Improved Distance Queries and Cycle Counting by Frobenius Normal Form

Piotr Sankowski¹ · Karol Węgrzycki¹

Published online: 21 November 2018
© The Author(s) 2018

Abstract

Consider an unweighted, directed graph G with the diameter D . In this paper, we introduce the framework for counting cycles and walks of given length in matrix multiplication time $\tilde{O}(n^\omega)$. The framework is based on the fast decomposition into Frobenius normal form and the Hankel matrix-vector multiplication. It allows us to solve the All-Nodes Shortest Cycles, All-Pairs All Walks problems efficiently and also give some improvement upon distance queries in unweighted graphs.

Keywords Graph algorithms · Frobenius normal form · All-nodes shortest cycles

1 Introduction

The *All-Pairs Shortest Paths* (APSP) problem asks to find distances between all pairs of vertices in a graph. For a directed graphs with weights in $\mathbb{R}$, there is a classical $O(n^3)$ time algorithm [11, 29]. Currently best upper bound for this problem is due to [30] who showed an $O\left(\frac{n^3}{2^{\Omega(\log n)^{0.5}}}\right)$ algorithm. It is asymptotically faster than $O(n^3 / \log^c n)$ for any $c > 0$ (see survey [6] for earlier algorithms). Showing any algorithm that would work in $O(n^{3-\epsilon})$ time for some $\epsilon > 0$ is a major open problem [30].

Extended Abstract of this paper was presented at STACS 2017 [22]

This article is part of the Topical Collection on *Special Issue on Theoretical Aspects of Computer Science (STACS 2017)*

✉ Piotr Sankowski
sank@mimuw.edu.pl

Karol Węgrzycki
k.wegrzycki@mimuw.edu.pl

¹ Institute of Informatics, University of Warsaw, Warsaw, Poland

If we consider unweighted, directed graphs there are subcubic algorithms that exploit fast matrix multiplication. For the undirected graph [24] presented the optimal $\tilde{O}(n^\omega)$ time algorithm, where $\omega < 2.373$ is the matrix multiplication exponent [16]. For the directed case [35] presented an $O(n^{2.575})$ time algorithm that is based on the fast rectangular matrix multiplication. Moreover, if we are interested in small integer weights from the set $\{-M, \dots, M\}$ we have $O(M^{0.68}n^{2.575})$ algorithm [35].

Because APSP in undirected graphs can be solved in $\tilde{O}(n^\omega)$, diameter, radius, shortest cycle, etc. can be determined in $\tilde{O}(n^\omega)$ time as well. It is surprising that for a directed case, where merely $O(n^{2.575})$ APSP is known there are also $\tilde{O}(n^\omega)$ algorithms for determining these properties. After a long line of improvements [9] showed an $\tilde{O}(Mn^\omega)$ time algorithms for finding minimum weight perfect matching, shortest cycle, diameter and radius (some of these results were already known [21]). Also, [9] showed an application of their techniques that improves upon [31] $\tilde{O}(Mn^\omega t)$ time algorithm for the following problem: *determine the set of vertices that lie on some cycle of length at most t* . Cygan et al. [9] managed to solve this problem in $\tilde{O}(Mn^\omega)$ time using Baur-Strassen's theorem.

All of these algorithm are effective merely in the case of a dense graphs. For graphs with the small number of edges there are more efficient algorithms (e.g., APSP in $\tilde{O}(|V||E|)$ time [27]). But these algorithms are $\Theta(n^3)$ when $|E| = \Theta(n^2)$.

1.1 Related Work

1.1.1 Distance Queries

Yuster and Zwick [33] considered the weighted, directed graphs with weights in $\{-M, \dots, M\}$. They showed an algorithm that needs $\tilde{O}(Mn^\omega)$ preprocessing time. After preprocessing each distance $\delta(u, v)$ in the graph can be computed exactly in $O(n)$ query time. In the special case $M = 1$ they showed $\tilde{O}(n^\omega)$ algorithm that solves *Single Source Shortest Paths* (SSSP). This is the best known algorithm for a dense, weighted graph.

We will match their bounds (up to the polylogarithmic factors) using Frobenius normal form. Next we will extend that approach so it will return more information about a graph in the same query/preprocessing time.

1.1.2 Counting Cycles

For a given graph G and k determining whether G contains a simple cycle of length exactly k is NP-hard (in particular determining whether a graph contains a Hamiltonian cycle is NP-complete). However, if we fix k to be a constant this problem can be solved in polynomial time.

Alon et al. [4] introduced a color coding technique. For a fixed k if a graph $G(V, E)$ contains a simple cycle of size exactly k then such a cycle can be found in $\tilde{O}(|V|^\omega)$ time. Unfortunately, their algorithm depends exponentially $2^{O(k)}$ on the length of the cycle and in consequence is inapplicable for large k . In the next years, [5] showed (using a different technique) that for $k \leq 7$ one can count the number of cycles of length exactly k in a graph in $\tilde{O}(|V|^\omega)$ time. In [32] it is shown that for

any even k , cycles of length k can be found in $O(|V|^2)$ time in undirected graphs (if they contain such a cycle). Alon et al. [5] showed more methods that depend solely on a number of edges in a graph. For example for odd k they showed $O\left(E^{2-\frac{2}{k+1}}\right)$ algorithm for finding cycles of length k . However, for dense graphs these results are worse than [4].

From the other hand, to detect whether a non-simple cycle of length exactly k exists one can use the folklore algorithm. It starts by taking the adjacency matrix A of a graph G . Subsequently, in $\tilde{O}(n^\omega)$ time compute A^k by repeated squaring. If $\text{Tr}[A^k] > 0$ then there exists a non-simple cycle of length k .¹

Yuster [31] considered the following problem: *for every vertex in a graph find a shortest cycle that contains it*. He called this problem *All-Nodes Shortest Cycle* (ANSC). He showed a randomized algorithm that solves ANSC for undirected graphs with weights $\{1, \dots, M\}$ in $\tilde{O}(\sqrt{M}n^{(\omega+3)/2})$ time. He noted that for simple digraphs (directed graphs with no anti-parallel edges) it reduces to All-Pairs Shortest Path problem. The fastest known APSP algorithm for unweighted, directed graphs runs in $O(n^{2.575})$ due to [35]. Here, we will show how to solve ANSC in $\tilde{O}(n^\omega)$ for general, unweighted, directed graphs. Unfortunately, our techniques will allow us merely to find the length of such a cycle. But we can return the set of points, that lie on some cycle of a given length. Independently to our work [3] proved that ANSC can be solved in $\tilde{O}(n^\omega)$ for unweighted, undirected graphs using a completely different technique.

Yuster [31] also considered following problem: *given a graph and an integer t . Let $S(k)$ denote the set of all vertices lying in a cycle of length $\leq k$. Determine $S(t)$* . He considered directed graphs with weights in $\{-M, \dots, M\}$ and showed $\tilde{O}(Mn^\omega t)$ algorithm. Recently, [9] improved his algorithm. They showed that for a fixed $t \in [0, nM]$ the set $S(t)$ can be computed in $\tilde{O}(Mn^\omega)$ randomized time. We show, that for unweighted ($M = 1$) directed graphs we can compute sets $S(1), S(2), \dots, S(D)$ in $\tilde{O}(n^\omega)$ time with high probability.

2 Preliminaries

Let $T(n)$ be the minimal number of algebraic operations needed to compute the product of $n \times n$ matrix by an $n \times n$ matrix. We say that ω is the exponent of square matrix multiplication. For now the best known upper bound on ω is due to [16]:

Theorem 1 [16] *For every $\epsilon > 0$, $T(n) < O(n^{\omega+\epsilon})$ where $\omega < 2.37287$.*

In this paper, we will omit ϵ in definition and will assume that $O(n^\omega)$ operations are needed to multiply two matrices. The best lower bound for the exponent of matrix multiplication is $\omega \geq 2$. For convenience in this paper we will assume that $\omega > 2$. The $\tilde{O}$ notation hides polylogarithmic factors in complexity. We will use it to emphasize that all our algorithms need polylogarithmic number of calls to the fast matrix multiplication algorithm.

¹ $\text{Tr}[A]$ denotes the trace of a matrix A , i.e., the sum of elements on the main diagonal.

In this paper we will consider the Word RAM model of computation with a word size $O(\log n)$. The $\mathbb{F}$ denotes a small finite field and often we will assume this field to be $\mathbb{Z}_p$ for a prime number p with $O(\log n)$ bits. Note, that in Word RAM model arithmetic operations in $\mathbb{Z}_p$ can be done in constant time. In some situations we will need to stress that we want to work on larger integers. We will use W to denote the upper bound on the largest number in the field. In this case the arithmetic operations in Word RAM model take $O(\log W)$ time.

For matrices $A \in \mathbb{F}^{n \times k}$ and $B \in \mathbb{F}^{n \times l}$ the $A \oplus B \in \mathbb{F}^{n \times (k+l)}$ is the concatenation of their columns. $C_{a,b} \in \mathbb{F}^{n \times (b-a)}$ denotes a matrix constructed by concatenating columns $c_a, c_{a+1}, \dots, c_b$ of matrix $C \in \mathbb{F}^{n \times m}$.

In this paper we will consider the unweighted, directed graph G unless stated otherwise. We will denote by $V(G)$ and $E(G)$ the sets of vertices and edges of graph G . The adjacency matrix $A(G) \in \{0, 1\}^{n \times n}$ of a directed graph G with n vertices is defined as:

$$A(G)_{i,j} = \begin{cases} 1 & \text{if } i \rightarrow j \in E(G) \\ 0 & \text{if } i \rightarrow j \notin E(G). \end{cases}$$

3 Introduction to Cyclic Subspaces and Connection to Frobenius Matrices

We will start with a motivational example of application of cyclic subspaces. We have a constant-recursive sequence

$$a_n = c_0 a_{n-1} + c_1 a_{n-2} + \dots + c_{r-1} a_{n-r}$$

of order r , where all c_i are constants and initial conditions $a_0, \dots, a_{r-1}$ are given. Perhaps, the most familiar example is the Fibonacci sequence $F_n = F_{n-1} + F_{n-2}$.

We can define the *companion matrix* of our general sequence as:

$$C = \begin{bmatrix} 0 & \dots & 0 & -c_0 \\ 1 & \ddots & & -c_1 \\ & \ddots & 0 & \vdots \\ 0 & & 1 & -c_{r-1} \end{bmatrix}.$$

The crucial property of the matrix C is that we can generate the next element of sequence with multiplication by C :

$$C^T \begin{bmatrix} a_{n-r} \\ \vdots \\ a_{n-2} \\ a_{n-1} \end{bmatrix} = \begin{bmatrix} a_{n-r+1} \\ \vdots \\ a_{n-1} \\ a_n \end{bmatrix}.$$

For example, for Fibonacci sequence we have:

$$\begin{bmatrix} 0 & 1 \\ 1 & 1 \end{bmatrix} \begin{bmatrix} F_k \\ F_{k+1} \end{bmatrix} = \begin{bmatrix} F_{k+1} \\ F_{k+2} \end{bmatrix}.$$

Now, if we want the a_{n+1} element of the series we can square the companion matrix and get:

$$C^T C^T \begin{bmatrix} a_{n-r} \\ \vdots \\ a_{n-2} \\ a_{n-1} \end{bmatrix} = C^T \begin{bmatrix} a_{n-r+1} \\ \vdots \\ a_{n-1} \\ a_n \end{bmatrix} = \begin{bmatrix} a_{n-r+2} \\ \vdots \\ a_n \\ a_{n+1} \end{bmatrix}.$$

And analogously to get the a_{n+k-1} element we need to compute the k -th power of the matrix C . In linear algebra such transformations are known as the *cyclic subspaces* generated by the vector a . For our purposes, we will restrict ourself to a conclusion, that some columns of a companion matrix occur in its powers. These properties are well known in the linear algebra theory (see [10, 13] for more cyclic properties).

3.1 Consequences of Frobenius Normal Form

Let $\mathbb{F}$ be a commutative field. For any matrix $A \in \mathbb{F}^{n \times n}$ there exists an invertible U over $\mathbb{F}$ such that:

$$U^{-1}AU = F = \begin{bmatrix} C_1 & & 0 \\ & C_2 & \\ & & C_3 \\ & & & \ddots \\ 0 & & & & C_k \end{bmatrix}.$$

and F is the Frobenius-canonical-form² of A . The diagonal block C_i is called the companion matrix:

$$C_i = \begin{bmatrix} 0 & \dots & 0 & -c_0 \\ 1 & 0 & & -c_1 \\ & 1 & \ddots & -c_2 \\ & & \ddots & 0 \\ & & & 1 & 0 & -c_{r-2} \\ 0 & & & & 1 & -c_{r-1} \end{bmatrix} \in \mathbb{F}^{r \times r}.$$

Each companion matrix corresponds to the monic polynomial $C_i(x) = x^r + c_{r-1}x^{r-1} + \dots + c_0 \in \mathbb{F}[x]$ (similarly to the sequence example) and this polynomial is called the *minimal polynomial* of A . Each minimal polynomial has a property that $C_i(A) = 0$. To guarantee that matrix has only one decomposition into Frobenius normal form we require that every polynomial must divide the next one, i.e., $C_i(x) | C_{i+1}(x)$. The final list of polynomials is called the *invariant factors* of matrix A [25]. Storjohann [25] proposed the deterministic algorithm to compute the Frobenius canonical-form efficiently.

²Sometimes this form is called rational-canonical form.

Theorem 2 [25] *The Frobenius canonical-form of a matrix can be computed deterministically using $\tilde{O}(n^\omega)$ field operations.*

Moreover, there are also probabilistic algorithms that compute this form in expected $\tilde{O}(n^\omega)$ time over small fields [10]. In this paper, all algorithms are deterministic if we the upper bound on the number of distinct walks is W . Then, due to the time of a single field operation we need additional $O(\log W)$ factor in the complexity. However, since we are mainly interested in determining if a cycle/walk of a given length exists in a graph, we can set a sufficiently small field $\mathbb{Z}_p$ (p has $O(\log n)$ bits). This way when algorithm returns nonzero we are sure that there exists some walk. If algorithm returns zero, then with high probability there will be no such walk.

3.2 Cyclic Subspaces

Frobenius decomposition can be used to get the desired power of a matrix (analogously to the diagonal decomposition):

$$A^k = (UFU^{-1})^k = U F(U^{-1}U) F \cdots F(U^{-1}U) F U^{-1} = U F^k U^{-1}.$$

Moreover, we will use the property that the power of block diagonal matrix F is block diagonal:

$$F^k = \begin{bmatrix} C_1^k & & & 0 \\ & C_2^k & & \\ & & C_3^k & \\ & & & \ddots \\ 0 & & & & C_l^k \end{bmatrix}.$$

Now, we need a property of companion matrices that will enable us to power them efficiently (Fig. 1).

Definition 1 (Cyclic Property) Let $v_1, \dots, v_n$ be the columns of a matrix $C \in \mathbb{F}^{n \times n}$. Let $v_{n+1}, \dots, v_{2n}$ be the columns of matrix C^{n+1} . If, for every $1 \leq k \leq n$ the columns of matrix C^k are $v_k, v_{k+1}, \dots, v_{k+n}$ then the C has a **cyclic property**.

Fig. 1 Visualisation of the cyclic property (Definition 1)

It turns out, that companion matrices have a cyclic property.

Theorem 3 (Folklore [13], see [17] for generalization) *Every companion matrix has a cyclic property.*

We will illustrate this property with an example:

$$C = \begin{bmatrix} 0 & 0 & 0 & 0 & 1 \\ 1 & 0 & 0 & 0 & 2 \\ 0 & 1 & 0 & 0 & 3 \\ 0 & 0 & 1 & 0 & 4 \\ 0 & 0 & 0 & 1 & 5 \end{bmatrix}, \quad C^2 = \begin{bmatrix} 0 & 0 & 0 & 1 & 5 \\ 0 & 0 & 0 & 2 & 11 \\ 1 & 0 & 0 & 3 & 17 \\ 0 & 1 & 0 & 4 & 23 \\ 0 & 0 & 1 & 5 & 29 \end{bmatrix}, \quad C^3 = \begin{bmatrix} 0 & 0 & 1 & 5 & 29 \\ 0 & 0 & 2 & 11 & 63 \\ 0 & 0 & 3 & 17 & 98 \\ 1 & 0 & 4 & 23 & 133 \\ 0 & 1 & 5 & 29 & 168 \end{bmatrix},$$

$$C^4 = \begin{bmatrix} 0 & 1 & 5 & 29 & 168 \\ 0 & 2 & 11 & 63 & 365 \\ 0 & 3 & 17 & 98 & 567 \\ 0 & 4 & 23 & 133 & 770 \\ 1 & 5 & 29 & 168 & 973 \end{bmatrix}, \quad C^5 = \begin{bmatrix} 1 & 5 & 29 & 168 & 973 \\ 2 & 11 & 63 & 365 & 2114 \\ 3 & 17 & 98 & 567 & 3284 \\ 4 & 23 & 133 & 770 & 4459 \\ 5 & 29 & 168 & 973 & 5635 \end{bmatrix}, \quad C^6 = \begin{bmatrix} 5 & 29 & 168 & 973 & 5636 \\ 11 & 63 & 365 & 2114 & 12243 \\ 17 & 98 & 567 & 3284 & 19019 \\ 23 & 133 & 770 & 4459 & 25824 \\ 29 & 168 & 973 & 5635 & 32634 \end{bmatrix}$$

The matrix C^i has 4 columns identical to matrix C^{i+1} . C has coefficients of order equal to dimension (dimension is 5 and maximum coefficient is 5). After powering to the 5th power, the coefficients can be of order 5^5 . Over a finite field $\mathbb{Z}_p$, all those coefficients will have $O(\log p)$ bits.

4 Matching Distance Queries on Directed Unweighted Graphs

In this section, we will present an algorithm that matches the best known upper bounds of [33] for distance queries in directed unweighted graphs and uses Frobenius matrices.

We take the adjacency matrix A of a graph G (i.e., $n \times n$ matrix with $a_{u,v} = 1$ when $(u, v) \in G$ and 0 otherwise). The k -th power of the adjacency matrix of the graph G holds the number of walks, i.e., an $a_{u,v}$ element of A^k is *the count of distinct walks from u to v of length k in the graph*.

Observation 1 (Folklore, [8]) Let $A \in \{0, 1\}^{n \times n}$ be the adjacency matrix of a directed graph G . The $(A^k)_{u,v}$ is the number of distinct walks from u to v of length exactly k in the graph G .

Hence, the shortest path between vertices u, v is the smallest k such that A^k has nonzero element $a_{u,v}$. This will allow us to forget about graph theory interpretation for a brief moment and focus only on finding such k with algebraic tools.

In this section we will prove the following Lemma.

Lemma 1 *Given a matrix $A \in \mathbb{F}^{n \times n}$. There exists an algorithm that after some preprocessing, can answer queries for any given pair of indices $i, j \in \{1, \dots, n\}$ and integer $k \in \{1, \dots, n\}$, such that:*

- query returns an element $(A^k)_{i,j}$,

- preprocessing takes $\tilde{O}(n^\omega)$ field operations and query takes $O(n)$ field operations.

The algorithm is deterministic.

To proof this Lemma we decompose matrix A into the Frobenius normal form. Storjohann [25] showed an algorithm that returns U and F deterministically in $\tilde{O}(n^\omega)$ field operations (note that matrix inverse can also be computed in $\tilde{O}(n^\omega)$ field operations).

To better explain the idea, in the next section we will consider a simple case when a number of invariant factors of A is exactly 1. Then in Section 4.2 we will show how to generalize it to multiple invariant factors.

4.1 Single Invariant Factor

In that situation, the matrix F is a companion matrix $C \in \mathbb{F}^{n \times n}$. First, we compute the $(n+1)$ -th power of the companion matrix F^{n+1} . This can be done by using $\tilde{O}(n^\omega)$ field operations by repeated squaring (compute $F, F^2, F^4, \dots, F^{n+1}$ with $O(\log n)$ matrix multiplications).

Let $v_1, \dots, v_n$ be the columns of matrix UF and $v_{n+1}, \dots, v_{2n}$ be the columns of matrix UF^{n+1} . Note, that because matrix F has a cyclic property, the columns $v_k, \dots, v_{k+n-1}$ construct UF^k (see Fig. 2).

This step took just two matrix multiplications, because we need to multiply U times F and F^{n+1} . The preprocessing phase takes only $\tilde{O}(n^\omega)$ field operations.

Now, if a query asks for the number of distinct walks from vertices u to v of length exactly k we:

- select u -th row of matrix UF^k (n numbers),
- select v -th column of matrix U^{-1} ,
- multiply them in by using $O(n)$ multiplications (dot product of two n -dimensional vectors).

This will give us the u, v element of matrix $UF^k U^{-1} = A^k$. To get the length of the shortest path (i.e., the minimal k such that $(A^k)_{u,v} > 0$), we will modify our matrices slightly to get the number of walks of length $\leq k$. At the end, we will use in $\tilde{O}(n)$ query time (by using binary search) and $\tilde{O}(n^\omega)$ preprocessing time.

Fig. 2 Construction of UF^k from matrices UF and UF^{n+1}

Basically, for a given k we need to get the u, v element of matrix $A + A^2 + \dots + A^k$. It suffices to add consecutive columns of matrix $UF \oplus UF^{n+1} = v_1 \oplus v_2 \oplus \dots \oplus v_{2n}$ in the following manner:³

$$M' = \begin{bmatrix} v_1 & v_1 + v_2 & v_1 + v_2 + v_3 & \dots & \sum_{i=1}^k v_i & \dots & \sum_{i=1}^{2n} v_i \end{bmatrix} \in \mathbb{F}^{n \times 2n}.$$

Now, to get $A + A^2 + \dots + A^k$ one would need to multiply $M'_{k,k+n-1} U^{-1}$ and subtract $M'_{1,n} U^{-1}$ for a balance.⁴

The naive algorithm can transform matrices U and F to matrix M' in $O(n^2)$ field operations during preprocessing. During query, we will need to compute two dot products (u -th row of $M'_{k,k+n-1}$ times v -th column of U^{-1} and u -th row of $M'_{1,n}$ times v -th column of U^{-1}) and subtract them.

We have an algorithm that for a given vertices $u, v \in G$ and integer $k \in \{1, \dots, n\}$ can answer: *how many walks from u to v of length less or equal k are in the graph G in $\tilde{O}(n)$ query time and $\tilde{O}(n^\omega)$ preprocessing time.*

Because the result of the query is increasing in k we can use binary search. We can determine the first k for which the query will answer nonzero value in $O(\log n)$ tries. Hence, in $\tilde{O}(n)$ we can find the length of the shortest path. This generalized query can also return the number of walks of length exactly k , i.e., $q(u, v, k) - q(u, v, k - 1)$.

We matched the result of [33] for unweighted graphs with a single invariant factor. In the next section, we will show how to generalize this technique for graphs with any number of invariant factors.

4.2 Multiple Invariant Factors

Now, we will consider a case when $k \geq 1$, i.e., matrix F has multiple invariant factors. First of all, we need to note that this generalization is not perfect and will allow to compute the number of walks of length up to D (the longest distance in a graph, i.e., diameter).

In a real world applications of our framework (detecting cycles, determining distance between nodes, etc.) one does not need to consider walks longer than the longest possible distance in a graph. It is natural that the diameter is considered to be a bound of an output in graph problems [1, 2, 7, 9].

4.2.1 Relation of the Graph Diameter and Frobenius Normal Form

We begin with relating the graph diameter to the Frobenius normal form. It turns out that the graph diameter is bounded by the degree of a smallest invariant factor.

³Operation $\oplus$ denotes concatenation.

⁴ $X_{a,b}$ denotes a matrix constructed by concatenating columns $x_a, x_{a+1}, \dots, x_b$ of a matrix X .

Lemma 2 [8] *Given a directed, unweighted graph G with a diameter D . Let μ denote the degree of the smallest invariant factor (i.e., the dimension of the smallest block in the Frobenius matrix F) of an adjacency matrix of the graph G . Then $D \leq \mu$.*

This theorem is well known in literature [8]. We include the proof of this theorem for completeness.

Proof For a contradiction assume that $D > \mu$ and let $u, v \in G$ be vertices such that $\delta(u, v) = D$. Let A be the adjacency matrix of G . We know, that there is the minimal polynomial of degree μ :

$$A^\mu = a_0 \mathbb{1} + a_1 A + a_2 A^2 + \dots + a_{\mu-1} A^{\mu-1}.$$

Term $a_{i,j}^k$ denotes the i, j element of the matrix A^k . Now, consider the elements u, v of each matrix. The diameter $D > \mu$ and $\delta(u, v) = D$, so for every $k \leq \mu$ the elements $a_{u,v}^k = 0$ (because there is no walk of length less than D from u to v). Now, if we multiply the minimal polynomial by the matrix A we get:

$$A^{\mu+1} = a_0 A + a_1 A^2 + a_2 A^3 + \dots + a_{\mu-1} A^\mu.$$

Hence $a_{u,v}^{\mu+1} = 0$, because every element in the sequence $a_{u,v}^k = 0$ for $k \leq \mu$. By repeating this reasoning, we get that for every $k > 0$ the element $a_{u,v}^k = 0$. So, for every achievable pair of vertices, there must be some $k \leq \mu$, such that $a_{u,v}^k \neq 0$ and diameter is bounded by μ . $\square$

The bounds of this inequality are tight. There are graphs with diameter $D = \mu$ and graphs with $\mu = n$ and arbitrary small diameter [8]. Our algorithms are able to return walks up to the length μ . We use the bound on D solely because it is easier to interpret diameter than the *smallest degree of the invariant factor*.

Fig. 3 Multiplication of the UFU^{-1} . Example for 3 invariant factors. We know that matrix F is block and consists of F_1, F_2, F_3 . We divide matrix U into strips U_1, U_2, U_3 that correspond to blocks of F . The observation is that we can compute $U_1 \cdot F_1, U_2 \cdot F_2$ and $U_3 \cdot F_3$ independently and concatenate them into matrix UF

4.2.2 Generalization to Multiple Invariant Factors

Let k denote the number of blocks in the Frobenius matrix F and μ be the number of columns of the smallest block. To multiply the matrix U by F we can start by multiplying strips of matrix U by appropriate blocks of F and concatenate them later (see Fig. 3).

We start by splitting the matrix U into k strips with rows corresponding to the appropriate blocks of F (strip U_i has as many columns as block F_i). Then we multiply UF and have k strips: $U_1F_1, U_2F_2, \dots, U_kF_k$ (each with at least μ columns). Next, we multiply UF^μ and also keep k strips: $U_1F_1^\mu, U_2F_2^\mu, \dots, U_kF_k^\mu$. Our goal is to get a data structure such that if we need UF^k , we can quickly choose appropriate columns and append them.

The matrix U_iF_i has l_i columns: $v_1, \dots, v_{l_i}$. Because F_i is a companion matrix, the $U_iF_i^\mu$ has the cyclic property (Definition 1). And the matrix $U_iF_i^\mu$ has columns: $v_\mu, \dots, v_{\mu+l_i}$. Note, that there are some duplicate columns in U_iF_i and $U_iF_i^\mu$, when $\mu < l_i$. Hence, we only need to keep columns $v_1, \dots, v_{\mu+l_i}$ for this strip. We do this for all strips $U_1F_1, \dots, U_kF_k$ (see Fig. 4).

We are left with a matrix that has at most $2n$ columns (because $l_1 + \mu + l_2 + \mu + \dots + l_k + \mu = k\mu + \sum_{i=1}^n l_i = n + k\mu \leq 2n$). To generate it we need to power F to μ and do multiplications $U \cdot F$ and $U \cdot F^\mu$. This can be computed in $O(n^\omega)$ field operations via fast matrix multiplication and repeated squaring.

4.2.3 Queries with Multiple Invariant Factors

When a query for the number of walks of length k from node u to v comes, we do:

1. For each strip i take the u -th row of $U_iF_i \oplus U_iF_i^\mu$ and concatenate them (see Fig. 5) into vector $\bar{u}$,
2. Take v -th column of U^{-1} matrix and denote it $\bar{v}$,
3. Return the dot product $\bar{u} \cdot \bar{v}$.

Because $l_1 + l_2 + \dots + l_k = n$ the vector $\bar{u} \in \mathbb{F}^n$. Query needs $O(n)$ field operations.

Finally, this dot product is $a_{u,v}$ element of the matrix UF^kU^{-1} , for a fixed $k \leq \mu$ because $\bar{u}$ is the concatenation of original vector u . Analogously to Section 4.1 one can extend this result to return the number of walks of length less or equal k . This matches (up to the polylogarithmic factor) the result of [33]. We will omit the details of this observation because in the next section we will extend this framework even further.

Fig. 4 Combining strips into a single matrix. The height of the matrix in the schema is scaled down. We computed U_1F_1 and $U_1F_1^\mu$. Now we noted that companion matrices have a cyclic property so some of the rows in the strips are repeated. So in the single strip we can store only subsequent columns

Fig. 5 Schema of obtaining vector $\tilde{u}$ (marked red) from 3 strips. We are given the row number and the power k and the lengths l_i of each strip. At the end we concatenate them

5 Almost Optimal Query

In the previous section, we showed how to preprocess a matrix A with $\tilde{O}(n^\omega)$ field operations in such a way that in query that uses $O(n)$ field operations we can return a number $(A^k)_{i,j}$. However, in linear time $O(n)$ we return only a single number. The goal of this section is to get far richer information in $\tilde{O}(n)$ query time (with extra factors from field operations).

Theorem 4 *Let $A \in \{0, 1\}^{n \times n}$ be a matrix such that the degree of smallest invariant factor is μ . There exists a deterministic algorithm that after some preprocessing can answer queries for any given $i, j \in \{1, \dots, n\}$:*

- *Query returns $\{a_k \mid 1 \leq k \leq \mu\}$, where $a_k = (A^k)_{i,j}$,*
- *Preprocessing takes $\tilde{O}(n^\omega \log W)$ and query takes $O(n \log W)$ time,*

where W is an upper bound on a_k for all $k \in \{1, \dots, \mu\}$.

Note, that this theorem has some immediate application in graph algorithms (see Section 6).

5.1 Hankel Matrix

Now, we will focus on the proof of Theorem 4. First, we need to introduce the Hankel matrix and its properties.

$$H = \begin{pmatrix} c_1 & c_2 & \dots & c_n \\ c_2 & c_3 & \dots & c_{n+1} \\ \vdots & \vdots & & \vdots \\ c_n & c_{n+1} & \dots & c_{2n-1} \end{pmatrix}$$

Hankel matrix is defined by its first row and last column ($2n - 1$ numbers define $n \times n$ Hankel matrix). The numbers from the previous row are left-shifted by one and the new number is added at the end. Hankel matrices have some similarities to Topelitz and Circulant matrices.

The basic property we need is that the product of Hankel matrix and vector can be computed in $O(n \log n)$ time (see [15, 26]) even though explicitly writing the Hankel matrix as $n \times n$ matrix takes $O(n^2)$ time. The algorithm takes $2n - 1$ parameters that define the Hankel matrix and n parameters that define the vector. The technique is based on the Fast Fourier Transformation [15, 26].

5.2 Using Hankel Matrices to Get Richer Query

To proof the Theorem 4 we will modify only the last step in Section 4.2.3. The algorithm from Section 4.2.3 concatenates the strips $U_i F_i$ and builds a single vector. Subsequently, that vector is multiplied by a column of matrix U^{-1} . But we can also do it in a different order: first we multiply the strip by a section of matrix U^{-1} and sum the results at the end. Thus, we perform k (number of strips) dot products of smaller vectors (see Fig. 6).

Consider the query for a number of walks of length exactly k . The strips in the matrix U^{-1} do not depend on k (vector $(u_0, \dots, u_l)$). However, the vector taken from $U_i F_i$ (vectors $(x_i, \dots, x_{i+l})$) will be left shifted if we want to compute the next one.

$$\begin{pmatrix} x_0 & x_1 & x_2 & \dots & x_l \\ x_1 & x_2 & \dots & x_l & x_{l+1} \\ x_2 & \dots & x_l & x_{l+1} & x_{l+2} \\ \vdots & & & \vdots & \\ x_\mu & \dots & & x_{\mu+l} \end{pmatrix} \times \begin{pmatrix} u_0 \\ \vdots \\ u_l \end{pmatrix}$$

As you can see, the subsequent rows can be written as the Hankel matrix (we need to add zeros to get a square matrix, but it will not influence asymptotic complexity since there will be at most $O(n)$ of them). By using the *fast Hankel matrix-vector multiplication* we can compute μ values for every strip i in time $O(l_i \log l_i)$ (l_i was defined as the length of i -th strip). At the end, we need to sum all results into a single array. Therefore, the number of operations is $O\left(\sum_{i=1}^k l_i \log l_i\right)$. Because $\sum_{i=1}^k l_i = n$ the algorithm needs $O(n \log n)$ field operations. This proves Theorem 4.

Here, we silently assumed that the number of walks is bounded by W . Note, that for large W , the algorithm needs to output $O(n^2 \log W)$ bits and the complexity of every arithmetic operation needs to be multiplied by $\log W$. If one is only interested in the deciding if an entry of some power of adjacency matrix is nonzero, we can use a standard randomization technique to eliminate $\log W$ factors from the running time.

Corollary 1 *Let $A \in \{0, 1\}^{n \times n}$ be a matrix such that the degree of smallest invariant factor is μ . There exists an algorithm that after some preprocessing can answer queries for any given $i, j \in \{1, \dots, n\}$:*

Fig. 6 Multiplication of strips by U^{-1} matrix. As you can see, matrix U^{-1} can be splitted into sections, that multiply only $U_i F_i$ strips

- Query returns $\{a_k \mid 1 \leq k \leq \mu\}$, where $a_k = 1$ if $(A^k)_{i,j}$ is nonzero,
- Preprocessing takes $\tilde{O}(n^\omega)$ and query takes $\tilde{O}(n)$ time.

The algorithm is randomized with one-sided bounded error.

Proof At the beginning we will randomly select a prime number with $O(\log n)$ bits. We can write the matrix A as a polynomial:

$$\tilde{A}_{i,j} = \begin{cases} x_{i,j} & \text{if } A_{i,j} = 1 \\ 0 & \text{otherwise,} \end{cases}$$

where $x_{i,j}$ are unique variables. Now we can apply Schwartz-Zippel Lemma [23, 34]. The $\tilde{A}_{i,j}^k$ is the polynomial of degree at most n (because $k \leq \mu \leq n$) and if we compute it over $\mathbb{Z}_p$ for $p = O(n^2)$ we can correctly determine if $\tilde{A}_{i,j}^k$ is a nonzero polynomial with a constant probability. We can repeat the above procedure $O(\log n)$ times to get a correct result in all entries with high probability. $\square$

6 Applications

In this section we will show how to use Theorem 4 to improve known algorithms on graphs. First we will develop a data structure that returns a number of distinct walks efficiently.

Lemma 3 *Let $G = (V, E)$ be a directed, unweighted graph with n vertices and a diameter D . There exists an deterministic algorithm that after some preprocessing can answer queries for any given $u, v \in V$:*

- Query returns $\{w_i \mid 1 \leq i \leq D\}$, where w_i is the number of distinct walks from u to v of length exactly i ,
- Preprocessing takes $\tilde{O}(n^\omega \log W)$ and query takes $\tilde{O}(n \log W)$ field operations.

where W is the upper bound on w_i for all $i \in \{1, \dots, D\}$.

Proof We encode the graph G as an adjacency matrix $A(G)$. We use the Theorem 4 to construct the data structure that given a query (u, v) outputs $(A^k)_{u,v}$ for all $1 \leq k \leq \mu$. Finally, we use Observation 1 to note, that $(A^k)_{u,v}$ is equal to the number of distinct walks from u to v of length exactly k . Moreover we use Lemma 2 to bound the number $D \leq \mu$, so we will always output more numbers (but we can truncate them in $O(n)$ time). Finally we note, that the preprocessing and query of Theorem 4 matches the statement and construction of adjacency matrix is $O(n^2)$. $\square$

This algorithm is a significant improvement over [33]:

- One can use Lemma 3 to find the distance between u, v by linearly scanning the array and returning the first k such that $w_k > 0$,
- Lemma 3 can count cycles. In contrast the [33] cannot, because the distance from u to u is always 0 (see Section 6.1),

- Lemma 3 is almost optimal, i.e., when $D = O(n)$ then query will need to output $O(n \log W)$ bits.

From the other hand, Lemma 3 is merely a functional improvement and it does not break the $\tilde{O}(n^\omega)$ of the *Single Source Shortest Paths* (SSSP) for dense, directed graphs.

Now we will show the application of Lemma 3. We begin with almost optimal algorithm to compute the number of all walks between all pairs of vertices. We are not aware of any other research concerning this problem.

Definition 2 (All-Pairs All Walk problem) Given a directed, unweighted graph G with a diameter D . The task is to return an array A , such that for every pair of vertices $u, v \in G$ and every $k \in \{1, \dots, D\}$ an element $A[u, v, k]$ is the number of distinct walks from u to v of length k .

The folklore solution to this needs $O(Dn^\omega \log W)$ time (where W is an upper bound on number of walks) and works as follows: *take the adjacency matrix A of graph G and save it in $A[u, v, 1]$. Then, square it to get A^2 and save it in $A[u, v, 2]$. Continue until you fill out complete table.* In the worst case this algorithm needs $D = O(n)$ matrix multiplications, thus it needs $\tilde{O}(Dn^\omega)$ field operations. At the first glance it is surprising that we can improve it to $\tilde{O}(n^3)$ field operations.

Theorem 5 *All-Pairs All Walk problem admits an $\tilde{O}(n^3 \log W)$ algorithm (where W is upper bound on number of walks between every pair of vertices).*

Proof We will apply the Lemma 3 algorithm. The preprocessing takes $\tilde{O}(n^\omega)$ time. Then, for every pair of vertices u, v ask a query. A single query takes $\tilde{O}(n \log W)$ time. Next we will save it in the table $A[u, v]$ (query gives D numbers $w_1, \dots, w_D$, such that w_i is the number of walks of length i and save it $A[u, v, i] := w_i$).

Because there are $O(n^2)$ pairs and for each pair we need $\tilde{O}(n \log W)$ time, the complexity of our solution is $\tilde{O}(n^3 \log W)$. The algorithm is almost optimal because the output in the worst case may be $O(n^3 \log W)$ (we may need $O(\log W)$ bits to encode a single entry in the table). $\square$

6.1 Counting and Determining the Lengths of Cycles

We will use Theorem 4 to solve All-Nodes Shortest Cycle (ANSC) problem efficiently.

Definition 3 (All-Nodes Shortest Cycles [31]) Given a directed, unweighted graph G . The problem *All-Nodes Shortest Cycle* asks to output for every vertex v the length of the shortest cycle that contains v .

Lemma 4 *There exists a deterministic algorithm that for a given unweighted, directed G with a diameter D :*

- For every vertex u returns D numbers: $c_u^1, c_u^2, \dots, c_u^D$, where c_u^k is the number of non-simple cycles of length exactly k , that contain vertex u ,
- Algorithm works in $\tilde{O}(n^\omega \log W)$ time (where W is an upper bound on c_u^k).

Proof We will use Theorem 4. We start by preprocessing the graph G in time $\tilde{O}(n^\omega \log W)$. Theorem 4 allows us to ask for a number of walks from u to v and receive D numbers: $w_{u,v}^k$. So, we ask for the number of walks from vertex u to the same vertex u . This is exactly the number of non-simple cycles of given length that contain vertex u .

Because we need to ask only n queries (it is the number of vertices in a graph) and each query takes $\tilde{O}(n \log W)$ time we have $\tilde{O}(n^\omega \log W + n^2 \log W) = \tilde{O}(n^\omega \log W)$ algorithm. $\square$

If we are only interested in deciding if the numbers c_u^i are nonzero, instead of Theorem 4 we can use Corollary 1. It introduces the one-sided randomization but allows us to shave $\log W$ factors in the running time.

Corollary 2 *There exists a randomized algorithm that for a given unweighted, directed G with a diameter D :*

- For every vertex u returns D numbers: $c_u^1, c_u^2, \dots, c_u^D$, where c_u^k is 1 if there exists a non-simple cycle of length exactly k , that contain vertex u or 0 otherwise,
- Algorithm works in $\tilde{O}(n^\omega)$ time with one sided bounded error.

Now we will show how to improve upon [31] $\tilde{O}(n^{(\omega+3)/2})$ algorithm with Corollary 2.

Theorem 6 *All-Nodes Shortest Cycles admits an $\tilde{O}(n^\omega)$ randomized time algorithm.*

Proof We use Lemma 2 to compute the table $S[v]$. For every vertex we search for the first nonzero element linearly. This with high probability is the length of the shortest cycle that contains it. Because the output contains $O(n^2)$ numbers the complexity is equal to the preprocessing time $\tilde{O}(n^\omega)$. $\square$

Also the Corollary 2 improves upon [9, Theorem 45] for unweighted graphs.

Corollary 3 *Given a directed, unweighted graph G with a diameter D . Let $S(c)$ denote the set of vertices that lie in the cycle of length exactly c . In $\tilde{O}(n^\omega)$ time we can return the sets $S(1), \dots, S(D)$ with constant probability of success.*

Proof Similarly to the proof of Theorem 6, we can scan the output to compute the set $S(c)$ that contains all vertices that lie on some cycle of length $\leq c$. Then, by linear scan we can return the sets $S(1), \dots, S(D)$. $\square$

7 Conclusion and Open Problems

We introduced the framework based on Frobenius normal form and used it to solve some problems on directed, unweighted graphs in matrix multiplication time. The main open question is to use this framework to prove that APSP on such graphs can be solved in $\tilde{O}(n^\omega)$ or at least $O(n^{2.5})$. The promising way is to use the algorithms that determine operators of matrices of polynomials (e.g., determinant, solving linear system [14, 18]). Additionally, algorithms for a black-box polynomial degree determination seem to be a promising path.

Another interesting problem is to use this framework to obtain additive approximation for APSP. Currently, the best additive approximation of APSP is due to [20]. However, no additive approximation of APSP is known that would work in $\tilde{O}(n^\omega)$ time.

Application in dynamic algorithm also seems to be a promising approach. Frandsen and Sankowski [12] showed an algorithm, that dynamically preserves Frobenius normal form in $O(kn^2)$ time. Our algorithms use fast Hankel matrix-vector multiplication that is based on FFT. Reif and Tate [19] presented an $O(\sqrt{n})$ time per request algorithm for FFT. Can we use these approaches to obtain a faster dynamic algorithm?

Finally, it remains open how to apply the Frobenius normal form in the weighted directed graphs with small, integer weights $\{-M, \dots, M\}$. Cygan et al. [9] took degree M polynomials and used [18] algorithms to get $\tilde{O}(Mn^\omega)$ time radius and diameter detection. We suspect that similar technique can be applied to Frobenius normal form framework.

Acknowledgment This work is part of a project TOTAL that has received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreement No 677651). Piotr Sankowski is supported by the Polish National Science Center, grant no 2014/13/B/ST6/00770. Karol Węgrzycki is supported by the grants 2016/21/N/ST6/01468 and 2018/28/T/ST6/00084 of the Polish National Science Center. Part of this work was done while authors were visiting the Simons Institute for the Theory of Computing.

Open Access This article is distributed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted use, distribution, and reproduction in any medium, provided you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made.

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

References

1. Abboud, A., Grandoni, F., Williams, V.V.: Subcubic equivalences between graph centrality problems, APSP and diameter. In: Indyk, P. (ed.) Proceedings of the 26th Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2015, pp. 1681–1697. SIAM, San Diego (2015). ISBN 978-1-61197-374-7. <https://doi.org/10.1137/1.9781611973730.112>
2. Abboud, A., Williams, V.V., Wang, J.R.: Approximation and fixed parameter subquadratic algorithms for radius and diameter in sparse graphs. In: Krauthgamer, R. (ed.) Proceedings of the 27th Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2016, pp. 377–391. SIAM, Arlington (2016). ISBN 978-1-61197-433-1. <https://doi.org/10.1137/1.9781611974331.ch28>

3. Agarwal, U., Ramachandran, V.: Fine-grained complexity for sparse graphs. In: Diakonikolas, I., Kempe, D., Henzinger, M. (eds.) *Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing*, STOC 2018, pp. 239–252. ACM, Los Angeles (2018). <https://doi.org/10.1145/3188745.3188888>
4. Alon, N., Yuster, R., Zwick, U.: Color-coding. *J. ACM* **42**(4), 844–856 (1995). <https://doi.org/10.1145/210332.210337>
5. Alon, N., Yuster, R., Zwick, U.: Finding and counting given length cycles. *Algorithmica* **17**(3), 209–223 (1997). <https://doi.org/10.1007/BF02523189>
6. Chan, T.M.: More algorithms for all-pairs shortest paths in weighted graphs. In: Johnson, D.S., Feige, U. (eds.) *Proceedings of the 39th Annual ACM Symposium on Theory of Computing*, pp. 590–598. ACM, San Diego (2007). ISBN 978-1-59593-631-8. <https://doi.org/10.1145/1250790.1250877>
7. Chung, F.R.K., Faber, V., Manteuffel, T.A.: An upper bound on the diameter of a graph from eigenvalues associated with its laplacian. *SIAM J. Discrete Math.* **7**(3), 443–457 (1994). <https://doi.org/10.1137/S0895480191217776>
8. Chung, F.R.K.: *Spectral graph theory*, volume 92 American Mathematical Soc. (1997)
9. Cygan, M., Gabow, H.N., Sankowski, P.: Algorithmic applications of baur-strassen’s theorem: shortest cycles, diameter, and matchings. *J. ACM* **62**(4), 28 (2015). <https://doi.org/10.1145/2736283>
10. Eberly, W.: Black box frobenius decompositions over small fields. In: [28], pp. 106–113. ISBN 1-58113-218-2. <https://doi.org/10.1145/345542.345596>
11. Floyd, R.W.: Algorithm 97: shortest path. *Commun. ACM* **5**(6), 345 (1962). <https://doi.org/10.1145/367766.368168>
12. Frandsen, G.S., Sankowski, P.: Dynamic normal forms and dynamic characteristic polynomial. In: *International Colloquium on Automata, Languages, and Programming*, pp. 434–446. Springer, Berlin (2008)
13. Gantmacher, F.R.: *The theory of matrices*, Vol. 1 Chelsea (1959)
14. Giesbrecht, M., Jacobson, M.J. Jr., Storjohann, A.: Algorithms for large integer matrix problems. In: Boztas, S., Shparlinski, I.E. (eds.) *14th International Symposium Applied Algebra, Algebraic Algorithms and Error-Correcting Codes, AAECC-14. Proceedings, volume 2227 of Lecture Notes in Computer Science*, pp. 297–307. Springer, Melbourne (2001). ISBN 3-540-42911-5. https://doi.org/10.1007/3-540-45624-4_31
15. Golub, G.H., Van Loan, C.F. *Matrix computations*, 3rd edn. Johns Hopkins University Press, Baltimore (1996). ISBN 978-0-8018-5414-9
16. Le Gall, F.: Powers of tensors and fast matrix multiplication. In: *Proceedings of the 39th International Symposium on Symbolic and Algebraic Computation*, pp. 296–303. ACM (2014)
17. Lim, A., Dai, J.: On product of companion matrices. *Linear Algebra Appl.* **435**(11), 2921–2935 (2011)
18. Mulders, T., Storjohann, A.: Rational solutions of singular linear systems. In: [28], pp. 242–249. ISBN 1-58113-218-2. <https://doi.org/10.1145/345542.345644>
19. Reif, J.H., Tate, S.R.: On dynamic algorithms for algebraic problems. *J. Algorithms* **22**(2), 347–371 (1997). <https://doi.org/10.1006/jagm.1995.0807>
20. Roditty, L., Shapira, A.: All-pairs shortest paths with a sublinear additive error. *ACM Trans. Algorithms* **7**(4), 45:1–45:12 (2011). ISSN 1549-6325. <https://doi.org/10.1145/2000807.2000813>
21. Roditty, L., Williams, V.V.: Minimum weight cycles and triangles: equivalences and algorithms. In: Ostrovsky, R. (ed.) *IEEE 52nd Annual Symposium on Foundations of Computer Science, FOCS 2011*, pp. 180–189. IEEE Computer Society, Palm Springs (2011). <https://doi.org/10.1109/FOCS.2011.27>
22. Sankowski, P., Wegrzycki, K.: Improved distance queries and cycle counting by frobenius normal form. In: Vollmer, H., Vallée, B. (eds.) *34th Symposium on Theoretical Aspects of Computer Science, STACS 2017, March 8-11, vol. 66 of LIPIcs*, pp. 56:1–56:14. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, Hannover (2017). ISBN 978-3-95977-028-6. <https://doi.org/10.4230/LIPIcs.STACS.2017.56>
23. Schwartz, J.T.: Fast probabilistic algorithms for verification of polynomial identities. *J. ACM* **27**(4), 701–717 (1980). ISSN 0004-5411. <https://doi.org/10.1145/322217.322225>
24. Seidel, R.: On the all-pairs-shortest-path problem. In: Rao Kosaraju, S., Fellows, M., Wigderson, A., Ellis, J.A. (eds.) *Proceedings of the 24th Annual ACM Symposium on Theory of Computing*, May 4-6, 1992, pp. 745–749. ACM, Victoria (1992). ISBN 0-89791-511-9. <https://doi.org/10.1145/129712.129784>

25. Storjohann, A.: Deterministic computation of the frobenius form. In: 42nd Annual Symposium on Foundations of Computer Science, FOCS 2001, 14–17 October 2001, pp. 368–377. IEEE Computer Society, Las Vegas (2001). ISBN 0-7695-1390-5. <https://doi.org/10.1109/SFCS.2001.959911>
26. Tang, Z., Duraiswami, R., Gumerov, N.A.: Fast algorithms to compute matrix-vector products for pascal matrices. Technical Reports from UMIACS UMIACS-TR-2004-08, 2004/03/25/. <http://drum.lib.umd.edu/handle/1903/1338> (2004)
27. Thorup, M.: Undirected single-source shortest paths with positive integer weights in linear time. *J. ACM* **46**(3), 362–394 (1999). <https://doi.org/10.1145/316542.316548>
28. Traverso, C. (ed.): Proceedings of the 2000 International Symposium on Symbolic and Algebraic Computation, ISSAC 2000. ACM, St. Andrews (2000). ISBN 1-58113-218-2. <http://dl.acm.org/citation.cfm?id=345542>
29. Warshall, S.: A theorem on boolean matrices. *J. ACM* **9**(1), 11–12 (1962). <https://doi.org/10.1145/321105.321107>
30. Williams, R.: Faster all-pairs shortest paths via circuit complexity. In: Shmoys, D.B. (ed.) Symposium on Theory of Computing, STOC 2014, pp. 664–673. ACM, New York (2014). ISBN 978-1-4503-2710-7. <https://doi.org/10.1145/2591796.2591811>
31. Yuster, R.: A shortest cycle for each vertex of a graph. *Inf. Process. Lett.* **111**(21–22), 1057–1061 (2011). <https://doi.org/10.1016/j.ipl.2011.07.019>
32. Yuster, R., Zwick, U.: Finding even cycles even faster. In: Abiteboul, S., Shamir, E. (eds.) 21st International Colloquium, Automata, Languages and Programming, ICALP94. Proceedings, volume 820 of Lecture Notes in Computer Science, pp. 532–543. Springer, Jerusalem (1994). ISBN 3-540-58201-0. <https://doi.org/10.1007/3-540-58201-0.96>
33. Yuster, R., Zwick, U.: Answering distance queries in directed graphs using fast matrix multiplication. In: Proceedings of 46th Annual IEEE Symposium on Foundations of Computer Science (FOCS 2005), 23–25 October 2005, pp. 389–396. IEEE Computer Society (2005). ISBN 0-7695-2468-0. <https://doi.org/10.1109/SFCS.2005.20>
34. Zippel, R.: Probabilistic algorithms for sparse polynomials. In: Ng, E.W. (ed.) Symbolic and algebraic computation, pp. 216–226. Springer, Berlin (1979). ISBN 978-3-540-35128-3
35. Zwick, U.: All pairs shortest paths using bridging sets and rectangular matrix multiplication. *J. ACM* **49**(3), 289–317 (2002). <https://doi.org/10.1145/567112.567114>