

Batch Dynamic Single-Source Shortest-Path Algorithms: An Experimental Study ^{*}

Reinhard Bauer and Dorothea Wagner

Karlsruhe Institute of Technology (KIT), Germany
{rbauer,wagner}@ira.uka.de

Abstract. A dynamic shortest-path algorithm is called a batch algorithm if it is able to handle graph changes that consist of multiple edge updates at a time. In this paper we focus on fully-dynamic batch algorithms for single-source shortest paths in directed graphs with positive edge weights. We give an extensive experimental study of the existing algorithms for the single-edge and the batch case, including a broad set of test instances. We further present tuned variants of the already existing SWSF-FP-algorithm being up to 15 times faster than SWSF-FP. A surprising outcome of the paper is the astonishing level of data dependency of the algorithms. More detailed descriptions and further experimental results of this work can be found in [1].

1 Introduction

The *single-source shortest-path problem* is a fundamental graph problem with many real-world applications, such as routing in road networks, routing/data harvesting in sensor networks and internet routing using link state protocols (for example OSPF and IS-IS). In these applications shortest-path trees are stored and have to be updated whenever the underlying graph undergoes changes [2–5].

Algorithms that update the trees without a full recomputation from scratch are called *dynamic single-source shortest-path algorithms*. Such algorithms slightly differ in the type of their output. Some store only the distances from the source, while others additionally store a shortest-path tree or the shortest-path subgraph. Some of the algorithms known in the literature are only able to cope with the update of one edge at a time, while others can perform *batch updates*, i.e. update the shortest-path information after multiple edges have simultaneously changed their weight.

We consider edge insertions and deletions as special cases of weight changes: Deletions correspond to weight increments to infinity, while insertions are weight decrements from infinity. An algorithm is called *fully dynamic* if both weight increases and decreases are supported, and *semi-dynamic* if only weight decreases or only increases are supported.

^{*} Partially supported by the Future and Emerging Technologies Unit of EC (IST priority – 6th FP), under contract no. FP6-021235-2 (project ARRIVAL) and the DFG (project WAG54/16-1).

In this paper we focus on fully-dynamic batch updates for directed graphs with positive edge weights. In order to compare the different approaches, the only requirement that we make regarding the tested algorithms is that they update the distance vector. We furthermore demand that the algorithms be able to cope with edge insertions and deletions. For our experimental study, we apply integer edge weights.

Related work. Ramalingam and Reps [6] introduce the batch algorithm SWSF-FP, Narvaez et al. [2] propose the NARVÁEZ-framework containing six single-edge update algorithms and a modification to the framework leading to the according batch algorithms. Pure single-edge update algorithms are RR [7] (due to Ramalingam and Reps) and FMN [8] (by Frigioni et al). Buriol et al [9] present a heuristic technique to speed up RR-like approaches. The technique is similar to techniques used in the NARVÁEZ-framework but does not support edge insertions or deletions. Furthermore, in [9] the RR algorithm is adapted to maintain a special (shortest-path) tree proposed in [10].

There is no algorithm known in the literature for which the worst case is asymptotically better than recomputing the new solution from scratch. In the original works the algorithms described in Section 3 are theoretically analyzed with respect to different measures. These measures mostly depend on the size of the subgraph for which the shortest-path subgraph changes.

There is some work on the variant of the problem where edge weights may also be negative. In [7] the algorithm RR is adapted to cope with the existence of negative cycles, in [11] the same is done for the algorithm FMN. In [12] Demetrescu gives some algorithms for that problem. These algorithms use the reweighting technique, which incorporates a complete Dijkstra run on the graph (with changed edge weights). Hence, this approach is impractical for the problem with non-negative edges.

A well-studied related problem is the *fully dynamic all-pairs shortest-path problem*, in which the distances between all pairs of nodes have to be maintained while the graph undergoes changes. See [13] for a survey on the problem.

There is only few experimental work on this topic, all concentrating on single-edge updates. In [14] the algorithms FMN, RR and a full recomputation from scratch are compared on two instance classes: Erdős-Rényi graphs, where updates are chosen uniformly at random and a graph representing the internet on the AS-level, where updates simulate the failure and recovery of the links. In [2] the algorithms of the NARVÁEZ-framework are evaluated on graphs originating from a generator. This generator randomly places nodes on a grid and connects them by edges with probability that exponentially decreases with the distance of the nodes. The generator does not seem to be available any more. In [15] the algorithms SWSF-FP, RR, FMN, NARVÁEZ and full recomputation from scratch using DIJKSTRA, BELLMAN FORD, D'ESOPPO PAPE are evaluated with single-edge updates on Erdős-Rényi-like graphs. In [4] one algorithm of the NARVÁEZ-framework is evaluated on random single-edge updates on a graph representing the road-network of Western Europe. In [9], the algorithm RR as well as seven variants thereof are evaluated on a real world AT&T IP network,

synthetic internet-related graphs and a large set of other synthetic instances, namely those of [16] with non-negative edge lengths.

Overview. This paper is organized as follows. Section 2 states basic definitions and formally introduces the problem. Section 3 reviews the existing algorithms. Section 4 presents our tuned variants of the SWSF-FP-algorithm, while an extensive experimental study of these algorithms on synthetic and real-world data is given in Section 5. The paper ends with a conclusion in Section 6.

2 Problem Statement

Let $G = (V, E)$ be a directed graph with n nodes and m edges and a non-negative length function $\text{len} : V \times V \rightarrow \mathbb{R}^+ \cup \{\infty\}$. Let $s \in V$ be an arbitrary but fixed *source*. With $d(v)$ we denote the length of a shortest s - v -path in G for any $v \in V$.

A *batch update* is a set of *edge modifications* on G which can be edge insertions, edge deletions, edge weight increases and edge weight decreases (that keep the length function non-negative). We want to maintain a distance vector $D[]$ containing $d(v)$ for each node v in a dynamic environment where G is undergoing batch updates. After each batch update, $D[]$ (and possible required auxiliary data needed by the recomputation algorithm) has to be updated accordingly.

Throughout the text, we will cope with the recomputation of $D[]$ and the auxiliary data when *one* concrete batch update is given (because of the recomputation of the auxiliary data the algorithms are able to handle following updates). We write len_{old} for length function and d_{old} for distance before the update. Accordingly we write len for length function and d for distance after the update. For notational convenience, we consider inserted or deleted edges to be existing in the original and the updated graph and set the edge length to infinity, if necessary.

Some of the following algorithms are designed to handle only one edge modification at a time. Obviously, repeated application of these algorithms also solves the batch case. We call such algorithms *iterative algorithms* while the others are called *batch algorithms*. Iterative algorithms can be split into two parts: the *incremental* part handles edge insertions and weight decreases while the *decremental* part handles edge deletions and weight increases. This terminology can be unintuitive on a first glance but originates from the point of view that the graph increases when edges are inserted.

3 Overview of Algorithms

In this section, we give an overview on the algorithms evaluated in our experimental study. Each algorithm includes a main phase in which a min-based priority queue Q is used to recompute the distances in a Dijkstra-like fashion but on a smaller subgraph.

RR. Ramalingam and Reps [7] describe the iterative algorithm RR that handles only edge insertions and deletions. It can be directly transferred to an algorithm that works with weight increases and decreases. We will use this variant.

FMN. The FMN-algorithm of Frigioni et al. [8] is an iterative algorithm similar to the algorithm RR that uses more complex auxiliary data to obtain better theoretical worst case bounds. The approach relies on the existence of a *k*-bounded *accounting function* on G , which is a mapping $K : E \rightarrow V$ such that for each edge (u, v) the node $K(u, v)$ is either u or v and such that for each node n , no more than k edges are n -valued. We use the constructive 2-approximation algorithm described in [11] for finding a *k*-bounded accounting function on G .

Narváez. Narvaez et al. [2] propose a batch algorithm incorporating two degrees of freedom. One degree of freedom is the choice of Q which does not necessarily need to be a priority queue but only has to maintain the operations INSERT and EXTRACT. Narvaez et al propose a FIFO queue (Bellman-Ford like approach), a heap (implemented as binary heap or linked list) and a D’Esopo-Pape like approach. The other degree of freedom consists of two different variants for the main phase of the algorithm which we will describe below. We will refer to the different variants as NAR{1st, 2nd}{HEAP, BF, PAP}. The main idea of the NARVÁEZ-framework is to early-propagate distance changes through the tentative shortest-path tree.

4 Tuning SWSF-FP

In this section we will review the algorithm SWSF-FP which is due to Ramalingam and Reps [6] and give some tuned variants of it. The input of the algorithms is the outdated distance vector $D[]$, the graph G , the original length function len_{old} , the batch update $U = (u_1, \dots, u_k)$ and some auxiliary data which will be described for each algorithm separately. The output is the updated distance vector $D[]$ and the updated auxiliary data.

Notation. Given the outdated distance vector $D[]$, we say we *relax* an edge (u, v) when we check if $D[v] > D[u] + \text{len}(u, v)$. We say we *relax and update* an edge (u, v) when we set $D[v] := \min\{D[v], D[u] + \text{len}(u, v)\}$. An edge (u, v) is said to be *consistent* if $D[v] = \text{len}(u, v) + D[u]$ and *underconsistent* if $D[v] > \text{len}(u, v) + D[u]$. The *consistent value* $\text{con}(v)$ of a node v is

$$\text{con}(v) := \begin{cases} \min_{(u,v) \in E} \{D[u] + \text{len}(u, v)\}, & v \neq s \\ 0, & v = s \end{cases}$$

A node is said to be *consistent* if $D[v] = \text{con}(v)$ and to be *over-consistent* if $D[v] > \text{con}(v)$. As convention, we use $\min \emptyset := \infty$.

SWSF-FP. For each node v , a label $d[v]$ is given. Initially, $d[]$ equals $D[]$ (in order to save time for the copy process we implemented $d[]$ as auxiliary data). We say we *adjust* an inconsistent node v when we set $d[v] := \text{con}(v)$ and insert v with priority $\min(D[v], d[v])$ in Q . In case v is already in Q we only update its priority. We adjust a consistent node v when we remove it from Q . If v is not in Q we do nothing.

Initially, we adjust each node which is target of an edge in U . *Main Phase.* While Q is not empty, we perform as follows: We extract and delete the minimum node w from Q . If $d[w] < D[w]$ we set $D[w] := d[w]$ and adjust each outgoing

neighbor of w . If $d[w] > D[w]$ we set $D[w] := \infty$ and adjust w and each of its outgoing neighbors.

Tuned SWSF. This algorithm basically works like the SWSF-FP-algorithm, but with less computational effort. When performing SWSF-FP we have to relax all incoming edges of a node n in order to compute $\text{con}(n)$. TUNED SWSF relaxes fewer of such incoming edges: When we adjust an outgoing neighbor v of a node w with $d[w] < D[w]$, we compute $\text{con}(v)$ by $\min\{d[w] + \text{len}(v, w), d[v]\}$. The same strategy works in the initialization phase when we compute $\text{con}(n)$ for a node n that is the target node of an edge with decreased edge weight. When we adjust an outgoing neighbor v of a node w with $d[w] > D[w]$, we set $D_{old} := D[w]$ and $D[w] := \infty$. We can skip v when $D_{old} + \text{len}(w, v) \neq d[v]$. The same strategy holds in the initialization phase for target nodes of edges with increased weight.

Tuned SWSF RR. This variant enhances the algorithm TUNED SWSF with a technique adapted from the RR-algorithm. For each node v , a label $\text{indegree}(v)$ is given indicating the number of edges (u, v) with $D[u] + \text{len}(u, v) = d[v]$. Further, for each edge (u, v) a boolean label $DAG(u, v)$ is given indicating if $D[u] + \text{len}(u, v) = d[v]$. The labels indegree and DAG are directly updated whenever len , $D[]$ or $d[]$ change. The algorithm performs like TUNED SWSF with the following difference: After a node v with $d[v] > D[v]$ is extracted from Q only those edges (v, w) have to be processed for which $\text{indegree}[w] = 0$.

Tuned SWSF NAR. This variant enhances the algorithm TUNED SWSF with a technique adapted from the NARVÁEZ-algorithm. For each node v that is not the source, a label $P(v)$ is given pointing at another node, such that $D[P(v)] + \text{len}(P(v), v) = d(v)$. At the beginning a shortest-path tree T on the original graph is given implicitly by this label. The main phase of the algorithm works like the main phase of TUNED SWSF. The initialization phase works as follows: First, we update the edge weights. We denote by A the set of all nodes that lie behind a target node of an updated edge. Then, we update the distances $D[]$ of nodes in T according to the new edge weights (but to the original shortest-path tree T). This can be implemented such that for each node $v \in A$, the distance $D[v]$ is updated at most once. Then, we set $d[v] = \text{con}(v)$ for each node v which either is contained in A or has a neighbor in A . Finally we insert each node with $d[v] \neq D[v]$ with priority $\min\{d[v], D[v]\}$ in Q .

5 Experiments

In this section, we present an experimental evaluation of the algorithms described above. Our implementation is written in C++ (using the STL at some points). Our tests were executed on one core of an AMD Opteron 2218, running SUSE Linux 10.3. The machine is clocked at 2.6 GHz, has 32 GB of RAM and 2 x 1 MB of L2 cache. The program was compiled with GCC 4.2, using optimization level 3.

For each experiment, 1000 update instances were generated. To properly measure the speed-ups, a full Dijkstra run is performed directly after each update and the speed-up compared to that run (i.e. the time needed by Dijkstra's

algorithm divided by the time needed by the update algorithm) is computed. Finally we compute the mean value of these speed-ups. Thus, measurement disturbances due to background processes etc are avoided as much as possible. For Tables 1-4 we showed in bold letters all algorithms whose performance was at least 85% of the best observed performance.

In our experiments we evaluated all previously described algorithms. We did not include the heuristic of Buriol et al [9] because it does not support edge insertions or deletions. Further, we did not include the D’Esopo-Pape variants of the NARVÁEZ-framework because pretests had revealed some instances with extremely bad performance with this approach. To gain further insights in the performance of the batch-algorithms (NARVÁEZ and TUNED SWSF), we executed these two times: one time with processing the edges in batch, as stated originally and one time with iteratively processing the edges one after another. We refer to these approaches as ITNAR and ITTUNED SWSF. Note that we refer to the NARVÁEZ-framework as a batch algorithm while it actually does not perform updates completely in a batch: its initialization phase handles edge updates iteratively but the following main phase handles all updates in a batch.

5.1 Graph Instances

UNIT DISK. Given n and m , a unit disk graph is generated by randomly assigning each of the n nodes to a point in the unit square of the Euclidean plain. Two nodes are connected by an edge in case their Euclidean distance is below a given radius. This radius is adjusted such that the resulting graph has approximately m edges. As edge weights we use the Euclidean distance to the power of 0 (hop length), 1 (Euclidean distance) and 2 (energy). All tested graphs consist of 15 000 nodes.

RAILWAY. The graph RAIL represents the condensed railway network of Europe, based on timetable information, provided by the company HaCon [17] for scientific use. Nodes represent stations while edges represent direct connections between the stations. The edge weight corresponds to the average travel time between two stations. The graph has 29 578 nodes and 159 914 edges.

AS-GRAPH. The graph AS-HOP represents the internet as of 2008/3/26 on the AS-level, i.e. each node corresponds to an autonomous system and edges represent connections between autonomous systems. This graph is taken from the Routeviews project page [18]. It has 27 909 nodes and 114 474 edges. The edge weight is 1 for each edge. The same graph with edge weights chosen uniformly at random from the interval $[1, 1000]$ is called AS-RAN.

CAIDA. This dataset represents the internet on the router level, i.e. nodes are routers and edges represent connections between routers. The network is taken from the CAIDA webpage [19] and has 190 914 nodes and 1 215 220 edges. The edge weight is 1 for each edge.

ROAD. We evaluate three road networks provided by the PTV AG [20]. DEU represents Germany with 4 378 447 nodes and 10 968 884 edges, NLD the Netherlands with 946 632 nodes and 2 358 226 edges and LUX represents Luxembourg

with 30 647 nodes and 75 576 edges. The edge weights are the corresponding travel times with speed profile ‘slow car’.

GRID. These are fully synthetic graphs based on two-dimensional square grids. The nodes of the graph correspond to the crossings in the grid. There is an edge between two nodes if these are neighbors on the grid. Edge weights are randomly chosen integer values between 1 and 1000. GRID 100 is a 100x100 grid graph while GRID 300 is a 300x300 grid graph.

5.2 Assessing the Performance of the Algorithms

Let $U = \{u_1, \dots, u_k\}$ be a set of updated edges. By $\Delta(G, U)$ we denote the number of vertices in V for which the distance from the source changes due to the update. The *expected speed-up* of an update is the number of vertices in the graph divided by $\Delta(G, U)$. This value is roughly the speed-up we expect from a good update algorithm. Of course, speed-ups can even be higher for special instances. It experimentally turned out that when the topology of the original shortest-path tree does not change, the propagation of the updated edge’s weights through the tree can gain a large speed-up.

When we want to measure the difficulty of an update for an iterative algorithm we consider $U = (u_1, \dots, u_k)$ to be ordered. We perform the updates u_i iteratively in the given ordering (always additional to the former updates) obtaining a sequence of graphs $G = G_0, G_1, \dots, G_k$. We write $\delta(G, (u_1, \dots, u_k)) := \sum_{i=0}^{k-1} \Delta(G_i, \{u_{i+1}\})$. We have the following hypothesis: the smaller the difference between $\Delta(G, U)$ and $\delta(G, U)$ is, the less do the contained single-edge updates interfere and it is reasonable to use an iterative algorithm for the update. If the difference is great, an iterative algorithm would change the distance of many nodes multiple times. Hence, it is more appropriate to use a batch algorithm. The experimental evaluation will support our hypothesis.

5.3 Space-Saving Implementation of RR

The algorithm RR needs to maintain the shortest-path subgraph. This subgraph is implicitly given by each edge (u, v) with $d[u] + \text{len}(u, v) = d[v]$. We implemented the algorithm doubly. One time with explicitly storing the subgraph (RR DAG) and one time with reconstructing it when needed (RR). It turned out that there are only small differences between both implementations, with no variant being clearly superior. We therefore only report the results for the space-saving implementation RR.

5.4 Single-Edge Update Experiments

We start our experimental study by single edge updates. Because of space restrictions and a different focus of our paper we do not carry out a separate analysis for the decremental and the incremental case. An update consists of choosing an edge uniformly at random and multiplying its weight by a random value in $(0, 2)$. The results can be seen in Table 1.

We observe that the algorithms of the NARVÁEZ-framework have only tiny differences in performance with Nar-1st BF being slightly (but not significantly) faster most times. There is no such uniform behavior for the SWSF-FP-like algorithms. TUNED SWSF is always faster (between 1.3 and 6 times) than SWSF-FP. The algorithm TUNED SWSF RR is always at least as fast as SWSF-FP and up to 5.5 times faster. The algorithm TUNED SWSF NAR seems to be very volatile being between half as fast and 4 times faster than SWSF-FP.

Comparing the different classes of algorithms, we find the algorithms to perform quite differently, but within the same order of magnitude. The algorithm FMN is most times much slower than the other ones. This is due to the overhead caused by maintaining and reading the priority queues used by this algorithm. The technique used in this algorithm can pay off in case nodes with high degree exists (for which many edge-relaxations can be saved). This is not the case for the test instances used. Exceptions are the INTERNET instances CAIDA, AS-HOP and AS-RAN. Here, the gap to the other algorithms is much smaller, (which meets the theoretical considerations). Hence, it is to be expected that there are dense graph classes for which FMN is the superior algorithm. On the ROAD and GRID instances, the NARVÁEZ-framework is superior. This is because the structure of the shortest-path tree stored by the algorithm hardly changes on these experiments. Therefore, the early-propagation of the weight change works well. On the INTERNET instances, RR is the fastest algorithm. Looking at the small value of $\Delta(G, U)$, we can see that updates hardly have any impact on these instances, which favors the RR-algorithm with its small computational overhead and the early detection of edge weight increases that do not change distances on the graph.

The achieved speed-ups vary greatly between the instances. This is mainly due to the different structure of the underlying graphs, which results in greatly differing expected speed-ups. It is interesting to see that in nearly all cases the best actual speed-ups are close to the expected speed-ups or even higher. This, in combination with the small absolute runtimes in the range of microseconds, makes us expect that there is not much space for further improvement for the single-edge update case.

	LUX	NLD	DEU	RAIL	CAIDA	AS-HOP	AS-RAN	GR100	GR300	UNIT H	UNIT E
FMN	42	1504	29087	151	22702	1624	2182	25	142	327	36
SWSF-FP	112	3759	65404	366	12429	416	691	59	351	1613	31
tun SWSF-FP	152	5140	84873	562	16406	893	3442	105	598	2436	186
tun SWSF-NAR	147	3354	70245	215	9306	614	695	94	523	748	129
tun SWSF-RR	118	3798	66068	412	26093	2148	3766	74	430	2096	102
RR	155	4666	74857	510	34586	2599	4057	103	568	2519	137
Nar-1st BF	284	5335	100944	357	6578	417	305	138	784	1176	20
$\Delta(G, U)$	130.42	140.52	70.72	30.68	0.21	0.41	0.74	59	113	0.01	93
expected speed-up	236	6762	62549	986	inf	inf	inf	169	804	inf	163

Table 1. Speed-ups of experiments with single-edge updates

5.5 Experiments on Batch Updates

Multiple Randomly Chosen Edges. In this experiment we chose 25 edges uniformly at random. For each edge, we chose uniformly at random a value from the interval $(0, 2)$ and multiplied the weight of the edge with that value. For each graph there is hardly any difference between $\Delta(G, U)$ and $\delta(G, U)$. Therefore, the single-edge updates did only interfere marginally with each other. Hence, not much news is to be expected by this setting regarding the comparison of the algorithms. This has been confirmed by the experiments.

However, we ran the batch-algorithms (NARVÁEZ and TUNED SWSF) twice. One time with processing the edges in batch as stated in the description and one time with iteratively processing the edges one after another. Nearly no runtime differences were observed between the iterative and the batch variants, which indicates a low overhead with batch updates.

Node Failure and Recovery. This update class uses the two parameters deg_{min} and deg_{max} . First, a node v with degree between deg_{min} and deg_{max} is chosen uniformly at random. The update consists of two steps. In the first step, v fails, i.e. the weights of all edges adjacent to v are set to infinity. In the second step, v recovers, i.e. the weights of all edges adjacent to v are reset to their original values. The results can be found in Tables 2 and 3.

degree	AS-HOP			AS-RAN			CAIDA		
	1-10	10-100	100-500	1-10	10-100	100-500	1-10	10-100	100-500
FMN	784	173	23	1368	129	3	7824	2284	185
ittun SWSF-FP	912	228	26	1320	235	11	12874	4212	382
SWSF-FP	273	68	8	389	28	1	9651	2203	128
tun SWSF-FP	967	250	24	1417	252	15	14042	4693	405
tun SWSF-NAR	407	92	9	410	50	4	9785	2187	122
tun SWSF-RR	1272	528	130	2475	433	21	12395	6839	969
RR	1438	576	142	2623	490	17	13915	7075	1163
Nar-1st Heap	53	21	9	86	59	16	4315	761	75
itNar-1st Heap	52	18	6	71	30	8	4060	573	63
$\delta(G, U)$	1.26	12.16	82.54	1.47	45.01	1365.85	1.97	7.4	90.99
$\Delta(G, U)$	1.07	8.59	71.28	1.1	34.6	712.3	1.45	5.7	85.73
expected speedup	27909	3489	393	27909	821	86	190914	38183	2246

Table 2. Speed-ups of experiments with node failure and recovery updates on INTERNET-instances

We now take a look at the INTERNET instances. The most remarkable result is the bad performance of the Narvaez-framework, which clearly is the inferior algorithm for that testset. One main reason for that is, that on this testset the edge-weight propagation in the initialization phase creates useless extra effort which gets overwritten later on. The gap between $\delta(G, U)$ and $\Delta(G, U)$ is small to mid-size, favoring RR with its small overhead, but big enough such that TUNED SWSF RR is nearly as fast. This difference also manifests in the small difference between ITTUNED SWSF and TUNED SWSF.

The situation is similar, but a bit clearer, for UNIT DISK graphs. When applying hop distance, $\delta(G, U)$ and $\Delta(G, U)$ are still quite near to each other, TUNED SWSF and RR are the best-performing algorithms (with RR being slightly better). When applying Euclidean or energy edge weights updates, the

difference between $\delta(G, U)$ and $\Delta(G, U)$ is much bigger, and TUNED SWSF clearly is the superior algorithm. We also observe the advantage of TUNED SWSF against SWSF-FP being between 2 and 15 times faster.

metric	hop			euclidean			energy		
	7	10	15	7	10	15	7	10	15
average degree									
FMN	30	40	55	27	21	24	12	14	20
ittun SWSF-FP	238	398	485	116	95	98	56	66	91
SWSF-FP	128	214	236	60	32	36	28	24	22
tun SWSF-FP	260	462	561	158	115	141	75	86	110
tun SWSF-NAR	106	116	147	101	77	97	57	61	67
tun SWSF-RR	223	395	527	105	75	89	49	54	67
RR	289	504	628	111	91	106	55	63	84
Nar-1st Heap	70	87	131	84	62	111	52	62	74
itNar-1st Heap	55	71	100	64	50	66	36	46	52
$\delta(G, U)$	19	8	6	86	107	99	194	174	132
$\Delta(G, U)$	18	7	5	54	79	55	128	119	98
expected speedup	833	2500	3750	283	190	273	117	126	153

Table 3. Speed-ups of experiments with node failure and recovery updates on UNIT DISK-instances

Traffic Jams. This update class models real-world traffic jams. It derives from the observation that traffic jams often occur along shortest paths. The number k of updated edges is given as a parameter. Initially, a node v is chosen uniformly at random. Then a shortest path SP ending at v and containing exactly k edges is chosen uniformly at random. The update consists of two steps: in the first step, the weights of edges in SP are multiplied by 10. In the second step, the edge weights are reset to their original values. The results can be found in Table 4.

edges	GRID			LUX			NLD			DEU		
	10	20	30	5	10	20	10	20	30	10	20	30
FMN	3	2	1	4	2	1	11	5	2	185	30	7
ittun SWSF-FP	15	9	5	15	7	2	39	17	6	755	100	23
SWSF-FP	13	10	6	15	9	5	75	32	12	873	173	40
tun SWSF-FP	23	16	9	20	12	6	107	44	17	1210	235	55
tun SWSF-NAR	22	16	9	22	13	7	107	41	17	1402	342	79
tun SWSF-RR	16	12	7	15	9	5	72	31	12	957	181	42
RR	17	10	5	20	9	3	43	18	6	924	149	36
Nar-1st Heap	16	9	5	20	10	4	37	15	5	1120	196	35
itNar-1st Heap	19	12	6	24	12	4	57	24	8	1231	219	54
$\delta(G, U)$	4367	7909	15552	1178	3052	8616	12910	32088	93725	7885	39260	142191
$\Delta(G, U)$	2591	3564	6412	821	1366	2567	4134	10899	26153	3884	13701	51252
expected speed-up	35	25	14	37	22	12	229	87	36	1127	320	85

Table 4. Speed-ups of experiments with traffic jam updates

We observe that this update class consists of strongly interfering single-edge updates: there is a big difference between $\delta(G, U)$ and $\Delta(G, U)$. TUNED SWSF and TUNED SWSF NAR are the best-performing algorithms for this testset. This is because pure batch algorithms avoid processing nodes many times. With an increasing number of edges, the interference between the updated edges increases and the advantage of these two algorithms grows.

For a small number of edges in the jam, the NARVÁEZ-framework is comparable to TUNED SWSF. The framework slows down with a growing number of

updated edges. This is because the initialization phase processes many nodes one time for each updated edge. It is astonishing to see that the NARVÁEZ-framework is not able to take advantage of the batch-character of the update. This can be seen through a comparison with ITNARVAEZ. The iterative variant is even faster than the batch one, which could be a hint at space for improvement. Again, FMN is much slower than the other algorithms, as its overhead does not pay off on these instances.

6 Conclusion

In this work we focused on the single-source shortest-path problem with non-negative weights. We gave the first experimental study evaluating the performance for single-edge updates that contains all current algorithms and incorporates a broad set of instance classes. It turned out that the algorithms perform quite differently, but within the same order of magnitude. Furthermore, the achieved speed-ups varied greatly between different instances. This can be explained by measuring the impact of the updates on the graphs.

Moreover, we presented the first experimental study at all for the case of multiple edge changes at a time. One experiment was to choose a set of edges uniformly at random. It turned out that this way the single-edge updates did almost not interfere. Therefore, the results deviated not much from the single-edge case. We also used two more realistic types of batch updates. One is the simulation of node failure and recovery, which affects all incident edges. The single-edge updates interfered for that class, but not very strongly. For internet instances, the best performing algorithms were RR and TUNED SWSF RR with RR being slightly faster. For UNIT DISK graphs, TUNED SWSF was the best algorithm with RR being slightly faster for hop distance. The other update class modelled traffic jams. The single-edge updates interfered greatly, TUNED SWSF and TUNED SWSF NAR were the superior algorithms there.

Furthermore, we presented tuned variants for the SWSF-FP-algorithm and evaluated their performance. For the tuned variants we observed speed-ups compared to SWSF-FP of up to 15. Finally, we gave a simple methodology (based only on Dijkstra’s algorithm) to decide if one should try a single-edge or a batch-update algorithm for a given instance class. We compared the ‘impact’ of the update when processed in batch with the ‘impact’ when processed iteratively. For updates with a big gap between both values, the algorithms TUNED SWSF or TUNED SWSF RR usually performed best. With a small gap, there was usually a better-performing iterative algorithm.

Concluding, we gave a first experimental overview on the different approaches for the problem, which can be used as a base for further research. The most important information that can be extracted from our experiments is the astonishing level of data dependency within the problem. It turned out that a proper assessment of an algorithm’s running time is not possible without full knowledge of the application it is used in. Further, a great amount of experiments is required to get the big picture of an algorithm’s efficiency.

References

1. Bauer, R., Wagner, D.: Batch Dynamic Single-Source Shortest-Path Algorithms: An Experimental Study. Technical Report 2009,6, ITI Wagner, Faculty of Informatics, Universität Karlsruhe (TH) (2009) <http://digbib.ubka.uni-karlsruhe.de/volltexte/1000010926>.
2. Narváez, P., Siu, K.Y., Tzeng, H.Y.: New Dynamic Algorithms for Shortest Path Tree Computation. *IEEE/ACM Transactions on Networking* **8** (2000) 734–746
3. Bruera, F., Cicerone, S., D’Angelo, G., Stefano, G.D., Frigioni, D.: Dynamic Multi-level Overlay Graphs for Shortest Paths. *Mathematics in Computer Science* (2008) To appear.
4. Delling, D., Wagner, D.: Landmark-Based Routing in Dynamic Graphs. In Demetrescu, C., ed.: *Proceedings of the 6th Workshop on Experimental Algorithms (WEA’07)*. Volume 4525 of *Lecture Notes in Computer Science.*, Springer (2007) 52–65
5. Wagner, D., Watternhofer, R., eds.: *Algorithms for Sensor and Ad Hoc Networks*. Volume 4621 of *Lecture Notes in Computer Science*. Springer (2007)
6. Reps, T., Ramalingam, G.: An Incremental Algorithm for a Generalization of the Shortest-Path Problem. *Journal of Algorithms* **21** (1996)
7. Ramalingam, G., Reps, T.: On the computational complexity of dynamic graph problems. *Theoretical Computer Science* **158** (1996)
8. Frigioni, D., Marchetti-Spaccamela, A., Nanni, U.: Fully Dynamic Algorithms for Maintaining Shortest Paths trees. *Journal of Algorithms* **34** (2000)
9. Buriol, L., Resende, M., Thorup, M.: Speeding Up Dynamic Shortest-Path Algorithms. *Inform Journal on Computing* **20** (2008)
10. King, V., Thorup, M.: A space saving trick for directed dynamic transitive closure and shortest path algorithms. In: *Proceedings of the 7th Annual International Conference on Computing Combinatorics (COCOON’01)*. Volume 2108 of *Lecture Notes in Computer Science.*, Springer (2001) 268–277
11. Frigioni, D., Marchetti-Spaccamela, A., Nanni, U.: Fully dynamic shortest paths in digraphs with arbitrary arc weights. *Journal of Algorithms* **49** (2003) 86–113
12. Demetrescu, C.: *Fully Dynamic Algorithms for Path Problems on Directed Graphs*. PhD thesis, Department of Computer and Systems Science (2001)
13. Demetrescu, C., Italiano, G.F.: Dynamic shortest paths and transitive closure: Algorithmic techniques and data structures. *Journal of Discrete Algorithms* **4** (2006)
14. Frigioni, D., Ioffreda, M., Nanni, U., Pasqualone, G.: Experimental Analysis of Dynamic Algorithms for the single Source Shortest Path Problem. *ACM Journal of Experimental Algorithmics* **3** (1998)
15. Taoka, S., Takafuji, D., Iguchi, T., Watanabe, T.: Performance Comparison of Algorithms for the Dynamic Shortest Path Problem. *IEICE Transactions on Fundamentals of Electronics, Communications and Computer Sciences* **E90-A** (2007)
16. Cherkassky, B.V., Goldberg, A.V., Radzik, T.: Shortest paths algorithms. *Mathematical Programming, Series A* **73** (1996) 129–174
17. HaCon - Ingenieurgesellschaft mbH: <http://www.hacon.de> (2008)
18. University of Oregon Routeviews Project: <http://www.routeviews.org/> (2008)
19. CAIDA: The Cooperative Association for Internet Data Analysis: <http://www.caida.org/> (2008)
20. PTV AG - Planung Transport Verkehr: <http://www.ptv.de> (2008)