

Engineering Fast Algorithms for the Bottleneck Matching Problem

Ioannis Panagiotas ✉

Neo4j, Malmö, Sweden

Grégoire Pichon ✉ 🏠 

Université Claude Bernard Lyon 1 and LIP, France

UMR5668, CNRS, ENS de Lyon, Inria, UCBL1, France

Somesh Singh ✉ 🏠 

CNRS and LIP, ENS de Lyon, France

UMR5668, CNRS, ENS de Lyon, Inria, UCBL1, France

Bora Uçar ✉ 🏠 

CNRS and LIP, ENS de Lyon, France

UMR5668, CNRS, ENS de Lyon, Inria, UCBL1, France

Abstract

We investigate the maximum bottleneck matching problem in bipartite graphs. Given a bipartite graph with nonnegative edge weights, the problem is to find a maximum cardinality matching in which the minimum weight of an edge is the maximum. To the best of our knowledge, there are two widely used solvers for this problem based on two different approaches. There exists a third known approach in the literature, which seems inferior to those two which is presumably why there is no implementation of it. We take this third approach, make theoretical observations to improve its behavior, and implement the improved method. Experiments with the existing two solvers show that their run time can be too high to be useful in many interesting cases. Furthermore, their performance is not predictable, and slight perturbations of the input graph lead to considerable changes in the run time. On the other hand, the proposed solver's performance is much more stable; it is almost always faster than or comparable to the two existing solvers, and its run time always remains low.

2012 ACM Subject Classification Theory of computation

Keywords and phrases bipartite graphs, assignment problem, matching

Digital Object Identifier 10.4230/LIPIcs.ESA.2023.87

Related Version *Full Version*: <https://inria.hal.science/hal-04146298>

Supplementary Material *Software*: <https://gitlab.inria.fr/bora-ucar/bottled>
archived at `swh:1:dir:c711aeb1f701657080e12c716a5476361b8099ae`

1 Introduction

A matching in a graph is a set of edges without any common vertices. A maximum cardinality matching in a graph has the largest number of edges among all matchings. We investigate algorithms for finding a maximum cardinality matching whose minimum edge weight is the maximum on bipartite graphs with edge weights. This is called the bottleneck matching problem or linear bottleneck assignment problem when all vertices can be matched [3, Section 6.2]. Formally, the bottleneck matching problem is to find a maximum cardinality matching \mathcal{M} which maximizes

$$\min_{(r_i, c_j) \in \mathcal{M}} w_{i,j}.$$

This problem can be solved in polynomial time.



© Ioannis Panagiotas, Grégoire Pichon, Somesh Singh, and Bora Uçar;
licensed under Creative Commons License CC-BY 4.0

31st Annual European Symposium on Algorithms (ESA 2023).

Editors: Inge Li Gørtz, Martin Farach-Colton, Simon J. Puglisi, and Grzegorz Herman; Article No. 87;
pp. 87:1–87:15



Leibniz International Proceedings in Informatics
Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

The bottleneck matching problem arises in different contexts [3, Section 6.2.8]. We are motivated by the Birkhoff–von Neumann (BvN) decomposition of doubly stochastic matrices [2], which arises in practical applications [1, 4, 16, 21]. In this case, the bipartite graphs associated with matrices have equal number of vertices on both sides and contain perfect matchings. A known heuristic for the BvN decomposition [10] repeatedly calls a bottleneck matching algorithm on the bipartite graph of a dynamically changing matrix.

■ **Table 1** Run time, in seconds, of MC64J2 and MC64J3 on 12 problems. \mathbf{A} is the original matrix, \mathbf{AP} is the same matrix with a random column permutation; $\mathbf{S} = \mathbf{D}P(\mathbf{A})\mathbf{E}$, where $P(\mathbf{A})$ is the 0-1 matrix with 1s at the nonzero positions of \mathbf{A} ; and \mathbf{D} and \mathbf{E} are positive diagonal matrices scaling $P(\mathbf{A})$ to be doubly stochastic.

matrix	\mathbf{A}		\mathbf{AP}		\mathbf{S}	
	J2	J3	J2	J3	J2	J3
atmosmodm	0.05	0.26	0.13	0.40	684.87	6487.07
CurlCurl_3	0.05	0.31	0.15	0.48	578.83	2958.07
ss	2.28	0.90	2317.46	1.73	810.10	16501.00
vas_stokes_2M	0.25	1.87	6310.44	727.57	1420.01	4571.14

The software MC64 [7, 8] is the state-of-the-art and implements two algorithms, denoted MC64J2 and MC64J3, for the bottleneck matching problem. To the best of our knowledge these are the only available codes that can handle bipartite graphs corresponding to large sparse matrices. Their worst-case run time are $O(n(m+n)\log_2 n)$ and $O(nm\log_2 n)$, on bipartite graphs with n vertices on each side and m edges [8]. MC64, especially the newer MC64J2 [8], is well engineered. It works very well for graphs corresponding to matrices from numerical applications. However, it does not have stable run time behavior in two senses. First, when run on the same bipartite graph twice with different edge weights the difference in run time can be in the order of hours. Second, on two equivalent problem instances, where one is obtained from the other by just reordering the vertices, the run time can change dramatically. We report the run time of MC64 on four matrices, from the SuiteSparse Matrix Collection [5], in Table 1 to explain this – more experiments of similar nature are in Section 4. Here, the set of rows and the set of columns of a matrix correspond to the two parts of the bipartite graph with an edge between two vertices if the corresponding entry in the matrix is nonzero, and the nonzero values are the edge weights. The table contains results for \mathbf{A} , for \mathbf{AP} where \mathbf{P} is a random permutation, and for the doubly stochastic matrix \mathbf{S} which is obtained by scaling the pattern of \mathbf{A} with Sinkhorn-Knopp algorithm [20].

The bipartite graphs of \mathbf{A} and \mathbf{AP} are the same apart from renumbering of the vertices in one part. While on \mathbf{A} both MC64J2 and MC64J3 are fast, both methods suffer on \mathbf{AP} ; the run time of MC64J2 is not acceptable for the last two instances, and that of MC64J3 is high for the last one in Table 1. The bottleneck matching problems on \mathbf{A} and \mathbf{AP} are essentially the same, as permuting the columns does not change the values, nor the bottleneck matching and its value. The bipartite graph of $\mathbf{S} = \mathbf{D}P(\mathbf{A})\mathbf{E}$ is the same as that of \mathbf{A} with different edge weights, hence the problems are not equivalent. Now, the run time of both methods is too much for all instances. The wildly varying run time of both MC64 routines on \mathbf{S} in comparison to those on \mathbf{A} further highlight the instability in their performance.

Our aim in this paper is to develop an algorithm for the bottleneck matching problem which is better than the state-of-the-art codes in MC64. For this purpose, we study an overlooked alternative from the literature. We make observations that pave the way for an efficient algorithm, implement and compare it against the codes from MC64. We conduct a large set of experiments to show that our approach is usually much faster than MC64 and in addition exhibits stable and robust performance.

Section 2 gives a brief background and a summary of the known algorithms. Section 3 contains the proposed algorithm. Section 4 presents the experimental results, and Section 5 concludes the paper. Appendix of the related version, which is available at <https://inria.hal.science/hal-04146298>, contains detailed description of the experiments and further experiments.

2 Background and related work

A matrix \mathbf{A} with $nnz(\mathbf{A})$ nonzero entries can be represented with a bipartite graph $G = (R \cup C, E)$ where each row of \mathbf{A} corresponds to a unique vertex in R , each column of \mathbf{A} corresponds to a unique vertex in C , and there is an edge (r_i, c_j) whenever $a_{ij} \neq 0$. The number m of edges in G is thus equivalent to $nnz(\mathbf{A})$. When the edges are weighted, the weight of the edge (r_i, c_j) is $|a_{ij}|$, that is the magnitude of the nonzeros of \mathbf{A} . For a vertex v , we use $\text{adj}(v)$ to denote the set of its neighbors.

A *matching* is a set of edges with no common vertices. A matching is of *maximum cardinality* if it has the largest number of edges. Given a matching \mathcal{M} , a vertex is *matched* if an edge from \mathcal{M} is incident on it and *free* otherwise. A matching is *perfect* if it matches all vertices. The *deficiency* of a matching \mathcal{M} is the difference between the maximum cardinality of a matching and $|\mathcal{M}|$. Given a matching \mathcal{M} in the graph G , a path in G is \mathcal{M} -*alternating* if its edges are alternately in \mathcal{M} . An \mathcal{M} -alternating path \mathcal{P} is \mathcal{M} -*augmenting* if the start and end vertices of \mathcal{P} are both free. A *vertex cover* is a set of vertices that includes at least one vertex from each edge. In a bipartite graph the maximum cardinality of a matching is equal to the minimum cardinality of a vertex cover [3, Th. 2.7].

Given a bipartite graph, any of its maximum cardinality matchings can be used to obtain a canonical decomposition called Dulmage-Mendelsohn (DM) decomposition [11]. Based on the DM decomposition, Pothen and Fan [18] describe algorithms to permute sparse matrices in a block upper triangular form (BTF):

$$\mathbf{A} = \begin{matrix} & H_C & S_C & V_C \\ \begin{matrix} H_R \\ S_R \\ V_R \end{matrix} & \begin{pmatrix} \mathbf{A}_H & * & * \\ O & \mathbf{A}_S & * \\ O & O & \mathbf{A}_V \end{pmatrix} \end{matrix}. \quad (1)$$

In a BTF, the submatrix \mathbf{A}_H has more columns than rows, and all rows in H_R are matched to a column in H_C in any maximum cardinality matching; the submatrix \mathbf{A}_S is square with at least one perfect matching; the submatrix \mathbf{A}_V has more rows than columns, and all columns in V_C are matched to a row in V_R . The rows/columns in each block are defined as follows

$$H_R = \{\text{row vertices reachable from free column vertices via alternating paths}\},$$

$$H_C = \{\text{free column vertices or column vertices reachable from free column vertices via alternating paths}\},$$

$$V_R = \{\text{free row vertices or row vertices reachable from free row vertices via alternating paths}\},$$

$$V_C = \{\text{column vertices reachable from free row vertices via alternating paths}\},$$

$$S_R = R \setminus (H_R \cup V_R), \text{ and}$$

$$S_C = C \setminus (H_C \cup V_C).$$

A standard BFS/DFS-based graph traversal algorithm will find these sets in linear time.

We make some observation on the BTF form of a matrix. First, the DM decomposition reveals a minimum cover [3, Alg. 3.1]. As all nonzeros in the BTF (1) are confined in the rows $H_R \cup S_R$ and the columns in V_C , the vertex set $\mathcal{C} = H_R \cup S_R \cup V_C$ is a cover. Since the cardinality of \mathcal{C} is equal to the maximum cardinality of a matching, it is a minimum cover. Second, if one adds new nonzeros to the diagonal blocks, upper diagonal blocks, or to the blocks (S_R, H_C) and (V_R, S_C) , the maximum cardinality of a matching does not change. This is so, as the new nonzeros cannot create augmenting paths.

Hall's theorem [13] states that for a bipartite graph G to have a column-perfect matching, the relation $|S| \leq |\bigcup_{c \in S} \text{adj}(c)|$ must hold for any subset S of columns. If a similar relation holds for all subsets of rows, then G will have perfect matchings.

An $n \times n$ matrix $\mathbf{A} \neq 0$ is called *doubly stochastic* if every entry is nonnegative and the sum of entries in each row and each column is equal to 1. Any nonnegative square matrix, whose bipartite graph has perfect matchings, can be scaled with two diagonal matrices to be doubly stochastic [20]. A *permutation matrix* is a square matrix where each row/column contains exactly one nonzero value equal to 1. A perfect matching in the bipartite graph representation of \mathbf{A} corresponds to a permutation matrix. The bipartite graphs of doubly stochastic matrices have perfect matchings.

Henceforth, we assume that the given bipartite graph contains perfect matchings. We comment on rectangular matrices and matrices without perfect matchings in Section 3.3.

2.1 Related work

We review three algorithms from the literature [3, 7, 8]. The first two are implemented in MC64, and to the best of our knowledge are currently the best practical algorithms. Burkard et al. [3, Section 6.2.4] describe two other algorithms [12, 19], which are more theoretical.

2.1.1 Shortest-augmenting path based algorithms

Algorithms based on shortest augmenting paths start with a matching which has the maximum bottleneck value for the currently matched vertices C' in one part, say C . In order to augment the matching, a shortest augmenting path from a free vertex c of C is found with a variant of Dijkstra's shortest path algorithm. Augmenting along the shortest paths maintains the invariant that the current matching has the maximum bottleneck value for any matching that matches the vertices $C' \cup \{c\}$. The process continues until a perfect matching is obtained.

The state-of-the-art implementation in MC64 [8], MC64J2, starts by computing an upper bound ω on the bottleneck value, which is the minimum of maximum in each column and row. It then computes a maximal matching on the graph containing only edge weights no smaller than ω , which is then improved by length-three augmenting paths in a preprocess step. Then, a shortest-augmenting path is sought from each free column vertex to solve the problem. MC64J2 implements an efficient adaptation of Dijkstra's algorithm to find these paths. Depending on the edge weights, the structure of the bipartite graph, or the visit order many edges and vertices may be visited while finding an augmenting path. As seen in Table 1, this can accumulate and result in very long run time.

2.1.2 Threshold-based algorithms

Let G be a weighted bipartite graph. For a value ω , let $G[\omega]$ contain only the edges of G with weight at least ω . Threshold-based algorithms find the largest ω for which $G[\omega]$ has a perfect matching. They do so by considering different values for ω , testing whether $G[\omega]$

has a perfect matching or not and tuning the next ω to a higher or lower value accordingly. MC64's implementation, MC64J3, discusses initialization and algorithmic choices to reduce the number of tests [7]. However, MC64 uses a depth-first-search algorithm to find the augmenting paths for each test which can slow it down, as seen in Table 1.

2.1.3 An algorithm based on duality

This algorithm is also threshold-based and uses the duality of matchings and coverings to find the next threshold. We explain this algorithm in more detail as ours improves upon it. Let ω be a value where $G[\omega]$ does not contain a perfect matching. Let \mathcal{M} be a maximum cardinality matching in $G[\omega]$. From \mathcal{M} one can define a minimum vertex cover $\mathcal{C} = H_R \cup S_R \cup V_C$ of $G[\omega]$ with vertex sets $I = H_R \cup S_R$ and $J = V_C$ [3, Section 6.2.3]. Since \mathcal{M} is not perfect, there must be an edge in G from a vertex in $\bar{I} = R \setminus I$ to another vertex in $\bar{J} = C \setminus J$. The value $\max_{i \in \bar{I}, j \in \bar{J}} a_{ij}$ thus cannot be smaller than the bottleneck matching value of G and can be used as the next threshold. This approach is shown in Algorithm 1 [3, Section 6.2.3].

■ **Algorithm 1** Duality-based algorithm.

```

Input :  $G$ , an edge weighted bipartite graph having perfect matchings
Output:  $\mathcal{M}$ , a bottleneck perfect matching

Let  $\omega$  be an upper bound on the bottleneck matching value
 $\mathcal{M} \leftarrow \emptyset$ 
while  $|\mathcal{M}| < n$  do
     $\mathcal{M} \leftarrow$  a maximum cardinality matching in  $G[\omega]$ 
    if  $|\mathcal{M}| < n$  then
        1  $I \subseteq R$  and  $J \subseteq C$  be the vertex sets of the associated cover of  $G[\omega]$ 
        2  $\omega \leftarrow \max_{i \in R \setminus I, j \in C \setminus J} w_{i,j}$  /* in  $G$ , not  $G[\omega]$  */

```

The maximum cardinality matching in $G[\omega]$ can be found in $O(\sqrt{n} \text{nnz}(\mathbf{A}[\omega]))$ time in the worst case [14]. Once such a matching is found, the associated minimum cover and the maximum uncovered value at Line 2 can be obtained in linear time. Therefore, the worst case time complexity of Algorithm 1 is $O(\sqrt{n} \text{nnz}(\mathbf{A}))$ times the number of iterations. The worst case run time for Algorithm 1 can hence be too high. This is so as a new edge, due to the reduced ω , does not mean one more edge in the maximum matching, and hence the while loop can even run for more than n iterations.

3 The proposed algorithm

Our algorithm is based on Algorithm 1 and integrates threshold techniques. Let $G = (R \cup C, E)$ be an edge weighted bipartite graph, with $w_{i,j}$ being the weight of the edge (r_i, c_j) . Let ω be a nonnegative value, $G[\omega]$ be as before, and the bottleneck matching value b^* be the largest ω for which $G[\omega]$ has a perfect matching. We call a value ω *safe*, when $\omega \geq b^*$. Algorithm 1 produces decreasing safe values that converge to b^* . We make a series of observations to find b^* faster. The first observation is that there are several minimum vertex covers associated with a given maximum cardinality matching. Using the BTF (1), let $\mathcal{C}' = H_R \cup S_C \cup V_C$. As all nonzeros are covered by \mathcal{C}' and $|\mathcal{C}'| = |H_R \cup S_R \cup V_C|$, \mathcal{C}' is also a minimum cover. If we choose this cover, the sets I and J at Line 1 of Algorithm 1 become H_R and $S_C \cup V_C$, in which case, the maximum value can be different. This leads to the following proposition.

► **Proposition 1.** Let $\mathcal{C}_1 = H_R \cup S_R \cup V_C$ and $\mathcal{C}_2 = H_R \cup S_C \cup V_C$ be two minimum covers of $G[\omega]$ revealed by the BTF (1). Let ω_1 and ω_2 be the maximum values defined in Line 2 of Algorithm 1 for \mathcal{C}_1 and \mathcal{C}_2 , respectively. Then, $\min(\omega_1, \omega_2)$ is safe.

Proof. Let ω be the current value, and $\omega_1 < \omega_2$ without loss of generality. This means that ω_2 is in $\mathbf{A}(V_R, S_C)$ and all entries in $\mathbf{A}(V_R, H_C)$ are smaller than ω_2 . The maximum cardinality of a matching in $G[\omega_2]$ cannot be larger than that in $G[\omega]$, as the set $\mathcal{C}_1 = H_R \cup S_R \cup V_C$ still covers all edges of $G[\omega_2]$, including those that arise in $\mathbf{A}(S_R, H_C)$. Hence the cover \mathcal{C}_1 can be used to get the next ω in Line 2, which concludes the proof. \blacktriangleleft

Based on Proposition 1, one can use the smaller of the largest uncovered element in $(S_R \cup V_R, H_C)$ and that in $(V_R, H_C \cup S_C)$. We propose to exploit the two identified covers as much as possible for a faster convergence of ω to b^* . As we reason in Lemma 2 below, one can find a tighter bound on b^* , depending on the deficiency of the current matching as well as the (original) adjacencies of the vertices in the identified covers.

► Lemma 2. *Let \mathcal{M} be a maximum cardinality matching in $G[\omega]$ with a deficiency of k in G ; $\mathbf{A}[\omega]$ and \mathbf{A} be, respectively, the matrices associated with $G[\omega]$ and G ; $\mathcal{C} = H_R \cup S_C \cup V_C$ be a minimum vertex cover of $G[\omega]$ associated with \mathcal{M} when $\mathbf{A}[\omega]$ is permuted in a BTF (1); Let ω_k^r be the k th largest element in $\bigcup_{i \in S_R \cup V_R} \max\{w_{i,j} : j \in H_C\}$, and ω_k^c the k th largest element in $\bigcup_{j \in H_C} \max\{w_{i,j} : i \in S_R \cup V_R\}$. Then, $\omega_k = \min(\omega_k^r, \omega_k^c)$ is safe.*

Proof. Consider first the set H_C of columns, and note that $|H_C| - |H_R| = k$ as all other columns are matched. By Hall's theorem, $|H_C| \leq |\bigcup_{c \in H_C} \text{adj}(c)|$ must hold in \mathbf{A} as there is a perfect matching. Among all rows in $\bigcup_{c \in H_C} \text{adj}(c)$, we have $|H_R|$ in the set H_R . Therefore there must be at least k other nonzero rows in $\mathbf{A}(S_R \cup V_R, H_C)$. The element ω_k^r from $\bigcup_{i \in S_R \cup V_R} \max\{w_{i,j} : j \in H_C\}$ is safe as any value greater than that will cover less than k rows and Hall's conditions cannot be satisfied. A similar argument applies to ω_k^c by considering the set $S_R \cup V_R$ of rows. In $\mathbf{A}[\omega]$ we have $\text{adj}(S_R \cup V_R) = S_C \cup V_C$, and $|S_R \cup V_R| - |S_C \cup V_C| = k$. The element ω_k^c must be safe since we need at least k nonzero columns in $\mathbf{A}(S_R \cup V_R, H_C)$. The value ω_k is thus safe as the minimum of two safe values. \blacktriangleleft

One can identify several minimum covers, collect the k th largest uncovered element with respect to each, and use the minimum of the collected elements as the next ω . As finding these minimum covers can be expensive, we propose using the two which are readily revealed by the BTF. That is, we use $\omega = \min(\omega_1, \omega_2)$ where

$$\begin{aligned} \omega_1^r, \omega_1^c &= k\text{-th largest row and column maximum entries in } \mathbf{A}(S_R \cup V_R, H_C), \\ \omega_1 &= \min(\omega_1^r, \omega_1^c), \end{aligned} \tag{2}$$

$$\begin{aligned} \omega_2^r, \omega_2^c &= k\text{-th largest row and column maximum entries in } \mathbf{A}(V_R, H_C \cup S_C), \\ \omega_2 &= \min(\omega_2^r, \omega_2^c). \end{aligned} \tag{3}$$

This corresponds to applying Hall's theorem to the sets H_C and $H_C \cup S_C$ of columns and to the sets V_R and $S_R \cup V_R$ of rows.

3.1 Putting it all together

The proposed algorithm BOTTLED is shown in Algorithm 2. The input is a sparse matrix represented in the compressed storage by columns (CSC) format. BOTTLED creates a compressed storage by rows (CSR) representation of the input matrix. It then sorts the nonzeros in each row and each column in non-increasing order of their values. Then the threshold ω is initialized as the minimum of the $2n$ nonzero values consisting of the maximum in each row and maximum in each column. The algorithm then updates the threshold ω in a while-loop as in Algorithm 1. In the while-loop there are three subroutines:

MaximumCardinalityMatching, SAP, and DM-dec. These subroutines correspond to a maximum cardinality matching algorithm, a shortest augmenting path-based method to match a column, and an algorithm obtaining the row and column blocks of the BTF (1).

■ **Algorithm 2** BOTTLED: The proposed bottleneck matching algorithm.

Input : \mathbf{A} , stored by columns
Output : \mathcal{M} , a bottleneck perfect matching

Create a CSR representation of \mathbf{A}
Sort the nonzeros in each row and in each column in non-increasing order of values

- 1 $\omega \leftarrow \min$ of the maximum in each row, maximum in each column
 $G[\omega] \leftarrow (R \cup C, \emptyset)$ and $\mathcal{M} \leftarrow \emptyset$
while $|\mathcal{M}| < n$ **do**
- 2 for each i , release new $a_{ij} \geq \omega$ and for each j release new $a_{ij} \geq \omega$ into $G[\omega]$
 if $|\mathcal{M}| = n - 1$ **then**
 $\mathcal{M} \leftarrow \text{SAP}(G, \mathcal{M}, c)$ with the last free vertex c
 else
- 3 $\mathcal{M}' \leftarrow \text{MaximumCardinalityMatching}(G[\omega], \mathcal{M})$
 if $|\mathcal{M}'| = n$ **then** $\mathcal{M} \leftarrow \mathcal{M}'$; **break**
 if $|\mathcal{M}'| = |\mathcal{M}|$ **then**
- 4 select a vertex c
 $\mathcal{M} \leftarrow \text{SAP}(G, \mathcal{M}', c)$
 else
 $\mathcal{M} \leftarrow \mathcal{M}'$
- 5 $(H_R, S_R, V_R, H_C, S_C, V_C) \leftarrow \text{DM-dec}(G[\omega], \mathcal{M})$
 $\omega \leftarrow \min(\omega_1, \omega_2)$ with ω_1 as in (2) and ω_2 as in (3)

Algorithm 2 stores the edges of $G[\omega]$ over the storage of \mathbf{A} in the CSC and CSR formats, without explicitly building adjacency lists. The start address of each row and column are the same as those of \mathbf{A} . For each row/column of $G[\omega]$, we keep an end-pointer which points to the smallest nonzero of \mathbf{A} in that row/column that is no smaller than ω . These end-pointers are initialized before the while-loop in $O(n)$ time, and incremented at Line 2 at each iteration of the while-loop. Therefore the total run time cost of building $G[\omega]$ s is $O(\text{nnz}(\mathbf{A}))$.

In Algorithm 2, $\text{SAP}(G, \mathcal{M}, c)$ refers to the algorithm summarized in Section 2.1.1. As stated before, SAP needs the current matching to have the bottleneck value among all matchings covering the same set of column vertices. The approach outlined in Algorithm 1 produces such matchings, that is why, at any point, one can resort to SAP. We invoke SAP in two cases: (i) when the deficiency is one; (ii) when an update of ω did not yield an increase in the cardinality of the current matching. The first case is straightforward. For the second case, we apply a simple heuristic to help the algorithm converge faster. As any free column vertex can be the start of an augmenting path, we choose $c \in C$ whose largest edge weight not included in the current $G[\omega]$ is minimum. Matching c will lead to a reduced ω , and the reduction will hopefully be large with this choice of c (some empirical results are in the appendix of the related version). The most common algorithms for the maximum cardinality matching problem take an initial matching as input and augment it. This is very suitable at Line 3 of Algorithm 2, as we have a maximum cardinality matching on a graph, we add new edges, and then ask for a maximum cardinality matching in the new graph. We have used the code-base of MatchMaker [6, 15] to implement this step in the implicit representation of $G[\omega]$.

At Line 5 of Algorithm 2, we use a binary heap with a limit k on its size. For ω_1^r , the nonzeros of each row in $S_R \cup V_R$ with value smaller than ω are visited, and the largest element is used as a key in the heap. When the heap is full, keys are added only if they are larger than the current minimum, which is then removed. The minimum of the heap is returned as

ω_1^r . Other quantities of (2) and (3) are computed similarly. All nonzeros of $\mathbf{A} - \mathbf{A}[\omega]$ in the rows $S_R \cup V_R$ and columns H_C can be visited and the heap operations can be performed in $O(\text{nnz}(\mathbf{A}(S_R \cup V_R, H_C)) + (|H_C| + |S_R \cup V_R|) \log k)$ time. A similar analysis holds for $\mathbf{A}(S_R \cup V_R, H_C)$. The run time of one iteration of the while loop is thus dominated by that of the maximum cardinality matching algorithm. We do not have an estimate on the number of iterations of the while-loop; in the empirical results in Section 4.3 the number is less than what a binary search approach would yield.

3.2 In the context of a BvN decomposition method

The Birkhoff–von Neumann (BvN) theorem [2] states that a doubly stochastic matrix \mathbf{A} can be written as $\mathbf{A} = \sum_{i=1}^{\ell} \alpha_i \mathbf{P}_i$ where each \mathbf{P}_i is a permutation matrix, and α_i s are positive coefficients summing up to one. Such a decomposition is not unique, and the problem of finding a decomposition with the smallest ℓ value is NP-Complete [10].

One heuristic [10] for obtaining a BvN decomposition of a double stochastic matrix \mathbf{A} works as follows. It finds the value b of a bottleneck matching whose pattern is the permutation matrix \mathbf{P} , replaces \mathbf{A} with $\mathbf{A} - b\mathbf{P}$, and continues until a zero matrix is obtained. A property of this heuristic is that the successive bottleneck values are in a non-increasing order [9]. Our bottleneck matching algorithm is very fitting in this case. One can create the CSR representation and sort each row and column once. Then, executing the while loop of Algorithm 2 will obtain a bottleneck matching for the current matrix. Once b and \mathbf{P} are obtained, replacing \mathbf{A} with $\mathbf{A} - b\mathbf{P}$ can be done by subtracting b from each matched entry and updating that entry’s position in the sorted list of both rows and columns in overall $O(n + \text{nnz}(\mathbf{A}))$ time. While doing so, one can update the end-pointers used for $G[\omega]$, and avoid the preprocessing in Algorithm 2 at subsequent invocations.

3.3 Rectangular matrices or matrices without perfect matchings

The case in which there are perfect matchings in the given bipartite graph is common in applications where the bipartite graphs correspond to sparse matrices. This is especially so in the BvN decomposition, which was our motivation. Nonetheless MC64’s J2 and J3 work for cases in which one part of the bipartite graph has more vertices than the other, where the smaller side can be perfectly matched. This corresponds to $n_R \times n$ matrices for $n_R > n$ that have column-perfect matchings. Our algorithm can handle this case either by initializing ω at Line 1 using only the column values, or by using those and only the n th maximum of the set of n_R maximum entries, one from in each row.

Consider now the most general case corresponding to $n_R \times n$ matrices, with $n_R \geq n$ and without column perfect matchings. In this case, MC64J2 returns a *maximum cardinality matching*, without necessarily finding the correct bottleneck value. That is so because not all vertices from which the shortest-augmenting paths are sought can be matched in a bottleneck maximum cardinality matching. We do not know of a suitable fix for this. MC64J3 and its equivalent THRESH on the other hand work correctly. The presented algorithm BOTTLED needs four minor modifications to handle this general case: (i) the maximum size of a matching n' should be computed at the beginning; (ii) the initialization should use the choose the smallest of the n' th maximum of n maximum entries, one from each column, and the n' th maximum of the n_R maximum entries, one from each row; (iii) at Line 5, the deficiency is $k = n' - |\mathcal{M}|$; and (iv) the shortest-augmenting path method should not be used.

4 Experiments

We observed in a preliminary set of experiments that MC64J2 is generally faster than MC64J3; the geometric mean of the ratios of the run time of the latter to that of the former was 4.6. As already highlighted in Table 1, the run time of MC64J3 can be too large to experiment (this happens more frequently than with MC64J2). That is why we have implemented a threshold-based approach, referred here as THRESH, using the same code base of BOTTLED, and use it in our experiments instead of MC64J3. THRESH uses the initialization procedure common to MC64J3 and BOTTLED to find a large initial matching and an upper bound ω on the bottleneck value. If this matching is not maximum, then the edge weights are sorted to find the bottleneck value by a binary search between the smallest edge weight and the initial value ω . The binary search uses only the available edge weights and each time half of the edges are discarded. Our codes, written in C, are available at <https://gitlab.inria.fr/bora-ucar/bottled>; the codes used in the experiments are elsewhere [17].

The next subsection describes the data set, and Section 4.2 conducts a performance analysis of BOTTLED. Section 4.3 compares BOTTLED with MC64J2 and THRESH. Last, some experiments using BOTTLED within a BvN decomposition heuristic are discussed in Section 4.4. Appendix of the related version contains detailed information about experiments.

4.1 Data set and measurements

We have experimented with all square matrices with perfect matchings, at least 100,000 rows, less than 250,000,000 nonzeros, and with no explicit zeros from SuiteSparse [5]. There were 113 such matrices at the time of experimentation. From each matrix, we created six types of problem instances, which are denoted as **A**, **DAE**, **DP(A)E**, **AP**, **DAPE**, and **DP(A)PE**. The **A**-type instance corresponds to the bipartite graph of the original matrix with the magnitudes of the entries as edge weights. The **DAE**-type and **DP(A)E**-type instances correspond, respectively, to the scaled version of the matrices and their patterns with 20 iterations of the Sinkhorn–Knopp [20] algorithm, and the other three types of instances are obtained from the first three by random column permutations. We discarded the instances **A** and **AP** for 0-1 matrices; for the same set of matrices **DAE**-type and **DP(A)E**-type instances are identical, and hence we kept only one of them. We discarded the instances in which the initialization algorithm found the bottleneck value, in which case all three methods are equivalent. We report the experiments with 14 **A**-type, 18 **DAE**-type, 58 **DP(A)E**-type instances, and the same number of instances with column permutations. For each instance, each algorithm is run five times and the geometric mean of the run time is reported as that algorithm’s performance on that instance; when column permutations are applied, these correspond to five different permutations. When comparing two algorithms’ run time, we do not include cases where both algorithms run in less than one second (as both are very small and the difference between the algorithms is insignificant). Appendix of the related version contains all the results.

We carry out the experiments on a machine having Intel(R) Xeon(R) CPU E7-8890 v4 with a clock-speed of 2.20 GHz, and 1.5TB memory. The machine runs Debian GNU/Linux 11 (64 bit). All the codes are compiled with GCC version 10.2.1, with option `-O3`. We ran MatchMaker [6, 15] to verify that there were perfect matchings. We used MatchMaker with options “no cheap matching”, “Push-Relabel + fairness” as the core algorithm. For the sake of completeness, we report that the maximum run time of MatchMaker for any matrix was 2.65 seconds for `vas_stokes_4M` and 28.19 for the column permuted version of the same matrix.

4.2 Performance analysis of BOTTLED

We first analyze the percentage of the total run time of BOTTLED spent in the preprocessing step – creating the CSR representation or sorting the entries. Table 2 summarizes the results for six different problem instances, where BOTTLED took more than one second for **A**-, **DAE**- and/or **DP(A)E**-type instances. This table presents the geometric mean of the percentage of the time spent creating the CSR representation and sorting (with respect to the total time of BOTTLED). The row “tTime” contains the geometric mean of the run time of BOTTLED on the instances **A**, **DAE** and **DP(A)E** in seconds, and for the instances **AP**, **DAPE** and **DP(A)PE** it contains the geometric mean of the ratio of the run time of BOTTLED on the permuted instances to the original ones. For example, the run time under the column **DP(A)PE** is obtained by taking the geometric mean of ratio of the run time of BOTTLED on the 33 **DP(A)PE** instances to that on the 33 **DP(A)E** instances.

As we can see from Table 2, the average run time of BOTTLED on any of the six problem instance types is below 10 seconds. A first observation is that for all instances, the two preprocessing steps account for a non-negligible part of the total run time. At two extremes, they take 7% and 37% of the total time in **DP(A)E** and **DAPE** instances, respectively. We further observe that the absolute run time of CSR and sorting increases for the permuted instances. The percentage of the total time spent in CSR also increases for the permuted instances while that of sorting either increases or remains the same. If the CSR representation is available, its creation can be skipped and one can reduce the run time considerably. As discussed in Section 3.2, for the targeted BvN application BOTTLED can skip not only the CSR creation, but also the sorting phase across different runs of the algorithm.

■ **Table 2** Percentage of the total time spent in creating a CSR matrix and sorting the nonzeros in BOTTLED; “tTime”: the geometric mean of the run time of BOTTLED on the instances **A**, **DAE** and **DP(A)E** in seconds, and the ratios of the others to their counterparts.

	A	AP	DAE	DAPE	DP(A)E	DP(A)PE
CSR	5%	13%	7%	25%	4%	17%
sort	5%	5%	10%	12%	3%	3%
tTime	6.54(s)	1.28	3.43(s)	1.02	3.40(s)	1.28

■ **Table 3** Breakdown of the total time (tTime) of BOTTLED in seconds.

matrix	instance-type	tTime	CSR	sort
vas_stokes_4M	A	14.81	6.70	3.17
	AP	32.36	20.13	5.51
	DAE	23.44	6.93	3.66
	DAPE	52.99	19.93	5.57
vas_stokes_2M	DP(A)E	11.67	2.96	2.01
	DP(A)PE	19.30	8.81	2.57

Another observation from Table 2 is that the run time of BOTTLED consistently increases for the permuted instances. To put this into perspective, we present the run time of BOTTLED on a few instances in Table 3. We see that the increase in run time for the permuted matrices can be attributed, in part, to the creation of the CSR. For example, for **vas_stokes_4M**, excluding CSR and sort times from the total time yields 4.94 and 6.72 seconds for the while loop for **A**- and **AP**-type respectively. A similar calculation for the pairs **DP(A)E** and **DP(A)PE** for **vas_stokes_2M** shows that the while loop takes 6.7 and 7.92 seconds, respectively. In the instance pairs **DAE** and **DAPE** for **vas_stokes_4M**, the increase in the CSR time is still a contributing factor. From these two tables, we conclude that BOTTLED’s preprocessing takes up a significant portion of the total run time.

We have also investigated the number of iterations of the while-loop of BOTTLED in different types of instances to see how stable and robust it is with respect to random column permutations. The number of iterations of the while loop were almost always the same for all 90 instances; in a few cases there was a difference of one in the number of iterations. The small changes in the number of iterations confirm the robustness of BOTTLED.


■ **Table 4** The geometric mean and the maximum of the run time of the three methods on six different instance types.

		A	AP	DAE	DAPE	DP(A)E	DP(A)PE
geomean	MC64J2	0.49	6.23	0.16	0.74	12.47	20.16
	THRESH	2.58	3.80	1.09	1.54	2.24	3.16
	BOTTLED	1.23	2.38	0.44	0.81	1.23	2.07
maximum	MC64J2	123.04	18054.00	31.42	9139.63	1419.99	6813.20
	THRESH	43.55	64.90	59.42	91.74	63.51	77.57
	BOTTLED	14.52	31.81	23.22	51.55	58.64	76.20


4.3 Comparison of different algorithms

We compared BOTTLED against MC64J2 and THRESH. We first provide a broad overview of the comparisons in Table 4, which lists the geometric mean of the run time of the three methods as well as their maximum run time in different problem instance types (all instances). As can be seen, BOTTLED exhibits the best performance overall, despite being slightly slower on average for **A**-, **DAE**-, **DAPE**- instance types than MC64J2– the margin is not large enough for MC64J2 to make up for the performance loss on the other instance types. Furthermore, BOTTLED’s maximum run time is consistently smaller than those of both MC64J2 and THRESH. In contrast, the maximum run time of MC64J2 is prohibitive in all but the **A** and **DAE**-type instances. While THRESH avoids the prohibitive run time of its equivalent MC64J3 (see Table 1) by using a better cardinality matching algorithm [15], its maximum run time is still noticeably larger than that of BOTTLED in all but two cases.

We now look more into the performance of BOTTLED versus that of MC64J2. Figure 1a shows the natural logarithm (in the y -axis) of the ratio of the run time of MC64J2 to that of BOTTLED for the instances on which either method has a run time greater than 1 second. Here, MC64J2 is faster than BOTTLED on 22 instances, and BOTTLED is faster on the remaining 117 instances. In all 139 instances, the geometric mean of the ratio of the run time of MC64J2 to that of BOTTLED is 8.5, confirming that BOTTLED is faster than MC64J2 in average.



(a) The natural logarithm of the ratio of the run time of MC64J2 to that of BOTTLED in the y -axis in sorted order on 139 problem instances in the x -axis.



(b) The run time behavior of MC64J2 on four **DP(A)E**-type instances measured at 20 uniform steps (x -axis), normalized to the total time (y -axis).

■ **Figure 1** Performance comparison between MC64J2 and BOTTLED, and investigation on MC64J2.

In order to put these numbers into a perspective with the run time, we present Table 5. Table 5 lists six matrices in which the ratio of the run time of MC64J2 to that of BOTTLED was the smallest for certain instance types with or without permutation. It next lists six

■ **Table 5** The run time of MC64J2 and BOTTLED on selected instances, where their ratios were the lowest or the highest in different problem instance types or the permuted versions.

matrix	type	MC64J2	BOTTLED	type	MC64J2	BOTTLED
c-73b	A	0.09	13.55	AP	0.10	9.42
c-73	A	0.19	14.52	AP	0.34	14.68
dielFilterV3real	A	0.33	11.72	AP	0.81	21.48
boyd2	DP(A)E	0.03	3.31	DP(A)PE	0.07	1.84
boyd2	DAE	0.03	2.43	DAPE	0.10	1.50
dielFilterV3clx	A	0.12	3.97	AP	0.49	7.19
atmosmodd	DP(A)E	555.76	7.01	DP(A)PE	1757.81	12.11
vas_stokes_1M	DP(A)E	452.36	4.20	DP(A)PE	2427.71	7.42
vas_stokes_2M	DP(A)E	1419.99	11.93	DP(A)PE	6813.20	18.37
ss	DP(A)E	824.91	4.66	DP(A)PE	2621.10	10.79
CurlCurl_3	DP(A)E	610.16	1.54	DP(A)PE	211.56	3.80
atmosmodj	DP(A)E	565.60	7.27	DP(A)PE	1683.32	11.97
vas_stokes_4M	DAE	0.83	23.22	DAPE	9139.63	51.55
ss	A	2.25	2.65	AP	2532.09	8.65
vas_stokes_4M	A	0.53	14.23	AP	18054.00	31.81

matrices with different instance types in which the ratio of the run time of MC64J2 to that of BOTTLED was the largest for the original matrices and three for the permuted ones (three others were already in the list). Here we see that MC64J2 can have equivalently good performance on both the **A**- and **AP**-type instances (see the first block in Table 5). Still, there were some **AP** instances where MC64J2 ran prohibitively long (see the last two rows with **ss** and **vas_stokes_4M**). MC64J2 also struggled on several **DP(A)E** and **DP(A)PE** instances as seen in the second block. The prohibitively high run time of MC64J2 highlight the issue in the approaches based on the shortest augmenting paths: one may visit many edges and vertices to find augmenting paths at different stages during the execution. As such a behavior is also seen for DFS-based-cardinality matching algorithms [6], it cannot be attributed on the particular implementation of the shortest path algorithm in MC64J2.


Figure 1b shows the run time behavior of MC64J2 on the four **DP(A)E**-type instances of Table 1. The total number of augmentations is divided by 20 to obtain a step size, the run time is measured after each step and normalized by the total time. As this figure shows, in some instances the augmentations took about the same time all throughout, whereas in others later augmentations took more time. As the greedy matching approach of MC64J2 does not always find a maximum matching, it can needlessly result in many additional augmentations. Those augmentations may lead to large run time, not solely because of their number. For example, in the **DP(A)E**-type instance of **CurlCurl_3**, a square matrix with $n = 1219574$ rows and $\text{nnz} = 13544618$ nonzeros, there are 20040 augmentations with a total run time of about 600 seconds, even though the maximum cardinality matching on $G[\omega_0]$ has a deficiency of one. Obviously, detecting this would lead to much better run time. On the other hand, for the **DP(A)E**-type instance of **rajat31**, a square matrix with $n = 4690002$ rows and $\text{nnz} = 20316253$ nonzeros, MC64J2 needs 1562500 augmentations (its greedy approach finds a maximum cardinality matching initially), and the whole run time is 48.39 seconds. As **rajat31** is much bigger than **CurlCurl_3** and needs more augmentations, its shorter run time attests that the process of augmenting one-by-one can take large time due to instance specific properties. It hence cannot constitute a reliable method for the bottleneck matching problem in general. In passing we note that MC64J2 is well-engineered and is faster than using our own SAP implementation for the augmentations.

We compare now BOTTLED and THRESH. In total, BOTTLED beat THRESH in 165 instances out of 180. The largest differences (in seconds) in favor of THRESH were in the **A**- and **AP**-type instances of **c-73b** and **c-73**. THRESH obtained a run time of 2.88 and 3.48 seconds


■ **Table 6** Run time, in seconds, of BOTTLED and the BvN decomposition heuristic, along with the number of permutation matrices in the decomposition.

matrix	instance	BOTTLED initialization	longest run time of BOTTLED in an iteration	BvN	
				num.perm	time
atmosmodm	DAE	0.22	58.73	46	452.14
	DP(A)E	0.21	14.43	50	200.80
CurlCurl_3	DAE	0.35	17.30	50	183.43
	DP(A)E	0.34	4.33	50	68.45
ss	DAE	1.15	13.37	50	228.92
	DP(A)E	1.46	39.77	50	186.31
vas_stokes_2M	DAE	5.04	4.74	50	78.32
	DP(A)E	5.10	4.87	50	74.36

on *c-73b*, and 6.35 and 5.11 seconds on *c-73*. On these instances BOTTLED was relatively close to THRESH (see Table 5) – the largest difference is 10.67 seconds on the **A**-type instance of *c-73b*. Since both methods utilize the same core matching algorithm, the superiority of BOTTLED over THRESH should come from doing fewer iterations. Figure 2a supports this reasoning by plotting the difference between the number of iterations of THRESH and that of BOTTLED in nondecreasing order. As seen in this figure, THRESH’s number of iterations is always larger than BOTTLED’s; the theoretical observations of Section 3 translate to practical gains. We compute the ratio of the run time of THRESH to that of BOTTLED for the instances on which either THRESH or BOTTLED has a run time greater than 1 second (123 instances), and present the natural logarithm of this ratio in the *y*-axis in Figure 2b. As seen in this figure, BOTTLED fares better than THRESH in the majority of cases.



(a) The number of iterations of THRESH minus that of BOTTLED in the *y*-axis in sorted order on all 180 instances.



(b) The natural logarithm of the ratio of the run time of THRESH to that of BOTTLED in the *y*-axis in sorted order on 123 instances in the *x*-axis.

■ **Figure 2** Performance comparison between THRESH and BOTTLED.

4.4 Inside a BvN decomposition method

Table 6 presents the run time of the BvN decomposition method on the **DAE**- and **DP(A)E**-type instances of four matrices (obtained with 2500 scaling iterations). We run the BvN decomposition method until 50 permutation matrices or a coefficient of 0.92 are obtained. As each permutation matrix is obtained by a call to BOTTLED, we show their number in the column “num.perm”. Further, the table also presents the time for the initial preprocessing of BOTTLED, and the maximum time taken in a call to BOTTLED subsequently for obtaining a

permutation matrix. We observe that the maximum run time of BOTTLED, in an iteration, multiplied by the number of calls is at least 2.92 (for **DAE** of **ss**) and up to 10.67 (for **DP(A)E** of **ss**) times the total BvN time. This suggests that the run time of the subsequent bottleneck matching calls reduces appreciably, and BOTTLED works well inside the decomposition method by avoiding the preprocessing.

5 Conclusion

We have investigated the problem of finding a maximum bottleneck matching in bipartite graphs. Existing implementations for the problem suffer from unpredictable run time that can get prohibitively large, i.e., requiring thousands or even tens of thousands of seconds to complete. We have proposed a new algorithm called BOTTLED that converts an inefficient, duality-based approach into an efficient one through theoretical findings. Experimental results show that BOTTLED is almost always faster than the state-of-the-art methods. Furthermore, its run time is reliable and always remains within reasonable time limits. We have also explored its use inside a heuristic for the Birkhoff–von Neumann decomposition of doubly stochastic matrices and experimentally established the suitability of the proposed algorithm for this purpose.

Currently the proposed approach resorts to an augmenting-path-based method in few corner cases and only when there are perfect matchings. We plan to explore the possibility to use them more effectively, along with potential data reduction rules.

References

- 1 M. Benzi and B. Uçar. Preconditioning techniques based on the Birkhoff–von Neumann decomposition. *Computational Methods in Applied Mathematics*, 17:201–215, 2017.
- 2 G. Birkhoff. Tres observaciones sobre el algebra lineal. *Univ. Nac. Tucumán Rev. Ser. A*, 5:147–150, 1946.
- 3 R. Burkard, M. Dell’Amico, and S. Martello. *Assignment Problems*. SIAM, Philadelphia, PA, USA, 2009.
- 4 C. Chang, W. Chen, and H. Huang. On service guarantees for input-buffered crossbar switches: A capacity decomposition approach by Birkhoff and von Neumann. In *Quality of Service, 1999. IWQoS ’99. 1999 Seventh International Workshop on*, pages 79–86, 1999.
- 5 T. A. Davis and Y. Hu. The University of Florida sparse matrix collection. *ACM Trans. Math. Softw.*, 38(1):1:1–1:25, 2011.
- 6 I. S. Duff, K. Kaya, and B. Uçar. Design, implementation, and analysis of maximum transversal algorithms. *ACM Transactions on Mathematical Software*, 38:13:1–13:31, 2011.
- 7 I. S. Duff and J. Koster. The design and use of algorithms for permuting large entries to the diagonal of sparse matrices. *SIAM Journal on Matrix Analysis and Applications*, 20(4):889–901, 1999.
- 8 I. S. Duff and J. Koster. On algorithms for permuting large entries to the diagonal of a sparse matrix. *SIAM Journal on Matrix Analysis and Applications*, 22:973–996, 2001. doi:10.1137/S0895479899358443.
- 9 F. Dufossé, K. Kaya, I. Panagiotas, and B. Uçar. Further notes on Birkhoff–von Neumann decomposition of doubly stochastic matrices. *Linear Algebra and its Applications*, 554:68–78, 2018. doi:10.1016/j.laa.2018.05.017.
- 10 F. Dufossé and B. Uçar. Notes on Birkhoff–von Neumann decomposition of doubly stochastic matrices. *Linear Algebra and its Applications*, 497:108–115, 2016. doi:10.1016/j.laa.2016.02.023.
- 11 A. L. Dulmage and N. S. Mendelsohn. Coverings of bipartite graphs. *Canadian Journal of Mathematics*, 10:517–534, 1958.

- 12 H. N. Gabow and R. E. Tarjan. Algorithms for two bottleneck optimization problems. *J. Algorithms*, 9(3):411–417, 1988. Hopcroft-Karp algorithm can be used as an approximation algorithm p415.
- 13 P. Hall. On representatives of subsets. *Journal of the London Mathematical Society*, s1-10(37):26–30, 1935. URL: <http://jlm.oxfordjournals.org/cgi/reprint/s1-10/37/26.pdf>.
- 14 J. E. Hopcroft and R. M. Karp. An $n^{5/2}$ algorithm for maximum matchings in bipartite graphs. *SIAM Journal on Computing*, 2(4):225–231, 1973.
- 15 K. Kaya, J. Langguth, F. Manne, and B. Uçar. Push-relabel based algorithms for the maximum transversal problem. *Computers & Operations Research*, 40(5):1266–1275, 2013.
- 16 J. Kulkarni, E. Lee, and M. Singh. Minimum Birkhoff-von Neumann decomposition. In *International Conference on Integer Programming and Combinatorial Optimization*, pages 343–354. Springer, 2017.
- 17 I. Panagiotas, G. Pichon, S. Singh, and B. Uçar. Bottled: Fast algorithms for the bottleneck matching problem, April 2023. doi:10.5281/zenodo.7871464.
- 18 A. Pothén and C.-J. Fan. Computing the block triangular form of a sparse matrix. *ACM Transactions on Mathematical Software*, 16(4):303–324, December 1990. doi:10.1145/98267.98287.
- 19 A. P. Punnen and K. Nair. Improved complexity bound for the maximum cardinality bottleneck bipartite matching problem. *Discret. Appl. Math.*, 55(1):91–93, 1994.
- 20 R. Sinkhorn and P. Knopp. Concerning nonnegative matrices and doubly stochastic matrices. *Pacific J. Math.*, 21:343–348, 1967.
- 21 V. Valls, G. Iosifidis, and L. Tassiulas. Birkhoff’s decomposition revisited: Sparse scheduling for high-speed circuit switches. *IEEE/ACM Transactions on Networking*, 29(6):2399–2412, 2021. doi:10.1109/TNET.2021.3088327.