AARHUS UNIVERSITY

Coversheet

This is the accepted manuscript (post-print version) of the article.

Contentwise, the accepted manuscript version is identical to the final published version, but there may be differences in typography and layout.

How to cite this publication

Please cite the final published version:

Wøhlk, S., & Laporte, G. (2017). Computational Comparison of Several Greedy Algorithms for the Minimum Cost Perfect Matching Problem on Large Graphs. *Computers & Operations Research*, 87, 107-113. <https://doi.org/10.1016/j.cor.2017.06.006>

Publication metadata

Title:	Computational Comparison of Several Greedy Algorithms for the Minimum Cost Perfect Matching Problem on Large Graphs
Author(s):	Sanne Wøhlk & Gilbert Laporte
Journal:	Computers & Operations Research
DOI/Link:	10.1016/j.cor.2017.06.006
Document version:	Accepted manuscript (post-print)

General Rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognize and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

If the document is published under a Creative Commons license, this applies instead of the general rights.

Computational Comparison of Several Greedy Algorithms for the Minimum Cost Perfect Matching Problem on Large Graphs

Sanne Wøhlk^{a,*}, Gilbert Laporte^b

^a*CORAL - Cluster for Operations Research and Logistics, Department of Economics and Business Economics, Aarhus University, Fuglesangs allé 4, DK-8210 Aarhus V, Denmark*

^b*GERAD and Canada Research Chair in Distribution Management, HEC Montréal, 3000, chemin de la Côte-Sainte-Catherine, Montréal, Canada H3T 2A7*

Abstract

The aim of this paper is to computationally compare several algorithms for the Minimum Cost Perfect Matching Problem on an undirected complete graph. Our work is motivated by the need to solve large instances of the Capacitated Arc Routing Problem (CARP) arising in the optimization of garbage collection in Denmark. Common heuristics for the CARP involve the optimal matching of the odd-degree nodes of a graph. The algorithms used in the comparison include the CPLEX solution of an exact formulation, the LEDA matching algorithm, a recent implementation of the Blossom algorithm, as well as six constructive heuristics. Our results show that two of the constructive heuristics consistently exhibit the best behavior compared with the other four.

Keywords: Perfect Matching Problem, Blossom algorithm, Capacitated Arc Routing Problem, Rural Postman Problem, Garbage collection

1. Introduction

The purpose of this paper is to provide a computational comparison of several algorithms for the Minimum Cost Perfect Matching Problem in complete graphs, simply referred to as the Matching Problem in what follows. The problem is
5 defined on an undirected complete graph $G(N, E)$, where N is the set of n nodes (n is even) and E is the set of edges $e = (i, j)$, where $i, j \in N$ and $i < j$. For each node $i \in N$, we denote by $\delta(i)$ the set of edges incident to i . Let $c_e = c_{ij}$

*Corresponding author

Email addresses: sanw@econ.au.dk (Sanne Wøhlk), gilbert.laporte@cirrelt.ca (Gilbert Laporte)

be a non-negative weight associated with edge $e = (i, j)$. In practice this weight represents an edge length or a travel cost along it. Given a node i , all other
10 nodes are called neighbors of i and we refer to the node closest to i as its nearest neighbor.

A perfect matching is defined as a set of edges $M \subset E$ for which $|M \cap \delta(i)| = 1$ for all $i \in N$. We denote by $c(M) = \sum_{e \in M} c_e$ the cost of matching M . The Matching Problem is to determine a minimum cost perfect matching M^* in G .

To formally describe the Matching Problem, we define binary variables x_e equal to 1 if and only if edge e belongs to the matching. The problem is then to

$$\begin{aligned} & \text{minimize} && \sum_{e \in E} c_e x_e \\ & \text{subject to} && \sum_{e \in \delta(i)} x_e = 1 \quad \forall i \in N \end{aligned} \tag{1}$$

$$x_e \in \{0, 1\} \quad \forall e \in E. \tag{2}$$

15 Our work is motivated by the need to solve large-scale instances of the Capacitated Arc Routing Problem (CARP) arising in the optimization of garbage collection routes in Denmark. An effective heuristic for the CARP is to first solve a Rural Postman Problem (RPP) by ignoring the vehicle capacity constraints, and then cut the RPP solution into feasible routes as was done, for example
20 in [1] and [2]. By partitioning the RPP solution optimally, an approximation algorithm with fixed ratio can be obtained for the CARP [3]. A good heuristic for the RPP consists of first identifying connected components of required edges (those with a positive demand), connecting these components by means of a minimum cost spanning tree, and then matching the odd-degree nodes of
25 the graph induced by the connected components and the spanning tree [4].

The Matching Problem can be solved exactly by CPLEX, by LEDA's matching algorithm [5], and by the Blossom algorithm [6]. The most recent publicly available implementations of the latter are [7, 8]. The matching problem has also been investigated from an approximation point of view and the most recent
30 algorithm has a worst-case performance ratio of $c(M)/c(M^*) \leq \log^2(n)$ [9]. Either of the exact algorithms can be time consuming. In some large-scale applications, such as ours, it is desirable to compute high quality solutions within short computing times. We have therefore implemented several greedy heuristics that meet this requirement. In this spirit, we did not develop post-optimization
35 procedures which would typically result in higher time complexity. We have compared the heuristics between themselves and with the exact algorithms. Our heuristics are described in Section 2, followed by comparative computational results in Section 3. Conclusions follow in Section 4.

2. Matching Heuristics

40 We now outline the algorithms we have used in our analysis. All algorithms start with an empty matching M and iteratively add a single edge to M until all nodes are matched. Throughout this section, we use S to denote the set of unmatched nodes. When multiple nodes satisfy the selection criteria, ties are broken by always selecting the node or edge that occurs first in the list of nodes
45 or edges.

In the *Greedy* heuristic, we add to the matching an edge of minimum cost that connects two unmatched nodes and repeat this until all nodes are matched. This is detailed in Algorithm 1. The time complexity of *Greedy* is $\mathcal{O}(n^3)$ in our implementation. It can be reduced to $\mathcal{O}(n^2 \log n)$ by using a priority queue.

Algorithm 1 *Greedy*

```

Set  $S = N$ ,  $M = \emptyset$ 
while  $S \neq \emptyset$  do
  Select  $(i, j) \in \arg \min \{c_{i'j'} | i', j' \in S, i' \neq j'\}$ 
   $M = M \cup \{(i, j)\}$ 
   $S = S \setminus \{i, j\}$ 
end while
return  $M$ 

```

50 The intuition behind the *Largest* heuristic is to first identify a match for the nodes that are most isolated in their location. To this aim, we first identify for every node i the distance $H(i)$ to the nearest neighbor, and we then consider the nodes in non-increasing order with respect to $H(i)$. If a node is not already matched, it will be matched to its nearest unmatched neighbor and the process
55 is reiterated until all nodes are matched. *Largest* is outlined in Algorithm 2 and has a time complexity of $\mathcal{O}(n^2)$.

Algorithm 2 *Largest*

```

Set  $S = N$ ,  $M = \emptyset$ 
Set  $H(i) = \min \{c_{ij} | j \in N, j \neq i\}$ ,  $\forall i \in N$ 
while  $S \neq \emptyset$  do
  Select  $i \in \arg \max \{H(i') | i' \in S\}$ 
  Select  $j \in \arg \min \{c_{ij'} | j' \in S, j' \neq i\}$ 
   $M = M \cup \{(i, j)\}$ 
   $S = S \setminus \{i, j\}$ 
end while
return  $M$ 

```

The heuristic *Largest** is a modified version of *Largest* in which the values of $H(i)$ are updated at each iteration, in such a way that the node that is furthest away from its nearest unmatched neighbor can be chosen as i at each iteration.

60 Note that by storing additional information regarding the nodes resulting in the $H(i)$ -values, this update needs only be performed for a subset of the nodes. However, in the worst case, the $H(i)$ -value must be updated for all $i \in S$. This results in a worst-case time complexity of $\mathcal{O}(n^3)$ in our implementation.

In the *Sum* heuristic, the nodes i are considered in non-increasing order with respect to the sum $\Delta(i)$ of their distance to all other nodes. If a node is unmatched, 65 it will be matched to the nearest unmatched neighbor. This is illustrated in Algorithm 3 which has a time complexity of $\mathcal{O}(n^2)$.

Algorithm 3 *Sum*

```

Set  $S = N$ ,  $M = \emptyset$ 
Calculate  $\Delta(i) = \sum_{e \in \delta(i)} c_e$ ,  $\forall i \in N$ 
while  $S \neq \emptyset$  do
    Select  $i \in \arg \max \{\Delta(i') | i' \in S\}$ 
    Select  $j \in \arg \min \{c_{ij'} | j' \in S, j' \neq i\}$ 
     $M = M \cup \{(i, j)\}$ 
     $S = S \setminus \{i, j\}$ 
end while
return  $M$ 

```

*Sum** is a modified version of *Sum* in which the values of $\Delta(i)$ are repeatedly updated to reflect the sum of the distances to all the unmatched neighbors and the next node to be matched to its nearest unmatched neighbor is the one where 70 this sum is highest. The update adds to the run time, but does not increase the worst-case time complexity which remains $\mathcal{O}(n^2)$.

The *Regret* heuristic outlined in Algorithm 4 requires some notation. For every node $i \in S$, let $m_1(i)$ denote the nearest neighbor and let $m_2(i)$ denote the second nearest, i.e. $m_1(i) \in \arg \min \{c_{ij} | j \in S, j \neq i\}$ and $m_2(i) \in 75 \arg \min \{c_{ij} | j \in S, j \neq m_1(i), j \neq i\}$. The regret value of i is then defined as $R(i) = c_{im_2(i)} - c_{im_1(i)}$ and reflects the extra cost incurred if i is not matched to its nearest neighbor.

The intuition behind *Regret* is to give priority to those nodes that we would 80 most regret not to match to their nearest neighbors and to try and match those nodes in the best possible way. Therefore, we repeatedly identify the unmatched node with highest regret value and match it to its nearest unmatched neighbor. During this process, the regret values are updated to only take unmatched nodes into account. As for *Largest**, the update can be time consuming in the worst case, but in practice we can decrease the time for this update significantly by 85 storing $m_1(k)$ and $m_2(k)$ and only update $R(k)$ if one of these are one of the nodes just matched. The worst-case time complexity of *Regret* is $\mathcal{O}(n^3)$.

Algorithm 4 *Regret*

```
Set  $S = N$ ,  $M = \emptyset$ 
Calculate  $R(i)$ ,  $\forall i \in N$ 
while  $S \neq \emptyset$  do
    Select  $i \in \arg \max\{R(i') | i' \in S\}$ 
    Set  $j = m_1(i)$ 
     $M = M \cup \{(i, j)\}$ 
     $S = S \setminus \{i, j\}$ 
    Update  $R(k)$ ,  $\forall k \in S$  with  $m_1(k) \in \{i, j\}$  or  $m_2(k) \in \{i, j\}$ 
end while
return  $M$ 
```

3. Results and Analyses

We now describe our test instances and we provide extensive computational
90 comparisons of our algorithms.

3.1. Test Instances

We have used two sets of data in our analysis. The first set is based on the 88
real-life graphs presented in [10]. These graphs were originally designed for arc
routing problems. For each graph, we have created four matching graphs by
95 selecting a subset of the nodes using different criteria. A complete graph is then
created for matching purposes by using shortest path distances between the
selected nodes in the original graphs. In the first graph, all nodes are selected
(if the number of nodes is odd, we leave out the one designated as “the depot”
in the data). In the second graph, all odd-degree nodes are selected, whereas
100 the odd-degree nodes with respect to the required edges are used in the third
graph. Finally, in the fourth graph, only nodes of degree one are selected (in
case of an odd number, the one listed first is left out). This procedure yields a
total of 352 graphs that we will call real graphs. The number of nodes range
from two to 11640.

105 The second set contains 352 randomly generated graphs with the same number
of nodes as those in the first set. For each of the graphs, the nodes are randomly
generated according to a discrete uniform distribution in the 1000×1000 square.
Based on these nodes, a complete graph is generated with Euclidean distances
between each pair of nodes, rounded up to the nearest integer. These graphs
110 are referred to as random graphs.

The algorithms were implemented in C++ and executed on an Intel Xeon CPU
with 12 cores running at 3.5 GHz and 64 GBs RAM. We used CPLEX 12.6.1 with
standard settings and allowed it to take advantage of the parallel processors.
All other algorithms were executed sequentially. For the Blossom algorithm, we

115 used the freely available Blossom V implementation [8, 11]. We also used LEDA
6.5 [5].

We have only run CPLEX on the 268 graphs of each set having less than 4000 nodes, with a time limit of one hour. Within this time limit, CPLEX could find an optimal matching for 148 real graphs and for 236 random graphs.

120 Blossom V crashed for graphs with $n > 4000$. We contacted the author of this algorithm, but neither he nor we have been able to find the explanation for this unintended behavior, but we believe it regards the interaction between MS Visual Studio and the algorithm. It was therefore only run on graphs with fewer than 4000 nodes. On these graphs, Blossom V successfully obtained an optimal
125 matching for 233 real graphs and for all 268 random graphs.

3.2. The Optimal Cost of Matchings

Figure 1: Optimal costs of matchings for random graphs.

We first consider the cost of optimal matchings for the randomly generated graphs. Figure 1 depicts the cost of the optimal matchings as a function of n and indicates a clear relationship between n and matching cost $f(n)$. We have

found the empirical relationship

$$f(n) = \frac{l\sqrt{n}}{3},$$

where $l = 1000$ is the size of the $l \times l$ -square in which we have generated the nodes. The determination coefficient is $R^2 = 99.95\%$. This estimate is consistent with the result presented in [12, 7]. It is also interesting to observe that the solution cost is proportional to $\sqrt{n}$, as in the Beardwood-Halton-Hammersley approximation formula for the Traveling Salesman Problem [13].

In Figure 2, we show the optimal matching costs of the real graphs as a function of n . The figure indicates that no clear relationship exists for the optimal matching cost for real graphs.

Figure 2: Optimal costs of matchings for real graphs.

3.3. Cost Analysis of Heuristics

We now investigate the six heuristics as regards their ability to obtain low cost matchings.

Figure 3 shows the cost obtained by each of the heuristics for the 352 random graphs as well as the optimal cost. Note that the curve for *Sum* is partially hidden behind the curve for *Sum**. This figure suggests that the two heuristics *Largest* and *Largest** are outperformed by the other heuristics. We also note that *Sum* and *Sum** perform equally well and appear to perform better than *Greedy* and *Regret* on random graphs.

To further analyze the performance of the heuristics, we show in Figure 4 the relative cost obtained by each heuristic as a percentage above the optimal cost. The left-hand side of the figure depicts the results for real graphs, whereas the right-hand side plots the results for random graphs.

Figure 3: Optimal costs and costs obtained by the six heuristics for random graphs.

Figure 4: Cost obtained by the heuristics displayed as percent above the optimum. Left: Real graphs. Right: Random graphs. Top: All heuristics. Bottom: Only the best two heuristics.

To support the figure, Tables 1–4 provide the mean percent above the best (Table 1 for real graphs and Table 2 for random graphs) and the standard deviation of the percentage (Table 3 for real graphs and Table 4 for random graphs), both for the full set of graphs and for different size groups. The first column in each table states the range of n in each group and the second column gives the number of graphs in the group.

Table 1: Mean percent above the optimal for real graphs

	Number	<i>Largest*</i>	<i>Largest</i>	<i>Greedy</i>	<i>Regret</i>	<i>Sum</i>	<i>Sum*</i>
0-500	104	51.7	39.8	22.0	19.6	13.7	12.4
500-1000	55	54.4	42.9	28.3	24.8	12.7	16.2
1000-2000	49	70.2	52.9	29.8	27.8	15.3	14.5
2000-4000	60	72.4	50.3	29.5	28.0	14.3	14.3
4000-6000	41	68.5	51.0	32.9	32.0	14.2	13.9
> 6000	43	60.0	45.2	33.4	30.2	11.2	11.0
All	352	61.2	45.9	28.0	25.7	13.6	13.6

Table 2: Mean percent above the optimal for random graphs

	Number	<i>Largest*</i>	<i>Largest</i>	<i>Greedy</i>	<i>Regret</i>	<i>Sum</i>	<i>Sum*</i>
0-500	104	57.1	46.1	25.1	20.5	15.0	14.3
500-1000	55	71.4	55.1	25.7	24.3	16.7	16.1
1000-2000	49	72.4	54.6	24.3	23.3	17.3	17.1
2000-4000	60	72.6	54.2	24.4	24.2	17.3	17.2
4000-6000	41	73.0	53.5	23.9	24.2	17.2	17.4
> 6000	43	71.9	53.2	23.0	24.2	17.2	17.0
All	352	67.8	51.8	24.6	23.0	16.5	16.1

Table 3: Standard deviation of percent above the optimal for real graphs

	Number	<i>Largest*</i>	<i>Largest</i>	<i>Greedy</i>	<i>Regret</i>	<i>Sum</i>	<i>Sum*</i>
0-500	104	23.8	15.5	12.0	10.7	8.0	7.4
500-1000	55	16.5	12.9	10.2	9.6	3.8	22.1
1000-2000	49	19.6	14.7	9.7	10.5	4.4	3.9
2000-4000	60	25.3	13.5	8.6	8.3	3.9	5.0
4000-6000	41	13.1	9.2	4.4	4.8	3.8	3.3
> 6000	43	14.6	10.7	6.0	6.0	2.8	2.4
All	352	22.2	14.4	10.4	10.1	5.5	10.2

Consider the top part of Figure 4 again. The first result to notice is how stable
the performance of the heuristics is for random graphs compared to the real
graphs. This is supported by Tables 3 and 4 in which the standard deviations
are reported. As an example, the standard deviation for *Largest** is 22.2% for
real graphs compared to 15.8% for random graphs. For both types of graphs,
the results vary more for small graphs than they do for the larger ones. Again,
this is more pronounced for random graphs, where the standard deviation of
*Largest** is 23.3% for $n < 500$, but only 4.3% for $n > 6000$. The corresponding
values for real graphs are 23.8% and 14.6%, respectively.

We now compare the performance of the six heuristics. When we consider the

Table 4: Standard deviation of percent above the optimal for random graphs

	Number	<i>Largest*</i>	<i>Largest</i>	<i>Greedy</i>	<i>Regret</i>	<i>Sum</i>	<i>Sum*</i>
0-500	104	23.3	17.5	8.1	9.8	6.6	6.2
500-1000	55	10.4	6.2	4.2	4.9	2.3	1.9
1000-2000	49	8.7	4.1	2.7	4.1	1.7	1.8
2000-4000	60	7.4	3.0	2.2	2.5	1.1	1.1
4000-6000	41	4.7	2.5	1.5	2.1	0.8	0.9
> 6000	43	4.3	1.9	1.3	1.9	0.8	0.7
All	352	15.8	10.7	5.0	6.3	3.9	3.8

Figure 5: Relative ranks of the heuristics. Left: Real graphs. Right: Random graphs.

top part of Figure 4, it is clear that the two heuristics *Largest** and *Largest* perform significantly worse than the other heuristics with a mean of 61.2% (67.8%) and 45.9% (51.8%), respectively, for all real (random) graphs. Both *Greedy* and *Regret* show an average performance with a mean of 28.0% (24.6%) and 25.7% (23.0%), respectively, for all real (random) graphs. The two heuristics that show the best performance are *Sum* and *Sum**, both having a mean of 13.6% for real graphs and of 16.5% and 16.1%, respectively, for random graphs. The lower part of Figure 4 provides the same plot as the top part, but only shows *Sum* and *Sum**. The plot clearly indicates that the two heuristics do indeed perform equally well and no clear winner among the two can be selected solely based on matching cost. We counted the number of times a heuristic was the best (1), worst (6), and so on. These results are shown in Figure 5 and further support the relative ranking. Even though this figure indicates that *Sum** is superior to *Sum* as it outperforms the other heuristics more often, Tables 1–4 show that it is only marginally better when considering the average performance and standard deviation. Indeed, *Sum* is 13.6 (16.5) percent above the optimal cost for real (random) graphs on average. The corresponding number for *Sum** is 13.6 (16.1) percent, confirming that we cannot identify one of these as the preferred one.

3.4. Run Time Analysis of the Heuristics

Finally, we consider the run time of the six heuristics as well as the Blossom V algorithm, LEDA, and CPLEX, which solve the problem optimally. Since no

Figure 6: Run time for real graphs. All graphs and all algorithms.

significant differences in run time between real graphs and random graphs are observed, we only show the plots for real graphs. This is done in Figures 6 – 9, where run times in seconds are shown as a function of n .

In Figure 6, we show all nine algorithms run on all graphs. It is clear from this plot that the run time for CPLEX rapidly increases and in fact, CPLEX often fails to solve the problem within the time limit of one hour. CPLEX could solve 148 graphs with $n < 4000$ (out of 268) to optimality within the time limit. For random graphs, 236 graphs were solved by CPLEX. This means that for practical purposes CPLEX is not the best tool for solving the Matching Problem. The remaining plots exclude CPLEX.

Figure 7: Run time for real graphs. Zoom on the graphs with $n < 4000$. All except CPLEX.

In Figure 7, we zoom in on the graphs with $n < 4000$. Here we clearly observe that even for relatively large graphs, the Blossom algorithm is quite fast and with this advanced implementation of a very complicated algorithm its run time is indeed very impressive and this algorithm is clearly preferred over CPLEX for solving matching problems with more than about 1000 nodes. Blossom V is significantly faster than the LEDA algorithm. We note that the longest run time observed for Blossom is 12 seconds, whereas the longest time for LEDA on these graphs is 57 seconds. However, we also see that the run time of Blossom tends to grow faster than that of the heuristics, hinting that Blossom may not be a realistic choice for large graphs even if the code could be fixed. It is, however, the best option if optimal matchings are needed.

Figure 8: Run time for real graphs. All graphs. The six heuristics.

The relative run time of the heuristics is most easily observed in Figure 8, which shows the runtime of the six heuristics for all graphs. This plot clearly shows that *Greedy* is slower than the other heuristics, particularly for large graphs. The two heuristics *Largest* and *Sum* show the best run time, while the remaining three are slightly slower. However, even the slowest of the five heuristics runs within 20 seconds on the largest graphs.

Finally, Figure 9 shows the run time of *Sum* and *Sum**, the two heuristics that had the best performance in terms of solution quality. From this plot, we see that even though both show acceptable run times, *Sum* is indeed faster than the version with updates. In fact, *Sum* is the fastest of all the heuristics tested.

4. Conclusions

We have compared several algorithms for the solution of the Matching Problem arising, namely, in the solution of the Capacitated Arc Routing Problem. These

Figure 9: Run time for real graphs. All graphs. Only *Sum* and *Sum**.

220 include CPLEX applied to a standard integer linear programming model, the
Kolmogorov implementation of the Blossom algorithm, and six new constructive
heuristics. For graphs involving fewer than 4000 nodes, the Blossom algorithm
outperforms CPLEX in terms of computing time. For larger instances, which
cannot be solved optimally, the constructive heuristics called *Sum* and *Sum**
225 consistently outperform the other four, but *Sum* is faster.

Acknowledgments

The authors thank Serge Bisailon for his technical support. This project
is funded by the Danish Council for Independent Research - Social Sciences.
Project “Transportation issues related to waste management” [grant number
230 4182-00021] and by the Natural Sciences and Engineering Research Council of
Canada [grant number 2015-06189]. This support is gratefully acknowledged.
Thanks are due to the referees for their valuable comments.

References

- [1] A. Hertz, G. Laporte, M. Mittaz, A tabu search heuristic for the capacited
235 arc routing problem, *Operations Research* 48 (2000) 129–135.
- [2] C. Prins, N. Labadi, M. Reghioui, Tour splitting algorithms for vehicle
routing problems, *International Journal of Production Research* 47 (2009)
507–536.
- [3] S. Wøhlk, An approximation algorithm for the capacitated arc routing
240 problem, *The Open Operational Research Journal* 2 (2008) 8–12.

- [4] G. N. Frederickson, Approximation algorithms for some postman problems, *Journal of the Association for Computing Machinery* 26 (1979) 538–554.
- [5] Algorithmic Solutions Software GmbH, <http://www.algorithmic-solutions.com>.
- 245 [6] J. Edmonds, Paths, trees, and flowers, *Canadian Journal of Mathematics* 17 (1965) 449–467.
- [7] W. Cook, A. Rohe, Computing minimum-weight perfect matchings, *INFORMS Journal on Computing* 11 (1999) 138–148.
- 250 [8] V. Kolmogorov, Blossom V: A new implementation of a minimum cost perfect matching algorithm, *Mathematical Programming Computation* 1 (2009) 43–67.
- [9] M. Wattenhofer, R. Wattenhofer, Fast and simple algorithms for weighted perfect matching, *Electronic Notes in Discrete Mathematics* 17 (2004) 285–291.
- 255 [10] L. Kiilerich, S. Wøhlk, New large-scale data instances for CARP and variations of CARP, *INFOR: Information Systems and Operational Research Accepted*. doi:10.1080/03155986.2017.1303960.
- [11] V. Kolmogorov, <http://pub.ist.ac.at/~vnk/software.html#blossom5>.
- 260 [12] C. H. Papadimitriou, The probabilistic analysis of matching heuristics, *Proceedings of the 15th Annual Allerton Conference on Communication, Control, and Computing* (1977) 368–378.
- [13] J. Beardwood, J. H. Halton, J. M. Hammersley, The shortest path through many points, *Proceedings of the Cambridge Philosophical Society* 55 (1959) 299–327.