

Witnesses for Boolean Matrix Multiplication and for Shortest Paths

Noga Alon ^{*} Zvi Galil [†] Oded Margalit [‡] Moni Naor [§]

Extended Abstract

Summary of results

The subcubic ($O(n^\omega)$ for $\omega < 3$) algorithms to multiply Boolean matrices do not provide the witnesses; namely, they compute $C = AB$ but if $C_{ij} = 1$ they do not find an index k (a witness) such that $A_{ik} = B_{kj} = 1$. We design a deterministic algorithm for computing the matrix of witnesses that runs in $\tilde{O}(n^\omega)$ time, where here $\tilde{O}(n^\omega)$ denotes $O(n^\omega(\log n)^{O(1)})$.

The subcubic methods to compute the shortest distances between all pairs of vertices also do not provide for witnesses; namely they compute the shortest distances but do not generate information for computing quickly the paths themselves. A witness for a shortest path from v_i to v_j is an index k such that v_k is the first vertex on such a path. We describe subcubic methods to compute such witnesses for several versions of the all pairs shortest paths problem. As a result, we derive shortest paths algorithms that provide characterization of the shortest paths in addition to the shortest distances in essentially the same time needed for computing the distances; namely $\tilde{O}(n^{(3+\omega)/2})$ in the directed case and $\tilde{O}(n^\omega)$ time in the undirected case.

We also design an algorithm that computes witnesses for the transitive closure in the same time needed to compute witnesses for Boolean matrix multiplication.

^{*}Department of Mathematics, Raymond and Beverly Sackler Faculty of Exact Sciences, Tel Aviv University, Tel Aviv, ISRAEL. Research supported in part by a USA Israeli BSF grant

[†]Department of Computer Science, Raymond and Beverly Sackler Faculty of Exact Sciences, Tel Aviv University, Tel Aviv, Israel and Department of Computer Science, Columbia University, New York, NY 10027, USA

[‡]Department of Computer Science, Raymond and Beverly Sackler Faculty of Exact Sciences, Tel Aviv University, Tel Aviv, Israel

[§]IBM Almaden Research Center, 650 Harry Road, San Jose, CA 95120, USA

1 Introduction

Consider a Boolean matrix multiplication: $C = AB$, $C_{ij} = \bigvee_{k=1}^n (A_{ik} \wedge B_{kj})$. The n^3 time method that evaluates these expressions gives for every i, j for which $C_{ij} = 1$ all the k 's for which $A_{ik} = B_{kj} = 1$. The subcubic methods on the other hand consider A and B as matrices of integers and do not provide any of these k 's. We call a k such that $A_{ik} = B_{kj} = 1$ a *witness* (for the fact that $C_{ij} = 1$). We want to compute in addition to the matrix C a matrix of witnesses. When there is more than one witness for a given i and j we are satisfied with one such witness.

We use $O(n^\omega)$ to denote the running time of some subcubic algorithm for Boolean matrix multiplication. All our algorithms can be derived from any such algorithm yielding a corresponding time bound as a function of w . The best asymptotic bound known at present is the one with the exponent $\omega < 2.376$ and is due to Coppersmith and Winograd [3].

For two functions $f(n)$ and $g(n)$ we let $g(n) = \tilde{O}(f(n))$ denote the statement that $g(n)$ is $O(f(n)(\log n)^{O(1)})$.

Several researchers, including Seidel [6], Karger (personal communication) and the first three authors, observed that there is a simple randomized algorithm that computes witnesses in $\tilde{O}(n^\omega)$ time. In Section 2 we describe a **deterministic** algorithm for computing the witnesses in $\tilde{O}(n^\omega)$ time. It is essentially a derandomization of a modified version of the simple randomized algorithm, and relies heavily on the known constructions of small sample spaces with almost independent random variables. We also outline an alternative approach that gives slightly worse running time but may still be useful for matrices of moderate size.

Our motivation for studying the computation of witnesses for Boolean matrix multiplication is related to our work on the all pair shortest paths problem. We use the following notation. $D = \{d_{ij}\}_{i,j=1}^n$ is the matrix of edge lengths, $d_{ij} = +\infty$ in case there is no edge from v_i to v_j . In the positive case $d_{ij} \in \{1, 2, \dots, M, +\infty\}$ and in the unrestricted case $d_{ij} \in \{0, \pm 1, \pm 2, \dots, \pm M, +\infty\}$. $D^* = \{d_{ij}^*\}$ is the matrix of shortest distances.

In an earlier paper [2] the first three authors designed subcubic algorithms for computing all pair shortest distances of directed graphs with integer edge lengths whose absolute value is bounded by M . We denote the problem and its time bound by $APSD(n, M)$, where n is the number of vertices in the graph. We showed that $APSD(n, M) = O((Mn)^\nu)$, where $\nu = (3 + \omega)/2$. For $\omega < 2.376$, we have $\nu < 2.688$. In a more recent work [4] the second and third authors have improved the dependence on M and obtained better bounds for undirected graphs, in which case $APSD(n, M) = O(M^{(\omega+1)/2} n^\omega \log n)$. A simple $O(n^\omega \log n)$ algorithm for undirected $APSD(n, 1)$ was discovered independently by Seidel [6], but it does not seem to be extendable to larger edge lengths. All these algorithms do not provide any subcubic deterministic way for finding the shortest paths themselves, only the shortest distances.

One cannot have a subcubic algorithm that computes the shortest paths between all pairs of vertices simply because in the example depicted in Figure 1 there are more than $n^3/27$ edges in all shortest paths.

One may use the following definition to obtain a more concise representation of all shortest paths: A *witness for a shortest path* from v_i to v_j is an index k such that v_k is the first vertex on such a path. This definition is certainly sufficient in case of positive edge lengths. A shortest path can be easily constructed from these witnesses.

This definition is insufficient in case of nonpositive cycles. If $d_{ij}^* = -\infty$ we want to be able to construct from the witnesses a simple path from v_i to v_j together with a vertex v_k on the path and a negative cycle containing v_k (i, j and k need not be distinct). This leads to the need to define witnesses for paths, not necessarily shortest paths.




Figure 1:

Consider the transitive closure of a directed graph. One could try to define a witness for a path identically to the definition of a witness for a shortest path: A witness for a path from v_i to v_j is an index k such that v_k is the first vertex on such a path. Any method for computing witnesses for Boolean matrix multiplication can be immediately used for computing these “witnesses”: compute witnesses for $A \cdot T$, where A is the incidence matrix and T the transitive closure. Unfortunately this definition is inappropriate as can be seen in Figure 2: v_k is a possible witness for the path from v_i to v_j , but it is a bad choice which leads to a cycle.




Figure 2:

We require a matrix of *witnesses for the transitive closure* to satisfy the following condition: If a path from v_i to v_j exists then such a path can be constructed by following the witnesses. Namely, there is a path $v_i = v_{i_0}, \dots, v_{i_k} = v_j$ and for $1 \leq r \leq k$, i_r is the witness for the path from $v_{i_{r-1}}$ to v_j . In Section 3 we give an $\tilde{O}(n^\omega)$ algorithm that computes witnesses for the transitive closure.

Coming back to the shortest paths problem we would like to compute in addition to the matrix D^* of shortest distances also

1. Witnesses for shortest paths.
2. A simple negative cycle for each i such that $d_{ii}^* = -\infty$.

Consequently, shortest paths of finite length can be easily obtained from 1. On the other hand, a shortest path of length $-\infty$ can be represented as a (possibly empty) path from 1 together with a negative cycle (from 2).

In Section 4 we first give an algorithm for computing witnesses for shortest paths when edge lengths are positive, then when edge lengths are nonnegative. Finally we give an algorithm that generates the characterization of shortest paths in the general case. Its running time is $\tilde{O}(n^{(3+\omega)/2})$.

Summarizing, we get the following bounds for $APSP(n, 1)$: $\tilde{O}(n^{(3+\omega)/2})$ in the directed case and $\tilde{O}(n^\omega)$ in the undirected case. Recall that the time bounds for $APSD(n, 1)$ are $O(n^{(3+\omega)/2})$ in the directed case and $O(n^\omega \log n)$ in the undirected case. Indeed, some of the Boolean matrix multiplications are now augmented to compute also witnesses, which explains our motivation to study the latter. (We believe that witnesses for Boolean matrix multiplication will be found useful

elsewhere as well.) The fact that the bounds for *APSP* are obtained from the bounds for *APSD* by adding polylogarithmic factors is not immediate. This would be a simple consequence of our algorithm for matrix multiplication with witnesses if the algorithms just added witnesses to each Boolean matrix multiplication. However, this reason is not the only one needed to explain this coincidence. More details are given in Section 4.

2 Boolean matrix multiplication with witnesses

All the matrices in this section are n by n matrices, unless otherwise specified. If M is such a matrix, we let M_{ij} denote the entry in its i th row and j th column. Let A and B be two matrices with $\{0, 1\}$ entries, and let C be their product over the integers. Our objective is to find witnesses for all the positive entries of C , i.e., for each entry $C_{ij} > 0$ of C we wish to find a k such that $A_{ik} = B_{kj} = 1$. This is clearly equivalent to the problem of finding witnesses in the Boolean case. As observed by several researchers (including Seidel, Karger and the first three authors) there is a simple randomized algorithm that solves this problem in expected running time $\tilde{O}(n^\omega)$. Here we consider *deterministic* algorithms for the problem. Our best algorithm, described in the next subsection, is, in a sense, a derandomized version of the simple randomized solution, and its running time is $\tilde{O}(n^\omega)$. The derandomization requires several modifications in the straightforward randomized algorithm together with an interesting application of the known constructions of [5] (or [1]) of almost c -wise independent random variables in small sample spaces.

2.1 The algorithm

The first simple observation is the fact that if E and F are two matrices with $\{0, 1\}$ entries and $G = EF$ then one multiplication of matrices with entries at most n suffices for finding witnesses for all the entries of G which are precisely 1. Indeed, simply replace every 1-entry in the k th row of F by k (for all $1 \leq k \leq n$) to get a matrix F' and compute $G' = EF'$. Now observe that if $G_{ij} = 1$ and $G'_{ij} = k$ then k is a witness for G_{ij} .

Define $c = \lceil \log \log n + 9 \rceil$ and $\alpha = \frac{8}{2^c}$. For two matrices E and F with $\{0, 1\}$ entries define $G = E \wedge F$ by $G_{ij} = E_{ij} \wedge F_{ij}$.

Here is an outline of the algorithm. Besides A, B and $C = AB$ it uses two additional matrices: R and D . The way to perform steps 3c and 3d will be described later.

- While not all witnesses are known
 1. Let L denote the set of all positive entries of C for which there are no known witnesses.
 2. Let R be the all 1 matrix.
 3. Perform the following $\lceil 1 + 3 \log_{4/3} n \rceil$ times:
 - (a) $D \leftarrow A \cdot (B \wedge R)$ (The matrix multiplication is over the integers)
 - (b) Let L' denote the set of all entries of D in L which are at most c .
 - (c) Find witnesses for all entries in L' .
 - (d) $R \leftarrow$ *good* matrix (see definition of good below).

A matrix R is *good* (in step 3d above) if the following two conditions hold:

- a) The total sum of the entries of $D = A \cdot (B \wedge R)$ in L is at most $3/4$ of what this sum was while using the previous matrix R . (Observe that this guarantees that after $1 + 3 \log_{4/3} n$ iterations all these entries of D will vanish.)

b) The fraction of entries of D in L that go from a value bigger than c to 0 is at most α .

Lemma 1 *If $R \leftarrow R \wedge S$ in step 3d where S is a random $0, 1$ matrix, then the new R is good with probability at least $1/6$.*

The lemma follows from the following three claims:

Claim 1 *The probability that the sum of entries of D in L goes down by at least a factor of $3/4$ is at least $1/3$.*

To see this, observe that the expected sum of entries of D in L goes down by $1/2$. Thus, the claim follows from Markov's Inequality. \square

Claim 2 *The probability that a fixed entry of D which is at least c drops down to 0 is at most $1/2^c$.*

This is obvious. Observe that the claim holds even if we only assume that every c entries of S are independent. \square

Claim 3 *The probability that more than a fraction α of the entries of D in L drop from at least c to 0 is at most $\frac{1}{2^c} \frac{1}{\alpha} = \frac{1}{8}$.*

This follows from Claim 2 by Markov's Inequality. Since $1/3 - 1/8 > 1/6$ the lemma follows. \square

Define $\epsilon = \frac{1}{2^{c+1}}$. The crucial point is to observe that the proof of the above lemma still holds, with almost no change, if the matrix S is not totally random but its entries are chosen from a c -wise ϵ -dependent distribution in the sense of [5], [1]. Recall that if m random variables whose range is $\{0, 1\}$ are c -wise ϵ -dependent then every subset of $i \leq c$ of them attains each of the possible 2^i configurations of 0 and 1 with probability that deviates from $1/2^i$ by at most ϵ .

Lemma 2 *If $R \leftarrow R \wedge S$ in step 3d where the entries of S are chosen as n^2 random variables that are c -wise ϵ -dependent, then the new R is good with probability at least $1/12 - 2\epsilon$.*

We note that in fact it is sufficient to choose only one column and copy it n times. The proof is by the following modified three claims, whose proof is analogous to that of the corresponding previous ones.

Claim 4 *The probability that the sum of entries of D in L goes down by at least a factor of $3/4$ is at least $1/3 - 2\epsilon$. \square*

Claim 5 *The probability that a fixed entry of D which is at least c drops down to 0 is at most $1/2^c + \epsilon$. \square*

Claim 6 *The probability that more than a fraction α of the entries of D in L drop from at least c to 0 is at most $(\frac{1}{2^c} + \epsilon) \frac{1}{\alpha} < \frac{2}{2^c} \frac{1}{\alpha} = 1/4$. \square*

The lemma follows from the above three claims. \square

As shown in [5] and in [1] there are explicit probability spaces with n^2 random variables which are c -wise ϵ -dependent, whose size is

$$(\log n \cdot c \cdot \frac{1}{\epsilon})^{2+o(1)},$$

which is less than, e.g., $O((\log n)^5)$. Moreover, these spaces can be easily constructed in time negligible with respect to the total running time of our algorithm. Now suppose that in step 3d all the matrices S defined by such a probability space are searched, until a good one is found. Checking whether a matrix is good requires only matrix multiplication plus $O(n^2)$ operations. Therefore the inner loop (starting at step 3) takes polylog n times matrix multiplication time. It is important to note that during the performance of step 3d, while considering all possible matrices S provided by our distribution, we can accomplish step 3c as well. This is true since c -wise ϵ -dependence guarantees that every entry in L' will drop to precisely 1 for some of the matrices S and hence, by the observation in the beginning of this subsection, if we replace each matrix multiplication in the search for a good S by two matrix multiplications as described in that observation, we complete steps 3c and 3d together.

In every iteration of the inner loop 3 at most α fraction of the entries of L are “thrown” (i.e. their witness will not be found in this iteration of the outer loop). Therefore at least $(1 - \alpha)^{1+3 \log_{4/3} n}$ fraction of the entries of D in L will *not* be thrown during the completion of these iterations. For those entries, which are at least $1/2$ of the entries in L , a witness is found. Therefore, only $O(\log n)$ iterations of the outer loop are required, implying the desired $\tilde{O}(n^\omega)$ total running time.

We have thus proved the following:

Theorem 1 *The witnesses for the Boolean multiplication of two n by n matrices can be found in deterministic $\tilde{O}(n^\omega)$ time.*

2.2 An alternative approach

The witnesses for Boolean matrix multiplication can be computed in a different manner. Although the running time obtained is slightly worse than that of our previous algorithm, it may give better performance for matrices of moderate size.

Here is a rough outline of the approach: We design a sequence of algorithms, the first algorithm ALG_0 , is the naive cubic way: test all the n possible witnesses for every positive entry C_{ij} . The next algorithm ALG_1 , is the following: Consider each of the two matrices A and B as an $L \times L$ block matrix where each block is of size $n/L \times n/L$. Multiply the two block matrices using the trivial L^3 time algorithm, and using fast matrix multiplication for any multiplication of two blocks. Now we know for each positive entry C_{ij} , a product of a block of A and a block of B which contains a witness. Use ALG_0 for finding witnesses inside that block. The running time is

$$O\left(L^3 \left(\frac{n}{L}\right)^\omega + L^2 \left(\frac{n}{L}\right)^3\right).$$

An appropriate choice of L gives an $O(n^{\frac{9-2\omega}{4-\omega}})$ time algorithm.

The sequence starting with these two algorithms can be extended, where each algorithm uses the previous one and the time complexity converges to $O(n^{\omega+O(\log^{-1/3}(n))})$. This requires several additional ideas including a generalization of the problem to that of finding witnesses for a prescribed subset of entries of the product of two rectangular matrices, given certain information on the location of these witnesses. The details are complicated and since the running time is inferior to that of our previous algorithm we do not include them. For any given problem, one can apply any of the algorithms from the sequence above. It seems that for certain possible sizes, one of the algorithms ALG_s for some small integer s may actually be faster than the algorithm in the previous subsection.

3 Computing witnesses for the transitive closure

In the introduction we explained why the immediate solution that computes witnesses for $A \cdot T$ does not work. Another simple solution is to add lengths to the edges and compute witnesses for shortest paths. However, the best time for computing only the distances in the directed case (even without the witnesses for the paths) is $O(n^{(\omega+3)/2})$.

The only reason that the immediate solution does not work are the cycles. So we first find the strongly connected components of G , then we contract them into new vertices. Now we can use the immediate algorithm to solve the new problem. Lastly, we “open” the contracted vertices and transform the solution to a solution for the original problem. More formally:

Algorithm

1. Compute the strongly connected components of the input graph $G = (V, E)$, where $V = \{v_1, \dots, v_n\}$. Denote by $V' = \{v'_1, v'_2, \dots, v'_m\}$ the set of strongly connected components of G , where $v'_i = \{v_{i1}, v_{i2}, \dots, v_{ir_i}\}$. We build the contracted graph $G' = (V', E')$, where $E' = \{(v'_i, v'_j) : \exists (v_{ix}, v_{jy}) \in E\}$. Each edge $(v'_i, v'_j) \in E'$ is arbitrarily associated with one edge $(v_{ix}, v_{jy}) \in E$. This can be done in $O(n^2)$ time.
2. Solve the transitive closure problem of the graph G' , denote the solution by T' . Compute witnesses for the Boolean matrix multiplication $T' \cdot A'$ by W' . This step can be done in $\tilde{O}(n^\omega)$ time.
3. For each strongly connected component find witnesses for the transitive closure (which is a clique). Denote the witnesses matrix for that problem by \hat{W} ; this matrix is defined only for pairs which are in the same strongly connected component. This can be done in $O(n^2)$ time, as described in Algorithm 3.1.
4. Expand the solution of the contracted problem into a solution for the whole problem. This can be done in $O(n^2)$ time as described in Algorithm 3.2.

Theorem 2 *The algorithm above computes the matrix of witnesses in time $\tilde{O}(n^\omega)$.*

3.1 Computing witnesses for a strongly connected graph

The algorithm has two stages. In the first, we perform breadth first search (BFS) from one of the vertices v_0 . In the process we generate a BFS tree T . For each edge $(u, v) \in T$ and every descendant w of v we set $W(u, w) \leftarrow v$. In the second stage, we use the reverse edges and perform another BFS from v_0 . We process a vertex when it is first visited. Assume we enter first u using edge (u, v) . We then consider all $w \in V$ and if $W(u, w)$ is undefined we set it to v .

Obviously, each stage takes $O(n^2)$ time. Correctness follows by induction. The induction hypothesis states that for every processed vertex u , and every w , starting with u and following W we obtain a simple path from u to w . The base is true because the first stage essentially processes v_0 . For the induction step, assume we process u . If $W(u, w) = z$ is defined, it was defined in the first stage and (u, z) is in the BFS tree of the first stage and following W we follow a path on the tree from u to w . If it is undefined, we set $W(u, w) \leftarrow v$, where v was processed before. The claim now follows from the induction hypothesis.

3.2 Joining solutions

Examine the solution for G' . Suppose that $W'(i, j) = k$; by the definition of a witness, there exists an edge (v'_i, v'_k) . Let (v_{ix}, v_{ky}) be the edge of G associated with it.

$$W(v_{ik_1}, v_{jk_2}) = \begin{cases} v_{ky} & k_1 = x \\ \hat{W}(v_{ik_1}, v_{ix}) & \text{otherwise.} \end{cases}$$

The time complexity of this algorithm is $O(n^2)$.

4 Finding paths

In this section we solve the $APSP(n, 1)$ problem. Solving the $APSP(n, M)$ problem is similar, since the treatment of ‘large’ edges is the same as in the $APSD$ problem. We first show what cannot be done. Then we solve the positive case (subsection 4.1). We solve the nonnegative case in subsection 4.2. The solution for the unrestricted case is complicated, and due to space limitations we only briefly mention our result for this case in subsection 4.3. In subsection 4.4 we consider the simpler special case of undirected graphs.

The naive way to represent the solution to the $APSP$ problem is to write down the full path for every one of the n^2 pairs. Figure 1 shows that this output can be $\Omega(n^3)$ long. In fact, this holds for an exponential number of input graphs. (Replace each cycle in Figure 1 by an arbitrary connected graph of $n/3$ vertices.)

As explained in the introduction, one way to avoid this cubic bottleneck is to compute witnesses. For each pair (i, j) we compute an index k of a first vertex on a shortest path from v_i to v_j . This certainly works in the positive case.

4.1 Positive $APSP$

Consider the algorithm for the positive $APSD(n, 1)$ problem [2]. It uses Boolean matrix multiplications to compute short distances. Computing Boolean matrix multiplication with witnesses gives witnesses for these shortest paths.

For computing large distances, we use the separator trick: We consider in turn each vertex as a source and the layered graph obtained by single source shortest paths. We take a block of consecutive layers, choose the smallest one and use it as a separator. Each path that goes beyond the separator must go through the separator. Hence we minimize over the choice of the vertex on the separator. This part is performed naively and provides witnesses: If a shortest path from v_i to v_j goes through v_k , where v_k belongs to the separator, then the i, j witness can be taken to be the i, k witness that has already been computed.

Consequently, to obtain the time bound we can substitute $\tilde{O}(n^\omega)$ (the time for witnessed Boolean matrix multiplication) for $O(n^\omega)$ in the bound for the positive $APSD(n, 1)$:

Theorem 3 *The positive $APSP(n, 1)$ problem can be solved in $\tilde{O}(n^{(3+\omega)/2})$ time.*

4.2 Nonnegative $APSP$

Consider the nonnegative $APSD(n, 1)$ problem. One way to solve it is to eliminate the zero length edges first. We first compute Z , the transitive closure of the zero edges. It gives us all zero (shortest) distances. Let D' be obtained from D by replacing the $+\infty$ entries by zero. Compute the Boolean

matrix multiplication $E \stackrel{\text{def}}{=} (Z + I)D'(Z + I)$. An edge in E corresponds to a path of length 1. Solve the *APSP* problem for E to obtain all shortest nonzero distances.

A first attempt to obtain witnesses for shortest paths is simply to combine three sets of witnesses: the ones obtained from the solution of the positive *APSP*, the witnesses implied by the definition of E (two per entry) and the witnesses of the transitive closure Z . This approach does not give witnesses for the shortest paths as the following example shows.




Figure 3:

Consider Figure 3. E contains a direct edge of length 1 from a to x and from b to x . Suppose that the witness which was chosen in the $(Z + I)D'(Z + I)$ multiplication for the $a \rightarrow x$ pair is b_1 and for the $b \rightarrow x$ pair is a_1 . Simply combining the three sets of witnesses would lead us to the infinite path a, b, a, b, \dots for the pair $a \rightarrow x$.

The difficulty above is caused by zero length cycles in G . But we overcome it as in Section 3. We find the strongly connected components of Z , compute witnesses for them and contract them in G , forming the contracted graph G' . We avoid multiple edges by always choosing a minimum length edge. We next find the witnesses for shortest paths in G' by the method described above. (In G' there are no zero length cycles.) Finally, we combine the two sets of witnesses as in Section 3 to obtain the witnesses for shortest paths of G . The full details are omitted.

Theorem 4 *The nonnegative APSP($n, 1$) problem can be solved in $\tilde{O}(n^{(3+\omega)/2})$ time.*

4.3 Unrestricted APSP

The subcubic algorithm for the unrestricted *APSP* requires familiarity with that for the unrestricted *APSD* [2]. The latter is a recursive algorithm that uses Boolean matrix multiplications, transitive closure computations and the separator trick. Some recursive calls construct a new graph and then operate on it. Our algorithm for the unrestricted *APSP*($n, 1$) problem is obtained by a very careful unwinding of the recursive calls in the *APSD*($n, 1$) algorithm, introducing a natural notion which we call the *negative components* of a graph and applying several additional ideas. The details are complicated and will be given in the full version. The running time obtained exceeds that of the algorithm for the corresponding *APSD*($n, 1$) problem only by a polylogarithmic factor.

Theorem 5 *The unrestricted APSP($n, 1$) problem can be solved in $\tilde{O}(n^{(3+\omega)/2})$ time.*

4.4 Undirected APSP

In the undirected case, zero edges or negative edges pose no problem and can be eliminated easily. The special algorithm for the undirected *APSP* uses essentially only Boolean matrix multiplications, which can easily be replaced by ones with witnesses. A *randomized* algorithm that solves the undirected *APSP*($n, 1$) problem in (expected) time $\tilde{O}(n^\omega)$ has been found by Seidel as well as by ourselves. Our result in Section 2 enable us to obtain a *deterministic* algorithm with essentially the same running time.

Theorem 6 *The undirected APSP($n, 1$) problem can be solved in $\tilde{O}(n^\omega)$ time.*

5 Open Problems

Of course one would like to improve the time bounds for computing various types of witnesses. In particular, can we compute the witnesses for Boolean matrix multiplication in time $O(n^\omega)$? Alternatively, one would like to improve the best algorithms for $APSD(n, M)$ and $APSP(n, M)$, for directed graphs and for undirected graphs, for the nonnegative case and for the unrestricted case, by improving the exponent (in the undirected case this is not possible without improving Boolean matrix multiplication) or by improving the dependence on M .

We are also looking for other problems for which one may need the witnesses in addition to Boolean matrix multiplication or to transitive closure. It seems very plausible that these will find additional applications in the future.

Given the shortest distances, how hard is it to compute witnesses for the shortest paths? Possibly, this can be solved in $O(n^2)$ time, but all our algorithms need additional matrix multiplications. In case of negative cycles we can prove the following negative result. Let P be the problem of testing whether a graph has a negative cycle and let P_1 be the problem of constructing a negative cycle given the matrix D^* of shortest distances.

Theorem 7 *The time complexity of P for graphs with n vertices, $T(n)$ satisfies $T(n) \leq T_1(2n) + O(n)$, where T_1 is the time complexity of P_1 .*

References

- [1] N. Alon, O. Goldreich, J. Hastad and R. Peralta, *Simple constructions of almost k -wise independent random variables*, Proc. 31st IEEE FOCS, St. Louis, Missouri, IEEE (1990), pp. 544-553.
- [2] N. Alon, Z. Galil and O. Margalit, *On the exponent of the All Pairs Shortest Path problem*, Proc. 32th IEEE FOCS, IEEE (1991), pp. 569-575.
- [3] D. Coppersmith and S. Winograd, *Matrix multiplication via arithmetic progressions*, Journal of Symbolic Computation 9(1990), pp. 251-280.
- [4] Z. Galil and O. Margalit, *A faster algorithm for the all pairs shortest path problem for undirected graphs*, August 1991.
- [5] J. Naor and M. Naor, *Small-bias probability spaces: efficient constructions and applications*, Proc. 22nd annual ACM STOC, ACM Press (1990), pp. 213-223.
- [6] R. Seidel, *On the All-Pairs-Shortest-Path Problem*, Proc. 24th ACM STOC, ACM Press (1992), to appear.