. *Algorithmica*, 74(2):947–967, 2016. doi:10.1007/s00453-015-9981-1.
- 12 Marek Cygan, Stefan Kratsch, and Jesper Nederlof. Fast hamiltonicity checking via bases of perfect matchings. In *Proceedings of the Forty-fifth Annual ACM Symposium on Theory of Computing, STOC’13*, pages 301–310, New York, NY, USA, 2013. ACM. doi:10.1145/2488608.2488646.
- 13 Marek Cygan and Marcin Pilipczuk. Faster exponential-time algorithms in graphs of bounded average degree. *Information and Computation*, 243:75–85, 2015. 40th International Colloquium on Automata, Languages and Programming (ICALP 2013). doi:10.1016/j.ic.2014.12.007.
- 14 R. A. DeMillo and R. J. Lipton. A probabilistic remark on algebraic program testing. *Information Processing Letters*, 7:193–195, 1978.
- 15 Peter Floderus, Andrzej Lingas, Mia Persson, and Dzmityr Sledneu. Detecting monomials with k distinct variables. *Information Processing Letters*, 115(2):82–86, 2015. doi:10.1016/j.ipl.2014.07.003.
- 16 Ariel Gabizon, Daniel Lokshtanov, and Michał Pilipczuk. Fast algorithms for parameterized problems with relaxed disjointness constraints. In *Algorithms – ESA 2015: 23rd Annual European Symposium, Patras, Greece, September 14-16, 2015, Proceedings*, pages 545–556, Berlin, Heidelberg, 2015. Springer Berlin Heidelberg. doi:10.1007/978-3-662-48350-3_46.
- 17 Ira M. Gessel and Richard P. Stanley. Algebraic enumeration. In R. L. Graham, M. Grötschel, and L. Lovász, editors, *Handbook of Combinatorics*, volume II, pages 1021–1061. North-Holland, Amsterdam, 1995.
- 18 Russell Impagliazzo and Ramamohan Paturi. On the complexity of k -sat. *Journal of Computer and System Sciences*, 62(2):367–375, 2001. doi:10.1006/jcss.2000.1727.
- 19 Richard M. Karp. Dynamic programming meets the principle of inclusion and exclusion. *Operations Research Letters*, 1(2):49–51, 1982. doi:10.1016/0167-6377(82)90044-X.
- 20 Samuel Kohn, Allan Gottlieb, and Meryle Kohn. A generating function approach to the traveling salesman problem. In *Proceedings of the Annual Conference (ACM’77), Association for Computing Machinery*, pages 294–300, 1977.
- 21 Ioannis Koutis and Ryan Williams. Limits and applications of group algebras for parameterized problems. *ACM Transactions on Algorithms*, 12:Art. 31, 18, 2016. doi:10.1145/2885499.
- 22 Jacob T. Schwartz. Fast probabilistic algorithms for verification of polynomial identities. *Journal of the ACM*, 27:701–717, 1980.

- 23 Joachim von zur Gathen and Jürgen Gerhard. *Modern Computer Algebra*. Cambridge University Press, Cambridge, 3rd edition, 2013.
- 24 Meirav Zehavi. Mixing color coding-related techniques. In *Algorithms ESA 2015: 23rd Annual European Symposium, Patras, Greece, September 14-16, 2015, Proceedings*, pages 1037–1049. Springer Berlin Heidelberg, 2015. doi:10.1007/978-3-662-48350-3_86.
- 25 Richard Zippel. Probabilistic algorithms for sparse polynomials. In *Proc. International Symposium on Symbolic and Algebraic Computation*, volume 72 of *Lecture Notes in Computer Science*, pages 216–226, Berlin, 1979. Springer.