Electronic Colloquium on Computational Complexity, Report No. 111 (2006)

Faster algorithms for finding lowest common ancestors in directed acyclic graphs*

Artur Czumaj† Mirosław Kowaluk‡ Andrzej Lingas§

Abstract

We present two new methods for finding a lowest common ancestor (LCA) for each pair of vertices of a directed acyclic graph (dag) on n vertices and m edges.

The first method is a natural approach that solves the all-pairs LCA problem for the input dag on n vertices and m edges in time $O(nm)$.

The second method relies on a novel reduction of the all-pairs LCA problem to the problem of finding maximum witnesses for Boolean matrix product. We solve the latter problem and hence also the all-pairs LCA problem in time $O(n^{2+\lambda})$, where λ satisfies the equation $\omega(1,\lambda,1) = 1 + 2\lambda$ and $\omega(1,\lambda,1)$ is the exponent of the multiplication of an $n \times n^{\lambda}$ matrix by an $n^{\lambda} \times n$ matrix. By the currently best bounds on $\omega(1,\lambda,1)$, the running time of our algorithm is $O(n^{2.575})$. Our algorithm improves the previously known $O(n^{2.688})$ time-bound for the general all-pairs LCA problem in dags by Bender *et al.*

Our additional contribution is a faster algorithm for solving the all-pairs LCA problem in dags of small depth, where the depth of a dag is defined as the length of the longest path in the dag. For all dags of depth at most $h \leq n^{\alpha}$, where $\alpha \approx 0.294$, our algorithm runs in time asymptotically the same as that of multiplying two $n \times n$ matrices, that is, $O(n^{\omega})$; we also prove that this running time is optimal even for dags of depth 1. For dags with depth $h > n^{\alpha}$, the running time of our algorithm is at most $O(n^{\omega} \cdot h^{0.468})$. This algorithm is faster than our algorithm for arbitrary dags for all values of $h \leq n^{0.42}$.

1 Introduction

The problem of finding a *lowest common ancestor* (LCA) in a tree, or more generally, in a *directed acyclic graph* (dag) is a basic problem in algorithmic graph theory. An LCA of vertices u and v in a dag is an ancestor of both u and v that has no descendant which is an ancestor of u and v, see

*This paper is the result of merging the ICALP'05 paper [14] with [10].

†Department of Computer Science, New Jersey Institute of Technology, Newark, NJ 07102, USA. Email: aczumaj@acm.org. Research supported in part by NSF ITR grant CCR-0313219.

‡Institute of Informatics, Warsaw University, Warsaw, Poland. Email: kowaluk@mimum.edu.pl. Research supported by KBN grant 4T11C04425.

§Department of Computer Science, Lund University, 22100 Lund, Sweden. Email: Andrzej.Lingas@cs.lth.se. Research supported in part by VR grant 621-2002-4049.

ISSN 1433-8092

Fig. 1 for example. We consider the problem of preprocessing a dag such that LCA queries can be answered quickly for any pair of vertices.

This problem has been very extensively studied in the literature in the context of (rooted) trees (see, e.g., [3, 12, 17]), where it appears naturally in various settings and where it found many applications in design of efficient algorithms and data structures (e.g., for the problem of computing maximum matching in graphs and for various string problems). Harel and Tarjan [12] were the first who showed that in rooted trees, the LCA queries can be answered in constant time after a linear time preprocessing of the input tree. This work has been later extended in many ways, including the simplified algorithm from [3] and the recent dynamic algorithm [5].

Recently, Bender *et al.* initiated investigations of the LCA problem for arbitrary directed acyclic graphs in [4]. They listed a number of natural applications of the LCA queries in dags in [4], e.g., in inheritance analysis in programming languages, analysis of genealogical data, and lattice operations in complex systems (for more details, see e.g., [9, 15, 16], and especially [4] and the references therein). Bender *et al.* also observed that the all-pairs LCA problem in dags is not simpler than transitive closure in arbitrary directed graphs and Boolean matrix multiplication [4]. On the other hand, they were first to design substantially subcubic-time solution to the all-pairs LCA problem in dags. Their algorithm for the all-pairs LCA problem in dags runs in time $O(n^{(3+\omega)/2}) = O(n^{2.688})$, where n is the number of vertices and $\omega < 2.376$ is the exponent of the fastest matrix multiplication algorithm [7].

New contributions. We present two new methods of efficiently preprocessing a directed graph on n vertices and m edges in order to answer an LCA query for any pair of vertices in constant time, subsuming the previously known best results from [4].

The first method is a natural approach that solves the all-pairs LCA problem for the input dag on n vertices and m edges in time $O(nm)$. For sparse dags, this method is optimal and substantially faster than the known $O(n^{\frac{\omega+3}{2}})$-time general method from [4].

The second method efficiently reduces the all-pairs LCA problem to the problem of finding maximum (index) witnesses for Boolean matrix product. We solve the latter problem and hence also the all-pairs LCA problem in time $O(n^{2+\lambda})$, where λ satisfies the equation $\omega(1,\lambda,1) = 1+2\lambda$ and $\omega(1,\lambda,1)$ is the exponent of the multiplication of an $n \times n^{\lambda}$ matrix by an $n^{\lambda} \times n$ matrix. By the currently best bounds on $\omega(1,\lambda,1)$, the running time of our algorithm is $O(n^{2.575})$. Our algorithm improves the previously known $O(n^{2.688})$ time-bound for the general all-pairs LCA problem in dags by Bender *et al.* [4].

In addition, we present a faster algorithm for solving the all-pairs lowest common ancestor problem in dags of small depth, where the *depth* of a dag is defined as the length of the longest path in the dag [1]. For all dags of depth at most $h \leq n^{\alpha}$, where $\alpha \approx 0.294$, our algorithm runs in time asymptotically the same as that of multiplying two $n \times n$ matrices, that is, $O(n^{\omega})$; we also prove that this running time is optimal even for dags of depth 1. For dags with depth $h > n^{\alpha}$, the running time of our algorithm is at most $O(n^{\omega} \cdot h^{0.468})$. This algorithm is faster than our method for arbitrary dags for all values of $h \leq n^{0.42}$.

[1] In [4], an equivalent notion of the *height* of a dag is used. We find however the term *depth* more intuitive.

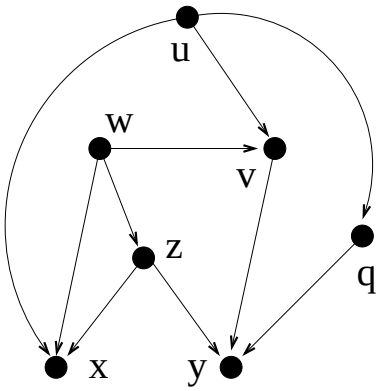

Figure 1: A dag with 7 vertices. The LCA of vertices x and y are both vertex u and vertex z; vertex w is a common ancestor of x and y but it is not the LCA of x and y. There is no common ancestor of vertices w and q.

Organization. Our paper is structured as follows. In Section 2, we present our natural $O(nm)$-time method for the all-pairs LCA problem in dags. In Section 3, we introduce several concepts and facts used by our matrix based methods for the all-pairs LCA problem in dags. In Section 4, we demonstrate the relationships between the problems of computing common ancestors and LCA in dags and those of finding witnesses and maximum witnesses for Boolean matrix product. In Section 5, we present our $O(n^{2.575})$-time method for the maximum witness problem and the all-pairs LCA problem for dags. In Section 6, we derive a more efficient solution to the all-pairs LCA problem in dags of small depth.

2 Optimal method for sparse dags

First, we shall describe preprocessing for answering queries about existence of a common ancestor for arbitrary pair of vertices in constant time.

For the input dag, we denote by n and m its number of vertices and edges, respectively. Also for a vertex v in the dag, $\mathsf{indeg}(v)$ and $\mathsf{outdeg}(v)$ stand respectively for the in-degree and out-degree of v. If $\mathsf{outdeg}(v) = 0$ then v is called a *terminal vertex* and if $\mathsf{indeg}(v) = 0$ then v is called a *source vertex*.

We may assume, without loss of generality, that the input dag is connected since otherwise we can decompose it into connected components and solve the problem for each component separately. For technical reasons, we also assume that every vertex is its own ancestor.

The following lemma immediately follows from the definition of a dag.

Lemma 1 *If two vertices have a common ancestor then there is a source vertex that is their common ancestor.*

In the first stage of the preprocessing, for each vertex of the input dag we form a table containing its descendants. In other words, we create the transitive closure of the dag which obviously can be done in time $O(nm)$.

Lemma 2 *The tables of descendants for all vertices can be formed in time* $O(nm)$.

In the second stage of the preprocessing, we determine for each vertex v the set of vertices that have a common ancestor with v. We proceed starting from source vertices. For the source vertices s, the sets are already computed: they are just the sets of descendants of s. Next, we iterate the following step: remove the vertices of in-degree 0 with incident edges and fill the tables for the new vertices v of in-degree 0 by merging the information from the tables associated with the removed direct ancestors of v. For each vertex v such an operation takes time $O(n) \times \mathsf{indeg}(v)$. Thus, for the entire graph it takes $O(nm)$ time.

For any vertex v, define the *depth* of v in a dag to be the length of the longest path from a source vertex to v in the dag. Note that the set of vertices having a common ancestor with a vertex v is the union of the sets of vertices having common ancestors with the ancestors of v (recall that v is also an ancestor of itself). Hence, we obtain the following lemma by induction on the depth of v.

Lemma 3 *For all vertices v, the tables of vertices having a common ancestor with v can be computed in time* $O(nm)$.

In order to answer LCA queries we need to refine the preprocessing slightly. During the second descending phase of the preprocessing we additionally enumerate the vertices in their visiting order. Since an ancestor is always visited before its descendant, we obtain the following lemma.

Lemma 4 *A vertex of a higher number cannot be an ancestor of a vertex of a lower number.*

For all vertices v, in the table keeping vertices w having a common ancestor with v, we keep also the maximum of the numbers assigned to the common ancestors of v and w. To achieve this, when we merge the information from the tables of direct ancestors of v, we pick the maximum number of a common ancestor of a direct ancestor of v and w. Clearly, the refinement can be accomplished within the same asymptotic time $O(mn)$. By induction, we obtain the following lemma.

Lemma 5 *For all vertices v, the tables of vertices w having a common ancestor with v with a pointer to a lowest common ancestor of v and w can be computed in time* $O(nm)$.

Hence, we obtain immediately the following theorem.

Theorem 6 *The all-pairs LCA problem for a dag on n vertices and m edges can be solved in time* $O(n(n+m))$.

Note that if $m = O(n)$ then our solution is optimal.

Shortest ancestral distance problem. One of the applications of the all-pairs LCA problem in dags discussed in [4] is the *shortest ancestral distance* problem in dag[2]. For two vertices u and v, Bender *et al.* define it as the minimum length of a path between u and v passing through a common ancestor of u and v (observe that the common ancestor is not necessarily the lowest one [3]). They showed that the all-pairs shortest ancestral distances can be computed in time $O(n^{2.575})$ [4].

We can modify our first method to obtain an $O(mn)$-time algorithm to compute the all-pairs shortest ancestral distances as follows. In the ascending phase, for each vertex v, and for each descendant u of v, we additionally compute the shortest directed distance between u and v. This can be easily accomplished within the same asymptotic time $O(mn)$. At the beginning of the descending phase, the previously computed shortest directed distances yield the shortest ancestral distances between sources and their descendants. While descending the shortest ancestral distances between the parents of the current vertex v and each other vertex u are increased by one. Next, the minimum of them and the shortest directed distance between v and u (it can be infinite) is taken as the shortest ancestral distance between v and u. In this way for all pairs of vertices v and u the shortest ancestral distance is computed.

Similarly, the so modified descending phase can be also implemented in time $O(mn)$. We conclude with the following theorem.

Theorem 7 *For a dag on n vertices and m edges, the all-pairs shortest ancestral distances can be computed in time $O(n(n+m))$.*

3 Preliminaries for the matrix based methods

In the remaining sections, we will present an efficient algorithm for the all-pairs LCA problems in dags that uses the matrix multiplication approach. We begin our study with introduction of several facts and concepts related to matrix multiplication, in particular Boolean matrix product.

We let ω to denote the exponent of square matrix multiplication, that is, the smallest constant for which the product of two $n \times n$ matrices can be computed in $O(n^{\omega})$ time. The best asymptotic upper bound on ω currently known is $\omega < 2.376$, by Coppersmith and Winograd [7].

The following fact that relates graph problems to fast matrix multiplication is folklore.

Fact 1 *The transitive closure of any directed graph with n vertices, in particular a dag, can be computed in time $\Theta(n^{\omega})$.*

We also need to define the concept of witnesses and maximum witness in Boolean matrix multiplication.

Definition 8 *If an entry $C[i,j]$ of the Boolean product of two Boolean matrices A and B is equal to 1 then any index k such that $A[i,k]$ and $B[k,j]$ are equal to 1 is a **witness** for $C[i,j]$. If k is the maximum possible witness for $C[i,j]$ then it is called the **maximum witness** for $C[i,j]$.*

[2] The identification of genes associated with genetic diseases was the original motivation for considering the shortest ancestral distance in sparse dags [9, 16]. It is now less relevant, see, e.g., [2].

[3] One can also consider the so called shortest ancestral LCA distance where the common ancestor is required to be lowest [4]

The following fact is due to Alon and Naor [1].

Fact 2 *The witnesses of the Boolean product of two $n \times n$ Boolean matrices can be computed in time*[4] $\widetilde{O}(n^\omega)$.

As it has been observed in [4], the concept of the level of a vertex in a dag plays an important role in the study of LCA in dags.

Definition 9 *The **level** i of a dag is the set of all its vertices of depth i. The **depth of a dag** is the maximum depth of its vertices, or equivalently, the number of its non-empty levels decreased by one.*

By using a variant of the Bellman-Ford algorithm for the single-source shortest path problem in dags running in time $O(n + m)$ (see, e.g., Chapter 24.2 in [8]), we obtain the following lemma.

Lemma 10 *For a dag G with n vertices and m edges, one can compute the partition of the vertices of G into the levels and the depth of G in time $O(n + m)$.*

4 Common ancestors versus Boolean matrix product witnesses

In this section we discuss the relationship between the problems of finding common ancestors and LCA for all pairs of vertices in a dag and the problems of computing witnesses and maximum witnesses of Boolean product of two Boolean $n \times n$ matrices.

We begin with a reduction of the problem of computing all-pairs common ancestors in a dag to the problem of computing witnesses of Boolean product of the corresponding Boolean matrices.

Theorem 11 *The problem of computing all-pairs common ancestors in a dag on n vertices can be reduced to the problem of computing witnesses of Boolean product of two Boolean $n \times n$ matrices in time $O(n^\omega)$. Similarly, the problem of computing all-pairs LCA in a dag on n vertices can be reduced to the problem of computing maximum witnesses of Boolean product of two Boolean $n \times n$ matrices in time $O(n^\omega)$.*

Proof. To prove the first part proceed as follows. Compute the transitive closure of the input dag. Number the vertices of the dag according to their topological ordering. Form the ancestor matrix A such that in its ith row there is 1 on the kth position iff the kth vertex is an ancestor of the ith vertex. Compute the witnesses of the Boolean product of the matrix A and its transpose A^T. For any (i, j) entry of the product matrix, if the entry is positive then its witness is the number of a common ancestor of the ith and jth vertex.

To prove the second part compute the maximum witnesses of the Boolean product of the matrix A and its transpose A^T instead of the ordinary ones. Next, observe that a common ancestor of vertices u and v having the largest number among common ancestors of u and v is an LCA of u

[4]Throughout the paper, the notation $\widetilde{O}(f(n))$ stands for $O(f(n)\log^c n)$ for some positive constant c.

and v. By the observation, for any (i, j) entry of the product matrix, if the entry is positive then its maximum witness is the number of an LCA of the ith and jth vertex. □

In turn, we consider a reduction of the problem of computing witnesses to the problem of computing all-pairs common ancestors in a dag of depth 1.

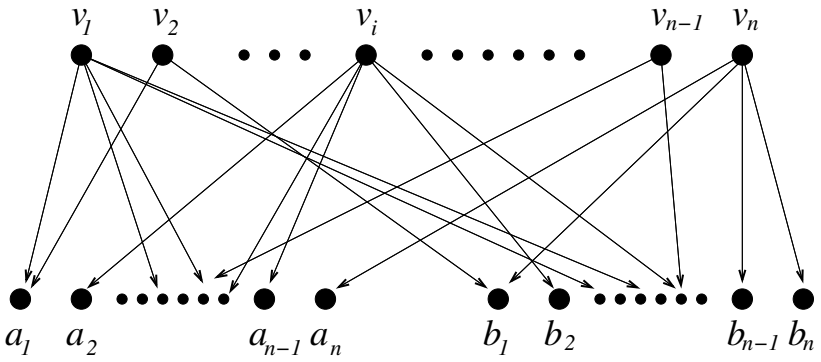

Figure 2: A scheme of the construction used in the proof of Theorem 12.

Theorem 12 *The problem of computing witnesses of Boolean product of two Boolean $n \times n$ matrices can be reduced to the problem of computing all-pairs common ancestors in a dag with $3n$ vertices and depth 1 in time $O(n^2)$.*

Proof. Let A and B be two Boolean $n \times n$ matrices. Construct a two-level dag G with vertices $v_1, v_2, \ldots, v_n$ on the zero level and vertices $a_1, a_2, \ldots, a_n$ and $b_1, b_2, \ldots, b_n$ on the first level. Create an edge (v_k, a_i) iff $A[i, k] = 1$. Analogously, form an edge (v_k, b_j) iff $B[k, j] = 1$.

Let C be the Boolean product of matrices A and B. By the definition of G, k is a witness of an entry $C[i, j]$ iff v_k is a common ancestor of the vertices a_i and b_j. □

Note that the all-pairs common ancestor problem (not the *lowest* common ancestor one) for an arbitrary dag can be solved in time $O(n^\omega)$ by computing the transitive closure of the dag, using Fact 1, and then computing the witnesses of the Boolean product of the transitive closure matrix and its transpose using Fact 2. This running time is optimal since Theorem 12 implies that the problem of computing all-pairs LCA in a dag with n vertices requires time $\Omega(n^\omega)$, even for dags of depth 1.

By slightly modifying the dag constructed in the proof of Theorem 12, we obtain the following theorem that relates computing *maximum witnesses* to computing all-pairs LCA in a dag.

Theorem 13 *The problem of computing maximum witnesses of Boolean product of two Boolean $n \times n$ matrices can be reduced to the problem of computing all-pairs LCA in a dag on $3n$ vertices in time $O(n^2)$.*

Proof. Assume the notation from the proof of Theorem 12. Additionally link the vertices $v_1, v_2, \ldots, v_n$ by the edges $(v_1, v_2), (v_2, v_3), \ldots, (v_{n-1}, v_n)$ in the dag G. (For an example, see Figure 3.)

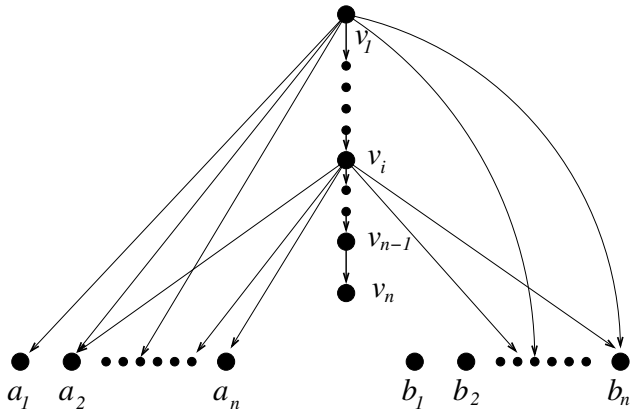

Figure 3: A scheme of the construction used in the proof of Theorem 13.

By the construction of the extended dag, k is a witness of an entry $C[i,j]$ iff v_k is a maximum common ancestor of the vertices a_i and b_j. □

By combining Theorem 11 with Theorems 12, 13, we obtain the following corollary.

Corollary 14 *The problem of computing all-pairs common ancestors in a dag on n vertices and the problem of computing witnesses of Boolean product of two Boolean $n \times n$ matrices are $O(n^\omega)$-time equivalent. Similarly, the problem of computing all-pairs LCA in a dag on n vertices and the problem of computing maximum witnesses of Boolean product of two Boolean $n \times n$ matrices are $O(n^\omega)$-time equivalent.*

Bender *et al.* [4] showed that the problem of computing all-pairs LCA in a dag with n vertices is not easier than that of computing transitive closure in a directed graph with n vertices. This implies the lower bound of $\Omega(n^\omega)$ for the all-pairs LCA in a dag with n vertices. However, our result on the relationship between the all-pairs LCA problem in dags and that of computing maximum witnesses of Boolean product of two Boolean matrices shows that these two problems have asymptotically identical complexity, and thus it yields a stronger relationship. Furthermore, our result (Theorem 12) shows that the problem of computing all-pairs LCA in a dag with n vertices requires time $\Omega(n^\omega)$ even for dags of depth 1.

5 An $O(n^{2.575})$-time method for maximum witnesses and the all-pairs LCA problem in dags

By Corollary 14, it is sufficient to compute maximum witnesses of Boolean product of two Boolean $n \times n$ matrices in order to solve the all-pairs LCA problem in dags. In this section, we present a substantially subcubic algorithm for the maximum witness problem using fast rectangular matrix multiplication.

Let ℓ be a positive integer smaller than n. Partition the matrix A into $n \times \ell$ sub-matrices A_p and the matrix B into $\ell \times n$ sub-matrices B_p, such that $1 \leq p \leq n/\ell$, and the sub-matrix A_p covers

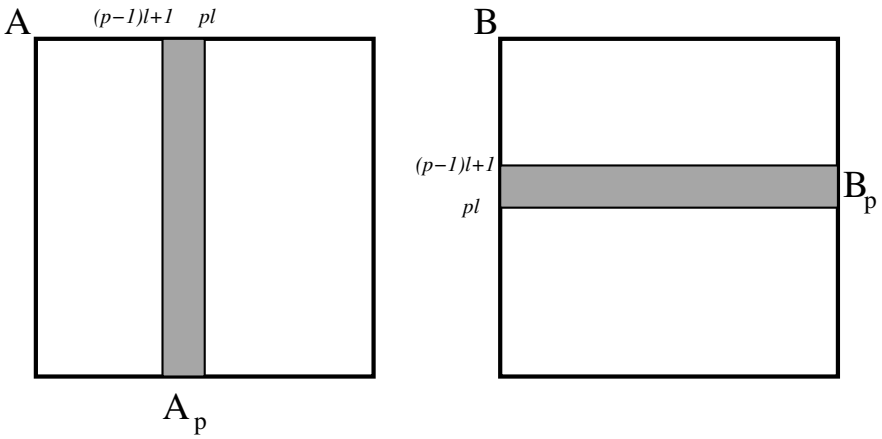

Figure 4: Rectangular matrices A_p and B_p.

the columns $(p-1)\,\ell+1$ through $p\,\ell$ of A whereas the sub-matrix B_p covers the rows $(p-1)\,\ell+1$ through $p\,\ell$ of B. (For an example, see Figure 4.)

For $p = 1, \ldots, n/\ell$, let C_p be the Boolean product of A_p and B_p. Then, $C_p[i,j] > 0$ iff there is an index k, $(p-1)\,\ell < k \leq p\,\ell$, such that $A[i,k] = B[k,j] = 1$. Therefore the following claim follows.

Claim 15 *Suppose that a* $C[i,j]$ *entry of the Boolean product* C *of* A *and* B *is positive. Let* p' *be the maximum value of* p *such that* $C_{p'}[i,j] > 0$*. The maximum witness of* $C[i,j]$ *belongs to the interval* $[(p'-1)\,\ell+1, p'\,\ell]$.

By this claim, after computing all the products $C_p = A_p \cdot B_p$, $1 \leq p \leq n/\ell$, we need $O(n/\ell+\ell)$ time per positive entry of C to find the maximum witness: $O(n/\ell)$ time to determine p' and then $O(\ell)$ time to locate the maximum witness.

Let $\omega(1,r,1)$ denote the exponent of the multiplication of an $n \times n^r$ matrix by an $n^r \times n$ matrix. It follows that the total time taken by our method for maximum witnesses is

$$O((n/\ell) \cdot n^{\omega(1,\log_n \ell,1)} + n^3/\ell + n^2\,\ell) \ .$$

By setting r to $\log_n \ell$ our upper bound transforms to $O(n^{1-r+\omega(1,r,1)} + n^{3-r} + n^{2+r})$. Note that by assuming $r \geq \frac{1}{2}$, we can get rid of the additive n^{3-r} term. Hence, by solving the equation $1-\lambda+\omega(1,\lambda,1) = 2+\lambda$ implying $\lambda \geq \frac{1}{2}$ by $\omega(1,\lambda,1) \geq 2$, we obtain our main result.

Theorem 16 *Let* λ *be such that* $\omega(1,\lambda,1) = 1+2\lambda$*. The maximum witnesses for all positive entries of the Boolean product of two* $n \times n$ *Boolean matrices can be computed in time* $O(n^{2+\lambda})$.

Coppersmith [6] and Huang and Pan [13] proved the following fact.

Fact 3 [6, 13] *Let* $\omega = \omega(1,1,1) < 2.376$ *and let* $\alpha = \sup\{0 \leq r \leq 1 : \omega(1,r,1) = 2+o(1)\} > 0.294$*. Then* $\omega(1,r,1) \leq \beta(r)$*, where* $\beta(r) = 2+o(1)$ *for* $r \in [0,\alpha]$ *and* $\beta(r) = 2+\frac{\omega-2}{1-\alpha}(r-\alpha)+o(1)$ *for* $r \in [\alpha,1]$.

Note that by Fact 3, the solution λ of the equation $\omega(1,\lambda,1) = 1 + 2\lambda$ is satisfied by $\lambda = \frac{1-\alpha(\omega-1)}{2(1-\alpha)-(\omega-2)} + o(1) < 0.575$. Therefore, we obtain the following concrete corollary.

Corollary 17 *The maximum witnesses for all positive entries of the Boolean product of two* $n \times n$ *Boolean matrices can be computed in time* $O(n^{2.575})$.

By combining Theorem 11 with Theorem 16 and Corollary 17, we obtain also an $O(n^{2.575})$-time solution to the all-pairs LCA problem in dags.

Theorem 18 *Let* $\omega(1,\lambda,1) = 1 + 2\lambda$. *The all-pairs LCA problem for an arbitrary dag with* n *vertices can be solved in time* $O(n^{2+\lambda})$, *which is bounded above by* $O(n^{2.575})$.

6 All-pairs LCA in dags of bounded depth

In this section we describe an algorithm for solving the all-pairs LCA problem for an arbitrary dag G of depth bounded by h. The algorithm has a similar flavor as that discussed in the previous section, but now we will additionally use the fact that the depth is bounded to speed up the process.

First, we compute the transitive closure of G and create the ancestor matrix A similarly as in the proof of Theorem 11. Next, using Lemma 10, we partition G into $h+1$ levels and extend the partial order induced by this partition to a linear order and number the vertices according to the linear order so the numbering is increasing with respect to vertex depth. These steps can be performed in time $O(n^{\omega})$.

Observe that the numbering naturally decomposes into $h+1$ continuous intervals in one-to-one correspondence with the levels of G. Our approach relies on the following generalization of the observation that the common ancestor of vertices u and v that has the highest number (in the topological ordering) is a lowest common ancestor of u and v.

Claim 19 [4] *Any common ancestor of vertices* u *and* v *that is of highest level among common ancestors of* u *and* v *is a lowest common ancestor of* u *and* v.

Claim 19 implies directly that the maximum witnesses of the product of the Boolean matrix of ancestors A and its transpose A^T yield the solution to the all-pairs LCA problem for a dag. However, since it is expensive to compute maximum witnesses, we modify this construction and reduce the problem to that of computing *witnesses* (instead of *maximum* witnesses).

Relying on Claim 19 and the bounded depth h, we decompose A and its transpose A^T into $h+1$ rectangular sub-matrices in one-to-one correspondence with the level intervals and compute witnesses for the products of the pairs of sub-matrices in A and A^T corresponding to the same level interval. Now, similarly as in the previous section, we observe that if vertices i and j have a common ancestor at level ℓ, then if we multiply the sub-matrix of A corresponding to level ℓ by the sub-matrix of A^T corresponding to level ℓ, then the resulting matrix C_ℓ will have $C_\ell[i,j] > 0$ and the witness for $C_\ell[i,j] > 0$ will be a vertex from level ℓ that is an ancestor of both i and j.

Let us estimate the time taken by finding the witnesses for the $h+1$ products of sub-matrices of A and A^T. Recall the definition of function β in Fact 3 due to Coppersmith, and Huang and

Pan. By using the straightforward reduction of Boolean matrix multiplication to the arithmetic one and by generalizing the derandomization method of Alon and Naor [1] for witnesses of Boolean matrix multiplication to include rectangular matrices, we obtain the following lemma.

Lemma 20 *The witnesses of the Boolean product of two Boolean matrices of sizes* $n \times n^r$ *and* $n^r \times n$ *can be computed in time* $\widetilde{O}(n^{\beta(r)})$.

For $k = 0, \ldots, h$, let ℓ_k be the number of vertices on level k in the dag G. By Lemma 20, the total time taken by computing the witnesses for $h+1$ sub-matrix products is

$$\widetilde{O}\left(\sum_{k=0}^{h} n^{\beta(\log_n \ell_k)}\right) \quad . \tag{1}$$

Finally, once we determined witnesses for every pair of vertices numbered i and j and for every level ℓ, it remains to find for each pair i and j, the maximum witness among the witnesses for the pair i, j in the products of the sub-matrices. It takes $O(h)$ time per pair i, j and hence, $O(n^2 h)$ time in total.

Next, by straightforward calculations, Jensen inequality implies that the value of (1) is maximized if the levels are of equal size $n/(h+1)$. (Indeed, we observe that the function $n^{\beta(x)-2}$ is concave and therefore $\sum_{k=0}^{h} n^{\beta(\log_n \ell_k)} = n^2 \sum_{k=0}^{h} n^{\beta(\log_n \ell_k)-2} \leq n^2 (h+1)\, n^{\beta(\log_n n/(h+1))-2} = (h+1)\, n^{\beta(\log_n n/(h+1))}$.) Hence, we obtain the following theorem.

Theorem 21 *The all-pairs LCA problem for an arbitrary dag with* n *vertices and depth* n^q *can be solved in time* $\widetilde{O}(n^\omega + n^{q+\beta(1-q)})$.

In particular, for dags of depth at most $n^\alpha \approx n^{0.294}$ *the running time is* $\widetilde{O}(n^\omega)$.

For larger values of the depth n^q*, the running time of this algorithm is*

$$\widetilde{O}\left(n^{\omega+q(1-\frac{\omega-2}{1-\alpha})+o(1)}\right) \quad \approx \quad O\left(n^{2.376+0.468\,q}\right) \quad .$$

Remark 22 *Note that for all values of* $q \leq \alpha \approx 0.294$*, the running time of our algorithm from Theorem 21 is asymptotically optimal due to our result in Theorem 12. Even for larger values of* q*, up to* $q \leq 0.42$*, our algorithm from Theorem 21 is faster than that general from Theorem 18.*

7 Final remarks

The problems of finding LCA are classical and central in the area of algorithms and data structures [4, 15, 17]. In spite of the long history of studies devoted to LCA problems, we have succeeded to design two quite natural methods for finding LCA in dags considerably subsuming the previously known best results [4].

Our second method relies on the presented close relationship between the problem of computing the maximum witnesses of Boolean matrix product and that of finding LCAs for all pairs of vertices in a dag and the use of fast algorithms for rectangular matrix multiplication.

It is an intriguing open problem whether the complexity gap between the problems of computing witnesses and computing maximum witnesses of Boolean matrix product, or alternatively, the problems of finding common ancestors and finding LCAs in a dag, can be further decreased.

The problem of finding maximum witnesses of Boolean matrix product seems to be of interest in its own rights. At first glance it seems that the recursive $O(n^{\omega+\epsilon})$-time method of Galil and Margalit [11] for witnesses of Boolean matrix product could be adapted to produce the maximum witnesses by considering the fragments containing maximum witnesses in the subproblems without substantially altering its asymptotic time. However, the aforementioned method may permute rows or columns in recursive steps which may disturb the search for maximum witnesses. Thus, the problem of whether or not our $O(n^{2.575})$-time method is optimal is open.

It is also an interesting question whether or not the instances of the problem of finding maximum witnesses of Boolean matrix product occurring in our reduction from the LCA problem in dags are computationally easier than the general ones.

8 Acknowledgments

The authors are grateful to Pavel Sumazin for inspiration, to Leszek Gąsieniec and Martin Wahlén for some discussions, and to the unknown referees for their comments.

References

[1] N. Alon and M. Naor. Derandomization, witnesses for Boolean matrix multiplication and construction of perfect hash functions. *Algorithmica*, 16: 434–449, 1996.

[2] A. Becker, D. Geiger, and A. A.Schäffer. Automatic selection of loop breakers for genetic linkage analysis. *Human Heredity*, 48:49–60, 1998.

[3] M. A. Bender and M. Farach-Colton. The LCA problem revisited. *Proc. 4th Latin American Symposium on Theoretical Informatics (LATIN'00)*, pp. 88–93, 2000.

[4] M. A. Bender, G. Pemmasani, S. Skiena and P. Sumazin. Finding least common ancestors in directed acyclic graphs. *Proc. 12th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA'01)*, pp. 845–853, 2001.

[5] R. Cole and R. Hariharan. Dynamic LCA queries in trees. *SIAM Journal on Computing*, 34(4): 894–923, 2005.

[6] D. Coppersmith. Rectangular matrix multiplication revisited. *Journal of Symbolic Computation*, 13: 42–49, 1997.

[7] D. Coppersmith and S. Winograd. Matrix multiplication via arithmetic progression. *Journal of Symbolic Computation*, 9: 251–290, 1990.

[8] T. H. Cormen, C. E. Leiserson, R. L. Rivest, and C. Stein. *Introduction to Algorithms.* 2nd edition, McGraw-Hill Book Company, Boston, MA, 2001.

[9] R. W. Cottingham Jr., R. M. Idury, and A. A. Shäffer. Genetic linkage computations. *American Journal of Human Genetics*, 53: 252–263, 1993.

[10] A. Czumaj and A. Lingas. Improved algorithms for the all-pairs lowest common ancestor problem in directed acyclic graphs. *Manuscript*, 2005.

[11] Z. Galil and O. Margalit. Witnesses for Boolean matrix multiplication and for transitive closure. *Journal of Complexity*, 9: 201–221, 1993.

[12] D. Harel and R. E. Tarjan. Fast algorithms for finding nearest common ancestors. *SIAM Journal on Computing*, 13(2): 338–355, 1984.

[13] X. Huang and V.Y. Pan. Fast rectangular matrix multiplications and applications. *Journal of Complexity*, 14: 257–299, 1998.

[14] M. Kowaluk and A. Lingas. LCA queries in directed acyclic graphs. *Proc. 32nd International Colloquium on Automata, Languages and Programming (ICALP'05)*, pp. 241–248, 2005.

[15] M. Nykänen and E. Ukkonen. Finding lowest common ancestors in arbitrarily directed trees. *Information Processing Letters*, 50(6): 307–310, 1994.

[16] A. A. Shäffer, S. K. Gupta, K. Shriram, and R. W. Cottingham Jr. Avoiding recomputation in linkage analysis. *Human Heredity*, 44: 225–237, 1994.

[17] R.E. Tarjan. Applications of path compression on balanced trees. *Journal of the ACM*, 26(4): 690–715, 1979.

[18] U. Zwick. All pairs shortest paths using bridging sets and rectangular matrix multiplication. *Journal of the ACM*, 49(3): 289–317, 2002.

ECCC ISSN 1433-8092

http://eccc.hpi-web.de/