

Single-Source Bottleneck Path Algorithm Faster than Sorting for Sparse Graphs

Ran Duan¹

Institute for Interdisciplinary Information Sciences, Tsinghua University, Beijing, China
duanran@mail.tsinghua.edu.cn

Kaifeng Lyu

Institute for Interdisciplinary Information Sciences, Tsinghua University, Beijing, China
lkf15@mails.tsinghua.edu.cn

Yuanhang Xie

Institute for Interdisciplinary Information Sciences, Tsinghua University, Beijing, China
xieyh15@mails.tsinghua.edu.cn

Abstract

In a directed graph $G = (V, E)$ with a capacity on every edge, a *bottleneck path* (or *widest path*) between two vertices is a path maximizing the minimum capacity of edges in the path. For the single-source all-destination version of this problem in directed graphs, the previous best algorithm runs in $O(m + n \log n)$ ($m = |E|$ and $n = |V|$) time, by Dijkstra search with Fibonacci heap [Fredman and Tarjan 1987]. We improve this time bound to $O(m\sqrt{\log n} + \sqrt{mn \log n \log \log n})$, which is $O(n\sqrt{\log n \log \log n})$ when $m = O(n)$, thus it is the first algorithm which breaks the time bound of classic Fibonacci heap when $m = o(n\sqrt{\log n})$. It is a Las-Vegas randomized approach. By contrast, the s-t bottleneck path has algorithm with running time $O(m\beta(m, n))$ [Chechik et al. 2016], where $\beta(m, n) = \min\{k \geq 1 : \log^{(k)} n \leq \frac{m}{n}\}$.

2012 ACM Subject Classification Theory of computation → Design and analysis of algorithms

Keywords and phrases Graph Algorithm, Bottleneck Path, Combinatorial Optimization

Digital Object Identifier 10.4230/LIPIcs.ICALP.2018.43

1 Introduction

The *bottleneck path* problem is a graph optimization problem finding a path between two vertices with the maximum flow, in which the flow of a path is defined as the minimum capacity of edges on that path. The bottleneck problem can be seen as a mathematical formulation of many network routing problems, e.g. finding the route with the maximum transmission speed between two nodes in a network, and it has many other applications such as digital compositing [11]. It is also the important building block of other algorithms, such as the improved Ford-Fulkerson algorithm [10, 12], and k-splittable flow algorithm [1]. The minimax path problem which finds the path that minimizes the maximum weight on it is symmetric to the bottleneck path problem, thus has the same time complexity.

¹ R. Duan is supported by a China Youth 1000-Talent grant.



1.1 Our Results

In a directed graph $G = (V, E)$ ($n = |V|, m = |E|$), we consider the single-source bottleneck path (SSBP) problem, which finds the bottleneck paths from a given source node s to all other vertices. In the comparison-based model, the previous best time bound for SSBP is the traditional Dijkstra’s algorithm [8] with Fibonacci heap [14], which runs in $O(m + n \log n)$ time. Some progress has been made for slight variants of the SSBP problem: When the graph is undirected, SSBP is reducible to minimum spanning tree [16], thus can be solved in randomized linear time [18]; for the single-source single-destination bottleneck path (s - t BP) problem in directed graphs, Gabow and Tarjan [15] showed that it can be solved in $O(m \log^* n)$ time, and this bound was subsequently improved by Chechik et al [5] to $O(m\beta(m, n))$. However, until now no algorithm is known to be better than Dijkstra’s algorithm for SSBP in directed graphs. And as noted in [14], Dijkstra’s algorithm can be used to sort n numbers, so a “sorting barrier”, $O(m + n \log n)$, prevents us from finding a more efficient implementation of Dijkstra’s algorithm.

In this paper, we present a breakthrough algorithm for SSBP that overcomes the sorting barrier. Our main result is shown in the following theorem:

► **Theorem 1.** *Let $G = (V, E)$ be directed graph with edge weights $w : E \rightarrow \mathbb{R}$. In comparison-based model, SSBP can be solved in expected $O(\sqrt{nm} \log n \log \log n + m\sqrt{\log n})$ time.*

An alternative way to state our result is: for $m \leq n \log \log n$, SSBP can be solved in $O(\sqrt{nm} \log n \log \log n)$ time; for $m > n \log \log n$, SSBP can be solved in $O(m\sqrt{\log n})$ time. In particular, when $m = O(n)$ the time bound is $O(n\sqrt{\log n} \log \log n)$. Our algorithm is inspired by previous works on the s - t BP problem: the $O(m \log^* n)$ -time algorithm by Gabow and Tarjan [15] and the $O(m\beta(m, n))$ -time algorithm by Chechik et al [5]. See Section 3 for our intuitions.

1.2 Related Works

A “sorting barrier” seemed to exist for the the Minimum Spanning Tree problem (MST) for many years [3, 17, 19], but it was eventually broken by [25, 6]. Fredman and Tarjan [14] gave an $O(m\beta(m, n))$ -time algorithm by introducing Fibonacci heap. The current best time bounds for MST include randomized linear time algorithm by Karger et al [18], Chazelle’s $O(m\alpha(m, n))$ -time deterministic algorithm [4] and Pettie and Ramachandran’s optimal approach [20].

The single-source single-destination version of the bottleneck path (s - t BP) problem is proved to be equivalent to the Bottleneck Spanning Tree (BST) problem (see [5]). In the bottleneck spanning tree problem, we want to find a spanning tree rooted at source node s minimizing the maximum edge weight in it. For undirected graph, the s - t BP can be reduced to the MST problem. For directed graph, Dijkstra’s algorithm [8] gave an $O(n \log n + m)$ -time solution using Fibonacci heap [14]. Then Gabow and Tarjan [15] gave an $O(m \log^* n)$ -time algorithm based on recursively splitting edges into levels. Recently, Chechik et al. [5] improved the time complexity of BST and BP to randomized $O(m\beta(m, n))$ time, where $\beta(m, n) = \min\{k \geq 1 : \log^{(k)} n \leq \frac{m}{n}\}$. All these algorithms are under comparison-based model. For word RAM model, an $O(m)$ -time algorithm has been found by Chechik et al. [5].

For the all-pairs version of the bottleneck path (APBP) problem, we can sort all the edges and use Dijkstra search to obtain an $O(mn)$ time bound. For dense graphs, it has been shown that APBP can be solved in truly subcubic time. Shapira et al. [21] gave an $O(n^{2.575})$ -time APBP algorithm on vertex-weighted graphs. Then Vassilevska et al. [23]

showed that APBP for edge-weighted graphs can be solved in $O(n^{2+\omega/3}) = O(n^{2.791})$ time based on computing (max, min)-product of real matrices, which was then improved by Duan and Pettie [9] to $O(n^{(3+\omega)/2}) = O(n^{2.686})$. Here $\omega < 2.373$ is the exponent of time bound for fast matrix multiplication [7, 24].

2 Preliminaries

For a directed graph G , we denote $w(u, v)$ to be the edge weight of $(u, v) \in E$. Without additional explanation, we use the symbol n to denote the number of nodes and m to denote the number of edges in G .

2.1 Bottleneck Path Problems

The *capacity* of a path is defined to be the minimum weight among traversed edges, i.e., if a path traverses $e_1, \dots, e_l \in E$, then the capacity of the path is $\min_{i=1}^l w(e_i)$. For any $u, v \in V$, a path from u to v with maximum capacity is called a *bottleneck path from u to v* , and we denote this maximum capacity by $b(u, v)$.

► **Definition 2.** The *Single-Source Bottleneck Path* (SSBP) problem is: Given a directed graph $G = (V, E)$ with weight function $w : E \rightarrow \mathbb{R}$ and a source $s \in V$, output $b(s, t)$ for every $t \in V$, which is the maximum path capacity among all the paths from s to t .

We use the triple (G, w, s) to denote a SSBP instance with graph G , weight function $w(\cdot)$ and source node s .

It is more convenient to present our algorithm on a slight variant of the SSBP problem. We shall call it *Arbitrary-Source Bottleneck Path with Initial Capacity* (ASBPIC) problem. We assume that the edge weight $w(e)$ of an edge $e \in E$ is either a real number or infinitely large $(+\infty)$. We say an edge e is *unrestricted* if $w(e) = +\infty$; otherwise we say the edge e is *restricted*. In the ASBPIC problem, an *initial capacity* $h(v)$ is given for every node $v \in V$, and the *capacity* of a path is redefined to be the minimum between the initial capacity of the starting node and the minimum edge weights in the path, i.e., if the path starts with the node $v \in V$ and traverses e_1, \dots, e_l , then its capacity is $\min(\{h(v)\} \cup \{w(e_i)\}_{i=1}^l)$. For any $v \in V$, a path ended with v with maximum capacity is called a *bottleneck path ended with v* , and we denote this maximum capacity as $d(v)$.

► **Definition 3.** The *Arbitrary-Source Bottleneck Path with Initial Capacity* (ASBPIC) problem is: Given a directed graph $G = (V, E)$ with weight function $w : E \rightarrow \mathbb{R} \cup \{+\infty\}$ and initial capacity function $h : V \rightarrow \mathbb{R} \cup \{\pm\infty\}$, output $d(v)$ for every $v \in V$, which is the maximum path capacity among all the paths ended with v .

We use the triple (G, w, h) to denote an ASBPIC instance with graph G , weight function $w(\cdot)$, and initial capacity function $h(\cdot)$.

Note that ASBPIC and SSBP are equivalent under linear-time reductions. Given an ASBPIC instance, we construct a new graph G' from G by adding a new node v_0 which has outgoing edges with weight $h(v)$ to all the nodes $v \in V$ having $h(v) > -\infty$, then it suffices to find bottleneck paths from v_0 to all other nodes in G' . On the other hand, a SSBP instance (G, w, s) can be easily reduced to the ASBPIC instance (G, w, h) , where $h(s) = \max_{e \in E} \{w(e)\}$ and $h(v) = -\infty$ for all $v \neq s$.

2.2 Dijkstra's Algorithm for SSBP and ASBPIC

SSBP can be easily solved using a variant of Dijkstra's algorithm [8]. In this algorithm, each node is in one of the three status: *unsearched*, *active*, or *scanned*. We associate each node $v \in V$ with a *label* $d'(v)$, which is the maximum path capacity among all the paths from s to v that only traverse scanned nodes or v .

Initially, all the nodes are unsearched except s is active, and we set $d'(s) = +\infty$ and $d'(v) = -\infty$ for all $v \neq s$. We repeat the following step, which we call the *Dijkstra step*, until none of nodes is active:

- Select an active node u with maximum label and mark u as scanned. For every outgoing edge (u, v) of u , update the label of v by

$$d'(v) \leftarrow \max\{d'(v), \min\{d'(u), w(u, v)\}\}, \quad (1)$$

and mark v as active if v is unsearched.

We use priority queue to maintain the order of labels for active nodes. This algorithm runs in $O(m + n \log n)$ time when Fibonacci heap [14] is used.

The algorithm we introduced above can also be adapted for solving ASBPIC. The only thing we need to change is that in the initial stage all nodes are active and $d'(v) = h(v)$ for every $v \in V$. The resulting algorithm again runs in $O(m + n \log n)$ time. We shall call these two algorithms as *Dijkstra's algorithm for SSBP* and *Dijkstra's algorithm for ASBPIC*, or simply call any of them *Dijkstra's algorithm* when no confusion can arise.

2.3 Weak and Strong Connectivity in Directed Graph

We also need some definitions about connectivity in graph theory in this paper. A directed graph is said to be *weakly-connected* if it turns to be a connected undirected graph when changing all of its directed edges to undirected edges. A directed graph is said to be *strongly-connected* if every pair of nodes can be reached from each other. A *weakly-* (or *strongly-*) *connected components* is defined to be a maximal weakly- (or strongly-) connected subgraph.

3 Intuitions for SSBP

If all the edge weights are integers and are restricted in $\{1, \dots, c\}$, then SSBP can be solved in $O(m + c)$ time using Dijkstra's algorithm with bucket queue. If the edge weights are not necessarily small integers but all the edges given to us are already sorted by weights, then we can replace the edge weights by their ranks and use Dijkstra's algorithm with bucket queue to solve the problem in $O(m)$ time. However, edges are not sorted in general. If we sort the edges directly, then a cost of $\Omega(m \log m)$ time is unavoidable in a comparison-based model, which is more expensive than the $O(m + n \log n)$ running time of Dijkstra's algorithm.

Our algorithm is inspired by previous works on the single-source single-destination bottleneck path problem (s - t BP): the $O(m \log^* n)$ -time algorithm by Gabow and Tarjan [15] and the $O(m\beta(m, n))$ -time algorithm by Chechik et al [5]. Gabow and Tarjan's algorithm for s - t BP consists of several stages. Let $b(s, t)$ be the capacity of a bottleneck path from s to t . Initially, we know that $b(s, t)$ is in the interval $(-\infty, +\infty)$. In each stage, we narrow the interval of possible values of $b(s, t)$. Assume that $b(s, t)$ is known to be in the range (l, r) . Let $m_{(l, r)}$ be the the number of edges with weights in the range (l, r) and k be a parameter. By applying the median-finding algorithm [2] repeatedly, we choose k thresholds $\lambda_1, \dots, \lambda_k$ to split (l, r) into $k + 1$ subintervals $(l, \lambda_1), [\lambda_1, \lambda_2), [\lambda_2, \lambda_3), \dots, [\lambda_{k-1}, \lambda_k), [\lambda_k, r)$ such that for each subinterval, there are $O(m_{(l, r)}/k)$ edges of weight in it. Gabow and Tarjan then show

that *locating* which subinterval contains $b(s, t)$ can be done in $O(m_{(l,r)})$ time by incremental search. Finally, the $O(m \log^* n)$ running time bound is achieved by setting k appropriately at each stage.

The algorithm by Chechik et al. is based on the framework of Gabow and Tarjan's algorithm, but instead of selecting the thresholds $\lambda_1, \dots, \lambda_k$ by median-finding repeatedly in $O(m_{(l,r)} \log k)$ time, in this algorithm we select the k thresholds by randomly sampling in edge weights, and sort them in $O(k \log k)$ time. These thresholds partition the edges evenly with high probability, but it requires $\Omega(m \log k)$ time to compute the partition explicitly. Then they show that actually we can locate which subinterval contains $b(s, t)$ in $O(m_{(l,r)} + nk)$ (or $O(m_{(l,r)} + n \log k)$) time, without computing the explicit partition. The time bound for the overall algorithm is again obtained by setting k appropriately at each stage.

We adapt Chechik et al.'s framework for the s - t BP problem to the SSBP problem. Our SSBP algorithm actually works on an equivalent problem called ASBPIC. In ASBPIC, there is no fixed source but every node has an initial capacity, and for all destination $t \in V$ we need to compute the capacity $d(t)$ of a bottleneck path ended with t (See Section 2.1 for details). Instead of locating the subinterval for a single destination t , our algorithm locates the subintervals for all destinations $t \in V$. Thus we adopt a divide-and-conquer methodology. At each recursion, we follow the idea from Chechik et al. [5] to randomly sample k thresholds. Then we *split* the nodes into $k + 1$ levels V_0, \dots, V_k , where the i -th level contains the nodes t that have $d(t)$ in the i -th subinterval ($0 \leq i \leq k$). For each level V_i of nodes, we compute $d(t)$ for every $t \in V_i$ by reducing to solve the SSBP on a subgraph consisting of all the nodes in V_i and some of the edges connecting them. We set k to be fixed in all recursive calls, and the maximum recursion depth is $O(\log n / \log k)$ with high probability.

The split algorithm becomes the key part of our algorithm. Note that at each recursion, we should reduce or avoid the use of operations that cost time $O(\log k)$ per node or per edge (e.g., binary searching for the subinterval containing a particular edge weight). This is because that, for example, if we execute an $O(\log k)$ -time operation for each edge at each recursion, then the overall time cost is $O(m \log k) \cdot O(\log n / \log k) = O(m \log n)$, which means no improvement comparing with the previous $O(m + n \log n)$ -time Dijkstra's algorithm. Surprisingly, we can design an algorithm such that whenever we execute an $O(\log k)$ -time operation, we can always find one edge that does not appear in any subsequent recursive calls. Thus total time complexity for such operations is $O(m \log k)$, which gives us some room to obtain a better time complexity by adjusting the parameter k .

4 Our Algorithm

Our algorithm for SSBP is as follows: Given a SSBP instance (G, w, s) , we first reduce the SSBP problem to an ASBPIC instance (G, w, h) , and then use a recursive algorithm (Figure 1) to solve the ASBPIC problem. The reduction is done by setting $h(s) = \max_{e \in E} \{w(e)\}$ and $h(v) = -\infty$ for all $v \neq s$ as described in the preliminaries.

For convenience, we assume that in the original graph, all edge weights are distinct. This assumption can be easily removed.

A high-level description of our recursive algorithm for ASBPIC is shown in Figure 1. For two set A and B , $A \uplus B$ stands for the union of A and B with the assumption that $A \cap B = \emptyset$. We use $E^{(r)}$ to denote the set of restricted edges in G , and similarly we use $E_i^{(r)}$ to denote the set of restricted edges in G_i for each ASBPIC instance (G_i, w_i, h_i) . When the thresholds $\lambda_0, \dots, \lambda_{k+1}$ are present, we define the *index* of x for every $x \in \mathbb{R}$ to be the unique index i such that $\lambda_i \leq x < \lambda_{i+1}$, and we denote it as $I(x)$. For $x = \pm\infty$, we define

Algorithm 1 Main Algorithm.

Input: Directed graph $G = (V, E)$, weight function $w(\cdot)$, initial capacity function $h(\cdot)$;
Parameter $k \geq 1$.

Output: For each $v \in V$, output $d(v)$, the capacity of a bottleneck path ended with v .

- 1: **if** G is not weakly-connected **then**
- 2: Compute $\{d(v)\}_{v \in V}$ in each weakly-connected component recursively.
- 3: Let $E^{(r)} = \{e \in E \mid w(e) < +\infty\}$ be the set of restricted edges.
- 4: **if** $|E^{(r)}| \leq 1$ **then**
- 5: Compute $\{d(v)\}_{v \in V}$ in linear time and exit. ▷ Section 4.2
- 6: Sample $l = \min\{k, |E^{(r)}|\}$ distinct edges from $E^{(r)}$ uniformly randomly.
- 7: Sort the sampled edges by weights, and let $\lambda_1 < \dots < \lambda_l$ be their weights.
- 8: Let $\lambda_0 = -\infty, \lambda_{l+1} = +\infty$.
- 9: Split V into $l + 1$ levels: $V_0 \uplus V_1 \uplus \dots \uplus V_l$, where $V_i = \{v \in V \mid \lambda_i \leq d(v) < \lambda_{i+1}\}$.
▷ Section 4.3
- 10: For every level V_i , reduce the computation of $\{d(v)\}_{v \in V_i}$ to a new ASBPIC instance (G_i, w_i, h_i) , where $G_i = (V_i, E_i)$ is a subgraph of G consisting of all the nodes in V_i and some of edges that connect them. Solve each (G_i, w_i, h_i) instance recursively.

$I(-\infty) = 0, I(+\infty) = l$. Note that all the subgraphs G_i at Line 10 are disjoint. We denote $r = |E| - \sum_{i=0}^l |E_i|$ to be the total number of edges in E that do not appear in any recursive calls of (G_i, w_i, h_i) . For an edge $(u, v) \in E$, if u and v belong to different levels, then we say that (u, v) is *cross-level*. If u and v belong to the same level V_i and $w(u, v) < \lambda_i$, then we say that (u, v) is *below-level*; conversely, if $w(u, v) \geq \lambda_{i+1}$ then we say that (u, v) is *above-level*.

Besides the problem instance (G, w, h) of ASBPIC, our algorithm requires an additional integral parameter k . At the top level of recursion, we set the parameter $k = 2^{\Theta(\sqrt{n \log n \log \log n / m})}$ if $m \leq n \log \log n$, or $k = 2^{\Theta(\sqrt{\log n})}$ if $m > n \log \log n$. We fix this value k all through our algorithm. The value of the parameter k does not affect the correctness of our algorithm, but it controls the number of recursive calls at each recursion.

At each recursion, our algorithm first checks if G contains only one weakly-connected component. If not, then our algorithm calls itself to compute d in each weakly-connected component recursively. Now we can assume that G is weakly-connected (so $n \leq m$).

If the number of restricted edges is no more than 1, we claim that we can compute $d(v)$ for all v in linear time. The specific algorithm will be introduced in Section 4.2.

► **Lemma 4.** *ASBPIC can be solved in $O(m)$ time if there is at most one restricted edge.*

If the number of restricted edges is more than 1, then our algorithm first sample $l = \min\{k, |E^{(r)}|\}$ distinct edges from $E^{(r)}$ uniformly randomly and sort them by weights, that is, if the number of restricted edges is more than k , then we sample k distinct restricted edges and sort them; otherwise we just sort all the restricted edges. Let λ_i be the weight of the edge with rank i ($1 \leq i \leq l$) and $\lambda_0 = -\infty, \lambda_{l+1} = +\infty$.

Next, we split V into $l + 1$ levels of nodes $V = V_0 \uplus V_1 \uplus \dots \uplus V_l$, where the i -th level of nodes is $V_i = \{v \in V \mid I(d(v)) = i\} = \{v \in V \mid \lambda_i \leq d(v) < \lambda_{i+1}\}$. The basic idea of the split algorithm is: we run Dijkstra's algorithm for ASBPIC on the graph produced by mapping every edge weight $w(e)$ and initial capacity $h(v)$ in G to their indices $I(w(e))$ and $I(h(v))$, and we obtain the final label value $d'(v)$ for each node $v \in V$ (Remember that $d'(v)$ is the label of v in Dijkstra's algorithm). It is easy to show that the final label value $d'(v)$ equals $I(d(v))$, so the nodes can be easily split into levels according to their final labels. The specific

split algorithm will be introduced in Section 4.3. The time complexity for a single splitting is given below. In Theorem 9 we show that this implies that the total time cost for splitting is $O(m \log n / \log k + n \log n \log \log k / \log k + m \log k)$.

► **Lemma 5.** *Splitting V into levels at Line 9 can be done in $O(m + n \log \log k + r \log k)$.*

Finally, for every level V_i , we compute $d(\cdot)$ for nodes in this level by reducing to a new ASBPIC instance (G_i, w_i, h_i) , where G_i is a subgraph of G consisting of all the nodes in V_i and some of edges that connect them. We solve each new instance by a recursive call. The construction of (G_i, w_i, h_i) is as follows:

- $G_i = (V_i, E_i)$, where V_i is the nodes at level i in G , and E_i is the set of edges which connect two nodes at level i and are not below-level, i.e., $E_i = \{(u, v) \in E \mid u, v \in V_i, w(u, v) \geq \lambda_i\}$;
- For any $e \in E_i$, $w_i(e) = +\infty$ if e is above-level; otherwise $w_i(e) = w(e)$;
- For any $v \in V_i$, $h_i(v) = \max(\{h(v)\} \cup \{w(u, v) \in E \mid u \in V_{i+1} \uplus \dots \uplus V_l\})$.

► **Lemma 6.** *We can construct all the new ASBPIC instances $(G_0, w_0, h_0), \dots, (G_l, w_l, h_l)$ in $O(m)$ time. For $v \in V_i$, the value of $d(v)$ in the instance (G_i, w_i, h_i) exactly equals to the value of $d(v)$ in the instance (G, w, h) .*

Proof. We can construct all these instances (G_i, w_i, h_i) for all $0 \leq i \leq l$ by linearly scanning all the nodes and edges, which runs in $O(m)$ time. We prove the correctness by transforming (G, w, h) to (G_i, w_i, h_i) step by step, while preserving the values of $d(v)$ for all $v \in V_i$.

By definition, $\lambda_i \leq d(v) < \lambda_{i+1}$ for all $v \in V_i$. We can delete all the nodes at level less than i and delete all the edges with weight less than λ_i , since no bottleneck path ended with a node in V_i can traverse them. Also, for every edge e with weight $w(e) \geq \lambda_{i+1}$, we can replace the edge weight with $+\infty$ since $w(e)$ is certainly not the minimum in any path ended with a node in V_i .

For every edge $e = (u, v)$ where $u \in V_{i+1} \uplus \dots \uplus V_l$ and $v \in V_i$, the edge weight $w(e)$ must be less than λ_{i+1} , otherwise $d(v) \geq \min\{d(u), w(u, v)\} \geq \lambda_{i+1}$ leads to a contradiction. Thus contracting all the nodes in $V_{i+1} \uplus \dots \uplus V_l$ to a single node v_0 with infinite initial capacity is a transformation preserving the values of $d(v)$ for all $v \in V_i$. Finally, our construction follows by taking $h_i(v)$ to be the maximum between the weight of incoming edges from v_0 and the initial capacity $h(v)$ for every $v \in V_i$. ◀

► **Remark.** In any subsequent recursive calls of (G_i, w_i, h_i) , neither cross-level nor below-level edges will appear, and all the above-level edges will become unrestricted. Also, it is easy to see that r is just the total number of cross-level and below-level edges (Recall that r is the number of edges that do not appear in any recursive calls).

4.1 Running Time

First we analyze the maximum recursion depth. The following lemma shows that randomly sampled thresholds evenly partition the restricted edges with high probability.

► **Lemma 7.** *Let $E^{(r)} = \{e_1, \dots, e_q\}$ be q restricted edges sorted by their weights in E . Let f_1, \dots, f_k be $k \geq 2$ random edges sampled from $E^{(r)}$ such that $w(f_1) < w(f_2) < \dots < w(f_k)$. Let $\lambda_i = w(f_i)$ for $i = 1, \dots, k$, and $\lambda_0 = -\infty$, $\lambda_{k+1} = +\infty$. Let $F_i = \{e \in E \mid \lambda_i \leq w(e) < \lambda_{i+1}\}$. Then for every $t > 0$, $\max_{0 \leq i \leq k} \{|F_i|\} < tq \log k / k$ holds with probability $1 - k^{-\Omega(t)}$.*

Proof. Let $M = tq \log k / k$. If $\max_{0 \leq i \leq k} \{|F_i|\} \geq M$, then there exists an edge e_p such that e_p is chosen but for any $p + 1 \leq j < p + M$, e_j is not chosen. Note that when p is given, this

event happens with probability $\leq k \cdot (1/q) \cdot \prod_{i=1}^{k-1} ((q-M-i)/(q-i)) \leq (k/q) \cdot (1-M/q)^{k-1}$. By the union bound for all possible p , we have

$$\Pr \left[\max_{0 \leq i \leq k} \{|F_i|\} \geq tq \log k/k \right] \leq k(1 - t \log k/k)^{k-1} \leq k^{-\Omega(t)},$$

which completes the proof. \blacktriangleleft

For our purposes it is enough to analyze the case that $k = 2^{\Omega(\sqrt{\log n})}$. The following lemma gives a bound for the maximum recursion depth using Lemma 7.

► **Lemma 8.** *For $k = 2^{\Omega(\sqrt{\log n})}$ where n is the number of nodes at the top level of recursion, the maximum recursion depth is $O(\log n / \log k)$ with probability $1 - n^{-\omega(1)}$.*

Proof. It is not hard to see that the total number of recursive calls of our main algorithm is $O(m)$. Applying Lemma 7 with $t = \Theta(\log n)$ and the union bound for all recursive calls, we know that with probability at least $1 - k^{-\Omega(t)} \cdot O(m) = 1 - n^{-\Omega(\log n)^{3/2}}$, after every split with $|E^{(r)}| > k$, the number of restricted edges in G_i is less than $(t \log k/k) \cdot |E^{(r)}|$ for every (G_i, w_i, h_i) . Thus after $O(\log m / \log(k/(t \log k))) = O(\log n / \log k)$ levels of recursion, every ASBPIC instance (G, w, h) has $|E^{(r)}| \leq k$, and this means that in any recursive call of (G_i, w_i, h_i) , the graph G_i has at most one restricted edge, which will be directly solved at Line 5. \blacktriangleleft

The overall time complexity of our algorithm is given by the following theorem:

► **Theorem 9.** *For $k = 2^{\Omega(\sqrt{\log n})}$, with probability $1 - n^{-\omega(1)}$, our main algorithm shown in Figure 1 runs in $O(n \log n \log \log k / \log k + m \log n / \log k + m \log k)$ time.*

Proof. Let $r = |E| - \sum_{i=0}^l |E_i|$, $r' = |E^{(r)}| - \sum_{i=0}^l |E_i^{(r)}|$. First we show that the running time in each recursive call is $O(m + n \log \log k + (r + r') \log k)$.

In each recursive call of our algorithm, the time cost for sorting at Line 7 is $O(l \log l)$. For the sample edge e_i with rank i , either V_i is not empty, or this edge is cross-level, below-level, or above-level. Let l_1 be the number of edges in the former case, and l_2 be the number of edges in the latter case. For the former case, note that we only run the split algorithm for weakly-connected graphs, so there are at least $l_1 - 1$ cross-level edges, which implies $l_1 \leq r + 1$. For the latter case, e_i becomes unrestricted or does not appear for every G_i , so $l_2 \leq r'$. Thus $l = l_1 + l_2 \leq r + r' + 1$ and the time cost for sorting is $O((r + r') \log k)$.

By Lemma 5, the split algorithm runs in $O(m + n \log \log k + r \log k)$ time in each recursive call. All other parts of our algorithm run in linear time. Thus the running time for each recursion is $O(m + n \log \log k + (r + r') \log k)$. Note that the recursion depth is $O(\log n / \log k)$ with probability $1 - n^{-\omega(1)}$. We can complete the proof by adding the time cost for all the $O(m)$ recursive calls together. \blacktriangleleft

Finally, we can get our main result by setting $k = 2^{\Theta(\sqrt{n \log n \log \log n / m})}$ for $m \leq n \log \log n$ and setting $k = 2^{\Theta(\sqrt{\log n})}$ for $m > n \log \log n$.

► **Theorem 10.** *For $m \leq n \log \log n$, SSBP can be solved in $O(\sqrt{nm \log n \log \log n})$ time with high probability; For $m > n \log \log n$, SSBP can be solved in $O(m \sqrt{\log n})$ time with high probability.*

► **Remark.** The above time bounds are also true for expected running time. It can be easily derived from the fact that the worst-case running time is at most $n^{O(1)}$.

In the rest of this section, we introduce the algorithm for ASBPIC with at most one restricted edge and the split algorithm.

4.2 Algorithm for the Graph with at Most One Restricted Edge

We introduce our algorithm for the graph with at most one restricted edge in the following two lemmas.

► **Lemma 11.** *For a given ASBPIC instance (G, w, h) , if there is no restricted edge in G , then the values of $d(v)$ for all $v \in V$ can be computed in linear time.*

Proof. G contains only unrestricted edges, so for every $v \in V$, $d(v)$ is just equal to the maximum value of $h(u)$ among all the nodes u that can reach v . If v_1 and v_2 are in the same strongly-connected component, then $d(v_1) = d(v_2)$. Thus we can use Tarjan's algorithm [22] to contract every strongly-connected component to a node. The initial capacity of a node is the maximum $h(u)$ for all u in the component, and the capacity of an edge between nodes is the maximum among edges connecting components. Then Dijkstra approach on DAG takes linear time. ◀

► **Lemma 12.** *For a given ASBPIC instance (G, w, h) , if there is exactly one restricted edge in G , then the values of $d(v)$ for all $v \in V$ can be computed in linear time.*

Proof. Let $e_0 = (u_0, v_0)$ be the only restricted edge in G . There are two kinds of paths in G :

1. Paths that do not traverse e . We remove e from G and use the algorithm in Lemma 11 to get $d(v)$ for every node $v \in V$.
2. Paths that traverse e . Note that $d(u_0)$ got in the previous step is the maximum capacity to u_0 through only restricted edges. Then we update $d(v)$ by $\max\{d(v), \min\{d(u_0), w(u_0, v_0)\}\}$ for every node v that can be reached from v_0 .

We output the values of $d(\cdot)$ after these two kinds of updates, then all the paths should have been taken into account. ◀

4.3 Split

Now we introduce the split algorithm at Line 9 in our main algorithm. As before, we use the notation $I(x)$ for the index of a value x and $d'(v)$ is the label of v in Dijkstra's algorithm. The goal of this procedure is to split V into $l + 1$ levels, $V = V_0 \uplus V_1 \uplus \dots \uplus V_l$, where $V_i = \{v \in V \mid I(d(v)) = i\}$. We need to show that this can be done in $O(m + n \log \log k + r \log k)$ time, where $r = |E| - \sum_{i=0}^l |E_i|$ is the total number of edges in E that do not appear in any (G_i, w_i, h_i) .

A straightforward approach to achieve this goal is to use Dijkstra's algorithm as described in Section 2.2. We map all the edge weights and initial capacities to their indices using binary searches, and run Dijkstra's algorithm for ASBPIC. The output should be exactly $d'(v) = I(d(v))$ for every v . However, this approach is rather inefficient. Evaluating the index $I(x)$ for a given x requires $\Omega(\log l)$ time in a comparison-based model, thus in total, this algorithm needs $\Omega(n \log l)$ time to compute indices, and this does not meet our requirement for the running time.

The major bottleneck of the above algorithm is the inefficient index evaluations. Our algorithm overcomes this bottleneck by reducing the number of index evaluations for both edge weights and initial capacities to be at most $O(r + n/\log l)$.

4.3.1 Index Evaluation for Edge Weights

First we introduce our idea to reduce the number of index evaluations for edge weights. Recall that in Dijkstra's algorithm for ASBPIC we maintain a label $d'(v)$ for every $v \in V$.

In every Dijkstra step, we extract an active node u with the maximum label, and for all edges $(u, v) \in E$ we compute $\min\{d'(u), I(w(u, v))\}$ to update $d'(v)$. In the straightforward approach we evaluate every $I(w(u, v))$ using binary search, but actually this is a big waste:

1. If $w(u, v) \geq \lambda_{d'(u)}$, then $\min\{d'(u), I(w(u, v))\} = d'(u)$, so there is no need to evaluate $I(w(u, v))$.
2. If $w(u, v) < \lambda_{d'(u)}$, then $\min\{d'(u), I(w(u, v))\} = I(w(u, v))$, so we do need to evaluate $I(w(u, v))$. However, it can be shown that (u, v) is either a cross-level edge or a below-level edge, so (u, v) will not appear in any subsequent recursive calls of (G_i, w_i, h_i) .

Using the method discussed above, we can reduce the number of index evaluations for edge weights to be at most $r = |E| - \sum_{i=0}^l |E_i|$ in Dijkstra's algorithm. Lemma 14 gives a formal proof for this.

4.3.2 Index Evaluation for Initial Capacities

Now we introduce our idea to reduce the number of index evaluations for initial capacities. Recall that in Dijkstra's algorithm, we need to initialize the label $d'(v)$ to be $I(h(v))$ for each $v \in V$, and maintain a priority queue for the labels of all active nodes. If we evaluate every $I(h(v))$ directly, we have to pay a time cost $\Omega(n \log l)$.

In our split algorithm, we first find a spanning tree T of G after replacing all the edges with undirected edges. Then we partition the tree T into $b = O(n/s)$ edge-disjoint (but not necessarily node-disjoint) subtrees, T_1, \dots, T_b , each of size $O(s)$. This partition can be found in $O(n)$ time using a slight variant of the topological partition algorithm in Frederickson's paper [13].

► **Theorem 13.** (see [13]) *Given a tree $T = (V, E)$ with n nodes and given an integer $1 \leq s \leq n$, there exists a linear time algorithm that can partition T into edge-disjoint subtrees, T_1, \dots, T_b , such that the number of nodes in each subtree is in the range $[s, 3s)$.*

We form b groups of nodes, U_1, \dots, U_b , where U_i is the group of nodes that are in the i -th subtree T_i . In the running of Dijkstra's algorithm, we divide the active nodes in V into two kinds:

1. **Updated node.** This is the kind of node v that has $d'(v)$ already been updated by Dijkstra's update rule (1), which means $d'(v) = \min\{d'(u), I(w(u, v))\} \geq I(h(v))$ for some u after a previous update. The value of $\min\{d'(u), I(w(u, v))\}$ is evaluated according to Section 4.3.1, so the value of $d'(v)$ can be easily known. We can store such nodes in buckets to maintain the order of their labels.
2. **Initializing node.** This is the kind of node v whose $d'(v)$ has not been updated by Dijkstra's update rule (1), so $d'(v) = I(h(v))$. However, we do not store the value of $I(h(v))$ explicitly. For each group U_i , we put the initializing nodes in U_i into a priority queue in descending order by their initial capacities $h(\cdot)$. We only compute the value of $I(h(v))$ when v is the maximum in its group, and use buckets to maintain the maximum values from all groups.

At each iteration in Dijkstra's algorithm, we extract the active node with the maximum label among the updated nodes and initializing nodes and mark it as scanned. For the case that the maximum node $v \in U_j$ is an initializing node, we remove v from the priority queue of U_j , and compute $d'(u) = I(h(u))$ for the new maximum node u . However, if we compute this value directly using an index evaluation for $h(u)$, then we will suffer a total cost $\Omega(n \log l)$, which is rather inefficient.

Algorithm 2 Split Algorithm.

Input: Directed graph $G = (V, E)$, weight function $w(\cdot)$, initial capacity function $h(\cdot)$, $l + 2$ thresholds $-\infty = \lambda_0 < \lambda_1 < \dots < \lambda_l < \lambda_{l+1} = +\infty$

Output: Split V into $V_0 \uplus V_1 \uplus \dots \uplus V_l$ where $V_i = \{v \in V \mid \lambda_i \leq d(v) < \lambda_{i+1}\}$.

- 1: Let $s = \min\{\lceil \log l \rceil, n\}$. Find a spanning tree T of G after replacing all the edges with undirected edges and partition T into $b = O(n/s)$ edge-disjoint subtrees T_1, \dots, T_b , each of size $O(s)$. Form b groups of nodes, U_1, \dots, U_b , where U_i is the group of nodes that are in the i -th subtree T_i .
 - 2: For nodes in each group, we maintain a priority queue implemented by binary heap by the order of $h(\cdot)$.
 - 3: Evaluate $I(h(v))$ only for the maximum node in each group, and initialize $l + 1$ buckets B_0, \dots, B_l which store the groups according to the index of $h(\cdot)$ of their maximum nodes.
 - 4: Initialize $l + 1$ buckets $C_0, \dots, C_l \leftarrow \emptyset$ which store the active nodes according to $d'(\cdot)$.
 - 5: **for** $i \leftarrow l, l - 1, \dots, 1, 0$ **do**
 - 6: **for all** $U \in B_i$ **do**
 - 7: **while** $U \neq \emptyset$ **and** the maximum node u in U has $I(h(u)) = i$ **do**
 - 8: Extract u from U and put u into C_i
 - 9: $d'(u) \leftarrow i$
 - 10: **while** $C_i \neq \emptyset$ **do**
 - 11: Extract a node u from C_i
 - 12: **for all** $(u, v) \in E$ **do**
 - 13: $\bar{w} \leftarrow \min\{d'(u), I(w(u, v))\}$. (Evaluate $I(w(u, v))$ only if $w(u, v) < \lambda_{d'(u)}$)
 - 14: **if** $\lambda_{\bar{w}} > h(v)$ **then**
 - 15: **if** $d'(v)$ does not exist **or** $\bar{w} > d'(v)$ **then**
 - 16: Delete v from $C_{d'(v)}$
 - 17: Delete v from the group U containing it (if any) and put v into $C_{\bar{w}}$
 - 18: $d'(v) \leftarrow \bar{w}$
 - 19: **while** B_i contains at least one non-empty group **do**
 - 20: Extract a non-empty group U from B_i
 - 21: Evaluate $I(h(u))$ for the maximum node u in U and put U into $B_{I(h(u))}$
 - 22: Split V according to d' and return
-

The idea to speed up is to check whether $d'(u) = d'(v)$ before performing an index evaluation. This can be done in $O(1)$ time since we can know the corresponding interval of $h(u)$ from the value of $d'(v)$ if $I(h(u)) = d'(v)$. We only actually evaluate $I(h(u))$ if $d'(u) \neq d'(v)$ after the Dijkstra step scans all nodes with level $d'(v)$. In this way, we can always ensure that the number of index evaluations in a group never exceeds the number of different final label values $d'(\cdot)$ in this group. Indeed, it can be shown that if there are c different final label values $d'(\cdot)$ in a group U_i , then there must be at least $c - 1$ cross-level edges in T_i , which implies that the number of index evaluations for initial capacities should be no greater than $r + b$. (Remember b is the number of groups in the partition.) Lemma 15 gives a formal proof for this.

4.3.3 The $O(m + n \log \log k + r \log k)$ -time Split Algorithm

Now we are ready to introduce our split algorithm in full details. A pseudocode is shown in Figure 2. During the search at Line 5 - 21, B_i may contain groups with maximum nodes not at the i -th level, e.g., when the maximum node in a group U is deleted at Line 17 and

$I(h(u))$ of the new maximum node u in U has not been evaluated yet. We have the following observations:

- For all $v \in V$, $d'(v)$ is non-decreasing, and at the end we have $d'(v) = I(d(v))$.
- All the binary heaps of U_i are deletion only (so the priority queue can also be implemented by a sorted linked list).
- Only at Line 3, 13 and 21 we need to evaluate the index of $h(\cdot)$ of a node. Each index evaluation costs $O(\log l)$ time.
- The numbers of executions of the while loops at Line 7 - 9, Line 10 - 18, Line 19 - 21 are all bounded by $O(n)$.

Our algorithm is an implementation of Dijkstra's algorithm, so the correctness is obvious. The running time analysis is based on the following lemmas:

► **Lemma 14.** *If we evaluate $I(w(u, v))$ at Line 13, then the edge (u, v) will not be in any recursive calls of (G_i, w_i, h_i) .*

Proof. We evaluate $I(w(u, v))$ only if $w(u, v) < \lambda_{d'(u)}$, so $\bar{w} = I(w(u, v))$ right after Line 13. Since u has already been scanned, $d'(u)$ here is just its final value $I(d(u))$. Note that $I(d(v)) \geq \min\{I(d(u)), I(w(u, v))\} = I(w(u, v))$. If $I(d(v)) > I(w(u, v))$, then $I(w(u, v))$ is smaller than both $I(d(u))$ and $I(d(v))$, thus (u, v) is a below-level edge or cross-level edge. If $I(d(v)) = I(w(u, v))$, then $I(d(v)) < I(d(u))$ and (u, v) is a cross-level edge. ◀

► **Lemma 15.** *At Line 3 and 21, if we evaluate $I(h(u))$ for c nodes u in some group U_i , then the number of different final label values $d'(\cdot)$ in U_i is at least c . Thus, the number of edges in the subtree T_i corresponding to U_i that do not appear in any recursive calls of (G_i, w_i, h_i) is at least $c - 1$.*

Proof. At line 21, $I(d(u))$ must be less than i ; otherwise, u should have been extracted at Line 11 before extracting u at Line 20, which is impossible. Also note that $d(u) \geq h(u)$, so $I(h(u)) \leq I(d(u))$ for all $u \in V$. Thus, if $u_1 \in U_i$ is evaluated at Line 3 and there are $c - 1$ nodes u_2, \dots, u_c extracted from U_i at Line 21, then $I(d(u_1)), I(d(u_2)), \dots, I(d(u_c))$ should be in c distinct ranges: $[I(h(u_1)), +\infty), [I(h(u_2)), I(h(u_1))), \dots, [I(h(u_c)), I(h(u_{c-1}))]$, which implies that the number of different final values of $d'(u)$ in U_i is at least c .

Suppose we remove all the cross-level edges in T_i , i.e., remove all the edges (u, v) in T_i whose final values of $d'(u)$ and $d'(v)$ differ. Then the tree should be decomposed into at least c components since there are at least c different final values of $d'(u)$ in U_i . Thus there are at least $c - 1$ cross-level edges in T_i . ◀

Finally, we can derive the $O(m + n \log \log k + r \log k)$ time bound for our split algorithm.

Proof for Lemma 5. By Lemma 14, the number of index evaluations at Line 13 is at most r . Let c_i be the number of index evaluations at Line 3 and 21 for nodes in the group U_i . Then by Lemma 15, $\sum_{i=1}^b c_i \leq r + b$. Thus the total number of index evaluations in our split algorithm can be bounded by $2r + b \leq O(r + n / \log l)$, which costs $O(r \log l + n)$ time.

The extraction operations of the maximum element at Line 8 and 21 and the deletion operation at Line 17 runs in $O(\log s) \leq O(\log \log l)$ time per operation. Note that there are $O(n)$ such operations, thus all these operations cost $O(n \log \log l)$ time in total.

It can be easily seen that all other parts of our split algorithm runs in linear time, so the overall time complexity is $O(m + n \log \log l + r \log l) \leq O(m + n \log \log k + r \log k)$. ◀

5 Discussion

We give an improved algorithm for solving SSBP in $O(m\sqrt{\log n} + \sqrt{mn \log n \log \log n})$ faster than sorting for sparse graphs. This algorithm is a breakthrough compared to traditional Dijkstra's algorithm using Fibonacci heap. There are some open questions remained to be solved. Can our algorithm for SSBP be further improved to $O(m\beta(m, n))$, which is the time complexity for the currently best algorithm for s - t BP? Can we use our idea to obtain an algorithm for SSBP that runs faster than Dijkstra in word RAM model?

References

- 1 Georg Baier, Ekkehard Köhler, and Martin Skutella. On the k -splittable flow problem. In *European Symposium on Algorithms*, pages 101–113. Springer, 2002.
- 2 Manuel Blum, Robert W. Floyd, Vaughan R. Pratt, Ronald L. Rivest, and Robert Endre Tarjan. Time bounds for selection. *J. Comput. Syst. Sci.*, 7(4):448–461, 1973.
- 3 O Boruvka. O jistém problému minimalním, praca moravske prirodovedecke spolocnosti 3, 1926.
- 4 B. Chazelle. A minimum spanning tree algorithm with inverse-Ackermann type complexity. *J. ACM*, 47(6):1028–1047, 2000.
- 5 Shiri Chechik, Haim Kaplan, Mikkel Thorup, Or Zamir, and Uri Zwick. Bottleneck paths and trees and deterministic graphical games. In *LIPICs-Leibniz International Proceedings in Informatics*, volume 47. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik, 2016.
- 6 D. Cheriton and R. E. Tarjan. Finding minimum spanning trees. *SIAM J. Comput.*, 5:724–742, 1976.
- 7 Don Coppersmith and Shmuel Winograd. Matrix multiplication via arithmetic progressions. *Journal of Symbolic Computation*, 9(3):251–280, 1990. Computational algebraic complexity editorial. doi:10.1016/S0747-7171(08)80013-2.
- 8 Edsger W Dijkstra. A note on two problems in connexion with graphs. *Numerische mathematik*, 1(1):269–271, 1959.
- 9 Ran Duan and Seth Pettie. Fast algorithms for (\max, \min) -matrix multiplication and bottleneck shortest paths. In *Proceedings of the twentieth annual ACM-SIAM symposium on Discrete algorithms*, pages 384–391. Society for Industrial and Applied Mathematics, 2009.
- 10 Jack Edmonds and Richard M Karp. Theoretical improvements in algorithmic efficiency for network flow problems. *Journal of the ACM (JACM)*, 19(2):248–264, 1972.
- 11 Elena Fernandez, Robert Garfinkel, and Roman Arbiol. Mosaicking of aerial photographic maps via seams defined by bottleneck shortest paths. *Oper. Res.*, 46(3):293–304, 1998. doi:10.1287/opre.46.3.293.
- 12 L. R. Ford and D. R. Fulkerson. Maximal flow through a network. *Canadian Journal of Mathematics*, 8(0):399–404, 1956. doi:10.4153/cjm-1956-045-5.
- 13 Greg N Frederickson. Data structures for on-line updating of minimum spanning trees, with applications. *SIAM Journal on Computing*, 14(4):781–798, 1985.
- 14 Michael L Fredman and Robert Endre Tarjan. Fibonacci heaps and their uses in improved network optimization algorithms. *Journal of the ACM (JACM)*, 34(3):596–615, 1987.
- 15 Harold N Gabow and Robert E Tarjan. Algorithms for two bottleneck optimization problems. *Journal of Algorithms*, 9(3):411–417, 1988.
- 16 T. C. Hu. The maximum capacity route problem. *Operations Research*, 9(6):898–900, 1961. URL: <http://www.jstor.org/stable/167055>.
- 17 Vojtech Jarník. O jistém problému minimálním. *Práce Moravské Přírodovědecké Společnosti*, 6:57–63, 1930.

- 18 David R Karger, Philip N Klein, and Robert E Tarjan. A randomized linear-time algorithm to find minimum spanning trees. *Journal of the ACM (JACM)*, 42(2):321–328, 1995.
- 19 Joseph B Kruskal. On the shortest spanning subtree of a graph and the traveling salesman problem. *Proceedings of the American Mathematical society*, 7(1):48–50, 1956.
- 20 S. Pettie and V. Ramachandran. An optimal minimum spanning tree algorithm. *J. ACM*, 49(1):16–34, 2002.
- 21 Asaf Shapira, Raphael Yuster, and Uri Zwick. All-pairs bottleneck paths in vertex weighted graphs. In *Proceedings of the eighteenth annual ACM-SIAM symposium on Discrete algorithms*, pages 978–985. Society for Industrial and Applied Mathematics, 2007.
- 22 Robert Tarjan. Depth-first search and linear graph algorithms. *SIAM journal on computing*, 1(2):146–160, 1972.
- 23 Virginia Vassilevska, Ryan Williams, and Raphael Yuster. All-pairs bottleneck paths for general graphs in truly sub-cubic time. In *Proceedings of the thirty-ninth annual ACM symposium on Theory of computing*, pages 585–589. ACM, 2007.
- 24 Virginia Vassilevska Williams. Multiplying matrices faster than Coppersmith-Winograd. In *Proceedings of the 44th symposium on Theory of Computing, STOC '12*, pages 887–898, New York, NY, USA, 2012. ACM. doi:10.1145/2213977.2214056.
- 25 A. C. Yao. An $O(|E| \log \log |V|)$ algorithm for finding minimum spanning trees. *Info. Proc. Lett.*, 4(1):21–23, 1975.