

# Algorithm Engineering of SSSP with Negative Edge Weights

Alejandro Cassis ✉

Saarland University and Max Planck Institute for Informatics,  
Saarland Informatics Campus, Saarbrücken, Germany

Andreas Karrenbauer ✉ 

Max Planck Institute for Informatics and Saarland University,  
Saarland Informatics Campus, Saarbrücken, Germany

André Nusser ✉ 

Université Côte d’Azur, CNRS, Inria, France

Paolo Luigi Rinaldi ✉ 

Max Planck Institute for Informatics and Saarland University,  
Saarland Informatics Campus, Saarbrücken, Germany

---

## Abstract

Computing shortest paths is one of the most fundamental algorithmic graph problems. It is known since decades that this problem can be solved in near-linear time if all weights are nonnegative. A recent break-through by [5] presented a randomized near-linear time algorithm for this problem. A subsequent improvement in [6] significantly reduced the number of logarithmic factors and thereby also simplified the algorithm. It is surprising and exciting that both of these algorithms are combinatorial and do not contain any fundamental obstacles for being practical.

We launch the, to the best of our knowledge, first extensive investigation towards a practical implementation of [6]. To this end, we give an accessible overview of the algorithm and discuss what adaptations are necessary to obtain a fast algorithm in practice. We manifest these adaptations in an efficient implementation. We test our implementation on a benchmark data set that is adapted to be more difficult for our implementation in order to allow for a fair comparison. As in [6] as well as in our implementation there are multiple parameters to tune, we empirically evaluate their effect and thereby determine the best choices. Our implementation is then extensively compared to one of the state-of-the-art algorithms for this problem [21]. On the hardest instance type, we are faster by up to almost two orders of magnitude.

**2012 ACM Subject Classification** Theory of computation → Shortest paths

**Keywords and phrases** Single Source Shortest Paths, Negative Weights, Near-Linear Time

**Digital Object Identifier** 10.4230/LIPIcs.SEA.2025.10

**Related Version** *Full Version*: <https://arxiv.org/abs/2502.11999>

**Supplementary Material** *Software (Source Code)*: [https://github.com/PaoloLRinaldi/negative\\_weight\\_SSSP](https://github.com/PaoloLRinaldi/negative_weight_SSSP); archived at [swh:1:dir:5507a6c96528f101c37559e498da241723527a7c](https://swh.1:dir:5507a6c96528f101c37559e498da241723527a7c)

**Funding** *André Nusser*: This work was supported by the French government through the France 2030 investment plan managed by the National Research Agency (ANR), as part of the Initiative of Excellence of Université Côte d’Azur under reference number ANR-15-IDEX-01.

**Acknowledgements** We want to thank Karl Bringmann and Danupon Nanongkai for their helpful input.



© Alejandro Cassis, Andreas Karrenbauer, André Nusser, and Paolo Luigi Rinaldi;  
licensed under Creative Commons License CC-BY 4.0

23rd International Symposium on Experimental Algorithms (SEA 2025).

Editors: Petra Mutzel and Nicola Prezza; Article No. 10; pp. 10:1–10:20

Leibniz International Proceedings in Informatics



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

## 1 Introduction

One of the most fundamental algorithmic problems on graphs is the single-source shortest path problem: given a source vertex  $s$  in the graph, the goal is to compute the distance from  $s$  to all other vertices. This is a problem with a very rich history that goes back to the early years of the field of modern computer science. This problem has two seemingly fundamentally different variants. Either all the edge weights are nonnegative, or we are in the general case where the edge weights are also allowed to be negative.<sup>1</sup>

The first variant, i.e., shortest paths on graphs with nonnegative edge weights, is very well understood. It is known since the 60s that the celebrated Dijkstra's algorithm solves the problem in near-linear time [14]. There is a long line of work (e.g., [32, 18, 1, 12]) giving theoretical improvements over Dijkstra's algorithm, culminating in a near-linear time algorithm with only  $\log \log$  factors for graphs with integer weights [31]. On the practical side, there also is a lot of interest in and work on the shortest path problems. In particular, the practically efficient computation of shortest paths in road networks is very well understood, and it can be considered as one of the great success stories of algorithm engineering, see [3] for an excellent overview on this topic.

For the second variant, i.e., shortest paths with negative edge weights, progress turned out to be significantly more challenging. Note that as opposed to the shortest path problem with non-negative edge weights, here we also have the problem of cycles of negative total weight, namely, negative cycles. If such a cycle is reachable from the source, we have shortest paths of arbitrary negative cost and we consequently have to detect these cases. A classical algorithm that solves the shortest path problem with negative edge weights is the famous Bellman-Ford algorithm [17, 4, 26], which was developed in the 50s and remained unbeaten for a long time. However, the problem continued to be of interest and saw steady progress on the theoretical side, e.g. [19, 29, 24]. Given its fundamental nature, the problem also drew interest from a practical perspective. We refer the reader to [28, 13, 27, 21] for some early experimental work. Some more recent works focused on negative cycle detection, also called the shortest-path feasibility problem, see e.g. [11, 8, 10]. We refer to [10] as an excellent resource on practical shortest path algorithms with negative edge weights and negative cycle detection. Note that shortest path algorithms can also be used as a subroutine for solving min-cost flow problems with the so-called successive shortest path approach [15].

Despite the progress in shortest paths with negative edge weights, until very recently all known algorithms were still far from achieving near-linear running time. In 2022, this barrier was overcome by [9], who gave an almost-linear time randomized algorithm for the more general min-cost flow problem. In an independent work, a near-linear time randomized algorithm for shortest paths with negative weights was discovered by [5]. Both of these results constitute stunning theoretical breakthroughs, see also [16]. One disadvantage of the min-cost flow algorithm of [9] is that it is very complex and relies on intricate tools from continuous optimization, which suggests that it is still far from being of practical use.<sup>2</sup> On the other hand, a notable feature of [5] is that it does not rely on complex technology. Indeed, the algorithm is purely combinatorial (it uses tools similar to classical graph decomposition techniques that are already known since the 80s, see e.g. [2]), and none of the techniques that it uses renders it a priori impractical. This gave rise to the hope of practical near-linear time

---

<sup>1</sup> We indeed disregard the interesting distinction between integer and real edge weights here.

<sup>2</sup> There was a very recent attempt [23] to implement the algorithm of [9]. This implementation is only partial, suggesting that indeed it is difficult and the practicability of the theoretical algorithm remains unclear.

algorithms for shortest paths with negative edge weights. This hope was further boosted by a result that reduced the number of logarithmic factors in the running time from 9 to 3 and thereby also resulted in a significantly simpler algorithm [6].<sup>3</sup> We note that in [6], the authors explicitly mention the goal of paving the way for a “comparably fast implementation”.

## Our Contribution

We initiate the transfer of the theoretical break-through on the problem of shortest paths with negative edge weights [5, 6] to the application domain and provide the first implementation inspired by the work of [6], to the best of our knowledge. In particular, our work primarily aims to demonstrate the potential of this novel approach, suggesting avenues for future research to establish its competitive advantage over classical Bellman-Ford based algorithms. We first present an overview of the algorithm of [6] to make it easily accessible to a broader audience and especially practitioners.<sup>4</sup> We discuss different adaptations that are necessary to obtain an implementation that is competitive with existing algorithms. We also show an improved instance-based upper bound on a crucial parameter that controls the recursion depth of the algorithm. We create a set of benchmark instances that is inspired by the work of [10], adapting their benchmark instances to make them more difficult to solve for our implementation in order to allow for a fair comparison. As the algorithm of [6] contains multiple constants that can be chosen in different ways and our modifications also allow for different parameter settings, we conduct experiments to empirically determine a good parameter choice. We then conduct extensive experiments to understand the running time and especially also the scaling behavior of our implementation compared to one of the state-of-the-art algorithms, namely GOR [21]. On the hardest instance type, we are faster by up to almost two orders of magnitude, suggesting that our implementation particularly thrives on hard instances. Indeed, the experiments also suggest that the running time of our implementation scales near-linearly.

We note that Li and Mowry [25] very recently and independent of our work further simplify the algorithm of [6]. However, they do not provide an implementation and conduct no experimental evaluation.

## 2 Preliminaries

We first introduce notation, then explain the type of graphs that we are working on, and finally explain the algorithm of [6] that we build our engineered solution on.

### 2.1 Setting

In this work, we are always given a directed, weighted graph  $G = (V, E, w)$  equipped with a weight function  $w : E \mapsto \mathbb{Z}$  that assigns (potentially negative) integer weights to the edges. With slight abuse of notation, we use  $|G|$  to denote the number of vertices of a graph  $G$ . Given a subset of vertices  $C \subset V$ , we denote with  $G[C]$  the induced subgraph of  $C$  in  $G$ . We sometimes consider the graph  $G_{\geq 0}$ , which we define as the graph  $G$  with the modified cost function  $w_{\geq 0}(e) = \max\{w(e), 0\}$  where all negative edge weights are set to zero. We

<sup>3</sup> We note that the algorithm was simplified, but the analysis is still challenging and out of the scope of this work to explain.

<sup>4</sup> Naturally, explaining the running time analysis of the algorithm of [6] is out of the scope of this paper and we merely focus on the algorithm.

also need access to the reverse graph. We hence denote with  $G^{\text{out}}$  the original graph  $G$ , and with  $G^{\text{in}}$  the graph  $G$  but where each edge is flipped in its orientation. We use the notation  $B_{\geq 0}^{\text{dir}}(v, r)$  to denote the set of vertices that have distance at most  $r$  from  $v$  in  $G_{\geq 0}^{\text{dir}}$ , for  $\text{dir} \in \{\text{in}, \text{out}\}$ . We say that  $B_{\geq 0}^{\text{dir}}(v, r)$  is an  $r$ -ball of  $v$ . Given such a ball  $B := B_{\geq 0}^{\text{dir}}(v, r)$ , we denote its boundary by  $\partial B$ , which comprises of all the edges that lead from nodes inside the ball to nodes outside the ball. Finally, given a direction  $\text{dir} \in \{\text{in}, \text{out}\}$ , we denote by  $\overline{\text{dir}}$  the opposite direction.

The goal of this work is to solve the following problem efficiently in practice:

► **Definition 1** (SSSP with Negative Weights). *Given a directed graph  $G = (V, E, w)$  equipped with a weight function  $w : E \mapsto \mathbb{Z}$  and a source vertex  $s \in V$ , compute the shortest path costs from  $s$  to all vertices in  $V$  w.r.t. the weights  $w$  or assert that  $G$  contains a negative cycle.*

SSSP with Negative Weights reduces to finding a *potential function*  $\phi : V \mapsto \mathbb{Z}$  such that the modified weights  $w(u, v) + \phi(u) - \phi(v)$  of all edges  $(u, v) \in E$  are nonnegative; we call such a potential function *valid*. This is also called “Johnson’s Trick” [22]. There are three important facts about such potential functions:

- a valid potential function  $\phi$  always exists iff  $G$  does not contain a negative cycle, and
- $\phi$  preserves the structure of the shortest paths in  $G$ , and
- given shortest path costs w.r.t.  $\phi$ , we can easily reconstruct the original costs.

Especially, if we found a valid potential function  $\phi$ , we can just apply Dijkstra’s algorithm [14] on  $G$  with the modified weights and hence obtain our shortest path distances in the graph with the original weights.

## 2.2 Restricted Graphs

Consider any cycle  $C = e_1, \dots, e_\ell$  in  $G$ , its *mean* is defined as  $\sum_{i \in [\ell]} w(e_i) / \ell$ . The *minimum cycle mean* is defined as the minimum mean any cycle in  $G$  has. In this work we focus on restricted graphs.

► **Definition 2.** *A graph  $G = (V, E, w)$  is restricted if  $w(e) \geq -1$  for all  $e \in E$ , and the minimum cycle mean of  $G$  is at least 1.*

Note that this definition also implies that  $G$  has no negative cycle. Furthermore, the original definition also includes the existence of a super source, due to technical reasons. We drop this requirement.

We mostly restrict to these types of graphs in this work (unless mentioned otherwise), due to the following reasons:

- The main contribution of [6] is an algorithm that runs on restricted graphs. How to generalize such an algorithm to the general case is technical, but comparably simple, and was already previously known.<sup>5</sup> Hence, restricting to these types of graphs enables for a cleaner evaluation of how their techniques perform in practice.
- Most of our benchmark graphs are derived from the hard instances described in [10]. These graphs are restricted or almost restricted and thus, even if we adapt our implementation for the general case, we expected the performance on these graphs to not significantly change.

---

<sup>5</sup> Our implementation can easily be adapted to general negative-cycle-free graphs by the scaling technique of [6]. On the other hand, the negative cycle detection of [6] is impractical as it works via budgeting – a technique that, very roughly speaking, given the theoretical worst-case running time of the algorithm on well-formed inputs, will stop if it is exceeded. We discuss negative cycles more in Section 3.

---

**Algorithm 1** LAZYDIJKSTRA.
 

---

```

1: procedure LAZYDIJKSTRA( $G, \phi$ )
2:    $d \leftarrow$  distance array used in the following
3:   for each  $v \in V$  do  $d(v) \leftarrow -\phi(v)$ 
4:    $Q \leftarrow$  priority queue containing each  $v \in V$  with its initial distance  $-\phi(v)$ 
5:   while  $Q \neq \emptyset$  do
6:     Dijkstra's alg. on edges in  $E$  with non-negative cost w.r.t.  $\phi$ , updating  $d$  and  $Q$ 
7:      $A \leftarrow$  vertices that were improved (or newly initialized)
8:     One round of Bellman-Ford from  $A$  using  $\phi$  and updating  $d$ 
9:      $Q \leftarrow$  vertices that were improved
10:  return  $\phi + d$ 

```

---

- The restriction does not have any effect on the correctness of our implementation: even if our implementation is given a non-restricted graph, it still outputs the correct result.<sup>6</sup> Our implementation is merely not optimized for these cases and might have higher running times. In particular, if the graph contains a negative cycle, it still will eventually be detected by our implementation.

## 2.3 Near-Linear Time Algorithm

In this subsection we give a description of [6, Algorithm 2]. For the readers' convenience, we refer to the statements and algorithms in the freely available full version [7] of [6]. Note that in the algorithm of [6] a designated source  $s \in V$  is added and connected to all vertices via zero-weight edges. In our description and our implementation, we replace the explicit source node by an implicit source. Computing a shortest path tree from this implicit source then gives us a valid potential function, and we can subsequently run Dijkstra's algorithm from any source with the modified weights, see Section 2.1.

### 2.3.1 Subroutines

Before describing the main algorithm, we have to describe several important subroutines. Note that most of these subroutines get a potential function  $\phi$  as input. This is due to the fact that the main algorithm recursively creates new potential functions, fixing more and more of the graph until we found a valid one.

#### LazyDijkstra

We first describe an algorithm that is efficient in the case where each path in the shortest path tree contains only few negative edges. This subroutine is called LAZYDIJKSTRA, see also [5, Lemma 3.3] and [7, Lemma 25]. The algorithm intuitively is Dijkstra's algorithm interleaved with Bellman Ford edge relaxations of all the relevant negative edges, see Algorithm 1.

#### FixDAGEdges

Now we describe the subroutine FIXDAGEDGES, see also [5, Lemma 3.2] and [7, Lemma 26], see Algorithm 2. This subroutine takes a set of strongly connected components  $C_1, \dots, C_\ell$  in their topological ordering according to  $G$  as input, and we are guaranteed that for each  $i \in [\ell]$

---

<sup>6</sup> This is the case as Line 9 in Algorithm 4 computes the correct result also on non-restricted graphs even with negative cycles.

■ **Algorithm 2** FIXDAGEDGES.

---

**Require:**  $C_1, \dots, C_\ell$  given in topological ordering, and edge costs are non-negative within all  $C_i$  w.r.t.  $\phi$

```

1: procedure FIXDAGEDGES( $G, C_1, \dots, C_\ell, \phi$ )
2:    $M \leftarrow \min_{(u,v) \in E} \{w(u,v) + \phi(u) - \phi(v), 0\} - 1$ 
3:   for  $i \leftarrow 1, \dots, \ell$  do
4:     for  $v \in C_i$  do  $\phi(v) \leftarrow \phi(v) + i \cdot M$ 
5:   return  $\phi$ 

```

---

the edges within  $C_i$  are non-negative with respect to  $\phi$ . This subroutine finds a new potential function such that all edges become non-negative. The idea to achieve this is simple: If we contract each  $C_i$  to a single vertex in  $G$ , then we obtain a DAG. Hence, we can simply assign large negative potentials that decrease (i.e., become more negative) the later the component is in the topological ordering.

### Decompose

This routine is at the heart of the near-linear time algorithm by [7], see [7, Section 3.1]. We provide the pseudo-code in Algorithm 3. At a high level, the goal of DECOMPOSE is to decompose the graph into small pieces by removing a small number of cut edges from the graph. The decomposition happens based on cutting out balls whose radii are expected to be in the order of  $\kappa/\log n$ , where  $\kappa$  is a parameter that we explain and set in Section 2.3.2. To achieve the goal, we want to grow the balls around vertices whose vicinity is relatively small. These are so-called light vertices, whose  $\frac{\kappa}{4}$ -ball contains at most a  $\frac{3}{4}$ -fraction of the vertices of  $G$ . However, classifying each vertex exactly is too costly, so instead we estimate the set of light vertices. This is the idea behind the set of vertices  $L$  that we compute in Lines 4 to 8 of Algorithm 3. To this end, we randomly sample vertices, growing a ball of radius  $\frac{\kappa}{4}$  around them in opposite direction, and marking all the vertices that are contained in this ball. This is done  $k$  times, where we choose  $k$  to be  $50 \log n$ . The set  $L$  then consists of the vertices that were marked less than  $\frac{3}{5}k$  times.

Note that this subroutine introduces randomization into the algorithm as the radius of the balls is sampled from a geometric distribution, where  $\text{Geom}(p)$  denotes the geometric distribution with mean  $1/p$ . Furthermore, in our description of DECOMPOSE, we use the constant factors of our implementation instead of the ones from [6].

### 2.3.2 Algorithm

We can now describe the main algorithm on restricted graphs from [6]. See Algorithm 4 for the pseudo code. Note that we use  $|G|$  here to denote the number of vertices of the graph  $G$ , that in the first call of RESTRICTEDSSSP the vector  $\phi$  is the all-zeros vector, and also that the algorithm is recursive and calls itself on the graphs of components induced by the separators found by Algorithm 3.

We focus here on the intuitive understanding of the described algorithm; for the in-depth description, see [7, Section 3]. The algorithm is recursive and we make progress in recursive calls by either reducing the component size or by halving the parameter  $\kappa$ , see Line 6 in Algorithm 4, the base case being  $\kappa \leq 2$ . We shed light on the role of  $\kappa$  in a moment. However, first consider a single recursive step. We can categorize the negative edges into three classes after the DECOMPOSE call: (i) edges within strongly connected components  $C_i$ , (ii) cut edges  $S$ , and (iii) other edges between the strongly connected components. Type (i) edges are handled recursively. Type (ii) edges could be hard to handle, but it follows from the

---

**Algorithm 3** DECOMPOSE.

---

```

1: procedure DECOMPOSE( $G, \kappa$ )
2:    $S \leftarrow \emptyset$ 
3:   for  $\text{dir} \in \{\text{in}, \text{out}\}$  do
4:      $k \leftarrow \lceil 50 \log |G| \rceil$ 
5:     for  $i = 1, \dots, k$  do
6:        $v \leftarrow$  random vertex from  $G$ 
7:       Mark all vertices in  $B_{\geq 0}^{\text{dir}}(v, \kappa/4)$ 
8:        $L \leftarrow$  vertices marked less than  $\frac{3}{5}k$  times
9:       while there is a  $v \in L$  do
10:         $B \leftarrow B_{\geq 0}^{\text{dir}}(v, \text{Geom}(20 \log(|G|)/\kappa))$ 
11:         $S \leftarrow S \cup \partial B$ ,  $L \leftarrow L \setminus B$ , and  $G \leftarrow G \setminus B$ 
12:    $C_1, \dots, C_\ell \leftarrow$  topol. sorted SCCs of  $G \setminus S$ 
13:   return  $C_1, \dots, C_\ell, S$ 

```

---

**Algorithm 4** Algorithm for restricted graphs from [6].

---

```

1: procedure RESTRICTEDSSSP( $G, \phi, \kappa$ )
2:   if  $\kappa \leq 2$  then
3:     return LAZYDIJKSTRA( $G, \phi$ )
4:    $C_1, \dots, C_\ell, S \leftarrow$  DECOMPOSE( $G, \kappa$ )
5:   for each component  $C_i$  do
6:     if  $|C_i| \geq \frac{3}{4}|G|$  then  $\kappa_i \leftarrow \kappa/2$  else  $\kappa_i \leftarrow \kappa$ 
7:      $\phi \leftarrow$  RESTRICTEDSSSP( $G[C_i], \phi, \kappa_i$ )
8:    $\phi \leftarrow$  FIXDAGEDGES( $G \setminus S, C_1, \dots, C_\ell, \phi$ )
9:    $\phi \leftarrow$  LAZYDIJKSTRA( $G, \phi$ )
10:  return  $\phi$ 

```

---

theoretical analysis that they are expected to be few, thus, we expect to find a valid potential function for them efficiently using LAZYDIJKSTRA. Type (iii) edges are DAG edges when considering the components as being contracted into a vertex. Hence, we easily find a valid potential function for them using the FIXDAGEDGES subroutine.

Let us now explain the crucial role of the parameter  $\kappa$ . The value  $\kappa(G)$  of a given restricted graph  $G$  is the largest number of negative edges in any simple path with non-positive cost. This intuitively measures the complexity of the shortest path problem incurred by the negative edges. The parameter  $\kappa$  is supposed to be an upper bound on  $\kappa(G)$ . If  $\kappa(G)$  is small, then we can efficiently solve the shortest path problem using LAZYDIJKSTRA. This is why halving  $\kappa$  intuitively makes progress – we expect to be in a less complex case.

### 3 Implementation

We now describe our implementation and different algorithmic choices.

#### 3.1 Our Algorithm

As the goal of this work is to understand the practicability of the algorithmic approach of [6] and [5], we use the structure described in Section 2.3, and in particular Algorithm 4, as basis and modify it to obtain a practically faster algorithm.

As a first preprocessing step, we decompose the graph into strongly connected components. This allows us to run our algorithm separately on each component and later obtain a valid potential function that also makes the edges in-between the components positive using `FIXDAGEDGES`. To this end, we also have to choose one of the many algorithms to compute strongly connected components. We use Kosaraju’s and Sharir’s algorithm [30], which also gives us a topological sorting of the strongly connected components as a by-product.

A subroutine that is pervasive in [6] is Dijkstra’s algorithm. In [6] the authors use Dijkstra’s algorithm with a priority queue by Thorup [31] with the goal to shave an additional logarithmic factor. We instead use a classical priority queue, more precisely, we use a 4-heap that has proven to be fast in practice in shortest path problems with non-negative edge weights, see e.g. [20].

A critical parameter that controls the size of the components that we recurse on and thereby also the recursion depth is  $\kappa$  – hence choosing  $\kappa$  well can have significant influence on the practical running time. In Section 4, we show that using  $\text{diam}(G_{\geq 0})$  is an upper bound to  $\kappa$  on any restricted graph  $G$ . As  $\text{diam}(G_{\geq 0})$  could be greater than  $n$ , we set  $\kappa = \min\{n, \text{diam}(G_{\geq 0})\}$  in our implementation. It is also possible that the graph contains only very few negative edges in total, so we could additionally have the total number of negative edges in  $G$  as third argument in the min above. However, as all the instances that we run our experiments on contain a significant amount of negative edges, we do not use this optimization in the experiments.

The computation of the estimated set of light vertices  $L$  in `DECOMPOSE` is an important factor in the algorithm. To reduce the running time, we can reduce the number of randomly sampled vertices that we use to estimate the lightness. While the choice in the work of [6] is aimed at ensuring a high probability of success, it is conceivable that on practical instances we need less vertices for a good estimation. To this end, we introduce a parameter  $K \in \mathbb{N}$  to our algorithm such that we choose  $1/K$ -times as many vertices to estimate the lightness in the `DECOMPOSE` routine.

We also have to decide at which point we want to invoke the base case and call `LAZYDIJKSTRA`. In [6] this is done when  $\kappa \leq 2$ , see Algorithm 4. While this is theoretically guaranteed to occur within a logarithmic number of recursive calls with sufficiently high probability, in practice this can slow down the implementation significantly. Hence, we settled on calling the base case if  $n + \kappa \leq 300$ . We use the sum as we want to only call the base case if both  $\kappa$  and  $n$  are small. Note that due to how we set  $\kappa$ , we also always have  $\kappa \leq n$ .

Finally, our implementation also performs negative cycle detection, however, this is not its strong suit. In [6] negative cycle detection is done via budgeting, i.e., when the algorithm exceeds a specific running time, then the existence of a negative cycle is reported; this technique is impractical and we do not use it. For an alternative approach, note that the only cases where negative cycles can occur in Algorithm 4 are the `LAZYDIJKSTRA` calls. Hence, we add negative cycle detection to `LAZYDIJKSTRA` by merely checking the number of distance improvements that any vertex sees. If this number exceeds the number of vertices, then we detected the existence of a negative cycle and we report this. How to adapt our implementation to also perform efficient negative cycle detection is an intriguing task for future work.<sup>7</sup>

Let us consider what asymptotic running time we can hope for our implementation. Note that we use the wording “hope for” on purpose, as the above modifications can potentially break the theoretical guarantees and our statements should be considered an intuitive

---

<sup>7</sup> We note that in a very recent parallel and independent work that appeared on arXiv, an improved cycle detection is presented, but this work does not present any implementation or experiments [25].

sketch. The theoretically expected running time for the version of Algorithm 4 in [6] is  $O((m + n \log \log n) \log^2 n)$ , where  $n$  and  $m$  are the number of vertices and edges in the input graph, respectively. We perform one crucial modification, that is, we use a 4-heap instead of the priority queue by [31]. This replaces the  $\log \log n$  factor in the above running time by a  $\log n$  factor. Hence, we can hope for an  $O((m + n \log n) \log^2 n)$  running time.

### 3.2 Other Algorithmic Choices

In our implementation, following [6], we use LAZYDIJKSTRA as algorithm to compute a valid potential function on graphs where each shortest path should only contain a small number of negative edges. This happens in the base case of Algorithm 4 as well as in the case of obtaining a valid potential function for the cut edges in-between components. Instead of LAZYDIJKSTRA we could actually employ any other shortest path algorithm for graphs with negative edge weights that is fast in practice. Especially, we can also use the baseline algorithm that we later compare to in Section 5 as subroutine.

## 4 Theoretical Insights

In Algorithm 4, we either make progress by reducing the component size or by reducing  $\kappa$ . In a restricted graph we know that  $\kappa \leq n$  as this is the maximal length of a simple path. This is sufficient in theory, but in practice we prefer a tighter and instance-dependent upper bound on  $\kappa$ . The following lemma establishes an upper bound on  $\kappa$  that comes from the diameter of  $G_{\geq 0}$ .

► **Lemma 3** (Diameter upper bound). *Given a restricted graph  $G$  that consists of a single strongly-connected component, we have that  $\kappa(G) \leq \text{diam}(G_{\geq 0})$ .*

**Proof.** We denote by  $w(P, G) := \sum_{e \in P} w_G(e)$  the cost of the path  $P$  in graph  $G$ . Let  $P$  be the path from  $u$  to  $v$  in  $G$  that realizes  $\kappa(G)$ . Hence,  $P$  has  $\kappa(G)$  negative edges and  $w(P, G) \leq 0$ . Let  $P'$  be the shortest path from  $v$  to  $u$  in  $G_{\geq 0}$ . Due to  $G$  being restricted (and hence having minimum cycle mean at least 1) and as  $P$  has  $\kappa(G)$  negative edges, we know that

$$w(P \circ P', G) \geq |P| + |P'| \geq \kappa(G).$$

Using the above statement and that  $w(P, G) \leq 0$ , we obtain that the diameter of  $G_{\geq 0}$  is at least

$$w(P', G_{\geq 0}) \geq w(P', G) \geq w(P \circ P', G) \geq \kappa(G),$$

hence  $\kappa(G) \leq \text{diam}(G_{\geq 0})$ . ◀

Note that it suffices to upper bound  $\kappa(G)$  on strongly connected components as we only ever use Algorithm 4 on strongly connected graphs, see Section 3.

## 5 Experiments

We now experimentally evaluate our implementation. To that end, we had the following challenges to solve:

- What algorithm(s) do we compare to? (Section 5.2)
- What data do we perform experiments on? (Section 5.3)
- What do we parameterize in our algorithm and how do we choose these parameters? (Section 5.4)

## 10:10 Algorithm Engineering of SSSP with Negative Edge Weights

In the following, we refer to our implementation as `OUR`. We conduct each experiment 5 times and present the average and the standard error of the mean, unless noted otherwise. For the quantitative evaluation of our experimental data, we use regression to the model  $a \cdot m^b$  by applying the common log-log-transform before an ordinary least squares fit from `statsmodels` (v.0.14.2) in Python. We report the errors of the slopes with a confidence interval of 95%. The error bars of single data points are reported with 1 standard error of the mean.

### 5.1 Code & Hardware

We wrote the entire experimental framework and all the algorithms in C++20. We compiled our code using GCC on Debian. We ran the experiments on a server with 48 Intel Xeon E5-2680 v3@2.50GHz CPUs with 256 GB of RAM. We note that we did not parallelize our implementation as our comparison is with a single-threaded implementation, but parallelization of our implementation is straight-forward (the recursive calls are independent).

### 5.2 Baseline Algorithm

We compare our implementation against a state-of-the-art algorithm by Goldberg and Radzik [21] called `GOR`. More specifically, we use the implementation in `SPLIB`<sup>8</sup>, which was developed by Cherkassky, Goldberg, and Radzik. We also implemented the algorithm called `BFCT` from [10] (a previous version was described in [11]). However, as our implementation was consistently outperformed by `GOR`, we only compare to `GOR` here. Our reason for choosing `GOR` is twofold:

- In [10], the authors show that `GOR` has overall the fastest average running times for scans of a single vertex and a good performance on all instances regarding the number of scans required to compute a shortest path tree, without large deviations.
- A highly-tuned implementation of `GOR` to compare to is publicly available.

### 5.3 Data

In the following, we present a brief overview of all the different types of instances we perform experiments on.

**BAD instances from [10].** These instances are worst-case instances that are taken from [10] and presented in Section A: `BAD-BFCT`, `BAD-DFS`, `BAD-GOR`, `BAD-RD`, and `BAD-RDB`. We take these instances as each of them is constructed to be pathological for a specific algorithm for shortest paths with negative edge weights. They are all DAGs, and the algorithms that they are adversarial constructions for are different variants of Bellman-Ford.

**Augmented BAD instances.** As all of the `BAD` instances are DAGs, and our implementation specifically exploits DAG structure<sup>9</sup>, a comparison merely on the `BAD` instances would not be reasonable and fair. However, these instances still reveal an interesting structure and the

---

<sup>8</sup> See <https://www3.cs.stonybrook.edu/~algorithm/implement/goldberg/distrib/splib.tar>. We modified it such that we can use it with our graph data structure and compile it within our implementation. We ensured that this did not incur a significant slowdown.

<sup>9</sup> We first decompose the graph into strongly connected components, which are singletons in the case of a DAG, and we then call `FIXDAGEDGES` on the set of components, which is a simple linear-time algorithm. Hence, our implementation degenerates to a simple algorithm to solve the shortest path problem on DAGs with negative edge weights.

work of [10] is the most comprehensive overview of shortest path algorithms on graphs with negative edge weights that we are aware of. Hence, we still use these instances, but to ensure that these instances are also interesting and difficult for our implementation, we perform two crucial modifications (denoting these instances by pre-fixing them with **AUG**, e.g., **AUG-GOR**).

- We randomly permute the vertex labels, as for many of the instances, the hardness for previous algorithms relies on the specific vertex ordering.
- We augment the graphs with 5 times as many edges as the original graph, which are selected uniformly at random using rejection sampling. We choose a constant factor as we aim to perform experiments on graphs with many vertices. Each of these new edges is assigned a large positive weight such that the minimum cycle mean remains at least 1, as we want to obtain restricted instances (see Section 2). As the initial graphs are DAGs, this weight is easy to determine: for any path  $P$  in  $G$ , we have to ensure that for the augmenting edge  $e$  that closes this path, we have

$$w(P \circ \{e\}) \geq |P| + 1 \iff w(e) \geq |P| + 1 - w(P),$$

and that all new edge weights are larger than the edge weights of the original graph. We simply set the edge weight of all augmenting edges to a single pre-determined value that ensures the above properties. Note that augmenting the graphs with these edges breaks the DAG structure while maintaining the shortest path structure, as desired.

Such augmentation introduces multiple strongly connected components inside the original instances, forcing all the subroutines of **OUR** to be run and guaranteeing executions that cover the implementation in its entirety. Note that the above modifications only ensure that the minimum cycle mean is at least 1, but we did not argue whether the instances fulfill the second condition of restrictedness (see Section 2), i.e., whether all edge weights are at least  $-1$ . Most of the instances are already restricted or require minimal modifications to be transformed into restricted instances. **BAD-BFCT** and **BAD-DFS** are restricted without modifications. The instances **BAD-RD** and **BAD-RDB** are originally not restricted, as both contain weight  $-2$  edges. However, since any weight  $-2$  edge is always preceded by a degree-2 vertex with only an incoming edge of weight 0, we set both of these edge weights to  $-1$ . After this modification both instances are restricted. The **BAD-GOR** instance is not restricted, however, the reason it is not restricted is only due to a single edge having weight  $-3\frac{n-1}{2}$ . We keep the instance unrestricted<sup>10</sup> as there is no straight-forward way to make it restricted, and it leads to insightful experimental results.

**SHIFT-GOR.** These instances result from studying the behavior of **OUR** with **GOR** used as subroutine instead of **LAZYDIJKSTRA** (see Section 3.2) on the **AUG-GOR** instances. Interestingly, we found that the potential shifts that our implementation creates (after the **FIXDAGEDGES** step) greatly deteriorates the performance of **GOR**, creating a novel, non-trivial, very hard instance for **GOR**. We extract the **SHIFT-GOR** instances as follows:

- Consider the largest SCC in the first recursion step of Algorithm 4, and consider the potentials before the execution of **LAZYDIJKSTRA** at the end of Algorithm 4.
- We add an explicit super source connected with zero weight edges to all vertices.
- We then change the edge weights according to the potentials, i.e.,  $w(u, v) \leftarrow w(u, v) + \phi(u) - \phi(v)$ .

<sup>10</sup> However, note that we still ensure that the minimum cycle mean is at least 1 after augmentation.

## 10:12 Algorithm Engineering of SSSP with Negative Edge Weights

Since the size of the SCCs depends on the allocation of the augmented edges during the generation of the AUG-GOR instance, the size of the SHIFT-GOR instances may vary slightly. Furthermore, note that the SHIFT-GOR instances are not restricted as they contain many edges with weight smaller than  $-1$ . Despite the non-restricted nature of these instances, we show that OUR outperforms GOR on them (see Section 5.5).

**RANDOM RESTRICTED instances.** We use these instances to analyze the behavior of OUR on random restricted graphs. How to generate such instances is not obvious as we want the graph to have edges with weight  $-1$  to make it interesting, while ensuring a minimum cycle mean of at least 1. We first generate the structure of the graph by inserting  $6n$  edges uniformly at random by rejection sampling. We set all edge weights to 2 initially, implying an initial minimum cycle mean of 2. We then iteratively run Dijkstra’s algorithm on yet unvisited vertices always choosing a random one as source. This procedure partitions the vertices into shortest path trees. We perform a potential shift by the distances, resulting in zero weight edges after potential application. We set the weight of the edges connecting any two shortest path trees to 0.<sup>11</sup> Finally, we decrease all the edge weights by 1. Now the minimum cycle mean is at least 1 and the minimum weights are at least  $-1$ .

**USA instances.** We also test our implementation on the “Full USA” distance graph instance from the 9th DIMACS challenge<sup>12</sup>. This is a graph with roughly 24 million nodes, 58 million edges, and all its edge weights are positive. To obtain a graph with negative edge weights for our experiments, we run a shortest path algorithm from the vertex with label zero and assign the distances as potentials to the vertices. Subsequently, we perform a random shift on the potentials by adding a random number in the range  $[0, W]$  to the potential of each vertex. We then apply the potentials to the edge weights. After the application, the edge weights on the shortest path tree are in the range  $[-W, W]$ ; for  $W = 1$  we obtain restricted instances.

### 5.4 Parameters

As discussed in Section 3, there are several choices to be made and parameters to be set. We list the interesting and non-obvious choices and parameters here and empirically evaluate them to understand what good choices are. Note that we set and evaluated the parameters in the order described here, i.e., the best choice of the first mentioned parameter is already used in the subsequent experiments. Hence, we chose the order below carefully and on purpose.

As described in Section 3.1, we can change the number of random vertices that we use to estimate the lightness in Algorithm 3. To that end, our implementation has a parameter  $K \in [1, \infty]$  such that we sample  $\max\{1, k/K\}$  random vertices, where  $k$  is the theoretical number of vertices as used in Algorithm 3. We tested OUR with  $K \in \{1, 5, 10, 20, 40, \infty\}$  on several large instances, see Table 1. The speed-up induced by this parameter is significant; it is up to almost a factor of 4. We notice that the worst choice is  $K = 1$ , i.e., the original value that is used in Algorithm 4. There is no clear best choice between  $K$  being 20, 40, or  $\infty$ . To strike a balance, we use  $K = 40$  in our subsequent experiments.

Next we consider whether we should use the improved upper bound on  $\kappa$  in our implementation, see Section 4 and Section 3.1. We call this method the *diameter upper bound* here. In almost all of our experiments we noticed a similar scaling behavior with the diameter

---

<sup>11</sup>This does not introduce negative cycles as these edges are always forward edges topologically.


<sup>12</sup><http://www.diag.uniroma1.it/~challenge9/>

upper bound and without it. On one hand, on small instances computing the diameter sometimes poses a significant running time overhead and makes it a bit slower in total. On the other hand, on the `RANDOM RESTRICTED` instances, using the diameter upper bound leads to a better scaling behavior and a speed-up of almost an order of magnitude, see Figure 1. Hence, despite small disadvantages on some instances, we decide to use it as it can massively speed-up some cases.

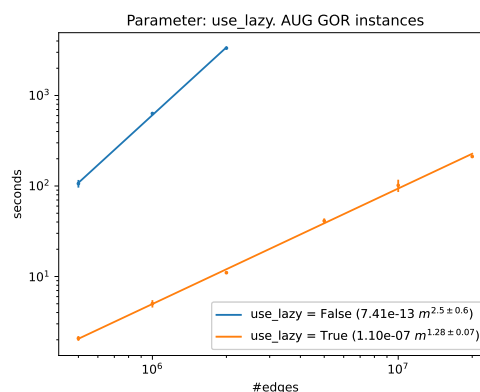
Finally, as discussed in Section 3.1, we can also replace `LAZYDIJKSTRA` by any other algorithm that performs shortest path computations on graphs with negative edge weights. The obvious replacement to evaluate is `GOR`. While the speed-up is minor when using `GOR` as replacement, the slowdown is massive on `AUG-GOR` instances, see Figure 2. Hence, we conclude that using `LAZYDIJKSTRA` as base algorithm is the more robust choice.

■ **Table 1**  $K$  on different instances with  $2 \cdot 10^7$  edges. Time is expressed in seconds and rounded to full seconds.

	$K = 1$	$K = 5$	$K = 10$	$K = 20$	$K = 40$	$K = \infty$
<code>BAD-DFS</code>	$124 \pm 5$	$48 \pm 3$	$56 \pm 2$	$39 \pm 2$	<b><math>35 \pm 1</math></b>	$38 \pm 2$
<code>BAD-BFCT</code>	$168 \pm 6$	$62 \pm 1$	$57 \pm 4$	<b><math>48 \pm 1</math></b>	$52 \pm 6$	$50 \pm 5$
<code>BAD-RD</code>	$164 \pm 5$	$62 \pm 1$	$64 \pm 2$	$48 \pm 2$	$49 \pm 3$	<b><math>43 \pm 2</math></b>
<code>BAD-RDB</code>	$126 \pm 4$	$51 \pm 3$	$47 \pm 3$	$40 \pm 1$	<b><math>37 \pm 1</math></b>	<b><math>37 \pm 1</math></b>
<code>BAD-GOR</code>	$260 \pm 20$	$270 \pm 20$	$250 \pm 30$	$230 \pm 30$	<b><math>203 \pm 3</math></b>	$240 \pm 20$



■ **Figure 1** Performance of OUR with and without the diameter upper bound on restricted random instances (see Section 5.3).




■ **Figure 2** Performance of OUR with `LAZYDIJKSTRA` and `GOR` on `AUG-GOR` (see Section 5.3). This plot is the reason why we investigated on `SHIFT-GOR`.

## 5.5 Comparison


In this section we compare the performance of `OUR` and `GOR` on several instances. The size of the instances that we use ranges from  $5 \cdot 10^5$  edges to  $2 \cdot 10^7$  edges. We consider the total running times but also the scaling behavior of both algorithms.

Let us first consider the behavior of `OUR` and `GOR` on the `RANDOM RESTRICTED` instances, see Figure 4. The scaling behavior of `OUR` is only slightly better than `GOR`, however `OUR` is consistently slower by a significant factor. We note, however, that the overall running

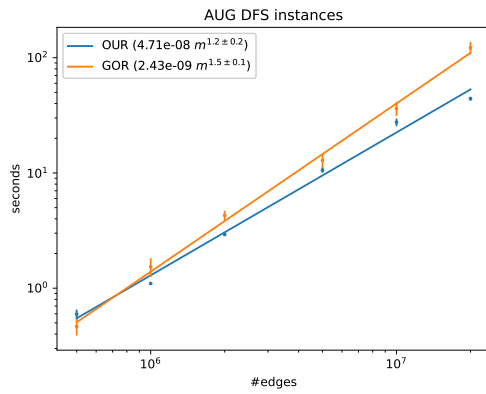
## 10:14 Algorithm Engineering of SSSP with Negative Edge Weights




(a) Performance of OUR vs GOR on AUG-GOR instances (see Section 5.3).



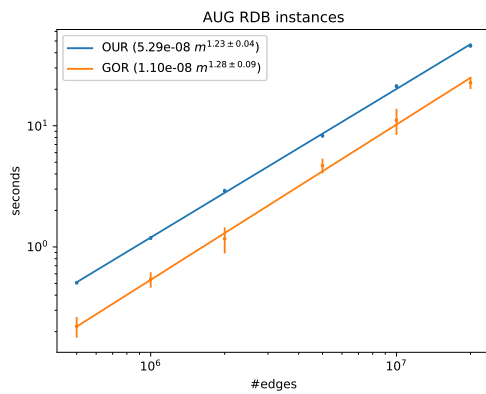
(b) Performance of OUR vs GOR on AUG-BFCT instances (see Section 5.3).




(c) Performance of OUR vs GOR on AUG-DFS instances (see Section 5.3).



(d) Performance of OUR vs GOR on AUG-RD instances (see Section 5.3).




(e) Performance of OUR vs GOR on AUG-RDB instances (see Section 5.3).



(f) Performance of OUR vs GOR on SHIFTED AUG-GOR instances (see Section 5.3).

■ **Figure 3** Our main experiments comparing the running time of OUR and GOR.



■ **Figure 4** Performance of OUR vs GOR on RANDOM RESTRICTED instances.

■ **Table 2** OUR vs GOR on BAD instances with  $10^6$  edges. Time is expressed in seconds.

	BAD-BFCT	BAD-DFS	BAD-GOR	BAD-RD	BAD-RDB
OUR	$1.14 \pm 0.03$	$0.71 \pm 0.03$	<b><math>0.95 \pm 0.02</math></b>	$0.90 \pm 0.03$	<b><math>0.87 \pm 0.02</math></b>
GOR	<b><math>0.014 \pm 0.003</math></b>	<b><math>0.016 \pm 0.002</math></b>	$720 \pm 20$	<b><math>0.028 \pm 0.002</math></b>	$260 \pm 3$

times for both algorithms on these instances are low compared to our other experiments on graphs with similar sizes. Let us now consider the DIMACS instance. We compare the running time of OUR and GOR on the unmodified instance, which we refer to as USA-0, and on the instance with a potential shift  $W$  (see Section 5.3) of 1, 10, and 100 (USA-1, USA-10, USA-100, respectively). On USA-0 OUR is 5.4 times slower. On USA-1 and USA-10 OUR is 7.1 times slower, and on USA-100 OUR is 6.1 times slower.

We now consider experiments on the BAD instances described in Section 5.3. As discussed in Section 5.3, a comparison on these instances is misleading as our implementation explicitly exploits the DAG structure while GOR does not. However, it is still enlightening to see for which instances GOR performs poorly. We show the results of these experiments in Table 2. We can see that OUR has a stable running time on these instances, as expected. Furthermore, we can clearly see the running time overhead of the data structures of OUR over the more lightweight GOR. However, the instances BAD-GOR and BAD-RDB are clearly pathological for GOR and it takes excessive time to solve these two instances, while the running time of OUR is similar to the one for other instances. We believe that an approach similar to ours that reduces the overhead induced by the data structures could become competitive to GOR on the instances where it is faster.

We now perform experiments on the AUG instances, see Figures 3a to 3e. First, note that on all AUG instances the scaling behavior of OUR is better than the scaling behavior of GOR. In particular, the regression results for the running times of OUR on all of the restricted instances (i.e., Figures 3b to 3e) are compatible with the theoretically expected near-linear running time: As described in Section 3.1, we could expect a running time of  $O((m + n \log n) \log^2 n)$ , which is  $O(n \log^3 n)$  for  $m \in O(n)$ . The same regression that we use for the running times applied to  $x \log^3(x)$  for sampled  $x \in [5 \cdot 10^5, 2 \cdot 10^7]$  leads to a polynomial factor of  $x^{1.20 \pm 0.01}$ , which is very close to the experimental results in the restricted cases. On the non-restricted augmented instance AUG-GOR, the running time of OUR is significantly worse than on the restricted instances, even though it is still significantly better than the running time of GOR.

When it comes to absolute running times, our experiments show a diverse picture. On the AUG-GOR and AUG-RD instances, OUR is always faster than GOR, see Figures 3a and 3d. However, on the AUG-BFCT and AUG-RDB instances, GOR is always faster than OUR, see Figures 3b and 3e. On AUG-DFS, OUR is slower on the smallest instance and consistently faster on all the other sizes, see Figure 3c. Summarizing, we want to stress that while GOR indeed is very fast on some instances, it also exhibits fluctuations of multiple orders of magnitudes on graphs of similar size due its adaptiveness; we do not see such fluctuations in our implementation when running it on restricted instances.

We now consider the experiments on the SHIFT-GOR instance, see Figure 3f. This instance is particularly hard for GOR, as it highlights its quadratic nature. Despite being far from restricted, OUR performs consistently better. OUR outperforms GOR by almost two orders of magnitudes on some of these non-trivial instances. It is to notice, however, that the running time of OUR is strongly unstable, probably because of the strong non-restricted nature of this instance. Each point in the plot in Figure 3f is the average of the running time of 5 different random generations of an instance of that size (and only one run per instance). This only marginally smoothed the unstable trend.

---

## References

- 1 Ravindra K. Ahuja, Kurt Mehlhorn, James B. Orlin, and Robert Endre Tarjan. Faster algorithms for the shortest path problem. *J. ACM*, 37(2):213–223, 1990. doi:10.1145/77600.77615.
- 2 Baruch Awerbuch, Andrew V. Goldberg, Michael Luby, and Serge A. Plotkin. Network decomposition and locality in distributed computation. In *30th Annual Symposium on Foundations of Computer Science, Research Triangle Park, North Carolina, USA, 30 October – 1 November 1989*, pages 364–369. IEEE Computer Society, 1989. doi:10.1109/SFCS.1989.63504.
- 3 Hannah Bast, Daniel Delling, Andrew Goldberg, Matthias Müller-Hannemann, Thomas Pajor, Peter Sanders, Dorothea Wagner, and Renato F. Werneck. Route planning in transportation networks. In Lasse Kliemann and Peter Sanders, editors, *Algorithm Engineering: Selected Results and Surveys*, pages 19–80. Springer International Publishing, Cham, 2016. doi:10.1007/978-3-319-49487-6\_2.
- 4 R. Bellman. On a routing problem. *Quarterly of Applied Mathematics*, 16:87–90, 1958.
- 5 Aaron Bernstein, Danupon Nanongkai, and Christian Wulff-Nilsen. Negative-weight single-source shortest paths in near-linear time. In *63rd IEEE Annual Symposium on Foundations of Computer Science, FOCS 2022, Denver, CO, USA, October 31 – November 3, 2022*, pages 600–611. IEEE, 2022. doi:10.1109/FOCS54457.2022.00063.
- 6 Karl Bringmann, Alejandro Cassis, and Nick Fischer. Negative-weight single-source shortest paths in near-linear time: Now faster! In *64th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2023, Santa Cruz, CA, USA, November 6-9, 2023*, pages 515–538. IEEE, 2023. doi:10.1109/FOCS57990.2023.00038.
- 7 Karl Bringmann, Alejandro Cassis, and Nick Fischer. Negative-weight single-source shortest paths in near-linear time: Now faster! *CoRR*, abs/2304.05279, 2023. doi:10.48550/arXiv.2304.05279.
- 8 Nitin Chandrachoodan, Shuvra S. Bhattacharyya, and K. J. Ray Liu. Adaptive negative cycle detection in dynamic graphs. In *Proceedings of the 2001 International Symposium on Circuits and Systems, ISCAS 2001, Sydney, Australia, May 6-9, 2001*, pages 163–166. IEEE, 2001. doi:10.1109/ISCAS.2001.922010.
- 9 Li Chen, Rasmus Kyng, Yang P. Liu, Richard Peng, Maximilian Probst Gutenberg, and Sushant Sachdeva. Maximum flow and minimum-cost flow in almost-linear time. In *63rd IEEE Annual Symposium on Foundations of Computer Science, FOCS 2022, Denver, CO, USA, October 31 – November 3, 2022*, pages 612–623. IEEE, 2022. doi:10.1109/FOCS54457.2022.00064.

- 10 Boris V. Cherkassky, Loukas Georgiadis, Andrew V. Goldberg, Robert Endre Tarjan, and Renato Fonseca F. Werneck. Shortest-path feasibility algorithms: An experimental evaluation. *ACM J. Exp. Algorithmics*, 14, 2009. doi:10.1145/1498698.1537602.
- 11 Boris V. Cherkassky and Andrew V. Goldberg. Negative-cycle detection algorithms. *Math. Program.*, 85(2):277–311, 1999. doi:10.1007/S101070050058.
- 12 Boris V. Cherkassky, Andrew V. Goldberg, and Craig Silverstein. Buckets, heaps, lists, and monotone priority queues. In Michael E. Saks, editor, *Proceedings of the Eighth Annual ACM-SIAM Symposium on Discrete Algorithms, 5-7 January 1997, New Orleans, Louisiana, USA*, pages 83–92. ACM/SIAM, 1997. URL: <http://dl.acm.org/citation.cfm?id=314161.314187>.
- 13 R. Dial, Fred W. Glover, David Karney, and Darwin Klingman. A computational analysis of alternative algorithms and labeling techniques for finding shortest path trees. *Networks*, 9(3):215–248, 1979. doi:10.1002/NET.3230090304.
- 14 Edsger W. Dijkstra. A note on two problems in connexion with graphs. *Numerische Mathematik*, 1:269–271, 1959. doi:10.1007/BF01386390.
- 15 Jack Edmonds and Richard M. Karp. Theoretical improvements in algorithmic efficiency for network flow problems. *J. ACM*, 19(2):248–264, 1972. doi:10.1145/321694.321699.
- 16 Chris Edwards. Historic algorithms help unlock shortest-path problem breakthrough. *Commun. ACM*, 66(9):10–12, August 2023. doi:10.1145/3607866.
- 17 L.R. Ford. *Network Flow Theory*. Number n° 923 in Network Flow Theory. Rand Corporation, 1956.
- 18 Michael L. Fredman and Robert Endre Tarjan. Fibonacci heaps and their uses in improved network optimization algorithms. *J. ACM*, 34(3):596–615, 1987. doi:10.1145/28869.28874.
- 19 Andrew V. Goldberg. Scaling algorithms for the shortest paths problem. *SIAM J. Comput.*, 24(3):494–504, 1995. doi:10.1137/S0097539792231179.
- 20 Andrew V. Goldberg, Haim Kaplan, and Renato F. Werneck. Reach for a\*: Shortest path algorithms with preprocessing. In Camil Demetrescu, Andrew V. Goldberg, and David S. Johnson, editors, *The Shortest Path Problem, Proceedings of a DIMACS Workshop, Piscataway, New Jersey, USA, November 13-14, 2006*, volume 74 of *DIMACS Series in Discrete Mathematics and Theoretical Computer Science*, pages 93–139. DIMACS/AMS, 2006. doi:10.1090/DIMACS/074/05.
- 21 Andrew V. Goldberg and Tomasz Radzik. A heuristic improvement of the bellman-ford algorithm. *Applied Mathematics Letters*, 6(3):3–6, 1993. doi:10.1016/0893-9659(93)90022-F.
- 22 Donald B. Johnson. Efficient algorithms for shortest paths in sparse networks. *J. ACM*, 24(1):1–13, 1977. doi:10.1145/321992.321993.
- 23 Nithin Kavi. Partial implementation of max flow and min cost flow in almost-linear time, 2024. doi:10.48550/arXiv.2407.10034.
- 24 Philip N. Klein, Satish Rao, Monika Rauch, and Sairam Subramanian. Faster shortest-path algorithms for planar graphs. In Frank Thomson Leighton and Michael T. Goodrich, editors, *Proceedings of the Twenty-Sixth Annual ACM Symposium on Theory of Computing, 23-25 May 1994, Montréal, Québec, Canada*, pages 27–37. ACM, 1994. doi:10.1145/195058.195092.
- 25 Jason Li and Connor Mowry. A bottom-up algorithm for negative-weight SSSP with integrated negative cycle finding. *CoRR*, abs/2411.19449, 2024. doi:10.48550/arXiv.2411.19449.
- 26 E.F. Moore. *The Shortest Path Through a Maze*. Bell Telephone System. Technical publications. monograph. Bell Telephone System., 1959.
- 27 Stefano Pallottino. Shortest-path methods: Complexity, interrelations and new propositions. *Networks*, 14(2):257–267, 1984. doi:10.1002/NET.3230140206.
- 28 U. Pape. Implementation and efficiency of moore-algorithms for the shortest route problem. *Math. Program.*, 7(1):212–222, 1974. doi:10.1007/BF01585517.
- 29 Piotr Sankowski. Shortest paths in matrix multiplication time. In Gerth Støltting Brodal and Stefano Leonardi, editors, *Algorithms – ESA 2005, 13th Annual European Symposium, Palma de Mallorca, Spain, October 3-6, 2005, Proceedings*, volume 3669 of *Lecture Notes in Computer Science*, pages 770–778. Springer, 2005. doi:10.1007/11561071\_68.

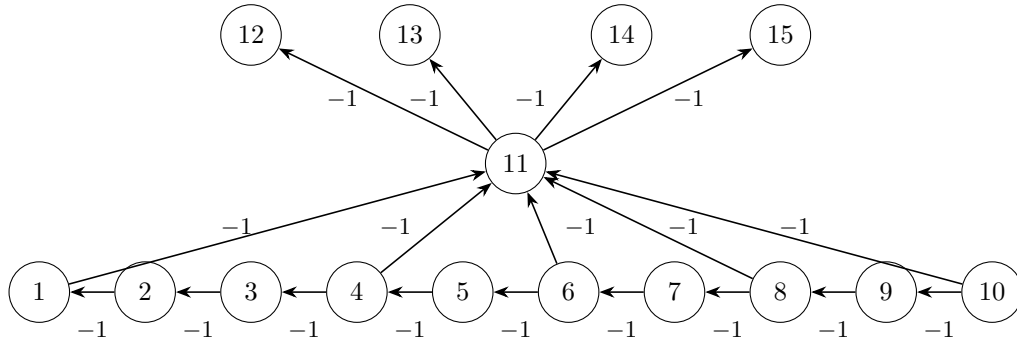
- 30 M. Sharir. A strong-connectivity algorithm and its applications in data flow analysis. *Computers & Mathematics with Applications*, 7(1):67–72, 1981. doi:10.1016/0898-1221(81)90008-0.
- 31 Mikkel Thorup. Integer priority queues with decrease key in constant time and the single source shortest paths problem. In Lawrence L. Larmore and Michel X. Goemans, editors, *Proceedings of the 35th Annual ACM Symposium on Theory of Computing, June 9-11, 2003, San Diego, CA, USA*, pages 149–158. ACM, 2003. doi:10.1145/780542.780566.
- 32 Peter van Emde Boas, R. Kaas, and E. Zijlstra. Design and implementation of an efficient priority queue. *Math. Syst. Theory*, 10:99–127, 1977. doi:10.1007/BF01683268.

**A BAD Instances**

In this section, we provide a description of the instances that we use in our experimental evaluation that are not yet described in Section 5. We note that many of these graphs were first described in [10] and we add the descriptions of these graphs here for completeness. In these cases, we closely follow the description of [10]. We refer to the vertices by their indices.

**BAD-BFCT**

The graph BAD-BFCT contains  $n = 4k - 1$  vertices and  $m = 5k - 3$  edges, where  $k$  is given as a parameter. The vertices 1 to  $3k - 2$  together with the edges  $(i + 1, i)$ ,  $1 \leq i \leq 3k - 3$ , form a path  $P$ . Every third vertex on  $P$  is connected to vertex  $3k - 1$ , that is, we have the edges  $(3(i - 1) + 1, 3k - 1)$  for  $1 \leq i \leq k$ . Finally, vertex  $3k - 1$  is connected to vertices  $3k$  to  $4k - 1$ . All edges in this graph have weight  $-1$ . Figure 5 gives an example for  $k = 4$ .




■ **Figure 5** Schematic representation of BAD-BFCT instances with 15 nodes and 18 edges.

**BAD-GOR**


The graph BAD-GOR consists of a path  $P$  of  $k$  vertices  $1, \dots, k$ , a vertex  $k + 1$  with  $k$  incoming and  $k$  outgoing edges, and  $k$  vertices  $k + 2, \dots, 2k + 1$ . The weights are  $w(1, 2) = -3k$ ,  $w(1, k + 1) = -1$ , and  $w(i, i + 1) = 1$  for  $2 \leq i \leq k - 1$ . Also,  $w(k + 1, k + 1 + i) = -1$  for  $1 \leq i \leq k$ , and  $w(i, k + 1) = 2(k - i)$  for  $2 \leq i \leq k$ . We have  $n = 2k + 1$  and  $m = 3k - 1$ . Figure 6 gives an example for  $k = 7$ .

**BAD-RD**

The graph BAD-RD consists of  $k$  edges  $(x_i, y_i) := (2i - 1, 2i)$  for  $1 \leq i \leq k$ , together with the edges  $(x_i, x_i + 1)$  and  $(y_i, x_i + 1)$  for  $1 \leq i < k$ . We set  $w(x_i, y_i) = 0$ ,  $w(y_i, x_i + 1) = -2$ , and  $w(x_i, x_i + 1) = -1$ . Figure 7 gives an example for  $k = 5$ .




■ **Figure 6** Schematic representation of BAD-GOR instances with 15 nodes and 20 edges.



■ **Figure 7** Schematic representation of BAD-RD instances with 10 nodes and 13 edges.

**BAD-RDB**

The graph BAD-RDB is obtained by modifying BAD-RD. We connect each even vertex  $y_i$  to a new vertex  $2k + 1$  and connect  $2k + 1$  to  $k$  new vertices  $2k + 2, \dots, 3k + 1$ . All new edges have weight  $-1$ . The new graph has  $n = 3k + 1$  vertices and  $m = 5k - 2$  edges. Figure 8 gives an example for  $k = 5$ .




■ **Figure 8** Schematic representation of BAD-RDB instances with 16 nodes and 23 edges.

**BAD-DFS**

The graph BAD-DFS is also obtained by modifying BAD-RD. We add the edges  $(y_i, y_i + 1)$  for  $1 \leq i < k$ . We set all edge weights to  $-1$ . We also relabeled nodes such that the index of the nodes below are all smaller than the indices above. Figure 9 gives an example for  $k = 5$ .

10:20 Algorithm Engineering of SSSP with Negative Edge Weights



■ **Figure 9** Schematic representation of BAD-DFS instances with 10 nodes and 17 edges.