Alcom-FT Technical Report Series
ALCOMFT-TR-01-18 (revised March 23, 2001)

A Heuristic for Dijkstra's Algorithm with Many Targets and its Use in Weighted Matching Algorithms*

Kurt Mehlhorn† Guido Schäfer‡

March 23, 2001

Abstract

We consider the single-source many-targets shortest-path (SSMTSP) problem in directed graphs with non-negative edge weights. A source node s and a target set T is specified and the goal is to compute a shortest path from s to a node in T. Our interest in the shortest path problem with many targets stems from its use in weighted bipartite matching algorithms. A weighted bipartite matching in a graph with n nodes on each side reduces to n SSMTSP problems, where the number of targets varies between n and 1.

The SSMTSP problem can be solved by Dijkstra's algorithm. We describe a heuristic that leads to a significant improvement in running time for the weighted matching problem; in our experiments a speed-up by up to a factor of 10 was achieved. We also present a partial analysis that gives some theoretical support for our experimental findings.

1 Introduction and Statement of Results

A matching in a graph is a subset of the edges no two of which share an endpoint. The weighted bipartite matching problem asks for the computation of a maximum weight matching in an edge-weighted bipartite graph $G = (A \dot\cup B, E, w)$ where the cost function $w : E \mapsto \mathbb{R}$ assigns a real weight to every edge. The weight of a matching M is simply the sum of the weights of the edges in the matching, i.e., $w(M) = \sum_{e \in M} w(e)$. One may either ask for a perfect matching of maximal weight (the weighted perfect matching problem or the assignment problem) or simply for a matching of maximal weight. Both versions of the problem can be solved by solving n, $n = \max(|A|, |B|)$, single-source many-targets shortest-path (SSMTSP) problems in a derived graph, see Section 4. We describe and analyse a heuristic improvement for the SSMTSP problem which leads to a significant speed-up in LEDA's weighted bipartite matching implementation, see Table 3.

In the SSMTSP problem we are given a directed graph $G = (V, E)$ whose edges carry a non-negative cost. We are also given a source node s. Every node in V is

*Partially supported by the IST Programme of the EU under contract number IST-1999-14186 (ALCOM-FT).

†Max-Planck-Institut für Informatik, Saarbrücken, Germany, www.mpi-sb.mpg.de/~mehlhorn.

‡Graduiertenkolleg (Graduate Studies Program) "Quality Guarantees for Computer Systems", Dept. of Computer Science, University of the Saarland, Germany, funded by the Deutsche Forschungsgemeinschaft (DFG), www.mpi-sb.mpg.de/~schaefer.

designated as either free or non-free. We are interested in finding the shortest path from s to a free node.

The SSMTSP problem is easily solved by Dijkstra's algorithm. Dijkstra's algorithm (see Section 2) maintains a tentative distance for each node and a partition of the nodes into settled and unsettled. At the beginning all nodes are unsettled. The algorithm operates in phases. In each phase, the unsettled node with smallest tentative distance is declared settled and its outgoing edges are relaxed in order to improve tentative distances of other unsettled nodes. The unsettled nodes are kept in a priority queue. The algorithm can be stopped once the first free node becomes settled. *We describe a heuristic improvement.* The improvement maintains an upper bound for the tentative distance of free nodes and performs only queue operations with values smaller than the bound. All other queue operations are suppressed. The heuristic significantly reduces the number of queue operations and the running time of the bipartite matching algorithm, see Tables 2 and 3.

This paper is structured as follows. In Section 2 we discuss Dijkstra's algorithm for many targets and describe our heuristic. In Section 3 we give an analysis of the heuristic for random graphs and report about experiments on random graphs. In Section 4 we discuss the application to weighted bipartite matching algorithms and present our experimental findings for the matching problem.

The heuristic was first used by the second author in his jump-start routine for the general weighted matching algorithm [Sch00, MS00]. When applied to bipartite graphs, the jump-start routine computes a maximum weight matching. When we compared the running time of the jump-start routine with LEDA's bipartite matching code [MN99, Section 7.8], we found that the jump-start routine is consistently faster. We traced the superiority to the heuristic described in this paper.

2 Dijkstra's Algorithm with Many Targets

It is useful to introduce some more notation. For a node $v \in V$, let $d(v)$ be the shortest path distance from s to v, and let $d_0 = \min\{d(v)\ ;\ v \text{ is free}\}$. If there is no free node reachable from s, $d_0 = +\infty$. Our goal is to compute (1) a node v_0 with $d(v_0) = d_0$ (or an indication that there is no such node), (2) the subset V' of nodes with $d(v) < d_0$, more precisely, $v \in V'$ if $d(v) < d_0$ and $d(v) \geq d_0$ if $v \notin V'$, and (3) the value $d(v)$ for every node $v \in \{v_0\} \cup V'$, i.e., a partial function $\tilde{d}$ with $\tilde{d}(v) = d(v)$ for any $v \in \{v_0\} \cup V'$. (Observe that nodes v with $d(v) = d_0$ may or may not be in V'.) We refer to the problem just described as the single-source many-targets shortest-path (SSMTSP) problem. It is easily solved by Dijkstra's algorithm as shown in Figure 1.

We maintain a priority queue PQ for the nodes of G. The queue is empty initially. For each node we compute a pair (*dist*,*pred*); $dist[v]$ is the tentative shortest path distance to v and $pred[v]$ is the edge into v defining $dist[v]$. For $v \neq s$, we have $pred[v] = nil$ iff $dist[v]$ is undefined. We initialize all *pred*-values to *nil*, set $dist[s]$ to zero and put the pair $(s, 0)$ into the priority queue. In the main loop, we delete a node u with minimum *dist*-value from the priority queue. If u is free, we are done: $v_0 = u$ and V' is the set of nodes removed in preceding iterations. Otherwise, we relax all edges out of u. Consider an edge $e = (u, v)$ and let $c = dist[u] + cost[e]$. If e is the first edge into v that is relaxed (this is the case iff $pred[v] = nil$ and

```
node_pq<int>  PQ(G);
node v; edge e;
dist[s] = 0;
PQ.insert(s,0);
forall_nodes(v,G) pred[v] = nil;
while ( !PQ.empty() )
{ node u = PQ.del_min();
  if ( is_free[u] ) break;
  forall_adj_edges(e,u)
  { v = G.target(e);
    int c = dist[u] + cost[e];
    if  ( pred[v] == nil && v != s ) PQ.insert(v,c);
    else if ( c < dist[v] ) PQ.decrease_p(v,c);
         else continue;
    dist[v] = c; pred[v] = e;
} }
```

Figure 1. Dijkstra's algorithm adapted for many targets. When the first free node is removed from the queue, the algorithm is stopped: v_0 is the node removed last and V' consists of all non-free nodes removed from the queue.

$v \neq s$), we insert (v, c) into PQ. Otherwise, we check whether c is smaller than the current tentative distance of v. If so, we decrease the priority of v in PQ. If a queue operation is performed, we also update $dist[v]$ and $pred[v]$.

Let T be the set of nodes removed from the queue. We relax the edges out of $|T| - 1$ nodes. Except for the initialization of the *pred*-values, the running time of the algorithm is

$$O(1 + \sum_{u \in T \setminus \{v_0\}} (1 + outdeg(u)) + Q) \ ,$$

where Q is the total cost of the queue operations performed. We perform $|T|$ removals from the queue, $|I|$ insertions, where I is the set of nodes reached by the shortest path computation, and some number of decrease priority operations.

The cost for initializing the *pred*-values is $O(n)$. If more than one shortest path computation is performed (which is the case for the application to weighted bipartite matchings), we can avoid the $O(n)$ cost by keeping track which *pred*-values were changed in the preceding shortest path computation and by initializing only these values.

We next describe a heuristic improvement of the scheme above. Let *upper_bound* be the smallest *dist*-value of a free node encountered by the algorithm; *upper_bound* $= +\infty$ initially. We claim that queue operations $PQ.op(v, c)$ with $c \geq$ *upper_bound* may be skipped without affecting correctness. This is clear, since the algorithm stops when the first free node is removed from the queue and since the *dist*-value of this node is certainly at least as small as *upper_bound*. Thus all dist-values less than $d(v_0)$ will be computed correctly. The modified algorithm may output a different node v_0 and a different set V'. However, if all distances are pairwise distinct the same node v_0 and the same set V' as in the basic algorithm are computed. The pruning heuristic can conceivably save on queue operations, since fewer insertions and decrease priority operations may be performed.

```
node_pq<int>  PQ(G);
node v; edge e;
dist[s] = 0;
PQ.insert(s,0);
forall_nodes(v,G) pred[v] = nil;
int upper_bound; bool upper_bound_is_defined = false;
while ( !PQ.empty() )
{ node u = PQ.del_min();
  if ( is_free[u] ) break;
  forall_adj_edges(e,u)
  { v = G.target(e);
    int c = dist[u] + cost[e];
    if ( upper_bound_is_defined && c >= upper_bound ) continue;
    if ( free[v] ) { upper_bound = c; upper_bound_is_defined = true; }
    if ( pred[v] == nil && v != s ) PQ.insert(v,c);
    else if ( c < dist[v] ) PQ.decrease_p(v,c);
         else continue;
    dist[v] = c; pred[v] = e;
} }
```

Figure 2. Dijkstra's algorithm for many targets with a pruning heuristic. An *upper_bound* for $d(v_0)$ is maintained and queue operations $PQ.op(v, c)$ with $c \geq$ *upper_bound* are not performed.

Figure 2 shows the algorithm with the heuristic added. In the program, we use a boolean flag *upper_bound_is_defined* to indicate whether *upper_bound* has value $+\infty$.

3 Analysis

We perform a partial analysis of the basic and the modified version of Dijkstra's algorithm for many targets. We use n for the number of nodes, m for the expected number of edges and f for the expected number of free nodes. We assume that our graphs are random graphs in the $B(n, p)$ model with $p = m/n^2$, i.e., each of the n^2 possible edges is picked independently and uniformly at random with probability p. We use c to denote $pn = m/n$. We also assume that a node is free with probability $q = f/n$ and that edge costs are random reals between 0 and 1. We could alternatively use the model in which all graphs with m edges are equally likely and in which the free nodes form a random subset of f nodes. The results would be similar. We are mainly interested in the case, where $p = c/n$ for a small constant c, say $2 \leq c \leq 10$, and q a constant, i.e., the expected number of free nodes is a fixed fraction of the nodes.

Deletions from the Queue: We first analyze the number of nodes removed from the queue. If our graph were infinite and all nodes were reachable from s, the expected number would be $1/q$, namely the expected number of trials until the first head occurs in a sequence of coin tosses with success probability q. However, our graph is finite (not really a serious difference if n is large) and only a subset of the nodes is reachable from s. Observe, that the probability that s has no outgoing edge

is $(1-p)^n \approx e^{-c}$. This probability is non-negligible. We proceed in two steps. We first analyze the number of nodes removed from the queue given the number R of nodes reachable from s and in a second step review results about the number R of reachable nodes.

Lemma 1 *Let R be the number of nodes reachable from s in G and let T be the number of iterations, i.e., in iteration T the first free node is removed from the queue. If there is no free node reachable from s, $T = R$. Then,*

$$\mathbf{Pr}\left(T = t \mid R = r\right) = \begin{cases} (1-q)^{t-1}q & \text{if } 1 \le t < r \ , \\ (1-q)^{t-1} & \text{if } t = r \ . \end{cases}$$

Moreover, for the expected number of iterations we have:

$$\mathbf{E}\left[T \mid R = r\right] = \frac{1}{q} - \frac{(1-q)^r}{q} \ .$$

Proof: Since each node is free with probability $q = f/n$ and since the property of being free is independent from the order in which nodes are removed from the queue, we have $\mathbf{Pr}\left(T = t \mid R = r\right) = (1-q)^{t-1}q$ and $\mathbf{Pr}\left(T \ge t \mid R = r\right) = (1-q)^{t-1}$, for $1 \le t < r$. If $t = r$, $\mathbf{Pr}\left(T = r \mid R = r\right) = (1-q)^{r-1} = \mathbf{Pr}\left(T \ge r \mid R = r\right)$.

The expected number of iterations is

$$\begin{aligned} \mathbf{E}\left[T \mid R = r\right] &= \sum_{t \ge 1} \mathbf{Pr}\left(T \ge t \mid R = r\right) = \sum_{1 \le t < r} (1-q)^{t-1} + (1-q)^{r-1} \\ &= \frac{1-(1-q)^r}{1-(1-q)} = \frac{1}{q} - \frac{(1-q)^r}{q} \ . \end{aligned}$$

□

The preceding Lemma gives us information about the number of deletions from the queue. The expected number of edges relaxed is $c\mathbf{E}\left[(T-1) \mid R = r\right]$ since $T-1$ non-free nodes are removed from the queue and since the expected out-degree of every node is $c = m/n$. We conclude that the number of edges relaxed is about $((1/q) - 1)(m/n)$.

Now, how many nodes are reachable from s? This quantity is analyzed in [ASE92, pages 149–155]. Let $\alpha > 0$ be such that $\alpha = 1 - \exp(-c\alpha)$, and let R be the number of nodes reachable from s. Then R is bounded by a constant with probability about $1 - \alpha$ and is approximately αn with probability about α. More precisely, for every $\epsilon > 0$ and $\delta > 0$, there is a t_0 such that for all sufficiently large n, we have

$$1 - \alpha - 2\epsilon \ \le \ \mathbf{Pr}\left(R \le t_0\right) \ \le \ 1 - \alpha + \epsilon$$

and

$$\alpha - 2\epsilon \ \le \ \mathbf{Pr}\left((1-\delta)\alpha n < R < (1+\delta)\alpha n\right) \ \le \ \alpha + 3\epsilon \ .$$

Table 1 indicates that small values of ϵ and δ work even for moderate n. For $c = 2$, we have $\alpha \approx 0.79681$. We generated 10000 graphs with $n = 1000$ nodes and 2000 edges and determined the number of nodes reachable from a given source

c	2	5	8	8
α	0.7968	0.993	0.9997	0.9997
MS	15	2	1	1
ML	714	981	996	1995
R	796.5	993	999.7	1999.3
F	7958	9931	9997	9995

Table 1. For all experiments (except the one in the last column) we used random graphs with $n = 1000$ nodes and $m = cn$ edges. For the last column we chose $n = 2000$ in order to illustrate that the dependency on n is weak. Nodes were free with probability q. The following quantities are shown; for each value of q and c we performed 10^4 trials.
α: the solution of the equation $\alpha = 1 - \exp(-c\alpha)$.
MS: the maximal number of nodes reachable from s when few nodes are reachable.
ML: the minimal number of nodes reachable from s when many nodes are reachable.
R: the average number of nodes reachable from s when many nodes are reachable.
F: the number of times many nodes are reachable from s.

node s. This number was either smaller than 15 or larger than 714. The latter case occurred in $7958 \approx \alpha \cdot 10000$ trials. Moreover, the average number of nodes reachable from s in the latter case was $796.5 \approx \alpha \cdot 1000 = \alpha n$.

For the sequel we concentrate on the case that $(1-\delta)\alpha n$ nodes are reachable from s. In this situation, the probability that all reachable nodes are removed from the queue is about

$$(1-q)^{\alpha n} = \exp(\alpha n \ln(1-q)) \approx \exp(-\alpha n q) = \exp(-\alpha f) \ .$$

This is less than $1/n^2$, if $c \geq 2$ and $f \geq 4 \ln n$,[1] an assumption which we are going to make. We use the phrase "R and f are large" to refer to this assumption.

Insertions into the Queue: We next analyze the number of insertions into the queue, first for the standard scheme.

Lemma 2 *Let IS be the number of insertions into the queue in the standard scheme. Then $\mathbf{E}[IS \mid T = t] = n - (n-1)(1-p)^{t-1}$ and*

$$\mathbf{E}[IS \mid R \text{ and } f \text{ are large}] = \frac{c(1-q)}{q + (1-q)c/n} - \frac{(1-q)c/n}{q + (1-q)c/n} + 1 + o(1)$$

Proof: In the standard scheme every node that is reached by the search is inserted into the queue. If we remove a total of t elements from the queue, the edges out of $t-1$ elements are scanned. A node v, $v \neq s$, is not reached if none of these $t-1$ nodes has an edge into v. The probability for this to happen is $(1-p)^{t-1}$ and hence the expected number $\mathbf{E}[IS \mid T = t]$ of nodes reached is $n - (n-1)(1-p)^{t-1}$. This is also the number of insertions into the queue under the standard scheme.

[1] For $c \geq 2$, we have $\alpha > 1/2$ and thus $\exp(-\alpha f) < \exp(-\frac{1}{2}f)$. Choosing $f \geq 4 \ln n$, we obtain: $\exp(-\alpha f) < 1/n^2$.

If R and f are large, we have

$$
\begin{aligned}
&\mathbf{E}\left[IS \mid R \text{ and } f \text{ are large}\right] \\
&\quad= \sum_{t=1}^{R} \mathbf{E}\left[IS \mid T = t \text{ and } R \text{ and } f \text{ are large}\right] \mathbf{Pr}\left(T = t \mid R \text{ and } f \text{ are large}\right) \\
&\quad= \sum_{t\geq 1} \left(n - (n-1)(1-p)^{t-1}\right)(1-q)^{t-1}q + \left(n - (n-1)(1-p)^{R-1}\right)(1-q)^{R-1} \\
&\qquad\quad - \sum_{t\geq R} \left(n - (n-1)(1-p)^{t-1}\right)(1-q)^{t-1}q \\
&\quad= \sum_{t\geq 1} \left(n - (n-1)(1-p)^{t-1}\right)(1-q)^{t-1}q + o(1) \\
&\quad= n - q(n-1)\sum_{t\geq 0}(1-q)^t(1-p)^t + o(1) \\
&\quad= n - q(n-1)\frac{1}{1-(1-p)(1-q)} + o(1) \\
&\quad= n - 1 - (n-1)\frac{q}{p+q-pq} + 1 + o(1) \\
&\quad= (n-1)\frac{p-pq}{p+q-pq} + 1 + o(1) \\
&\quad= \frac{c(1-q)}{q+(1-q)c/n} - \frac{(1-q)c/n}{q+(1-q)c/n} + 1 + o(1) \\
&\quad\approx \frac{c}{q} - c + 1 + o(1)\ .
\end{aligned}
$$

□

The final approximation is valid if $c/n \ll q$. The approximation makes sense intuitively. We relax the edges out of $1/q - 1$ nodes and hence relax about c times as many edges. There is hardly any sharing of targets between these edges, if n is large. We conclude that the number of insertions into the queue is $\frac{c}{q} - c + 1$.

Observe that the standard scheme makes about c/q insertions into but only $1/q$ removals from the queue. This is where the refined scheme saves. Let *INRS* be the number of nodes which are **i**nserted into the queue but **n**ever **r**emoved in the **s**tandard scheme. Then, by the above,

$$\mathbf{E}\left[INRS \mid R \text{ and } f \text{ are large}\right] \quad\approx\quad \frac{c}{q} - c + 1 - \frac{1}{q} \approx \frac{c-1}{q}\ .$$

The standard scheme also performs some *decrease_p* operations on the nodes inserted but never removed. This number is small since the average number of incoming edges scanned per node is small.

We turn to the refined scheme. We have three kinds of savings.

- Nodes that are removed from the queue may incur fewer queue operations because they are inserted later or because some distance decreases do not

lead to a queue operation. This saving is small since the number of distance decreases is small (recall that only few incoming edges per node are scanned)

- Nodes that are never removed from the queue in the standard scheme are not inserted in the refined scheme. This saving is significant and we will estimate it below.
- Nodes that are never removed from the queue in the standard scheme are inserted in the refined scheme but fewer decreases of their distance labels lead to a queue operation. This saving is small for the same reason as in the first item.

We concentrate on the set of nodes that are inserted into but never removed from the queue in the standard scheme. How many of these *INRS* insertions are also performed in the refined scheme? We use *INRR* to denote their number. We compute the expectation of *INRR* conditioned on the event E_l, $l \in \mathbb{N}$, that in the standard scheme there are exactly l nodes which are inserted into the queue but not removed.

Let $e_1 = (u_1, v_1), \ldots, e_l = (u_l, v_l)$ be the edges whose relaxations lead to the insertions of nodes that are not removed, labeled in the order of their relaxations. Then, $d(u_i) \leq d(u_{i+1})$, $1 \leq i \leq l-1$, since nodes are removed from the queue in non-decreasing order of their distance values.

Node v_i is inserted with value $d(u_i) + w(e_i)$; $d(u_i) + w(e_i)$ is a random number in the interval $[d(t), d(u_i) + 1]$, where t is the target node closest to s, since the fact that v_i is never removed from the queue implies $d(u_i) + w(e_i) \geq d(t)$ but reveals nothing else about the value of $d(u_i) + w(e_i)$.

In the refined scheme e_i leads to an insertion only if $d(u_i) + w(e_i)$ is smaller than $d(u_j) + w(e_j)$ for every free v_j with $j < i$. The probability for this event is at most $1/(k+1)$, where k is the number of free v_j preceding v_i. The probability would be exactly $1/(k+1)$ if the values $d(u_h) + w(e_h)$, $1 \leq h \leq i$, were all contained in the same interval. Since the upper bound of the interval containing $d(u_h) + w(e_h)$ increases with h, the probability is at most $1/(k+1)$.

Thus (the expectation is conditioned on the event E_l)

$$\begin{aligned}
\mathbf{E}\left[INRR \mid E_l\right] &\leq \sum_{1 \leq i \leq l} \sum_{0 \leq k < i} \binom{i-1}{k} q^k (1-q)^{i-1-k} \frac{1}{k+1} \\
&= \sum_{1 \leq i \leq l} \frac{1}{iq} \sum_{0 \leq k < i} \binom{i}{k+1} q^{k+1} (1-q)^{i-(k+1)} \\
&= \sum_{1 \leq i \leq l} \frac{1}{iq} \sum_{1 \leq k \leq i} \binom{i}{k} q^k (1-q)^{i-k} \\
&= \sum_{1 \leq i \leq l} \frac{1}{iq} \left(1 - (1-q)^i\right)
\end{aligned}$$

where the first equality follows from $\binom{i-1}{k}\frac{1}{k+1} = \frac{1}{i}\binom{i}{k+1}$. The final formula can also be interpreted intuitively. There are about iq free nodes preceding v_i and hence v_i is inserted with probability about $1/(iq)$.

In order to estimate the final sum we split the sum at a yet to be determined index i_0. For $i < i_0$, we estimate $(1 - (1-q)^i) \leq iq$, and for $i \geq i_0$, we use

$(1-(1-q)^i) \leq 1$. We obtain

$$\mathbf{E}\left[INRR \mid E_l\right] \leq k + \frac{1}{q} \sum_{k \leq i \leq l} \frac{1}{i} \approx k + \frac{1}{q} \ln \frac{l}{k} \ .$$

For $k = 1/q$ (which minimizes the final expression[2]) we have

$$\mathbf{E}\left[INRR \mid E_l\right] \leq \frac{1}{q} \cdot (1 + \ln(lq)) \ .$$

Since $\ln(lq)$ is a convex function of l (its first derivative is positive and its second derivative is negative), we obtain an upper bound on the expectation of *INRR* conditioned on R and f being large, if we replace *INRS* by its expectation. We obtain

$$\begin{aligned} \mathbf{E}\left[INRR \mid R \text{ and } f \text{ are large}\right] &\leq \frac{1}{q} \cdot (1 + \ln(q\mathbf{E}\left[INRS \mid R \text{ and } f \text{ are large}\right])) \\ &\approx \frac{1}{q} \cdot \left(1 + \ln\left(q\frac{c-1}{q}\right)\right) \\ &= \frac{1}{q} \cdot (1 + \ln(c-1)) \ . \end{aligned}$$

We can now finally lower bound the number S of queue operations saved. By the above the saving is at least *INRS* − *INRR*. Thus

$$\begin{aligned} \mathbf{E}\left[S \mid R \text{ and } f \text{ are large}\right] &\geq \frac{c-1}{q} - \frac{1}{q}(1 + \ln(c-1)) \\ &\approx \frac{c}{q}\left(1 - \frac{2 + \ln c}{c}\right) \ . \end{aligned}$$

We have a guaranteed saving of $\frac{2+\ln c}{c}$. Moreover, if $\frac{2+\ln c}{c} < 1$ we are guaranteed to save a constant fraction of the queue operations. For example, if $c = 8$, we will save at least a fraction of $1 - \frac{2+\ln 8}{8} \approx 0.49$ of the queue operations. The actual savings are higher, see Table 2. Also, there are substantial savings, even if the assumption of R and f being large does not hold (e.g., for $c = 2$ and $q = 0.02$).

It is interesting to observe how our randomness assumptions were used in the argument above. G is a random graph and hence the number of nodes reachable from s is either bounded or very large. Also, the expected number of nodes reached after t removals from the queue has a simple formula. The fact that a node is free with fixed probability gives us the distribution of the number of deletions from the queue. In order to estimate the savings resulting from the refined scheme we use that every node has the same chance of being free and that edge weights are random. For this part of the argument we do not need that our graph is random.

4 Bipartite Matching Problems

As already mentioned in the introduction, the starting point for our investigation was the observation that the pruning heuristic described in Section 2 significantly

[2] Take the derivative with respect to k ... .

c	2	2	2	5	5	5	8	8	8	8
q	0.02	0.06	0.18	0.02	0.06	0.18	0.02	0.06	0.18	0.18
D	49.60	16.40	5.51	49.33	16.72	5.50	50.22	16.79	5.61	5.53
D^*	50.00	16.67	5.56	50.00	16.67	5.56	50.00	16.67	5.56	5.56
IS	90.01	31.40	10.41	195.20	73.71	22.98	281.30	112.90	36.45	36.52
IS^*	90.16	31.35	10.02	197.60	73.57	23.25	282.30	112.30	36.13	36.77
$INRS$	40.41	15.00	4.89	145.80	56.99	17.49	231.00	96.07	30.85	30.99
$INRS^*$	40.16	14.68	4.46	147.60	56.90	17.69	232.30	95.60	30.57	31.22
$INRR$	11.00	4.00	1.00	35.00	12.00	4.00	51.00	18.00	5.00	5.00
$INRR^*$	39.05	14.56	4.34	104.10	37.13	11.99	126.80	45.78	15.03	15.15
DP_s	1.42	0.19	0.02	13.78	1.90	0.19	36.55	5.28	0.56	0.28
DP_r	0.71	0.09	0.01	2.63	0.31	0.03	4.60	0.50	0.05	0.03
Q_s	140.00	46.98	14.94	257.30	91.33	27.67	367.00	133.90	41.62	41.34
Q_r	110.40	36.12	11.52	134.50	45.33	13.97	154.40	50.85	16.00	15.77
S	29.58	10.86	3.42	122.80	46.00	13.69	212.70	83.08	25.62	25.57
S^*	1.12	0.13	0.12	43.47	19.77	5.70	105.50	49.82	15.54	16.07
P	21.12	23.11	22.87	47.74	50.37	49.50	57.94	62.03	61.55	61.85

Table 2. For all experiments (except the one in the last column) we used random graphs with $n = 1000$ nodes and $m = cn$ edges. For the last column we chose $n = 2000$ in order to illustrate that the dependency on n is weak. Nodes were free with probability q. The following quantities are shown; for each value of q and c we performed 10^4 trials. Trials where only a small number of nodes were reachable from s were ignored, i.e., about $(1-\alpha)n$ trials were ignored.

D: the number of deletions from the queue.
$D^* = 1/q$: the predicted number of deletions from the queue.
IS: the number of insertions into the queue in the standard scheme.
$IS^* = \frac{c(1-q)}{q+(1-q)c/n} - \frac{(1-q)c/n}{q+(1-q)c/n} + 1$: the predicted number of insertions into the queue.
$INRS$: the number of nodes inserted but never removed.
$INRS^* = IS^* - D^*$: the predicted number.
$INRR$: the number of extra nodes inserted by the refined scheme.
$INRR^* = \frac{1}{q} \cdot (1 + \ln(qN^*))$: the predicted number.
DP_s: the number of decrease priority operations in the standard scheme.
DP_r: the number of decrease priority operations in the refined scheme.
Q_s: the total number of queue operations in the standard scheme.
Q_r: the total number of queue operations in the refined scheme.
$S = Q_s - Q_r$: the number of saved queue operations.
S^*: the lower bound on the number of saved queue operations.
$P = S/Q_s$: the percentage of queue operations saved.

improved the running time of LEDA's bipartite matching algorithm [MN99, Section 7.8].

The input for a weighted bipartite matching algorithm is an edge-weighted bipartite graph $G = (A \dot\cup B, E, w)$ where the cost function $w : E \mapsto \mathbb{R}$ assigns a real weight to every edge. One may either ask for a perfect matching of maximal weight (the weighted perfect matching problem or the assignment problem) or simply for a matching of maximal weight. Essentially the same algorithms apply to both problems; we discuss the assignment problem.

A popular algorithm for the assignment problem follows the primal dual paradigm [AMO93, Section 12.4], [MN99, Section 7.8], [Gal86]. The algorithm constructs a perfect matching and a dual solution simultaneously. A dual solution is

simply a function $\pi : V \mapsto \mathbb{R}$ that assigns a real potential to every node. We use V to denote $A \cup B$. The algorithm maintains a matching M and a potential function π with the property that

(a) $w(e) \leq \pi(a) + \pi(b)$ for every edge $e = (a, b)$,

(b) $w(e) = \pi(a) + \pi(b)$ for every edge $e = (a, b) \in M$ and

(c) $\pi(b) = 0$ for every free[3] node $b \in B$.

Initially, $M = \emptyset$, $\pi(a) = \max_{e \in E} w(e)$ for every $a \in A$ and $\pi(b) = 0$ for every $b \in B$. The algorithm stops when M is a perfect matching[4] or when it discovers that there is no perfect matching. The algorithm works in phases. In each phase the size of the matching is increased by one (or it is determined that there is no perfect matching).

A phase consists of the search for an augmenting path of minimum reduced cost. An augmenting path is a path starting at a free node in A, ending at a free node in B and using alternately edges not in M and in M. The reduced cost of an edge $e = (a, b)$ is defined as $\overline{w}(e) = \pi(a) + \pi(b) - w(e)$; observe that edges in M have reduced cost zero and that all edges have non-negative reduced cost. The reduced cost of a path is simply the sum of the reduced costs of the edges contained in it. There is no need to search for augmenting paths from all free nodes in A; it suffices to search for augmenting paths from a single arbitrarily chosen free node $a_0 \in A$.

If no augmenting path starting in a_0 exists, there is no perfect matching in G and the algorithm stops. Otherwise, for every $v \in V$, let $d(v)$ be the minimal reduced cost of an alternating path from a_0 to v. Let $b_0 \in B$ be a free node in B which minimizes $d(b)$ among all free nodes b in B. We update the potential function according to the rules (we use π' to denote the new potential function):

(d) $\pi'(a) = \pi(a) + \max(d(b_0) - d(a), 0)$ for all $a \in A$,

(e) $\pi'(b) = \pi(b) - \max(d(b_0) - d(b), 0)$ for all $b \in B$.

It is easy to see that this change maintains (a), (b), and (c) and that all edges on the least cost alternating path p from a_0 to b_0 become tight[5]. We complete the phase by switching the edges on p: matching edges on p become non-matching and non-matching edges become matching edges. This increases the size of the matching by one.[6]

A phase is tantamount to a SSMTSP problem: a_0 is the source and the free nodes are the targets. We want to determine a target (= free node) b_0 with minimal distance from a_0 and the distance values of all nodes v with $d(v) < d(b_0)$. For nodes

[3] A node is free if no edge in M is incident to it.

[4] It is easy to see that M has maximal weight among all perfect matchings. Observe that if M' is any perfect matching and π is any potential function such that (a) holds then $w(M') \leq \sum_{v \in V} \pi(v)$. If (b) also holds, we have a pair (M', π) with equality and hence the matching has maximal weight (and the node potential has minimal weight among all potentials satisfying (a)).

[5] An edge is called tight if its reduced cost is zero.

[6] The correctness of the algorithm can be seen as follows. The algorithm maintains properties (a), (b), and (c) and hence the current matching M is optimal in the following sense. Let A_m be the nodes in A that are matched. Then M is a maximal weight matching among the matchings that match the nodes in A_m and leave the nodes in $A \setminus A_m$ unmatched. Indeed if M' is any such matching then $w(M') \leq \sum_{a \in A_m} \pi(a) + \sum_{b \in B} \pi(b) = w(M)$, where the inequality follows from (a) and (c) and the equality follows from (b) and (c).

C	n	m	LEDA$^-$	MS$^-$	LEDA$^+$	MS$^+$	LEDA*	MS*
U	10000	40000	24.31	8.98	24.84	9.10	26.28	9.38
U	10000	60000	33.25	6.63	34.36	6.69	36.33	6.58
U	10000	80000	35.81	4.90	36.59	4.54	38.10	4.53
U	20000	80000	83.19	30.17	88.12	30.58	91.00	31.73
U	20000	120000	114.65	21.45	119.10	21.70	124.96	22.63
U	20000	160000	131.01	14.29	134.42	14.03	140.73	14.86
U	40000	160000	287.28	107.38	304.81	108.43	317.21	108.70
U	40000	240000	420.63	78.71	445.44	78.52	461.37	78.21
U	40000	320000	458.53	46.28	481.13	45.34	502.95	46.19
R	10000	40000	1.53	1.09	1.38	0.88	1.03	0.69
R	10000	60000	5.11	2.66	4.63	2.36	3.73	1.94
R	10000	80000	15.60	6.91	14.89	6.75	13.43	6.12
R	20000	80000	3.35	2.29	2.65	1.84	2.02	1.46
R	20000	120000	11.06	5.76	10.16	5.16	8.22	4.26
R	20000	160000	36.60	15.30	36.50	14.76	33.36	13.47
R	40000	160000	7.33	4.99	5.79	4.01	4.41	3.20
R	40000	240000	24.51	12.65	22.58	11.36	18.38	9.49
R	40000	320000	86.86	35.62	83.29	33.70	75.85	30.65
L	10000	40000	10.19	7.30	10.26	7.33	9.43	7.01
L	10000	60000	15.25	9.27	14.32	8.94	14.04	8.63
L	10000	80000	17.33	10.30	17.21	10.06	16.14	9.70
L	20000	80000	27.09	19.86	28.19	20.22	27.07	19.35
L	20000	120000	46.65	28.43	45.98	28.22	42.26	26.77
L	20000	160000	57.13	31.74	56.03	31.89	53.14	29.78
L	40000	160000	89.37	60.92	85.00	59.24	83.48	60.50
L	40000	240000	151.47	87.55	150.83	87.04	148.62	88.47
L	40000	320000	174.69	97.19	175.29	96.40	171.17	95.57

Table 3. Effect of the pruning heuristic for the maximum-weight bipartite matching algorithm. LEDA stands for LEDA's bipartite matching algorithm (up to version LEDA-4.2) as described in [MN99, Section 7.8] and MS stands for a modified implementation using the pruning heuristic. We created random graphs with n vertices on each side of the bipartition and m edges inbetween. U, R and L denote that random weights were chosen out of the range [1], $[1,\ldots,1000]$ and $[1000,\ldots,1005]$, respectively. $^-$, $^+$ and * indicates whether no, a greedy or a refined heuristic for constructing an initial matching was used. The running time is stated in CPU-seconds and is an average of 5 trials.

v with $d(v) \geq d(b_0)$, there is no need to know the exact distance. It suffices to know that the distance is at least $d(b_0)$.

Table 3 shows the effect of the pruning heuristic for the bipartite matching algorithm. LEDA stands for LEDA's bipartite matching algorithm (up to version LEDA-4.2) as described in [MN99, Section 7.8] and MS stands for a modified implementation with the pruning heuristic. We timed our algorithms in combination with three heuristics for finding an initial matching: starting with an empty matching, using a greedy matching or using a refined greedy heuristic (see [MN99, Section 7.8] for details). The heuristics used to construct an initial matching have little influence on the running time. The improved code will be part of LEDA Version 4.3.

References

[AMO93] R.K. Ahuja, T.L. Magnanti, and J.B. Orlin. *Network Flows.* Prentice Hall, 1993.

[ASE92] N. Alon, J.H. Spencer, and P. Erdös. *The Probabilistic Method*. John Wiley & Sons, 1992.

[Gal86] Z. Galil. Efficient algorithms for finding maximum matching in graphs. *ACM Computing Surveys*, 18(1):23–37, 1986.

[MN99] K. Mehlhorn and S. Näher. *The LEDA Platform for Combinatorial and Geometric Computing*. Cambridge University Press, 1999. 1018 pages.

[MS00] K. Mehlhorn and G. Schäfer. Implementation of $O(nmlogn)$ weighted matchings in general graphs: The power of data structures. In *Workshop on Algorithm Engineering (WAE)*, Lecture Notes in Computer Science, to appear, 2000. `www.mpi-sb.mpg.de/~mehlhorn/ftp/WAE00.ps.gz`.

[Sch00] G. Schäfer. Weighted matchings in general graphs. Master's thesis, Fachbereich Informatik, Universität des Saarlandes, Saarbrücken, Germany, 2000.