

Shortest Distances as Enumeration Problem^{*}

Katrin Casel^a, Tobias Friedrich^a, Stefan Neubert^{a,1,*}, Markus L. Schmid^{b,2}

^a*Hasso Plattner Institute, University of Potsdam, Prof.-Dr.-Helmert-Str.
2-3, Potsdam, 14482, Germany*

^b*Humboldt University, Unter den Linden 6, Berlin, 10099, Germany*

Abstract

We investigate the single source shortest distance (SSSD) and all pairs shortest distance (APSD) problems as enumeration problems (on unweighted and integer weighted graphs), meaning that the elements $(u, v, d(u, v))$ – where u and v are vertices with shortest distance $d(u, v)$ – are produced and listed one by one without repetition. The performance is measured in the RAM model of computation with respect to preprocessing time and delay, i. e., the maximum time that elapses between two consecutive outputs. This point of view reveals that specific types of output (e. g., excluding the non-reachable pairs (u, v, ∞) , or excluding the self-distances $(u, u, 0)$) and the order of enumeration (e. g., sorted by distance, sorted row-wise with respect to the distance matrix) have a huge impact on the complexity of APSD while they appear to have no effect on SSSD.

In particular, we show for APSD that enumeration without output restrictions is possible with delay in the order of the average degree. Excluding non-reachable pairs, or requesting the output to be sorted by distance, in-

*Declarations of interest: none

*Corresponding author

Email addresses: Katrin.Casel@hpi.de (Katrin Casel), Tobias.Friedrich@hpi.de (Tobias Friedrich), Stefan.Neubert@hpi.de (Stefan Neubert), MLSchmid@MLSchmid.de (Markus L. Schmid)

URL: <https://hpi.de/friedrich/people/katrin-casel.html> (Katrin Casel), <https://hpi.de/friedrich/people/tobias-friedrich.html> (Tobias Friedrich), <https://hpi.de/friedrich/people/stefan-neubert.html> (Stefan Neubert), <http://www.mlschmid.de> (Markus L. Schmid)

¹Supported by the HPI Research School on Data Science and Engineering.

²Supported by the German Research Foundation – project number 416776735 (gefördert durch die Deutsche Forschungsgemeinschaft (DFG) – Projektnummer 416776735)

creases this delay to the order of the maximum degree. Further, for weighted graphs, a delay in the order of the average degree is also not possible without preprocessing or considering self-distances as output. In contrast, for SSSD we find that a delay in the order of the maximum degree without preprocessing is attainable *and* unavoidable for any of these requirements.

Keywords: Enumeration, shortest paths, APSP, fine-grained complexity

1. Introduction

Computing the distance (and a corresponding shortest path) for each two vertices of a given (directed or undirected, weighted or unweighted) graph is a fundamental computational problem. Its theoretical research spans over more than half a century and many of the main results are nowadays part of a typical undergraduate computer science curriculum (see [1]). Nevertheless, even for such well-researched problems, a slight change in perspective can lead to new algorithmic challenges and reveal new interesting aspects. As an example, the setting where the underlying graph is subject to constant change (e.g. insertions or deletions of edges) and the resulting task of maintaining the shortest distance information without recomputing it from scratch has given rise to an important and well-established research area within the field of dynamic graph algorithms (see, e.g., the surveys [2, 3]).

Our paradigm shift with respect to computing shortest distances lies in treating it as enumeration problem: we want to enumerate all *distance triples* $(u, v, d(u, v))$ (where u and v are vertices and $d(u, v)$ their distance) one by one, without repetitions. The performance of algorithms solving this task is measured with respect to *delay* (i.e., worst-case time between any two outputs, where the end-of-computation is also seen as an output) and *preprocessing time* (i.e., additional time needed before the enumeration starts). While this general algorithmic framework has practical motivation and is heavily applied in other areas (see discussion of related work further below), we primarily use it in order to reveal new complexity aspects of shortest distance problems. More specifically, we consider the problems *single source shortest distance* (SSSD) and *all pairs shortest distance* (APSD).

Many algorithms for SSSD or APSD work by using some already computed distances in order to obtain new ones. So, although not explicitly formulated as enumeration algorithms, such strategies nevertheless internally perform some kind of enumeration in the sense that at any point of the computation

some distances are already known, while others are yet to be determined. It is well known that the order of such hidden enumeration for shortest distances can have a crucial impact. In particular, Dijkstra’s algorithm fixes distances in increasing order and therefore is, in the comparison-model, limited by the lower bounds on sorting.

Explicitly treating SSSD and APSD as enumeration problems provides a fine-grained tool to reveal new and interesting aspects of computing shortest distances; allowing for a taxonomy of shortest distance algorithms according to their enumeration behavior. For example, our results show that even in the unweighted case the maximum degree of the graph is an upper and (unconditional) lower bound for the delay for SSSD, even if no enumeration order is required. This bound consequently also holds for APSD if the distances are required to be enumerated “row-wise” (partially ordered by their first vertex). However, if we drop this requirement for APSD, we can achieve algorithms with a delay in the order of the average degree of the graph. While we explicitly chose to enumerate distances instead of paths, all our algorithms can be altered, without changing their performance, to also give pointers to the second vertex on a shortest path to enable recovering shortest paths from the output.

1.1. Related Work

We use n and m to denote the number of vertices and edges of the input graph, respectively.

Most algorithms for SSSD on integer-weighted graphs, although they aim for a best possible total time, internally fix distances one by one. The famous Dijkstra algorithm [4], and all its many subsequent improvements (see [5] for a survey), fix distances one by one, in increasing order. A sequence of papers achieved significant improvements by explicitly breaking from this order constraint: Thorup developed a linear time (i.e., $O(m+n)$) algorithm for SSSD on undirected graphs with integer weights in [6], and later generalized this to an almost linear time algorithm for directed graphs in [7]. These ideas were later also used to develop algorithms for APSD which beat n runs of Dijkstra’s algorithm even for real-valued weights by Pettie and Ramachandran [8] for undirected, and Pettie [9] for directed graphs. In a sense, these results have been achieved by changing the enumeration order. In our enumeration point of view (which is, to the best of our knowledge, a new perspective w.r.t. shortest distance problems), we therefore also put a focus on the role of orders.

For APSD, there seem to be more algorithmic approaches that do not fix distances one at a time. Probably the first thing that comes to mind in this regard is the classical Floyd-Warshall algorithm [10, 11], followed by the algorithms based on fast matrix multiplication by Seidel [12] and Alon et al. [13]. But also other approaches are more successful with non-trivial steps before fixing distances. Chan [14], for example, shows how to solve APSD in unweighted, undirected graphs in $o(nm)$ by first heavily transforming the instance. Hagerup [15] also uses a non-trivial structure he calls component tree. For APSD enumeration, such strategies, at least at first glance, result in large preprocessing time. We will however see a curious case where fast matrix multiplication may be used to improve the delay.

On the side of lower bounds, we like to mention the tight $\Theta(nm)$ bound by Karger et al. [16] for APSD with integer weights and so called “path-comparison based” algorithms, that only use the edge weight function to compare path weights. Further, the lack of progress for APSD with general weights has lead to assuming the impossibility of significant improvement (i. e., a truly subcubic algorithm) in the form of the *APSP conjecture* [17].

Finally, let us discuss some aspects of the general enumeration framework adopted here. In the literature, the term *enumeration* is also used for the task of listing *all* solutions to an instance of a combinatorial (optimization) problem, e. g., listing all maximal independent sets in a graph [18]. Our understanding of the term enumeration is rather “enumerating *all parts of a solution*” instead of “enumerating *all solutions*”. This perspective is mainly motivated by practical scenarios where a user is constantly *querying* solutions for rapidly changing problem instances and where the need for seeing the (potentially large) solution in its entirety is negligible compared to the need for repeatedly accessing small portions of it very fast.

This also explains why in query evaluation in database theory this enumeration perspective is quite common and, over roughly the last decade, has even become a standard algorithmic setting, see [19, 20, 21] for more details. Another interesting observation is that despite extensive work on enumeration problems in database theory, lower bounds on the delay are rather scarce in this area. In fact, the few conditional lower bounds that do exist follow directly from the fact that some problems with lower bound assumptions can be encoded as getting the first element only, or computing the whole solution set; see [19, Section 6]. In this regard, it seems noteworthy that in our study of shortest distances as enumeration problems, we are able to show an unconditional lower bound. Note, however, that in the database

theory setting we usually assume a linear time preprocessing phase as a prerequisite for the enumeration algorithms. Whilst, in some cases, we are able to lower the delay by making use of a preprocessing phase, we are in this work primarily interested in enumeration algorithms without preprocessing and our lower bound is tailored to this setting.

1.2. Our Contribution

In this paper we analyze enumeration of APSD for graphs with unweighted or non-negative integer weighted edges. We present algorithmic upper and lower bounds for five output types:

- *Unconstrained APSD enumeration*: Algorithms have to output all elements $(u, v, d(u, v))$ with $u, v \in V$. This precisely mimics the standard APSD problem.
- *Row-wise-APSD enumeration*: The elements $(u, v, d(u, v))$ have to be enumerated sorted by the source vertices u . This is essentially equivalent to consecutive runs of SSSD enumeration.
- *No-self-APSD enumeration*: The elements $(u, u, 0)$ have to be omitted in the enumeration.
- *Reachable-APSD enumeration*: The elements (u, v, ∞) have to be omitted in the enumeration. Note that such elements only exist if the graph is not (strongly) connected.
- *Sorted-APSD enumeration*: The elements $(u, v, d(u, v))$ have to be enumerated sorted by the distance $d(u, v)$ in increasing order.

We analyze both the five individual types as well as all their combinations. The resulting upper bounds are summarized in Table 1. Likewise, Table 2 gives an overview of our lower bounds, some of which are unconditional whilst others are based on popular conjectures or on limited computational models. All bounds are given in terms of average $(\bar{\Delta})$ and maximum (Δ) vertex degree of the input graph. Combinations not listed explicitly are included in one of the other types: Adding the restrictions of the latter three output types to row-wise-APSD does not have any impact on the problem's complexity. Furthermore, sorted-APSD is a strictly stronger requirement than reachable-APSD.

Output Type	Unweighted Graphs		Non-negative Edge Weights	
	time	space	time	space
unconstrained	$(\overline{\Delta}, 1)$	$(-, n)$	$(\overline{\Delta} + \log(n), 1)$	$(-, n)$
row-wise (SSSD)	$(\underline{\Delta}, 1)$	(n, n)	$(\Delta + \log(n), 1)$	(n, n)
no-self	$(\overline{\Delta}, 1)$	$(-, n)$	$(\overline{\Delta} + \log(n), n)$	$(-, n)$
reachable	$(\Delta, 1)$	(n, n)	$(\Delta + \log(n), 1)$	(n, n)
reachable & no-self	(Δ, n)	$(-, n)$	$(\Delta + \log(n), n)$	$(-, n)$
sorted	$(\Delta, 1)$	(n^2, n^2)	$(\Delta + \log(n), 1)$	(n^2, n^2)
sorted & no-self	(Δ, \mathbf{n})	(n^2, n^2)	$(\Delta + \log(n), \mathbf{n} + \mathbf{m})$	(n^2, n^2)

Table 1: Upper bounds for APSD enumeration with different output types, where *row-wise* also transfers to SSSD. For each type we present time bounds (order of delay, order of preprocessing) and space bounds (order of lazy-initialized memory [or “-” if none is needed], order of overall space complexity). Bounds in **bold** are complemented with lower bounds, listed in more detail in Table 2. Results are clickable references to the respective theorem.

Output Type	Unweighted Graphs	Non-negative Edge Weights
unconstrained	-	$\Omega(\overline{\Delta})$ preprocessing or delay (p-c)
row-wise (SSSD)	$\Omega(\Delta)$ preprocessing or delay (u)	
no-self	-	$\Omega(\Delta)$ preprocessing or delay (u)
reachable (& no-self)	$\omega(\overline{\Delta})$ delay with preprocessing in $O(n + m)$ (BMM)	
sorted	$\omega(\overline{\Delta})$ delay with preprocessing in $O(n + m)$ (BMM)	
sorted & no-self	$\Omega(n)$ preprocessing or delay (u)	$\Omega(n + m)$ preprocessing or delay (u)

Table 2: Lower bounds for APSD enumeration with different output format, where *row-wise* also transfers to SSSD. Results are clickable references to the respective theorem. (p-c: path-comparison-based [16], u: unconditional, BMM: unless breakthrough in Boolean matrix multiplication)

In traditional asymptotic complexity analysis, the time for initializing memory for data structures, that are fully read or written later, is negligible and thus mostly ignored. This no longer applies if for computing the first output a data structure is needed, that cannot be initialized within preprocessing and one unit of delay. For those data structures we make use of lazy initialization (see, e.g., the textbook [22] and the detailed description in section 2). All our upper bounds specify the exact amount of lazy-initialized memory along with the overall space complexity. Lazy-initialized memory could be replaced with balanced search trees; the additional cost of $\log(n)$ steps for each read or write operation roughly increases the delay by a factor of $\log(n)$. All our lower bounds are independent of available memory.

Interestingly, whether or not edges are directed does not seem to alter

the difficulty of these enumeration problems; all our results hold for both undirected and directed graphs. In this regard, observe also that if we redefined the output of APSD enumeration to be distances for *unordered* pairs of vertices, the delay would only increase by a factor of at most 2. We discuss the details of this adapted definition in Appendix A.

We show a tight delay of $\Theta(\Delta)$ for unweighted row-wise-APSD enumeration without preprocessing by an unconditional adversary-based lower-bound, complemented by an algorithm meeting this delay. In contrast, for unconstrained- and no-self-APSD enumeration we achieve a delay in $O(\overline{\Delta})$. This means that deviation from the order in which a breadth-first search (BFS) fixes distances cannot be used to improve SSSD enumeration but does have an effect on APSD. For sorted- and reachable-APSD, an equal improvement would be surprising, as this would provide a quadratic-time algorithm for Boolean matrix multiplication (BMM).

Generally, we solve the weighted problems with additional $O(\log(n))$ cost on the delay compared to their unweighted counterparts, which is due to the cost of minimum extraction on standard priority queues. We discuss possible speedups through faster integer priority queues along with the respective algorithmic results in subsection 3.2 and Theorem 9.

In three considered cases it seems favorable to make use of additional preprocessing time. Two cases (sorted- & no-self- APSD enumeration) are backed up by a matching unconditional lower bound on either preprocessing or delay. The third case (weighted no-self-APSD enumeration) leaves a gap between the $\Omega(\Delta)$ lower bound and the $O(n)$ preprocessing time required for our algorithm. This and further open questions are summarized in section 7.

We start with the results for row-wise APSD enumeration (i. e., SSSD enumeration), in section 3, as these nicely illustrate the techniques and change of perspective required to approach APSD in the new light of enumeration. Then we drop the row-wise requirement (i. e., we consider actual APSD enumeration), which allows more freedom to distribute the overall computational effort over the enumerated elements in a clever way and therefore leads to better delays. We exemplify this idea in section 4 by using the trivial self-distances to initially buy a head start of computation time. In section 5, we extend this general approach to the technically more involved results for the constrained output types (including APSD enumeration without self-distances). Finally, the results on sorted APSD enumeration are discussed in section 6.

Before diving into the first results, we give some notation and the computational model.

2. Algorithmic Model and Basic Definitions

For all problems in this paper the input is a graph $G = (V, E)$ with n vertices $V = \{1, \dots, n\}$ and $|E| = m$ edges. The degree of a vertex v is denoted by $\deg(v)$ and we write Δ for the maximum and $\bar{\Delta}$ for the average degree in G . All proofs work on both undirected and directed graphs; by *degree* we always refer to the number of outgoing edges.

Unless otherwise stated, graphs are unweighted. In this case the shortest s - t -distance $d(s, t)$ is the minimum number of (directed) edge hops to reach t from s . For graphs with edge weights given by a function $w: E \rightarrow \{0, \dots, n^c\}$ for some constant $c \in \mathbb{N}$, $d(s, t)$ is the minimum sum of edge weights on a (directed) path from s to t . For a vertex t that is not reachable from vertex s we define $d(s, t) = \infty$. Note that, for both weighted and unweighted graphs, we have $d(s, s) = 0$ for every $s \in V$. We call the elements $(s, s, d(s, s))$ *self-distances*.

We study algorithms that *enumerate* APSD in the sense that for each s - t -distance the triple $(s, t, d(s, t))$ is emitted exactly once. We say that such an algorithm has *delay* in $O(f)$, if both the time to compute the first or next output and the time to detect that no further output exists, is in $O(f)$. On top of that we sometimes allow for additional *preprocessing time* the algorithm can spend before the first delay starts. Our analysis only considers preprocessing times asymptotically smaller than the total time of the algorithm, as otherwise the enumeration trivially degrades to solving the problem during the preprocessing phase and reading out the finished solution afterwards.

In order to minimize the worst-case delay, algorithms ought to spread out their computational effort as evenly as possible over all produced outputs. For this, algorithms are not required to emit an output immediately after it is computed. Quite the contrary: Holding back outputs until the current delay expires is crucial to minimize the worst-case delay of an algorithm. We analyze this behavior with a variant of the accounting method: The initial account balance is in the order of the preprocessing time plus one unit of delay. Every emitted output raises the account balance by at most the aspired delay; we say the output is *charged credit* for this increase; the unit of credit being some implementation specific constant. This credit is used to pay for all computation steps and must never become negative. As we expect the algorithm to actively emit each output after holding it back for the appropriate time, the delay (and thus all credit changes) must be computable

for the algorithm.

We perform our analyses in the word RAM model with word size in $\Omega(\log(n))$. This allows for all read/write and basic arithmetic operations such as comparison and addition of edge weights to be performed in constant time. Graphs are encoded as read-only adjacency lists with constant time access to the number of vertices and the degree of each vertex. This is for example supported by the *standard representation* by Kammer and Sajenko [23].

Algorithms can “lazily” reserve uninitialized memory within constant time. To be able to distinguish written memory cells and uninitialized ones in a data structure A , one level of indirection is introduced: The algorithm allocates auxiliary memory B of twice the size of A along with a counter c , starting at 0. B will hold pairs of pointers and actual values, the pairs being written from left to right, whilst c keeps track of the number of pairs in B . Instead of writing and reading from A directly, memory is used as follows: An initialized cell $A[x]$ is a pointer to a pair in B , which in turn holds a pointer to $A[x]$ and the actual value associated with $A[x]$. Uninitialized memory cells in A either point outside the valid area of B (as specified by c) or do not match with the reverse pointer stored in the referenced pair.

3. Single Source Shortest Distance Enumeration

In this section, we consider the following enumeration problem:

Problem Definition (Single Source Shortest Distance (SSSD) Enumeration).

Input: Graph $G = (V, E)$, vertex $s \in V$.

Output: For each $t \in V$ the shortest distance solution $(s, t, d(s, t))$.

Note that running an SSSD enumeration algorithm with delay t_d and preprocessing $t_p \in O(n \cdot t_d)$ for each vertex solves row-wise-APSD enumeration with the same delay, preprocessing and memory.

Conversely, consider an algorithm for row-wise-APSD enumeration with delay t_d and preprocessing t_p . We can stop a modified version of this algorithm after it produced the complete first row of the distance matrix to solve SSSD enumeration with the same delay and preprocessing: Normally, the first row of the APSD distance matrix includes the distances from vertex 1 to all vertices in the graph. However, vertex names are arbitrary. So by introducing one level of abstraction over the input data that, for each data access, swaps all

occurrences of vertex 1 and the given source vertex s , we effectively rename those two vertices with only constant-time overhead per access to the input.

As the above reductions transfer all lower and upper bounds between these two problems, we study SSSD enumeration instead of analyzing row-wise APSD.

3.1. Unweighted Graphs

We first consider the case of unweighted graphs, for which a single breadth-first search (BFS) started in vertex s solves SSSD in optimal total time $O(n + m)$ and with $\Theta(n)$ space. It turns out that a BFS can also be carried out in such a way that it produces the distance triples $(s, t, d(s, t))$ with a delay bounded by the maximum degree of the graph (Theorem 1). Moreover, we also show that this upper bound on the delay is essentially optimal (Theorem 2).

Theorem 1. *Unweighted SSSD enumeration can be solved with delay in $O(\Delta)$, without preprocessing, with $\Theta(n)$ lazy-initialized memory and space complexity in $\Theta(n)$.*

Proof. Our modified BFS version works as follows: It maintains a solution queue Q to collect computed distances for later output. Whenever the search visits a vertex t , it iterates over all its neighbors and, when indicated, updates their distance and notes them down to be visited later. Then, the algorithm appends $(s, t, d(s, t))$ to Q . After the BFS is finished, the algorithm iterates over the vertices and appends for all unvisited vertices v the solution (s, v, ∞) to Q . (Note that this final pass takes $\Theta(n)$ time while perhaps not producing further output if there are no unreachable vertices, which is why we actually need the solution queue.)

During the breadth-first search, the algorithm tracks the maximum node degree seen so far. We denote the observed maximum node degree after the BFS visited k vertices by Δ_k . For some implementation specific constant c and the respective current Δ_k , the algorithm emits the next solution from Q after $c \cdot \Delta_k \in O(\Delta)$ steps. Whenever the algorithm is required to emit an output, Q must not be empty. To prove this, we apply the accounting method, effectively tracking how many solutions are in Q and how much computation time can be spent before a new output must be emitted. For each credit unit, the algorithm can perform a constant number of steps. As long as the credit stays positive, Q is not empty.

By induction over the number k of vertices visited by the BFS we show that (1) the account balance is always non-negative, (2) it is at least $\Delta + k$

after the k -th vertex has been visited and (3) we achieve this by charging each output with at most $\Delta + 1$ credit. For $k = 0$ this is true, as the initial balance is Δ as specified in section 2. Let the claims be true for some $0 \leq k \leq n - 1$ and let t be the $(k + 1)$ st vertex visited by the BFS. It takes $O(\deg(t))$ steps to process the neighbors of t and to put the solution $(s, t, d(s, t))$ into Q , which is paid for with $\deg(t) \leq \Delta$ units. Recall that the algorithm tracks the maximum node degree Δ_{k+1} seen so far. The newly produced solution is charged $\Delta_{k+1} + 1$ units. As $\Delta_{k+1} \geq \deg(t)$ holds, the new credit balance is at least $\Delta + k + 1$ as claimed.

When the last vertex was visited by the BFS, the account balance is at least $\Delta + r$, where r is the number of vertices reachable from s . In the final iteration over all vertices, the algorithm uses this balance to skip the r already seen vertices. The remaining $n - r$ solutions with distance ∞ directly pay for the rest of the iteration.

As each output is charged at most $\Delta + 1$ credit and the account balance is never negative, the algorithm enumerates all solutions to SSSD with delay in $O(\Delta)$.

The BFS uses $\Theta(n)$ lazy-initialized memory to keep track of the distance, predecessor and visited state of the vertices. Both the internal queue of the BFS and the solution queue Q hold at most n entries. Thus, the overall space complexity of the algorithm is in $\Theta(n)$. \square

Before presenting the lower bound, we note that (on a connected graph) a BFS spends on average $O(\bar{\Delta})$ time per discovered vertex. Consequently, the question arises whether $O(\bar{\Delta})$ can also be achieved as the delay in the enumeration setting. The following lower bound, however, shows that, for no or little preprocessing time, this is not possible and that the delay of $O(\Delta)$ achieved by the algorithm of Theorem 1 is optimal. The intuitive idea of the lower bound is that any algorithm for SSSD enumeration can be forced to fully process first a potentially densely connected neighborhood of s with many edges but few solutions, before progressing to areas of the graph where less effort is needed to produce many solutions.

Theorem 2. *Unweighted SSSD enumeration cannot be solved with both preprocessing and delay in $o(\Delta)$, even if the graph is connected and has constant average degree.*

Proof. Assume some algorithm A was able to enumerate the distances with delay in $o(\Delta)$ and consider the following adversarial setup that is equivalent to

receiving the input graph as adjacency lists: A is allowed to ask the adversary two types of queries for a vertex v :

neighbor query Returns the next neighbor of v . (Where “next” refers to the adjacency list of v .)

degree query Returns the degree of v .

The adversary will, for arbitrary k , construct a graph consisting of a k -clique $C = \{v_1, \dots, v_k\}$ with the start vertex s in it. One edge of C is replaced by a path of length $k^2 + 1$ with k^2 inner vertices $P = \{u_1, \dots, u_{k^2}\}$. Figure 1 depicts the graph structure. Note, that the graph has $|V| = k + k^2$ vertices and $\frac{k(k-1)}{2} + k^2$ edges and thus constant average degree.

When asked for a vertex degree, the adversary consistently reports $\deg(v) = k - 1$ for all $v \in C$ and $\deg(u) = 2$ for all $u \in P$. For the first at least $k^2 - 3$ neighbor queries on a vertex in P , the adversary answers with yet another path vertex that extends the seen sub-path but is not adjacent to C . More precisely, the adversary can guarantee that the algorithm has only seen a continuous sub-path P' of P that has no connection to C . Whenever the algorithm makes a neighbor query on one of the end vertices of the P' , the adversary answers with a vertex that extends P' . For a neighbor query on a previously unseen path vertex, the adversary answers with an edge to an end vertex of P' . This way, after $i < k^2 - 3$ neighbor queries on vertices in P , P' has length at most $i + 1$ and there are at least two path vertices with unseen neighborhood. Let p_1, p_2 be two such vertices. Similarly, to the first at least $(k - 1)^2 - 1$ neighbor queries on a vertex in C , the adversary answers with neighbors that also lie in C . Until those queries are all completed, there are at least 2 clique vertices left, for which the algorithm has not queried the complete neighborhood. In particular, the algorithm is missing information on the potential connection between these clique vertices and p_1, p_2 . Let c_1, c_2 be two such clique vertices. In summary, the algorithm has to make at least $(k - 1)^2$ queries until it sees the first edge between one of c_1, c_2 and one of p_1, p_2 .

Now assume the algorithm emits for any $t \in P$ a shortest distance solution $(s, t, d(s, t))$ before having made at least $(k - 1)^2$ queries. As no adjacent edge of neither p_1 nor p_2 was fixed yet, the adversary is free to position the two vertices as shown in Figure 1 or instead move p_1 in-between c_2 and p_2 or move p_2 in-between c_1 and p_1 . For at least one of those three configurations, $(s, t, d(s, t))$ is not a correct solution: If $d(s, t)$ is the correct solution for

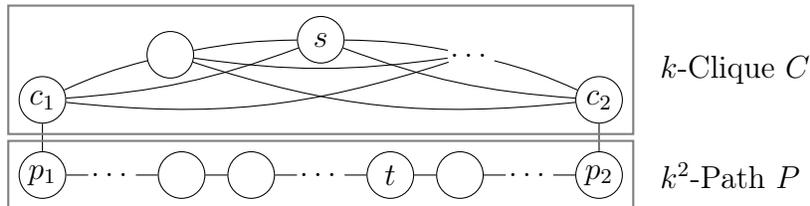


Figure 1: The adversary for Theorem 2 constructs a clique with one edge replaced by a long path.

the configuration shown in Figure 1 and a shortest path from s to t w.l.o.g. enters P through the vertices c_1, p_1 , then moving p_1 in-between c_2 and p_2 decreases the shortest distance by 1 and thereby renders the solution invalid. In the second case, $d(s, t)$ is incorrect for the arrangement in Figure 1 and the adversary can simply choose this configuration.

Thus, the algorithm needs $\Omega(k^2)$ queries to the adversary before being able to correctly determine a distance to any vertex from P . As there are only k vertices outside P (and hence only k outputs), this results in an enumeration delay of at least $\Omega(k)$. The constructed graph has maximum degree $\Delta = k - 1$, thus the delay is in $\Omega(\Delta)$. \square

3.2. Non-negative Edge Weights

Next, we consider graphs with non-negative integer edge weights. In this setting, Dijkstra's algorithm [4] using a Fibonacci Heap as priority queue [24] solves the SSSD problem in $O(m + n \cdot \log(n))$ with $\Theta(n)$ space. Similarly as we modified BFS for Theorem 1, we can extend Dijkstra's algorithm to an algorithm for SSSD enumeration for weighted graphs.

Theorem 3. *SSSD enumeration on graphs with edge weights in $\{0, \dots, n^c\}$, for some constant c , can be solved with delay in $O(\Delta + \log(n))$, without preprocessing, with $\Theta(n)$ lazy-initialized memory and space complexity in $\Theta(n)$.*

Proof. Our enumeration algorithm replaces BFS with Dijkstra's algorithm in Theorem 1. By using a priority queue such as Strict Fibonacci Heaps [25] with $O(\log(n))$ worst case time for EXTRACTMIN and $O(1)$ worst case time for MAKEHEAP, INSERT and DECREASEKEY, we get a delay of $O(\log(n))$ for determining the next vertex t and its shortest distance $d(s, t)$, and additionally $\Delta \cdot O(1)$ delay for processing the neighborhood of t . \square

We next argue that this is the best one can hope for when basing an enumeration algorithm without preprocessing on Dijkstra’s algorithm with a Fibonacci Heap. First, note that the lower bound in Theorem 2 also applies for weighted graphs. Consequently, if $\Delta \in \Omega(\log(n))$, then the delay $O(\Delta + \log(n)) = O(\Delta)$ is optimal. If, on the other hand, the maximum degree is smaller, then our enumeration version of Dijkstra’s algorithm has delay $O(\log(n))$, but also the number of edges is in $O(n \log(n))$, which means that in the non-enumeration setting, Dijkstra’s algorithm also needs on average $O(\log(n))$ time per computed distance triple.

There are several algorithms for SSSD with non-negative integer weights that outperform Dijkstra’s algorithm in the theoretical analysis, such as Thorup’s algorithm with a total runtime of $O(m + n(\log \log(n))^{1+\epsilon})$ in our setting. However, it is not clear whether such improvements indicate the existence of enumeration algorithms for SSSD with smaller delay. As they are usually based on priority queues with much higher initialization overhead, these improvements seem to at least require much higher preprocessing time. In contrast, such algorithms *can* be used to enumerate APSD faster without additional preprocessing time, as explained in subsection 4.2.

3.3. Constrained SSSD Enumeration

Any standard SSSD algorithm would produce as output an array of length n that includes the shortest distance from s to all vertices of the input graph, including s itself and all vertices unreachable from s . For graphs with only non-negative edge weights the former distance is always 0 (independent of the input graph) and can therefore be considered unnecessary information. The distances to unreachable vertices are all ∞ and, while the actual unreachable vertices obviously do depend on the input graph, they may be completely omitted from the output, since they are nevertheless implicitly given by omission (i. e., if we know that distances to unreachable vertices are not part of the output, they are all still uniquely represented by the output). In addition to omitting redundant information in the output, we could also require the distance triples to be sorted with respect to the distance. The question arises how these modifications influence the upper and lower delay bounds shown so far.

In fact, the next corollary states that the above mentioned modifications do not change our upper delay bounds (and this also holds for all possible combinations of them): Regarding the omission of self-distances or of distance

triples of unreachable vertices, our algorithms can just skip the single self-distance by doubling the delay and simply stop enumerating when all reachable vertices have been fully processed. Moreover, the requirement of producing the distance-triples sorted by the distance is already fulfilled by both BFS and Dijkstra’s algorithm. With the terminology defined in subsection 1.2 this can be summarized as follows.

Corollary 4. *No-self-, reachable- and sorted-SSSD enumeration as well as all combinations of the three output types can be solved with delay in $O(\Delta)$ for unweighted graphs and $O(\Delta + \log(n))$ for graphs with edge weights in $\{0, \dots, n^c\}$, for some constant c . Both results use no preprocessing, $\Theta(n)$ lazy-initialized memory and have an overall space complexity of $\Theta(n)$.*

With respect to lower bounds, we can first observe that Theorem 2 still applies if we omit self-distances, or distances of unreachable vertices, or both. For sorted enumeration on graphs with non-negative edge weights, we inherit an $\Omega(\log(n))$ lower bound on the delay in the addition-comparison model from the $\Omega(n \log(n))$ lower bound for comparison-based sorting, as enumerating the distances of a star center to its spikes sorted by distance is equivalent to sorting the edge weights.

Corollary 5. *In the addition-comparison model, no algorithm can solve sorted-SSSD enumeration on graphs with non-negative edge weights with preprocessing in $o(n \log(n))$ and delay in $o(\log(n))$.*

In summary, omitting the redundant self-distances or distances to unreachable vertices, or requiring sorted enumeration has no influence on SSSD enumeration, but, as we shall see in section 5 and section 6, the situation is quite different for APSD enumeration.

4. Unconstrained All Pairs Shortest Distance Enumeration

For APSD enumeration, we require that all entries of the distance matrix are listed:

Problem Definition (All Pairs Shortest Distance (APSD) Enumeration).

Input: Graph $G = (V, E)$.

Output: For each $s, t \in V$ the shortest distance solution $(s, t, d(s, t))$.

Again, we start with unweighted graphs and then extend to non-negative edge weights.

4.1. Unweighted Graphs

Since solving the APSD problem reduces to solving SSSD for each vertex $s \in V$ as start vertex, it is obvious that Theorem 1 extends to APSD, i. e., unconstrained APSD enumeration is possible without preprocessing and with delay in $O(\Delta)$. However, such an algorithm just performs n breadth-first searches in total time $O(n^2 + nm)$ and therefore spends on average $O(\bar{\Delta})$ time per produced distance triple. Hence, just like discussed for SSSD after the proof of Theorem 1, the question arises whether $O(\bar{\Delta})$ can also be achieved as delay for APSD enumeration. While Theorem 2 shows that this is impossible for SSSD, we can actually achieve delay $O(\bar{\Delta})$ for APSD enumeration (this means that for APSD enumeration, we can do better than just running the SSSD enumeration algorithm n times).

Theorem 6. *Unweighted APSD enumeration can be solved with delay in $O(\bar{\Delta})$, without preprocessing, and with space complexity in $\Theta(n)$.*

Proof. Our algorithm maintains a queue Q of known solutions from which it emits, one by one, a shortest distance whenever the current delay comes to an end. As first step, the algorithm iterates over all vertices, appending to Q all solutions with distance 0. All longer distances are found and appended to Q through breadth-first searches on all vertices. Because the algorithm initially does not know the average degree $\bar{\Delta}$, it emits the first $\frac{n}{2}$ solutions in Q with constant delay and in parallel computes $\bar{\Delta}$ by summing up node degrees and dividing by n . From then on, all further entries in Q will be emitted with $\Theta(\bar{\Delta})$ delay.

We again apply the accounting method to prove that Q is never empty when a solution has to be emitted. Clearly, the first $\frac{n}{2}$ solutions (distance 0) can be produced and emitted with constant delay. We assign a payment of 3 to each of those solutions, whilst the actual cost is 1, leaving a credit of n which is used to compute $\bar{\Delta}$ (recall that we assume constant access to the degree of a vertex, thus $\bar{\Delta}$ can be computed in $O(n)$ time). The remaining $\frac{n}{2}$ trivial solutions are also produced with an actual cost of 1 but are in addition charged $4 \cdot (\bar{\Delta} + 1)$, resulting in $2n(\bar{\Delta} + 1)$ credit. All further solutions produced by the breadth-first searches are assigned a cost of $\bar{\Delta}$.

We now show that the algorithm can run BFS on all vertices without running out of credit. For this we claim that, after the algorithm has performed a breadth-first search for each start vertex in the set $U \subseteq V$, the credit balance is at least $B(U) = |V \setminus U| \cdot (\bar{\Delta} + 1) + m + n$.

Recall that $m = n \cdot \bar{\Delta}$, thus the claim is true for $U = \emptyset$. Assuming the claim holds for any $U \subset V$ we inspect a breadth-first search on any vertex $s \in V \setminus U$. As the search costs $m + n$ and we have $B(U) > m + n$, the BFS can be performed without the solution queue running empty in-between. The search produces $n - 1$ new solutions which are each charged with a cost of $\bar{\Delta} + 1$. (Additionally, it skips one solution with distance 0 that has already been put into Q during initialization.) Thus we get $B(U \cup \{s\}) = B(U) - (m + n) + (n - 1) \cdot (\bar{\Delta} + 1)$ as claimed.

As each solution is assigned a cost in $O(\bar{\Delta})$, the algorithm solves APSD enumeration with delay in $O(\bar{\Delta})$.

By alternating between two sets of data structures for the breadth-first searches, only $\Theta(n)$ memory is needed for all BFS runs. During each BFS the algorithm can prepare the memory needed for the subsequent search; same holds for the first BFS whose memory can be prepared during the initial pass over the vertices. Note that the solution queue Q never has to hold more than $O(n)$ solutions, as for all $U \subseteq V$ we have $B(U)/\bar{\Delta} \in O(n)$. Thus, the algorithm has a space complexity of $\Theta(n)$. \square

The key observation for the proof of Theorem 6 is, that we can use the distance triples $(v, v, 0)$ in order to get a head start: These trivial solutions can be charged for the regular delay but only need constant time to be produced. This allows us to carry out the BFSs in such a way that the delay of $O(\bar{\Delta})$ is maintained. But just like the distance triples $(v, v, 0)$ can be used for this purpose, we could also use the distance triples $(u, v, 1)$ that are directly given by the edges $(u, v) \in E$. Assuming that we aim for maintaining a delay bound of D , this would give us a head start of $\Theta(m \cdot D)$ computation time. Now for dense graphs, m may be large enough such that we can fully solve the APSD problem within time $O(m \cdot D)$ by means of fast matrix multiplication (FMM), for example with the algorithm by Seidel [12], and afterwards emit the solutions with distance > 1 with constant delay. This observation is formally stated in the next remark.

Remark 7. Unweighted APSD enumeration can be solved without preprocessing and with delay in $O(\min(\bar{\Delta}, \frac{n^\omega \log(n)}{m}))$, where $\omega < 2.37286$ is the matrix multiplication constant [26].

4.2. Non-negative Edge Weights

We can tackle non-negative edge weights by again replacing BFS with Dijkstra's algorithm. In comparison to Theorem 6, we only need to adjust

the credit payments assigned to the solutions to match the aspired delay.

Theorem 8. *APSD enumeration on graphs with edge weights in $\{0, \dots, n^c\}$, for some constant c , can be solved with delay in $O(\bar{\Delta} + \log(n))$, without preprocessing, and with space complexity in $\Theta(n)$.*

Proof. If BFS is replaced by Dijkstra’s algorithm in Theorem 6, the analysis needs to be changed as follows:

The second half of the trivial solutions is charged, in addition to the constant cost for producing it, with $4(\bar{\Delta} + \log(n))$, resulting in $2n(\bar{\Delta} + \log(n))$ units on the account. We now have $B(U) = |V \setminus U| \cdot (\bar{\Delta} + \log(n)) + m + n \log(n)$. For $U \subset V$ this credit is enough to fully run Dijkstra’s algorithm with cost $m + n \log(n)$, producing $n - 1$ new solutions which are each charged with $\bar{\Delta} + \log(n)$ units. Thus, $B(U)$ is non-negative for all $U \subseteq V$ and the algorithm has delay in $O(\bar{\Delta} + \log(n))$. The space analysis from Theorem 6 applies without changes, resulting in $\Theta(n)$ space complexity. \square

In case $\bar{\Delta} \in \Omega(\log(n))$ holds, this delay is optimal in the path-comparison model, in which the edge weights must not be used for anything but comparing path weights, as it matches the $\Omega(mn)$ total time lower bound shown by Karger et al. [16]. For other models one can use improved SSSD algorithms instead of Dijkstra’s algorithm to achieve a better delay. If one allows the full power of word RAM *including* multiplication this goes as far as using Thorup’s $O(m)$ shortest paths algorithm for non-negative integer weights [6], which gives a delay of $O(\bar{\Delta})$ for APSD enumeration on undirected graphs.

Corollary 9. *Let T be the runtime of an SSSD algorithm such that T is computable in $O(n)$ time. Then APSD enumeration on graphs with edge weights in $\{0, \dots, n^c\}$, for some constant c , can be solved without preprocessing and with delay in $O(\frac{1}{n} \cdot T)$.*

5. Constrained All Pairs Shortest Distance Enumeration

Standard APSD algorithms produce a complete distance matrix and our previous enumeration algorithms mimic this by enumerating the values of all matrix cells in the form of all distance triples $(u, v, d(u, v))$. In the following, we consider the output restrictions discussed at the beginning of subsection 1.2, i. e., no-self-APSD enumeration (where the self-distances $(u, u, 0)$ are omitted) and reachable-APSD enumeration (where distance triples (u, v, ∞) are omitted).

5.1. No-Self-APSD Enumeration

In contrast to the SSSD enumeration with delay in $\Omega(\Delta)$, the APSD enumeration algorithms of section 4 buy some head start time by emitting all entries on the main diagonal of the matrix; which are the solutions that indicate the distance of a vertex to itself. For graphs with negative edge weights these entries are important, as they indicate whether the graph contains a cycle of negative total weight. With none or only non-negative weights however, they do not depend on the input graph and are thus somewhat uninteresting for the receiver of the enumerated solutions. So, what happens, if only distances between pairwise different vertices are to be emitted – in other words, what delay can we achieve for no-self-APSD enumeration?

In the unweighted case, each edge $(u, v) \in E$ of the input graph corresponds to a shortest distance solution $(u, v, 1)$; we have used this before in Theorem 7 for solving unweighted APSD enumeration via FMM. As it takes $\Omega(n)$ steps to find an edge in the worst case, this is not enough to ensure a delay in $O(\overline{\Delta})$. However, vertices with (out-)degree 0 admit for even more easy-to-compute solutions: For each $s \in V$ with $\deg(s) = 0$ and each $t \in V \setminus \{s\}$ there is no shortest path from s to t and therefore (s, t, ∞) is a solution. By modifying the algorithm of Theorem 6 to use these solutions instead of self-distances, we achieve the same delay in the constrained case.

Theorem 10. *Unweighted no-self-APSD enumeration can be solved with delay in $O(\overline{\Delta})$, without preprocessing, and with space complexity in $\Theta(n)$.*

Proof. The algorithm maintains a solution queue Q and a queue S of start vertices for later breadth-first searches. In the first step, the algorithm iterates over all vertices and edges, computes $\overline{\Delta}$ and distinguishes two cases for the current vertex s :

1. If $\deg(s) = 0$, append for each $t \in V \setminus \{s\}$ the solution (s, t, ∞) to Q .
2. Otherwise, append for each $(s, t) \in E$ the solution $(s, t, 1)$ to Q and append s to S .

Afterwards, the algorithm finds all remaining distances by running BFS on each vertex in S and from each such search emitting all solutions with distance > 1 . The first $\frac{n}{2}$ solutions are emitted with constant delay, all remaining ones with delay in $\Theta(\overline{\Delta})$.

During the first pass over all vertices and edges, the algorithm produces at least $\max(m, n)$ solutions, each within a constant number of steps. We again

assign a payment of 3 to the first $\frac{n}{2}$ solutions and pay for the computation of $\bar{\Delta}$ with the surplus n units of credit. The remaining at least $\max(m, n) - \frac{n}{2} \geq \frac{\max(m, n)}{2}$ initial solutions are assigned an additional payment of $4 \cdot (\bar{\Delta} + 1) + 2 \cdot (\bar{\Delta} + 1) + 4$ units. We claim that, after the algorithm has run a BFS for each start vertex in the set $U \subseteq S$, the credit balance is at least

$$B(U) = \sum_{v \in V \setminus U} \deg(v) \cdot (\bar{\Delta} + 1) + |V \setminus U| \cdot (\bar{\Delta} + 1) + m + n.$$

For $U = \emptyset$ this is true: Recall that, by the handshaking lemma, $2m = \sum_{v \in V} \deg(v)$ holds. Assuming the claim holds for any $U \subset V$ we inspect a breadth-first search with cost $m+n$ on any vertex $s \in S \setminus U$. As $B(U) > m+n$, the BFS can be performed without the solution queue running empty during the search. The BFS computes the distance to all vertices in the graph, producing $n - 1 - \deg(s)$ new solutions which are each charged with a cost of $\bar{\Delta} + 1$. (Additionally, the BFS computes the self-distance and $\deg(s)$ solutions with distance 1 that have already been put into Q during initialization.) Thus we get $B(U \cup \{s\}) = B(U) - (m+n) + (n - 1 - \deg(s)) \cdot (\bar{\Delta} + 1)$ as claimed.

Note that this algorithm produces up to $\Theta(n^2)$ solutions during the first iteration over all vertices and edges which would require $\Theta(n^2)$ memory for the solution queue. However, these solutions can be produced on demand whenever the queue runs empty which only requires a single pointer to track the iteration progress. Thus, as before in Theorem 6, the algorithm has an overall space complexity in $O(n)$. \square

For graphs with non-negative weights we can apply a similar idea: Not all edges form a shortest path, but at least for each vertex the outgoing edge with smallest weight does. It takes $\Omega(\Delta)$ time in the worst case to find such an edge, which is also the minimum time any algorithm must spend before it can possibly emit the first solution.

Theorem 11. *No algorithm can solve no-self-APSD enumeration on a weighted graph with both delay and preprocessing in $o(\Delta)$, even for connected graphs with constant average degree.*

Proof. We extend the adversarial setup from Theorem 2 with the graph structure shown in Figure 1 with the following additions:

- The first k vertices the algorithm receives any information on (through neighborhood or degree queries) are set to be from the clique C .

- With each query for a neighbor v of some vertex u , the algorithm also receives the edge weight $w(u, v)$.
- The first $k - 2$ edge weights reported by the adversary are set to 3.

Assume, an algorithm reports (s, t, d) as solution after at most $k - 2$ queries to the adversary. Clearly, the algorithm cannot emit a shortest distance solution containing some unseen vertex, thus $s, t \in C$ must hold. If $d \geq 3$, then the adversary sets all remaining edge weights to 1. As there must be a common neighbor v of both s and t of which the algorithm has not queried a single edge, the correct shortest distance would have been $(s, t, 2)$ with the shortest path (s, v, t) . For $d \leq 2$ all remaining edge weights are also set to 3 and $d(s, t) \geq 3$ holds, again making the output wrong. Therefore every algorithm needs to ask at least k queries before it can emit the first shortest distance. As $\Delta = k - 1$, an algorithm needs to have delay or preprocessing time in $\Omega(\Delta)$. \square

Whilst not fully closing the gap to this $\Omega(\Delta)$ lower bound on delay plus preprocessing, we can give an algorithm that uses $O(n)$ preprocessing time and then achieves the same delay as previous algorithms without output constraints: During preprocessing, the algorithm computes $\bar{\Delta}$ and sorts the vertices by increasing degree. In this order, the algorithm can afterwards find for each vertex the adjacent edge with minimum weight and put the corresponding solution into the solution queue. This way, the algorithm does not run into vertices with high degree early on, but can build some head start with prepared solutions first.

Theorem 12. *No-self-APSD enumeration on graphs with edge weights in $\{0, \dots, n^c\}$, for some constant c , can be solved with delay in $O(\bar{\Delta} + \log(n))$, preprocessing in $O(n)$ and with space complexity in $\Theta(n)$.*

Proof. During preprocessing, the algorithm sorts the vertices by increasing node degree and computes $\bar{\Delta}$. Again, two queues are maintained, Q for solutions and S for start vertices for Dijkstra's algorithm. In the enumeration phase the algorithm iterates over the sorted vertices and distinguishes two cases for the current vertex s :

1. If $\deg(s) = 0$, append for each $t \in V \setminus \{s\}$ the solution (s, t, ∞) to Q .
2. Else, find $t = \arg \min_{v \in V} \{ w(s, v) \mid (s, v) \in E \}$, append $(s, t, w(s, t))$ to Q and s to S .

The remaining distances are found by running Dijkstra’s algorithm on each vertex in S . In each such run, the respective edge from case 2 can be found again and the corresponding distance solution can be skipped. All solutions are emitted with delay in $\Theta(\bar{\Delta})$.

The preprocessing can be done in $O(n)$ time with $O(n)$ memory as follows. Initialize n empty lists, one for each possible degree. Scan through all vertices once, place vertices with degree i in list i . Going through these lists once then creates the required list of all vertices, sorted by their degree. As the algorithm then iterates over the vertices in order of their degree, the initial solutions can be produced with delay in $O(\bar{\Delta})$: In the first case, the algorithm produces $n - 1$ solutions with constant delay, in the second case it takes $O(\deg(s))$ time to compute one solution for vertex s . Let v_1, v_2, \dots, v_n be the vertices sorted by increasing (out-)degree. Then, for each $1 \leq j \leq n$, the inequality $\sum_{i=1}^j \deg(v_i) \leq j \cdot \bar{\Delta}$ holds. Thus, by charging each solution with $\bar{\Delta}$ units of credit, the computation effort for the expensive solutions can always be paid for; incorporating the ∞ -solutions only increases the head start.

Similar to Theorem 8 there are at least n initial solutions; each of those is charged additional $3(\bar{\Delta} + \log(n))$ units. This results in a credit balance of at least $B(U) = 2|V \setminus U| \cdot (\bar{\Delta} + \log(n)) + m + n \log(n)$. Each of the $n - 2$ new solutions per run of Dijkstra’s algorithm is charged $\bar{\Delta} + \log(n)$ units. Again, $B(U)$ is non-negative for all $U \subseteq V$ and the algorithm enumerates the solutions with delay in $O(\bar{\Delta} + \log(n))$.

By using the same trick as in Theorem 10, the initial solutions can be produced on demand, without memory-overhead. The space complexity of the remaining algorithm is $\Theta(n)$ as before in Theorem 8. \square

5.2. Reachable-APSD Enumeration

Both BFS and Dijkstra’s algorithm primarily deal with the vertices reachable from the start vertex. Thus, for the SSSD algorithms based on BFS and Dijkstra’s algorithm, the distance ∞ to unreachable and thus unvisited vertices is produced merely as a byproduct. However, restricting APSD enumeration to distances $< \infty$ (i. e., *reachable-APSD*, as defined in subsection 1.2) has a big impact on the delay. In fact, solving this restriction with delay in $O(\bar{\Delta})$ would imply a surprising breakthrough in Boolean Matrix Multiplication (BMM), as we can reduce BMM to reachable-APSD enumeration: We augment the standard graph representation of the $d \times d$ BMM instance with additional $O(d^2)$ vertices that reduce the average degree to a constant.

Enumerating the $O(d^2)$ reachable distances with delay in $O(\overline{\Delta})$ then solves BMM in $O(d^2)$ total time (since the average degree is constant).

Theorem 13. *No algorithm can solve reachable-APSD enumeration with delay in $O(\overline{\Delta})$ and preprocessing in $O(n+m)$ unless $d \times d$ BMM can be solved in $O(d^2)$ total time.*

Proof. We reduce Boolean Matrix Multiplication to reachable-APSD enumeration. Given two $d \times d$ Boolean matrices A and B we construct the standard graph representation of the BMM instance: The graph consists of three vertex partitions I, J, K , each with vertices $\{1, \dots, d\}$. Connect two vertices $i \in I, j \in J$ iff $A[i, j] = 1$ and two vertices $j \in J, k \in K$ iff $B[j, k] = 1$. Thus we have, for all $i \in I, k \in K$, the equivalence $(A \cdot B)[i, k] = 1 \Leftrightarrow d(i, k) = 2$.

We augment this graph with $2d^2$ isolated vertices. Note that the resulting graph has $n = 2d^2 + 3d$ vertices, $m \leq 2d^2$ edges and therefore constant average degree. Building the graph can be done in $O(d^2)$ time. The restricted APSD enumeration produces $O(d^2)$ solutions from which the resulting matrix $A \cdot B$ can be constructed in $O(d^2)$ time. Thus, an enumeration algorithm for reachable-APSD with delay in $O(\overline{\Delta})$ and preprocessing in $O(n+m)$ yields an $O(d^2)$ BMM algorithm. \square

Recall that by Theorem 4, reachable-(no-self)-SSSD can be solved with delay in $O(\Delta)$ on unweighted graphs and with delay in $O(\Delta + \log(n))$ on graphs with non-negative edge weights. Repeatedly applying these SSSD enumeration algorithms solves the corresponding APSD-enumeration with the same delay bounds. With the no-self restriction it may take up to $\Theta(n)$ steps to find the first edge.

Corollary 14. *Reachable-APSD enumeration can be solved with delay in $O(\Delta)$ for unweighted graphs and $O(\Delta + \log(n))$ for graphs with edge weights in $\{0, \dots, n^c\}$, for some constant c . Both results need no preprocessing, use $\Theta(n)$ lazy-initialized memory, and have an overall space complexity of $\Theta(n)$. Without self-distances, $O(n)$ additional preprocessing is needed, which in turn eliminates the need for lazy-initialized memory.*

Both algorithms have cubic total time for graphs with high maximum degree. The fastest so-called *combinatorial* algorithm to date multiplies $n \times n$ matrices in $O(n^3 / \log(n)^3)$ time (ignoring $\log \log$ factors) [27]. As BMM can be solved with FMM in $O(n^\omega)$ it is open whether reachable-APSD enumeration admits for a non-combinatorial algorithm with truly subcubic preprocessing time and delay in $o(\Delta)$.

6. Sorted-APSD Enumeration

Although the presented algorithms for SSSD enumeration happen to already produce output sorted by distance, running these algorithms in sequence as presented in section 4 does not preserve this output property and thus does not solve sorted-APSD enumeration.

As first observation, we note that a delay in $O(\overline{\Delta})$ (as we have for unweighted unconstrained APSD enumeration) would be surprising, as the conditional lower bound of Theorem 13 for reachable-APSD enumeration transfers to the sorted-APSD enumeration problem: By enumerating the solutions to APSD sorted by increasing distance and stopping the enumeration as soon as the first infinite distance is emitted, one solves reachable-APSD with a delay at least as good as sorted-APSD.

Corollary 15. *No algorithm can solve sorted-APSD enumeration with delay in $O(\overline{\Delta})$ and preprocessing in $O(n + m)$ unless $d \times d$ BMM can be solved in $O(d^2)$ total time.*

One can solve sorted-APSD enumeration by running the respective SSSD algorithms in parallel. In the unweighted case, the individual breadth-first searches are run in a cyclic fashion, each search being interrupted as soon as the reported next lowest hop distance increases.

Theorem 16. *Unweighted sorted-APSD enumeration can be solved with delay in $O(\Delta)$, without preprocessing, with $\Theta(n^2)$ lazy-initialized memory and space complexity in $\Theta(n^2)$.*

Proof. Similar to Theorem 6, the algorithm maintains a solution queue Q and begins by computing the trivial n solutions with distance 0 along with the maximum degree Δ ; emitting the first $\frac{n}{2}$ solutions in constant delay and the second half with delay in $O(\Delta)$. Afterwards it initializes on every vertex v a BFS B_v that maintains its private solution queue Q_v ; each BFS B_v is initially executed up to the point that it would emit the solution $(v, v, 0)$ from this queue. These solutions are discarded instead, as they have already been added to Q .

The algorithm keeps track of the next expected hop distance d , which is initialized to $d = 1$. All BFS instances are put in a queue B , which is processed as follows: As long as B is not empty, select the head entry B_s . If Q_s has an entry, move it to the end of Q and set d to the distance of this solution. Then, advance B_s to produce its next solution $(s, t, d(s, t))$

and put it in its private solution queue Q_s . If the distance is $d(s, t) = d$, then do nothing. Should $d(s, t)$ be larger, then move B_s to the end of B . If B_s is finished and thus has not produced a new solution, remove it from B . Afterwards, the loop over B continues.

Note that each BFS B_s is effectively executed until it emits its first solution of the next higher hop distance; this respective solution is held back until all other searches produced all remaining solutions with the current hop distance and B_s arrives again at the head of B . Thus, the algorithm enumerates all solutions sorted by increasing distance.

We use accounting to prove that this procedure indeed emits solutions from Q with delay in $O(\Delta)$. The initial n solutions are produced with constant delay, the first half being also emitted with constant delay. These solutions are used to pay for the degree computation. The second set of $\frac{n}{2}$ solutions is also produced with constant delay but is additionally charged $4 \cdot \Delta$ units, resulting in $2n\Delta$ credit. From that, $n\Delta$ credit is used to pay for the initialization of n BFS instances, the queue B and the distance tracker d .

Recall from Theorem 1 and Theorem 4 that unweighted SSSD can be solved without preprocessing and with delay in $O(\Delta)$ with an algorithm that produces the solutions sorted by increasing distance. Thus, each loop iteration costs Δ credit. Each respective first time a BFS B_s is the head entry of B , its private solution queue Q_s will be empty. These n iterations are paid for with the remaining $n\Delta$ credit from the initial solutions. For all other iterations, the solution being moved to Q is charged Δ which pays for all computation in this iteration.

As each solution is assigned a cost in $O(\Delta)$, the algorithm solves unweighted sorted-APSD enumeration with delay in $O(\Delta)$. \square

Without trivial self-distances, the first solution to be emitted corresponds to an arbitrary edge. As it takes up to $\Omega(n)$ time to find the first edge, this results in a lower bound on the delay or preprocessing of algorithms for sorted-no-self-APSD.

Theorem 17. *No algorithm can solve unweighted sorted-no-self-APSD enumeration with both delay and preprocessing in $o(n)$.*

Proof. An adversary can construct a graph in which two vertices are connected by an edge and the remaining $n - 2$ vertices are isolated. For the first $n - 2$ queries of the algorithm, the adversary only returns isolated vertices. Should the algorithm emit a solution $(s, t, 1)$ before having made at least $n - 2$

queries, the adversary can claim for at least one of s or t , that this vertex is isolated and thus must not be part of any solution. If the algorithm emits a solution with distance greater or smaller 1, it violates the sorted- or the no-self-requirement. Therefore, the algorithm has to make at least $\Omega(n)$ queries to the adversary before it can safely produce the first output. \square

This bound is tight in the sense that with $O(n)$ preprocessing one can achieve the same delay as for sorted-APSD *with* self-distances. Basically the preprocessing time is used to locate all edges; afterwards producing and emitting solutions works roughly as for Theorem 16.

Theorem 18. *With $O(n)$ preprocessing, unweighted sorted-no-self-APSD enumeration can be solved with delay in $O(\Delta)$, with $\Theta(n^2)$ lazy-initialized memory and space complexity in $\Theta(n^2)$.*

Proof. Within $O(n)$ preprocessing time the algorithm can divide the vertices into those with (out-)degree 0 and those with at least one outgoing edge; additionally it computes Δ . Afterwards it transforms the edges into solutions with hop distance 1 within constant time per solution, but charges additional $4 \cdot \Delta$ units for each. Of this credit, $\min(m, n) \cdot \Delta$ units are used to initialize BFS instances on all vertices with at least one outgoing edge; $3m\Delta$ units pay for the at most $\min(m, n) + 2m$ iterations in which these instances produce (unwanted or duplicate) solutions with distance ≤ 1 . When all BFS instances are fully finished, the algorithm emits all ∞ -distances of the vertices with (out-)degree 0 within constant time per solution. The rest of the analysis remains unchanged, therefore the algorithm solves unweighted sorted-no-self-APSD enumeration with $O(n)$ preprocessing and delay in $O(\Delta)$. \square

If the graph has non-negative edge weights, we again use Dijkstra's algorithm instead of BFS to compute the distances. Running n instances of Dijkstra's algorithm in a cyclic fashion (as done with BFS in the unweighted case) would not guarantee sorted output; instead a priority queue is used to choose the respective Dijkstra instance to be advanced next.

Theorem 19. *Sorted-APSD enumeration on graphs with edge weights in $\{0, \dots, n^c\}$, for some constant c , can be solved with delay in $O(\Delta + \log(n))$, without preprocessing, with $\Theta(n^2)$ lazy-initialized memory and space complexity in $\Theta(n^2)$.*

Proof. To derive this result, the algorithm and analysis from Theorem 16 are adjusted: The BFS instances are replaced with Dijkstra’s algorithm using Strict Fibonacci Heaps as data structure, increasing the time for both the initialization and each loop iteration to $O(\Delta + \log(n))$. Accordingly, every Δ units paid by solutions before are replaced with $\Delta + \log(n)$ units, thereby again covering all computation costs.

We replace the queue of BFS instances with a priority queue to select in each loop iteration the respective instance of Dijkstra to be advanced. The initial key of each priority queue entry is the minimum edge weight adjacent to the start vertex of this Dijkstra instance. In each loop iteration, the current instance is determined by EXTRACTMIN. If the instance produces a new solution, the instance is re-inserted into the priority queue with the distance of the new solution as key. Therefore, the key of each instance always represents the next solution this instance will contribute to the overall solution queue. Repeatedly selecting the minimum among all keys ensures that the algorithm enumerates the solutions sorted by increasing distance. By using a priority queue with both EXTRACTMIN and INSERT in $O(\log(n))$, those operations are dominated by the existing computation time per loop iteration.

As each solution is assigned a cost in $O(\Delta + \log(n))$, the algorithm solves sorted-APSD enumeration with non-negative edge weights with delay in $O(\Delta + \log(n))$. \square

Without trivial self-distances, the time until an algorithm can safely report the first solution increases to $\Omega(m + n)$, as this solution corresponds to the edge with minimum weight.

Theorem 20. *No algorithm can solve sorted-no-self-APSD enumeration on graphs with edge weights in $\{0, \dots, n^c\}$, for some constant c , with both delay and preprocessing in $o(m + n)$.*

Proof. An adversary can construct an arbitrary graph with n vertices and m edges. If there are isolated vertices, the adversary presents those first before revealing the first edge to the algorithm. For the first $m - 1$ neighborhood queries of the algorithm that result in an actual edge, the adversary returns 3 as edge weight. Thus, should the algorithm emit a solution (s, t, d) before having made at least $m + n$ queries, the adversary is free to choose the remaining edge weight as either 1 or 2 and thereby falsify the algorithm’s output: If $d \geq 2$, the adversary chooses edge weight 1 and the algorithm violates the sorted-requirement. For $d = 1$, the adversary sets the last edge

weight to 2 and thus s and t have at least distance 2. Therefore, the algorithm has to make at least $\Omega(m + n)$ queries to the adversary before it can safely produce the first solution. \square

Again we can match the delay upper bound for sorted-APSD while only using the minimum amount of preprocessing, as $O(m + n)$ preprocessing time is enough for the algorithm to perform all initialization steps and afterwards run the loop of Dijkstra instances unchanged.

Theorem 21. *With $O(m + n)$ preprocessing, sorted-no-self-APSD enumeration on graphs with edge weights in $\{0, \dots, n^c\}$, for some constant c , can be solved with delay in $O(\Delta + \log(n))$, with $\Theta(n^2)$ lazy-initialized memory and space complexity in $\Theta(n^2)$.*

Proof. We adapt the algorithm from Theorem 19 by moving all initialization to the $O(m + n)$ preprocessing phase while skipping self-distance solutions. Afterwards, the loop of Dijkstra instances runs unchanged and produces solutions with delay in $O(\Delta + \log(n))$. \square

7. Open Problems

In this paper we only started investigating what an enumeration point of view can reveal about the problems SSSD and APSD. Improvements on our results, different model assumptions, as well as other input types remain wide open for future research.

In Theorem 9 we explained how to transfer improvements of SSSD algorithms for weighted graphs to unconstrained-APSD enumeration. It is however open, whether such improvements could be applied to speed up the delay for row-wise-APSD/SSSD enumeration on weighted graphs. With the lower bound of Theorem 2 in mind this leaves the question, whether one can get rid of the additional $O(\log(n))$ steps per delay.

For weighted no-self-APSD enumeration we showed a lower bound of $\Omega(\Delta)$ on either preprocessing or delay in Theorem 11, but only manage to solve the problem with preprocessing in $O(n)$ in Theorem 12. Closing this gap remains an interesting open problem.

If one lifts the restriction that input-data is read-only, there might be room for improvements in the space complexity, as Kammer and Sajenko recently showed how to solve BFS in-place in linear time [23].

At last, we have not at all addressed SSSD or APSD for graphs with negative edge weights. Observe that in this general case, even emitting the first distance answers the question if a strongly connected graph contains a negative cycle. Whilst this problem seems to be easier than computing distances, state of the art algorithms require about as much time to decide the presence of negative cycles as solving SSSD or even APSD completely (see e.g. the prior work section in [28]).

References

- [1] T. H. Cormen, C. E. Leiserson, R. L. Rivest, C. Stein, Introduction to Algorithms, third edition, 3rd Edition, The MIT Press, Cambridge, Mass, 2009.
- [2] C. Demetrescu, G. F. Italiano, Dynamic shortest paths and transitive closure: Algorithmic techniques and data structures, *Journal of Discrete Algorithms* 4 (3) (2006) 353–383. doi:10.1016/j.jda.2005.12.003.
- [3] D. P. Martin, Dynamic Shortest Path and Transitive Closure Algorithms: A Survey, arXiv:1709.00553 [cs] (Oct. 2017). arXiv:1709.00553. URL <http://arxiv.org/abs/1709.00553>
- [4] E. W. Dijkstra, A note on two problems in connexion with graphs, *Numerische Mathematik* 1 (1) (1959) 269–271. doi:10.1007/BF01386390.
- [5] U. Zwick, Exact and Approximate Distances in Graphs — A Survey, in: F. M. auf der Heide (Ed.), *Algorithms — ESA 2001*, Lecture Notes in Computer Science, Springer, Berlin, Heidelberg, 2001, pp. 33–48. doi:10.1007/3-540-44676-1_3.
- [6] M. Thorup, Undirected Single-source Shortest Paths with Positive Integer Weights in Linear Time, *Journal of the ACM* 46 (3) (1999) 362–394. doi:10.1145/316542.316548.
- [7] M. Thorup, Integer priority queues with decrease key in constant time and the single source shortest paths problem, *Journal of Computer and System Sciences* 69 (3) (2004) 330–353. doi:10.1016/j.jcss.2004.04.003.

- [8] S. Pettie, V. Ramachandran, A Shortest Path Algorithm for Real-Weighted Undirected Graphs, *SIAM Journal on Computing* 34 (6) (2005) 1398–1431. doi:10.1137/S0097539702419650.
- [9] S. Pettie, A new approach to all-pairs shortest paths on real-weighted graphs, *Theoretical Computer Science* 312 (1) (2004) 47–74. doi:10.1016/S0304-3975(03)00402-X.
- [10] R. W. Floyd, Algorithm 97: Shortest Path, *Communications of the ACM* 5 (6) (1962) 345–. doi:10.1145/367766.368168.
- [11] S. Warshall, A Theorem on Boolean Matrices, *Journal of the ACM* 9 (1) (1962) 11–12. doi:10.1145/321105.321107.
- [12] R. Seidel, On the All-Pairs-Shortest-Path Problem in Unweighted Undirected Graphs, *Journal of Computer and System Sciences* 51 (3) (1995) 400–403. doi:10.1006/jcss.1995.1078.
- [13] N. Alon, Z. Galil, O. Margalit, On the Exponent of the All Pairs Shortest Path Problem, *Journal of Computer and System Sciences* 54 (2) (1997) 255–262. doi:10.1006/jcss.1997.1388.
- [14] T. M. Chan, All-pairs Shortest Paths for Unweighted Undirected Graphs in $o(mn)$ Time, *ACM Transactions on Algorithms* 8 (4) (2012) 34:1–34:17. doi:10.1145/2344422.2344424.
- [15] T. Hagerup, Improved Shortest Paths on the Word RAM, in: U. Montanari, J. D. P. Rolim, E. Welzl (Eds.), *Automata, Languages and Programming*, Vol. 1853 of *Lecture Notes in Computer Science*, Springer, Berlin, Heidelberg, 2000, pp. 61–72. doi:10.1007/3-540-45022-X_7.
- [16] D. R. Karger, D. Koller, S. J. Phillips, Finding the Hidden Path: Time Bounds for All-Pairs Shortest Paths, *SIAM Journal on Computing* 22 (6) (1993) 1199–1217. doi:10.1137/0222071.
- [17] V. V. Williams, R. R. Williams, Subcubic Equivalences Between Path, Matrix, and Triangle Problems, *Journal of the ACM* 65 (5) (2018) Article 27. doi:10.1145/3186893.
- [18] D. S. Johnson, M. Yannakakis, C. H. Papadimitriou, On generating all maximal independent sets, *Information Processing Letters* 27 (3) (1988) 119–123. doi:10.1016/0020-0190(88)90065-8.

- [19] C. Berkholz, F. Gerhardt, N. Schweikardt, Constant delay enumeration for conjunctive queries: A tutorial, *ACM SIGLOG News* 7 (1) (2020) 4–33. doi:10.1145/3385634.3385636.
- [20] K. Casel, M. L. Schmid, Fine-Grained Complexity of Regular Path Queries, in: K. Yi, Z. Wei (Eds.), 24th International Conference on Database Theory (ICDT 2021), Vol. 186 of Leibniz International Proceedings in Informatics (LIPIcs), Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl, Germany, 2021, pp. 19:1–19:20. doi:10.4230/LIPIcs.ICDT.2021.19.
- [21] L. Segoufin, Constant Delay Enumeration for Conjunctive Queries, *ACM SIGMOD Record* 44 (1) (2015) 10–17. doi:10.1145/2783888.2783894.
- [22] B. M. E. Moret, H. D. Shapiro, *Algorithms from P to NP (Vol. 1): Design and Efficiency*, Benjamin-Cummings Publishing Co., Inc., USA, 1991.
- [23] F. Kammer, A. Sajenko, Linear-Time In-Place DFS and BFS on the Word RAM, in: P. Heggernes (Ed.), *Algorithms and Complexity*, Lecture Notes in Computer Science, Springer International Publishing, Cham, 2019, pp. 286–298. doi:10.1007/978-3-030-17402-6_24.
- [24] M. L. Fredman, R. E. Tarjan, Fibonacci Heaps and Their Uses in Improved Network Optimization Algorithms, *Journal of the ACM* 34 (3) (1987) 596–615. doi:10.1145/28869.28874.
- [25] G. S. Brodal, G. Lagogiannis, R. E. Tarjan, Strict fibonacci heaps, in: *Proceedings of the 44th Symposium on Theory of Computing - STOC '12*, ACM Press, New York, New York, USA, 2012, p. 1177. doi:10.1145/2213977.2214082.
- [26] J. Alman, V. V. Williams, A Refined Laser Method and Faster Matrix Multiplication, in: *Proceedings of the 2021 ACM-SIAM Symposium on Discrete Algorithms (SODA)*, Proceedings, Society for Industrial and Applied Mathematics, 2021, pp. 522–539. doi:10.1137/1.9781611976465.32.
- [27] H. Yu, An improved combinatorial algorithm for Boolean matrix multiplication, *Information and Computation* 261 (2018) 240–247. doi:10.1016/j.ic.2018.02.006.

- [28] K. Bringmann, T. Dueholm Hansen, S. Krinninger, Improved Algorithms for Computing the Cycle of Minimum Cost-to-Time Ratio in Directed Graphs, in: DROPS-IDN/7439, Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 2017. doi:10.4230/LIPIcs.ICALP.2017.124.

Appendix A. A Note on Undirected Graphs

When applied to an undirected graph, an APSD enumeration algorithm would emit the distance between each unordered pair of vertices $u \neq v$ twice. We can however always adjust the algorithm to instead, for all vertices u, v , emit exactly one of $(u, v, d(u, v))$ and $(v, u, d(v, u))$, without compromising the asymptotic delay: The adjusted algorithm only emits the solution found first and skips the second. By emitting all solutions with twice the delay, each first solution pays for the omission of the respective second occurrence.

Whilst such tie-breaking can easily always be done with $\Theta(n^2)$ lazy-initialized memory in which the produced solutions are tracked, all algorithms presented in this paper allow for a constant-time tie-breaking without additional memory. This is essentially due to the fact that our APSD enumeration algorithms can control the order in which the SSSD enumeration algorithms they build on are executed:

- For the unconstrained algorithms presented in section 4 and the algorithms for reachable-APSD enumeration of Theorem 14, selecting the start vertices for the SSSD algorithms with increasing vertex id and only emitting solutions (u, v, \cdot) with $u \leq v$ suffices, as apart from the trivial self-distances all solutions are produced by the SSSD algorithms.
- The algorithm of Theorem 10 for unweighted no-self-APSD enumeration can be adjusted the same way; additionally the initial pass over all vertices and edges has to respect the same vertex order.
- Weighted no-self APSD enumeration according to Theorem 12 changes the vertex order according to their node degree. As our sorting of vertices by degree is stable, all following operations iterate over vertices sorted by node degree first and vertex id second. If tie-breaking uses the same ordering, duplicate solutions can again be identified in constant time.
- The algorithms for sorted APSD enumeration in section 6 executes SSSD algorithms in parallel instead of consecutively, but can again apply tie-breaking by vertex id to select the respective SSSD algorithm instance to advance next.