

EGERVÁRY RESEARCH GROUP  
ON COMBINATORIAL OPTIMIZATION



TECHNICAL REPORTS

TR-2013-04. Published by the Egerváry Research Group, Pázmány P. sétány 1/C,  
H-1117, Budapest, Hungary. Web site: [www.cs.elte.hu/egres](http://www.cs.elte.hu/egres). ISSN 1587-4451.

---

# Minimum-cost flow algorithms: An experimental evaluation

Péter Kovács

---

September 2013  
Revised: May 2015

# Minimum-cost flow algorithms: An experimental evaluation<sup>\*</sup>

Péter Kovács<sup>\*\*</sup>

**Keywords:** network flows; minimum-cost flows; efficient algorithms; experimental study; LEMON

## Abstract

An extensive computational analysis of several algorithms for solving the minimum-cost network flow problem is conducted. Some of the considered implementations were developed by the author and are available as part of an open-source C++ optimization library called LEMON (<http://lemon.cs.elte.hu/>). These codes are compared to other publicly available solvers: CS2, MCF, RelaxIV, PDNET, MCFSimplex, as well as the corresponding components of the IBM ILOG CPLEX Optimization Studio and the LEDA C++ library. This evaluation, to the author's knowledge, is more comprehensive than earlier studies in terms of the range of considered implementations as well as the diversity and size of problem instances. The presented results demonstrate that the primal network simplex and cost-scaling algorithms are the most efficient and robust in general. Other methods, however, can outperform them in particular cases. The network simplex code of the author turned out to be far superior to the other implementations of this method, and it is the most efficient solver on the majority of the considered problem instances. In contrast, the cost-scaling algorithms tend to show better asymptotic behavior, especially on sparse graphs, and hence they are typically faster than simplex-based methods on huge networks.

---

<sup>\*</sup>The final version of this paper is published in Optimization Methods and Software [69]

<sup>\*\*</sup>Department of Algorithms and Applications and MTA-ELTE Egerváry Research Group, Eötvös Loránd University, Pázmány Péter sétány 1/c, H-1117 Budapest, Hungary; and ChemAxon Ltd., Záhony utca 7, H-1031 Budapest, Hungary. Research was partially supported by grants (no. CNK 77780 and no. CK 80124) from the National Development Agency of Hungary, based on a source from the Research and Technology Innovation Fund.

# 1 Introduction

The *minimum-cost flow* (MCF) problem plays a fundamental role within the area of network optimization. It is to find a minimum-cost transportation of a specified amount of flow from a set of supply nodes to a set of demand nodes in a directed network with capacity constraints and linear cost functions defined on the arcs. This problem has a remarkably wide range of applications in various fields, for instance, telecommunication, network design, transportation, routing, scheduling, resource planning, and manufacturing. Furthermore, it often arises as a subtask of more complex optimization problems. Comprehensive studies of network flow theory and algorithms can be found in, for example, [3, 93, 68]. A recent survey related to minimum-cost network flows is presented in [95].

Over the last five decades, numerous algorithms have been devised for solving the MCF problem, and they have been thoroughly studied both from theoretical and from practical aspects. It has become evident that the asymptotic worst-case time complexity of an algorithm does not determine its actual performance on real-life problem instances. Therefore, many researchers have worked on efficient implementation and experimental analysis of MCF algorithms, for example, [45, 19, 58, 5, 56, 13, 14, 16, 17, 15, 90, 48, 46, 74, 21, 86, 88, 40]. Dominant contributions in this area were attained during the First DIMACS Implementation Challenge [33]. Furthermore, Brunsch et al. [20] recently applied the smoothed analysis model to an MCF algorithm to obtain better explanation of its practical behavior.

A previous paper [65] presents our implementations of several MCF algorithms along with an experimental evaluation. This study improves on that work by focusing on the fastest algorithms and providing a more comprehensive computational analysis including three additional MCF solvers and more test instances. The contribution of the presented results is twofold. First, a great number of implementations are compared within the same benchmark framework using a wide variety of test instances. Second, we consider larger networks than that in previous studies, which turned out to be essential to draw appropriate conclusions. The author is not aware of other recent works that provide an empirical study of this extent.

Fourteen MCF solvers are evaluated in this paper. Seven implementations are due to the author: three cycle-canceling algorithms; implementations of the successive shortest path and capacity-scaling methods; a cost-scaling algorithm; and a primal network simplex code. All of them are available with full source code as part of the LEMON [73] optimization library under permissive license terms. These codes are compared to seven other publicly available solvers: CS2 code of Goldberg and Cherkassky (see [46, 26]); the corresponding function of the LEDA library [71]; MCF solver of Löbel [74, 75]; the NETOPT component of the IBM ILOG CPLEX Optimization Studio [25]; MCFSimplex code of Bertolini, Frangioni, and Gentile [9]; RelaxIV [9], a C++ translation of the original Fortran code due to Bertsekas and Tseng [14, 10]; and PDNET solver of Portugal et al. [88, 87].

The benchmark tests were conducted on an extensive collection of large-scale networks (up to millions of nodes and arcs). These instances were created either using standard generators NETGEN, GRIDGEN, GOTO, and GRIDGRAPH (see [33]), or

based on networks arising in real-life problems.

The presented results provide a profound and reliable survey of the efficiency and robustness of the considered codes. Implementations of the cost-scaling and primal network simplex algorithms turned out to have the best overall performance. On special problem instances, however, other methods (for instance, the augmenting path algorithms and the relaxation method) can outperform them. Among the four primal network simplex codes evaluated in this study, our implementation turned out to be the most efficient and robust. Moreover, it is consistently superior to all other solvers on relatively small network instances (up to many thousands of nodes). However, the cost-scaling codes are substantially faster on the largest sparse networks due to their better asymptotic behavior in terms of the number of nodes. The most robust implementations of this algorithm are CS2 and the one developed by the author. The latter implementation, to the author's knowledge, is the first to apply Goldberg's partial augment-relabel technique [47] in the MCF context, which turned out to be a considerable improvement in accordance with Goldberg's results related to the maximum flow problem.

The rest of this paper is organized as follows. Section 2 provides a brief overview of the MCF problem and its solution methods. Section 3 describes the implementations used in our experiments. Section 4 presents the computational results. Finally, the conclusions are drawn in Section 5.

## 2 The MCF problem

We assume familiarity with network flow theory and standard notations (see, e.g., [3, 93, 68]). Here we discuss only the most important definitions and results.

Let  $G = (V, A)$  be a weakly connected directed graph consisting of  $n = |V|$  nodes and  $m = |A|$  arcs. We associate with each arc  $ij \in A$  a *capacity* (upper bound)  $u_{ij} \geq 0$  and a *cost*  $c_{ij}$ . Each node  $i \in V$  has a signed *supply* value  $b_i$ . The single-commodity, linear minimum-cost network flow (MCF) problem is defined as

$$\min \sum_{ij \in A} c_{ij} x_{ij}, \quad (1a)$$

subject to

$$\sum_{j:ij \in A} x_{ij} - \sum_{j:ji \in A} x_{ji} = b_i \quad \forall i \in V, \quad (1b)$$

$$0 \leq x_{ij} \leq u_{ij} \quad \forall ij \in A. \quad (1c)$$

We refer to (1b) as *flow conservation constraints* and (1c) as *capacity constraints*. We assume that all data are integer and wish to find an integer-valued optimal solution. Without loss of generality, we may further assume that all arc capacities are finite, all arc costs are nonnegative, and the problem has a feasible solution (see [3]). It also implies that  $\sum_{i \in V} b_i = 0$ .

A *pseudoflow* is a function  $x$  defined on the arcs that satisfies the capacity constraints (1c), but may violate (1b). A feasible flow is also a pseudoflow. The *excess* value of a node  $i$  with respect to a pseudoflow  $x$  is defined as

$$e_i = b_i + \sum_{j:ji \in A} x_{ji} - \sum_{j:ij \in A} x_{ij}. \quad (2)$$

Node  $i$  is referred to as an *excess node* if  $e_i > 0$  and as a *deficit node* if  $e_i < 0$ . Note that  $\sum_{i \in V} e_i = \sum_{i \in V} b_i = 0$ .

Given a pseudoflow  $x$ , the corresponding *residual network*  $G_x = (V, A_x)$  is defined as follows. For each original arc  $ij \in A$ , corresponds a *forward arc*  $ij \in A_x$  with residual capacity  $r_{ij} = u_{ij} - x_{ij}$  and cost  $c_{ij}$  if  $r_{ij} > 0$ ; and a *backward arc*  $ji \in A_x$  with residual capacity  $r_{ji} = x_{ij}$  and cost  $-c_{ij}$  if  $r_{ij} > 0$ .

The linear programming dual solution of the MCF problem is represented by *node potentials*  $\pi_i$  ( $i \in V$ ). The *reduced cost* of an arc  $ij$ , with respect to a potential function  $\pi$ , is defined as  $c_{ij}^\pi = c_{ij} + \pi_i - \pi_j$ .

Optimality criteria can be defined for the MCF problem as follows. These results are discussed in [39, 49, 3, 93, 68].

**Theorem 2.1** (Negative cycle optimality conditions). *A feasible solution  $x$  of the MCF problem is optimal if and only if the residual network  $G_x$  contains no directed cycle of negative total cost.*

An equivalent formulation can be stated using node potentials and reduced costs.

**Theorem 2.2** (Reduced cost optimality conditions). *A feasible solution  $x$  of the MCF problem is optimal if and only if for some node potential function  $\pi$ ,  $c_{ij}^\pi \geq 0$  holds for each arc  $ij$  in the residual network  $G_x$ .*

The concept of *approximate optimality* is also essential for some algorithms. For a given  $\epsilon \geq 0$ , a pseudoflow  $x$  is called  $\epsilon$ -*optimal* if for some node potential function  $\pi$ ,  $c_{ij}^\pi \geq -\epsilon$  holds for each arc  $ij$  in the residual network  $G_x$ . If the arc costs are integer and  $\epsilon < 1/n$ , then an  $\epsilon$ -optimal feasible flow is indeed optimal. Furthermore, for a non-optimal feasible solution  $x$ , the smallest  $\epsilon \geq 0$  for which  $x$  is  $\epsilon$ -optimal equals to the negative of the minimum-mean cost of a directed cycle in the residual network  $G_x$  (see, e.g., [3, 41]).

The MCF problem and its solution methods have been the object of intensive research for more than fifty years. Dantzig was the first to solve a special case of the problem, the so-called transportation problem, by specializing his well-known simplex method. Later, he also applied this approach directly to the MCF problem and devised the network simplex algorithm (see [29]). Ford and Fulkerson [39] developed the first combinatorial algorithms by generalizing Kuhn's remarkable Hungarian Method [70].

Various algorithms were proposed in the next few years, but they do not run in polynomial time. Edmonds and Karp [36] introduced the scaling technique and developed the first weakly polynomial-time algorithm. The first strongly polynomial method is due to Tardos [99]. These results were followed by many other algorithms of improved running time bounds, most of which rely on the scaling technique.

Efficient implementation and profound computational evaluation of MCF algorithms have also been of high interest (see, e.g., [45, 19, 58, 5, 56, 13, 14, 16, 17, 15, 90, 48, 46, 74, 21, 86, 88, 40]). The network simplex algorithm became quite popular when spanning tree labeling techniques were developed to improve its practical performance. Later, implementations of relaxation and cost-scaling algorithms also turned out to be highly efficient.

Table 1 provides a complexity survey of MCF algorithms, which was made according to [3], [93], and the papers referenced in the table.

## 3 Implementations of MCF algorithms

In this section, we first introduce the MCF algorithms implemented by the author in the LEMON library, then we briefly discuss all other solvers used in our experiments. For a comprehensive study of MCF algorithms, see the book of Ahuja, Magnanti, and Orlin [3].

### 3.1 MCF algorithms in the LEMON library

LEMON [73] is an open-source C++ template library with focus on combinatorial optimization tasks related mainly to graphs and networks. Its name is an abbreviation of *Library for Efficient Modeling and Optimization in Networks*. LEMON provides highly efficient implementations of graph algorithms and related data structures, which can be combined easily to solve complex real-life optimization problems. The LEMON project is maintained by the MTA-ELTE Egerváry Research Group on Combinatorial Optimization (EGRES) [37] at the Department of Operations Research, Eötvös Loránd University, Budapest, Hungary. The library is also a member of the COIN-OR initiative [24], a collection of open-source projects related to operations research. LEMON can be used in both commercial and non-commercial software development under very permissive license terms. For more information about LEMON, the readers are referred to the introductory paper [32] and to the web site of the library: <http://lemon.cs.elte.hu/>.

LEMON provides implementations of several MCF algorithms, which are presented in this section along with the applied data structures and heuristics. For a more detailed discussion, see [65].

#### 3.1.1 SCC – Simple cycle-canceling algorithm

Cycle-canceling is a general primal method for solving the MCF problem, originally proposed by Klein [66]. A feasible solution  $x$  is first established, which can be carried out by a maximum flow computation. After that, the algorithm throughout maintains feasibility and gradually decreases the total cost of the solution. At each iteration, a directed cycle of negative cost is found in the current residual network  $G_x$ , and it is canceled by augmenting flow along it. When  $G_x$  contains no negative-cost cycle, the algorithm terminates with an optimal solution found (see Theorem 2.1).

	$O(n^4CU)$	Ford and Fulkerson (1958) [38]
	$O(m^3U)$	Yakovleva (1959) [106], Minty (1960) [78], Fulkerson (1961) [42] <i>out-of-kilter</i>
	$O(n^2mU)$	Jewell (1958) [61], Busacker and Gowen (1960) [22], Iri (1960) [60] <i>successive shortest path</i>
	$O(nm^2CU)$	Klein (1967) [66] <i>cycle-canceling</i>
*	$O(nU \cdot \text{SP}_+(n, m, nC))$	Edmonds and Karp (1970) [35, 36], Tomizawa (1971) [102] <i>successive shortest path (using potentials)</i>
*	$O(m \log U \cdot \text{SP}_+(n, m, nC))$	Edmonds and Karp (1972) [36] <i>capacity-scaling</i>
	$O(n \log C \cdot \text{MF}(n, m, U))$	Röck (1980) [92], Bland and Jensen (1985) [16, 17] <i>cost-scaling</i>
	$O(m^2 \log n \cdot \text{MF}(n, m, U))$	Tardos (1985) [99]
	$O(n^2m \log(nC))$	Goldberg and Tarjan (1987) [50, 54] <i>cost-scaling (generic version)</i>
	$O(n^3 \log(nC))$	Goldberg and Tarjan (1987) [50, 54], Bertsekas and Eckstein (1988) [11] <i>cost-scaling</i>
	$O(n^{5/3}m^{2/3} \log(nC))$	Goldberg and Tarjan (1987) [50, 54] <i>cost-scaling</i>
	$O(nm \log n \log(nC))$	Goldberg and Tarjan (1987) [50, 54] <i>cost-scaling (with dynamic trees)</i>
	$O(n^2m^2 \min\{\log(nC), m \log n\})$	Goldberg and Tarjan (1988) [51, 53] <i>minimum-mean cycle-canceling</i>
	$O(n^2m \min\{\log(nC), m \log n\})$	Goldberg and Tarjan (1988) [51, 53] <i>cancel-and-tighten</i>
		Orlin (1997) [83] <i>primal network simplex</i>
	$O(nm \log n \min\{\log(nC), m \log n\})$	Goldberg and Tarjan (1988) [51, 53] <i>cancel-and-tighten (with dynamic trees)</i>
		Tarjan (1997) [101] <i>primal network simplex (with dynamic trees)</i>
*	$O(m \log n \cdot \text{SP}_+(n, m, nC))$	Orlin (1988) [81, 82] <i>enhanced capacity-scaling</i>
		Vygen (2000) [103, 104]
*	$O(nm \log \log U \log(nC))$	Ahuja, Goldberg, Orlin, and Tarjan (1988) [1, 2] <i>double scaling</i>
*	$O((m^{3/2}U^{1/2} + mU \log(mU)) \log(nC))$	Gabow and Tarjan (1989) [43]
*	$O((nm + mU \log(mU)) \log(nC))$	Gabow and Tarjan (1989) [43]
	$O(n^3 \min\{\log(nC), m \log n\})$	Goldberg and Tarjan (1990) [54] <i>generalized cost-scaling</i>
*	$O(nm \log(n^2/m) \min\{\log(nC), m \log n\})$	Goldberg and Tarjan (1990) [54] <i>generalized cost-scaling (with dynamic trees)</i>
	$O(m(m + n \log n) \min\{\log(nU), m \log n\})$	Orlin, Plotkin, and Tardos (1993) [84] <i>dual network simplex</i>
	$O(nm \log n(m + n \log n))$	Armstrong and Jin (1997) [4] <i>dual network simplex</i>

$n$  and  $m$  denote the number of nodes and arcs, respectively;  $U$  denotes the maximum of supply values and arc capacities;  $C$  denotes the largest arc cost.  $\text{SP}_+(n, m, K)$  and  $\text{MF}(n, m, K)$  denote the time complexity of solving the single-source shortest path problem and the maximum flow problem, respectively, in a directed graph with  $n$  nodes,  $m$  arcs, and nonnegative integer arc lengths or capacities, each at most  $K$ .

\* indicates an asymptotically best complexity bound.

Table 1: Complexity survey of MCF algorithms

Many variants of this approach have been developed by applying different rules for cycle selection (see, e.g., [105, 6, 53, 97]). These algorithms have quite different theoretical and practical behavior; some of them run in polynomial or even strongly polynomial time. Moreover, the primal network simplex algorithm (see Section 3.1.7) can also be viewed as a particular variant of the cycle-canceling method.

LEMON provides three cycle-canceling algorithms. One of them, which we call *simple cycle-canceling* (SCC), identifies negative cycles using the Bellman–Ford algorithm with successively increased limit on the number of its iterations. The worst-case time complexity of this method is  $O(nm^2CU)$  (using the notations defined in Section 2).

### 3.1.2 MMCC – Minimum-mean cycle-canceling algorithm

Goldberg and Tarjan [51, 53] devised this famous special variant of the cycle-canceling method, which selects a negative cycle of minimum mean cost at each iteration. It is the simplest strongly polynomial-time MCF algorithm; it performs  $O(nm^2 \log n)$  iterations for arbitrary real-valued arc costs and  $O(nm \log(nC))$  iterations for integer arc costs (see, e.g., [53, 49, 3, 93, 68]).

The MMCC algorithm relies on finding minimum-mean directed cycles in a graph. This optimization problem has also been studied for a long time, and various algorithms have been developed for solving it. Comprehensive experimental studies can be found in [31, 30, 44]. In the MMCC implementation, we use a combination of Howard’s policy-iteration algorithm (see [59, 30]) and the algorithm of Hartmann and Orlin [57]. This combined method is quite efficient in practice and also ensures the best strongly polynomial time bound  $O(nm)$ . Therefore, the overall complexity of the MMCC algorithm is  $O(n^2m^2 \min\{\log(nC), m \log n\})$  for the MCF problem with integer data.

### 3.1.3 CAT – Cancel-and-tighten algorithm

This algorithm is also due to Goldberg and Tarjan [51, 53]. It improves on the MMCC method by maintaining node potentials to make the detection of negative residual cycles easier and faster. The CAT algorithm proceeds by ensuring  $\epsilon$ -optimality of the current solution for successively smaller values of  $\epsilon \geq 0$  until it becomes optimal. At every iteration, two main steps are performed. In the *cancel* step, directed cycles consisting entirely of arcs with negative reduced cost are successively canceled until no such cycle exists. In the *tighten* step, the node potentials are modified in order to introduce new arcs of negative reduced cost and to decrease  $\epsilon$  to at most  $(1 - 1/n)$  times its former value.

We implemented the cancel step using a straightforward method, which is based on a depth-first search and has a worst-case complexity of  $O(nm)$ . Although it could be improved to run in  $O(m \log n)$  time using dynamic tree data structures (see [53] and [96]), we did not implement this variant since the cycle-canceling algorithms turned out to be slow in practice.

The tighten step can be carried out in  $O(m)$  time using a topological ordering of the nodes. However, this method does not ensure strongly polynomial bound on



the number of iterations. Therefore, after every  $k$  iterations, the parameter  $\epsilon$  is set to its smallest possible value using a minimum-mean cycle computation. Goldberg and Tarjan [53] suggest to use  $k = n$  (to preserve the amortized running time  $O(m)$  of the tighten step), but we use  $k = \lfloor \sqrt{n} \rfloor$  since it turned out to be more efficient in practice. The minimum-mean cycle computations are carried out using the same combined method that was applied in the MMCC algorithm.

The CAT algorithm performs  $O(n \min\{\log(nC), m \log n\})$  iterations, and thus our implementation runs in  $O(n^2 m \min\{\log(nC), m \log n\})$  time for the MCF problem with integer data.

### 3.1.4 SSP – Successive shortest path algorithm

This is another fundamental MCF algorithm, which solves the problem by a sequence of shortest path computations. The initial version of this method was devised independently by Jewell [61], Iri [60], and Busacker and Gowen [22]. Later, Edmonds and Karp [35, 36] and Tomizawa [102] independently suggested the utilization of node potentials to ensure nonnegative arc costs throughout the algorithm, which greatly improves its performance both in theory and in practice.

The SSP algorithm applies a dual approach. It maintains an optimal pseudoflow along with corresponding node potentials and attempts to achieve feasibility. At each iteration, it selects a node  $v$  with  $e_v > 0$  and augments flow from node  $v$  to a node  $w$  with  $e_w < 0$  along a shortest path in the residual network with respect to the reduced arc costs. The node potentials are then increased by the computed shortest path distances to preserve the reduced cost optimality conditions (see Theorem 2.2). The algorithm terminates after  $O(nU)$  iterations.

The practical performance of the SSP algorithm mainly depends on the shortest path computations. Our implementation uses Dijkstra’s algorithm with binary heap data structure (other heap structures were also tested, but they did not turn out to be more robust in practice). Therefore, an iteration is performed in  $O(m \log n)$  time and the overall worst-case complexity of our SSP code is  $O(nmU \log n)$ . An essential improvement of the algorithm is also applied as it is discussed in [3]. At every iteration, Dijkstra’s algorithm is terminated once it permanently labels the deficit node  $w$ , and the potentials are modified accordingly. Furthermore, we use an efficient representation of the residual network that allows fast traversal of the outgoing arcs of a node, which is crucial for the shortest path computations.

### 3.1.5 CAS – Capacity-scaling algorithm

This is an improved version of the SSP method, which was developed by Edmonds and Karp [36] as the first weakly polynomial-time MCF algorithm. This method performs capacity-scaling phases to ensure that the path augmentations carry sufficiently large amounts of flow, which often reduces the number of iterations. In a  $\Delta$ -scaling phase, each path augmentation delivers  $\Delta$  units of flow from a node  $v$  with  $e_v \geq \Delta$  to a node  $w$  with  $e_w \leq -\Delta$ . When no such augmenting path is found, the value of  $\Delta$  is halved, and the algorithm proceeds with the next phase. At the end of the phase with

$\Delta = 1$ , an optimal solution is found.

The CAS algorithm performs  $O(m \log U)$  iterations under the additional assumption that a directed path of sufficiently large capacity exists between each pair of nodes. Although this condition can easily be achieved by a simple extension of the underlying network, we decided to avoid this transformation. As a result, more units of excess may remain at the end of a  $\Delta$ -scaling phase, but we allow the path augmentations of subsequent phases to deliver more units of flow than the current value of  $\Delta$ . The polynomial running time is thereby not guaranteed, but our experiments showed that this version does not perform more path augmentations and runs significantly faster in practice.

Our implementation uses a scaling factor  $\alpha$  other than two, that is,  $\Delta$  is initially set to  $\alpha^{\lceil \log_\alpha U \rceil}$  and divided by  $\alpha$  at the end of each phase. We use  $\alpha = 4$  by default, as it turned out to provide the best overall performance. The CAS algorithm has much in common with the SSP method, hence similar data structures and algorithmic improvements are applied in its implementation.

### 3.1.6 COS – Cost-scaling algorithm

The cost-scaling technique was first proposed independently by Röck [92] and Bland and Jensen [16, 17]. Goldberg and Tarjan [50, 54] improved on these methods by also utilizing the concept of  $\epsilon$ -optimality and developed an algorithm that has turned out to be among most efficient ones, both in theory and in practice.

The COS algorithm performs  $\epsilon$ -phases for successively smaller values of  $\epsilon$  until  $\epsilon < 1/n$ . In each phase, a *refine* procedure is applied to transform an  $\epsilon$ -optimal solution into an  $(\epsilon/\alpha)$ -optimal solution for a given scaling factor  $\alpha > 1$ .

The refine procedure first converts the  $\epsilon$ -optimal flow  $x$  to an optimal pseudoflow. The value of  $\epsilon$  is then divided by  $\alpha$  and the pseudoflow  $x$  is gradually transformed into a feasible flow again, but in a way that preserves  $\epsilon$ -optimality for the new value of  $\epsilon$ . This is achieved by performing a sequence of local *push* and *relabel* operations similarly to the well-known push-relabel algorithm for the maximum flow problem, which is also due to Goldberg and Tarjan [52].

Given a pseudoflow  $x$  and a potential function  $\pi$ , a node  $i$  is called *active* if  $e_i > 0$  and a residual arc  $ij$  is called *admissible* if  $c_{ij}^\pi < 0$ . A basic operation selects an active node  $v$  and either *pushes* flow on an admissible residual arc  $vw$  or if no such arc exists, *relabels* node  $v$ . It means that the potential of node  $v$  is decreased by the largest possible amount that does not violate the  $\epsilon$ -optimality conditions. This operation introduces new admissible outgoing arcs at node  $v$  and thereby allows subsequent push operations to carry the remaining excess of this node. The refine procedure terminates when no active node remains in the network, and hence an  $\epsilon$ -optimal feasible solution is obtained (since the basic operations preserve  $\epsilon$ -optimality).

The generic version of the refine procedure takes  $O(n^2m)$  time. As the number of  $\epsilon$ -scaling phases is  $O(\log(nC))$ , the COS algorithm runs in weakly polynomial time  $O(n^2m \log(nC))$ . Goldberg and Tarjan [54] also developed special versions of improved running time bounds by applying particular selection rules of the basic operations and using complex data structures. The best variant runs in strongly polynomial time

$O(nm \log(n^2/m) \min\{\log(nC), m \log n\})$ .

We implemented three variants of the COS algorithm that perform the refine procedure rather differently. We first discuss the standard push-relabel implementation, and only the differences are presented for the other two variants.

**Push-relabel method** This implementation is based on the generic version of the COS algorithm and hence performs local push and relabel operations. Its practical performance highly depends on the application of various effective heuristics and other improvements. Such ideas are discussed in [48, 46, 21]. In fact, most of these improvements are analogous to similar techniques devised for the push-relabel maximum flow algorithm (see, e.g., [23]). In the followings, we briefly present the improvements we applied.

The bottleneck of the COS algorithm corresponds to searching for admissible arcs for the basic operations. Therefore, we apply the same representation of the residual network as for the SSP and CAS algorithms. Furthermore, we maintain a current arc pointer for each node, which is a standard technique for making these arc searches even faster.

We also apply a well-known technique to overcome the issue that the generic COS algorithm performs internal computations with non-integer values of  $\epsilon$  and non-integer node potentials. All arc costs are multiplied by  $\alpha n$  for a given integer scaling factor  $\alpha \geq 2$ , and  $\epsilon$  is also scaled accordingly. Initially,  $\epsilon$  is set to  $\alpha^{\lceil \log_\alpha(anC) \rceil}$  and is divided by  $\alpha$  in each phase until  $\epsilon = 1$ . Thereby, it is ensured that all computations operate solely on integer numbers, which improves the performance of the operations and avoids potential numerical issues.

The optimal value of  $\alpha$  depends on the heuristics applied in the algorithm and on the problem instance, as well. Our experiments showed that its optimal value is usually between 8 and 24, and the differences are typically moderate, thus we use  $\alpha = 16$  by default.

The strategy for selecting an active node for the next basic operation is essential. We experimented with a few variants, including the so-called wave implementation, but the simple FIFO selection rule turned out to be the most robust (as it is also observed in [46, 21]).

In addition to these improvements, we also applied three complex heuristics out of the four proposed in [48, 46].

The *potential refinement* (or *price refinement*) heuristic introduces an additional step at the beginning of each phase to check if the current solution is already  $(\epsilon/\alpha)$ -optimal. This step attempts to adjust the potentials to satisfy the  $(\epsilon/\alpha)$ -optimality conditions, but without modifying the flow. If it succeeds, the refine procedure is skipped and the next phase begins. Our results also verified that this heuristic substantially improves the overall performance of the algorithm in most cases.

Another variant of this heuristic performs a minimum-mean cycle computation at the beginning of each phase to determine the smallest  $\epsilon$  for which the current flow is  $\epsilon$ -optimal and computes corresponding node potentials. However, our experiments showed that an efficient implementation of the former variant of the heuristic is clearly superior to this one, which contradicts the results of [21].

The *global update* heuristic performs relabel operations on several nodes at once by iteratively applying the following operation. Let  $S \subset V$  denote a set of nodes such that all deficit nodes are in  $S$ , but at least one active node is in  $V \setminus S$ . If no admissible arc enters  $S$ , then the potential of every node in  $S$  can be increased by  $\epsilon$  without violating the  $\epsilon$ -optimality conditions. This heuristic turned out to greatly improve the efficiency of the COS algorithm on some problem classes, although it does not help to much or even slightly worsens the performance on other instances.

The *push-look-ahead* heuristic attempts to avoid pushing flow from node  $i$  to node  $j$  when a subsequent push operation is likely to send this amount of flow back to node  $i$ . To achieve this, only a limited amount of flow is allowed to be pushed into a node  $j$ , and the relabel operation is extended to those nodes at which this limitation is applied regardless of their current excess values. In most cases, this heuristic significantly decreases the number of push operations.

Another heuristic, the *speculative arc fixing*, is also proposed in [48, 46]. Although this heuristic would most likely improve the performance of our implementation, as well, we did not implement it until now, mainly because it seems to be rather involved and sensitive to parameter settings.

**Augment-relabel method** This variant of the COS algorithm performs path augmentations instead of local push operations, but relabeling is heavily used to find augmenting paths. At each step of the refine procedure, this method selects an active node  $v$  and performs a depth-first search on the admissible arcs to find an augmenting path to a deficit node  $w$ . Whenever the search process has to step back from a node  $i$ , then this node is relabeled.

The flow augmentation on the admissible paths can be performed in two natural ways. We can either push the same amount of flow on each arc of the path or push the maximum possible amount of flow on each arc independently. According to our experiments, the latter variant performs slightly better, thus it is applied in our implementation.

Note that these path search and flow augmentation procedures correspond to particular sequences of local push and relabel operations. However, the actual push operations are carried out in a delayed and more guided manner, in aware of an admissible path to a deficit node. This helps to avoid such problems for which the push-look-ahead heuristic is devised (see above), but a lot of work may be required to find augmenting paths, especially if they are long.

In this implementation, we used the same data representation, improvements, and heuristics as for the push-relabel variant except for the push-look-ahead heuristic.

**Partial augment-relabel method** The third implementation of the COS algorithm can be viewed as an intermediate approach between the other two variants. It is based on the partial augment-relabel algorithm recently proposed by Goldberg [47] as a novel variant of the push-relabel algorithm for solving the maximum flow problem. As this method turned out to be highly efficient and robust in practice, Goldberg also suggests the utilization of the idea in the MCF context, but he does not

investigate it. Our implementation of the COS algorithm, to the author's knowledge, is the first to incorporate this technique.

The partial augment-relabel implementation of the COS algorithm is quite similar to the augment-relabel variant, but it limits the length of the augmenting paths. The path search process is stopped either if a deficit node is reached or if the length of the path reaches a given parameter  $k \geq 1$ . In fact, the push-relabel and augment-relabel variants are special cases of this approach with  $k = 1$  and  $k = n - 1$ , respectively. Our code uses  $k = 4$  by default, just like Goldberg's maximum flow implementation, as it results in a quite robust algorithm in the MCF context, as well.

Apart from the length limitation for the augmenting paths, this variant is exactly the same as the augment-relabel method. (Actually, they have a common implementation but with different values of the parameter  $k$ .) However, the partial augment-relabel technique attains a good compromise between the former two approaches and turned out to be superior to them. We refer to this implementation as COS in the followings.

### 3.1.7 NS – Network simplex algorithm

The primal network simplex (NS) algorithm is one of the most popular solution methods for the MCF problem. It is a specialized version of the well-known linear programming (LP) simplex method that exploits the network structure of the MCF problem and performs the basic operations directly on the graph representation. The LP variables correspond to the arcs of the graph, and the LP bases are represented by spanning trees.

The NS algorithm was devised by Dantzig [28, 29], the inventor of the LP simplex method. Although the generic version of the algorithm does not run in polynomial time, it turned out to be rather effective in practice. Therefore, subsequent research has focused on efficient implementation of the NS algorithm (see [98, 45, 19, 8, 64, 56, 74]). Furthermore, researchers also developed particular variants of both the primal and the dual network simplex methods that run in polynomial time (see [80, 85, 100, 55, 84, 4, 83, 101]). Detailed discussion of the NS method can be found in, for example, [3] and [63].

The NS algorithm is based on the concept of *spanning tree solutions*. Such a solution is represented by a partitioning of the arc set  $A$  into three subsets  $(T, L, U)$  such that each arc in  $L$  has flow fixed at zero, each arc in  $U$  has flow fixed at the capacity of the arc, and the arcs in  $T$  form an (undirected) spanning tree of the network with flow values not restricted to the lower or upper bound. If an instance of the MCF problem has an optimal solution, then it also has an optimal spanning tree solution, which can be found by successively transforming a spanning tree solution to another one. These solutions actually correspond to the LP basic feasible solutions of the problem.

The LP simplex method maintains a basic feasible solution and gradually improves its objective function value by small transformations, known as pivots. Accordingly, the primal NS algorithm throughout maintains a spanning tree solution of the MCF problem and successively decreases the total cost of the flow until it becomes optimal. Node potentials are also maintained such that the reduced cost of each arc in the

spanning tree is zero. At each iteration, a non-tree arc violating the optimality conditions is added to the current spanning tree, which uniquely determines a negative-cost residual cycle. This cycle is canceled by augmenting flow on it and a tree arc corresponding to a saturated residual arc is removed from the tree. The tree structure is then updated, and the node potentials are adjusted accordingly. This whole operation transforming a spanning tree solution to another one is called pivot. If no suitable entering arc can be found, the current flow is optimal and the algorithm terminates.

In fact, the NS algorithm can also be viewed as a particular variant of the cycle-canceling algorithm. Due to the sophisticated method of maintaining spanning tree solutions, however, a negative cycle can be found and canceled much faster (in  $O(m)$  time).

Similarly to the LP simplex method, degeneracy is also a crucial issue for the NS algorithm. In many cases, a cycle of zero residual capacity is found in a pivot step, and hence the flow cost cannot be decreased, only the spanning tree is modified. Degeneracy not only worsens the performance of the algorithm, but its finiteness is even not guaranteed. A simple and popular technique to overcome these issues is to ensure that the spanning tree solution is strongly feasible at each iteration. It means that a positive amount of flow can be sent from each node to a designated root node of the spanning tree along the tree path. For the details, see [27, 7, 3].

**Spanning tree data structures** The representation of spanning tree solutions is essential to implement the NS algorithm efficiently. Several storage schemes have been developed for this purpose along with efficient methods for updating them during the pivot operations (see [62, 98, 45, 8, 64, 3]).

We implemented the primal NS algorithm with two spanning tree storage techniques. The first one is usually referred to as the ATI (*Augmented Threaded Index*) method and represents a spanning tree as follows. The tree has a designated root node, and three indices are stored for each node: the depth of the node in the tree, the index of its parent node, and a so-called thread index that is used to define a depth-first traversal of the spanning tree. It is a quite popular method, which is discussed in detail in [63] and [3].

The ATI technique has an improved version, which is due to Barr, Glover, and Klingman [8] and is usually referred to as the XTI (*eXtended Threaded Index*) method. The XTI scheme replaces the depth index by two indices for each node: the number of successors of the node in the tree and the last successor of the node according to the traversal defined by the thread indices. This modification allows faster update process since a tree alteration of a single pivot usually modifies the depth of several nodes in subtrees that are moved from a position to another one, while the set of successors is typically modified only for much fewer nodes.

We implemented the XTI scheme with an additional improvement that a reverse thread index is also stored for each node to represent the depth-first traversal as a doubly-linked list. This modification turned out to substantially improve the performance of the update process. In fact, the inventors of the XTI technique also discussed this extension [8], but they did not apply it in order to reduce the memory requirements of the representation.

Although the XTI labeling method is not as widely known and popular as the simpler ATI method, our experiments showed that it is much more efficient on all problem instances, so the final version of our code only implements the XTI technique.

**Initialization** The NS algorithm also requires an initial spanning tree solution to start with. It is possible to transform any feasible solution to a spanning tree solution, and the required tree indices can also be computed by a depth-first traversal of the tree arcs. However, artificial initialization is more common in practice. It means that the underlying network is extended with an artificial root node and additional arcs connecting this node and the original nodes in a way that a strongly feasible spanning tree solution can easily be constructed. We apply this method as it is much easier and usually provides better overall performance than starting with a non-artificial solution.

**Pivot rules** One of the most crucial aspects of the NS algorithm is the selection of entering arcs for the pivot operations. The applied method affects the “goodness” of the entering arcs and thereby the number of iterations as well as the average time required for selecting an entering arc, which is a dominant part of each iteration.

We implemented five pivot rules, which are briefly discussed in the followings. A non-tree arc  $ij$  is called *eligible* if it violates the optimality conditions, that is, either  $x_{ij} = 0$  and  $c_{ij}^\pi < 0$  or  $x_{ij} = u_{ij}$  and  $c_{ij}^\pi > 0$  with respect to the current potential function  $\pi$ . We refer to  $|c_{ij}^\pi|$  as the *violation* of arc  $ij$ .

The two simplest pivot rules are the *best eligible* and the *first eligible* strategies. The former one always selects an eligible arc of maximum violation to enter the tree, while the latter one examines the arcs cyclically and selects the first eligible arc at each iteration. Neither of them provides good overall performance in practice.

The *block search pivot rule* was proposed by Grigoriadis [56]. This method cyclically examines certain subsets (blocks) of the arcs and selects the best eligible candidate among these arcs at each iteration. The block size  $B$  is an important parameter of this method. In fact, the previous two rules are special cases of this one with  $B = m$  and  $B = 1$ , respectively. Several sources suggest to set  $B$  proportionally to the number of arcs, for example, between 1% and 10% [56, 63]. However, our experiments showed that a much robust implementation can be achieved using  $B = \lfloor \sqrt{m} \rfloor$ .

The *candidate list pivot rule* is another well-known approach, which was proposed by Mulvey [79]. In a so-called major iteration, this method examines the arcs cyclically to build a list containing at most  $L$  eligible arcs. This list is then used by at most  $K$  subsequent iterations to select an arc of maximum violation among the candidates. If an arc becomes non-eligible, it is removed from the list. We obtained the best average running time using  $L = \lfloor \sqrt{m}/4 \rfloor$  and  $K = \lfloor L/10 \rfloor$ .

We also developed an improved version of the previous method, which we call *altering candidate list pivot rule*. It also maintains a candidate list, but it attempts to extend this list at each iteration and keeps only the several best candidates for the next one. At least one arc block of size  $B$  is examined at every iteration to extend the list with new eligible arcs. After that, an arc of maximum violation is selected to

enter the basis, and the list is truncated to contain at most  $H$  of the best candidate arcs. According to our measurements,  $B = \lfloor \sqrt{m} \rfloor$  and  $H = \lfloor B/100 \rfloor$  result in a rather robust implementation.

Our experimental results showed that the block search and the altering candidate list pivot rules are the most efficient for all classes of problem instances. Since the block search rule is simpler and turned out to be slightly more robust, it is our default pivot strategy, and we refer to this variant as NS in the followings.

## 3.2 Other MCF solvers

In the followings, we briefly introduce the MCF solvers that are compared to the implementations provided by the LEMON library. The MCFClass project [9] was of great help for us in these experiments. This project features a common and flexible C++ interface for several MCF solvers, which are all considered in this paper. However, MCFClass does not support the latest versions of CS2 and MCFZIB, thus we used these two solvers directly.

### 3.2.1 CS2

CS2 is a highly robust, authoritative implementation of the cost-scaling push-relabel algorithm. It was written in C language by A.V. Goldberg and B. Cherkassky applying the improvements and heuristics described in [48, 46]. CS2 has been widely used as a benchmark for solving the MCF problem for a long time (see, e.g., [46, 74, 40]). We used the latest version, CS2 4.6, which is available from the IG Systems, Inc. [26]. CS2 software can be used for academic research and evaluation purposes free of charge, but commercial use requires a license.

### 3.2.2 LEDA

LEDA is a comprehensive commercial C++ library providing efficient implementations of various data types and algorithms [71, 77]. It is widely used in many application areas, such as telecommunication, scheduling, GIS, traffic planning, VLSI design, and computational biology. The `MIN_COST_FLOW()` procedure of the LEDA library implements the cost-scaling push-relabel algorithm similarly to CS2. We used LEDA 5.0 in our experiments. In fact, LEDA 5.1.1 was also tested, but its MCF method turned out to be slower than that of version 5.0.

### 3.2.3 MCFZIB

MCF is a network simplex code written in C language by A. Löbel [74] at the Zuse Institute Berlin (ZIB). We denote this solver as MCFZIB in order to differentiate it from the problem itself (similarly to the MCFClass project [9]). This solver features both a primal and a dual network simplex implementation, from which the former one is used by default as it tends to be more efficient. It applies a usual data structure for representing the spanning tree solutions along with an improved version of the candidate list pivot rule, which is called *multiple partial pricing* (see [74] for the



details). This implementation has been shown to be rather efficient (see, e.g., [74, 40]). We used the latest version, MCF 1.3, which is available for academic use free of charge, see [75].

### 3.2.4 CPLEX

IBM ILOG CPLEX Optimization Studio [25], usually referred to simply as CPLEX, is a well-known and powerful software package aimed at large-scale optimization problems. It provides efficient methods for solving different kinds of mathematical programming problems, including linear, mixed integer, and convex quadratic programming. The NETOPT module of CPLEX implements the primal network simplex algorithm for solving the MCF problem. We used this component through a “wrapper class” called MCFComplex, which is provided by the MCFClass project [9]. MCFComplex implements the common interface defined in this project using the CPLEX Callable Library. Although CPLEX is a comprehensive commercial software, it is available free of charge for academic use and for non-commercial research. We used its latest version, 12.4.

### 3.2.5 MCFSimplex

MCFSimplex is a recent, open-source software written by A. Bertolini and A. Frangioni in C++ language. It is available as part of the MCFClass project [9] and can be used under the flexible LGPL license. This code implements both the primal and the dual network simplex algorithms, from which the former one is the default, as it is usually faster and more robust. MCFSimplex is reported to be reasonably efficient compared to other network simplex implementations, moreover, it can solve separable quadratic MCF problems, as well.

### 3.2.6 RelaxIV

RelaxIV is an efficient, authoritative implementation of the relaxation algorithm [13]. The original FORTRAN code was written by D.P. Bertsekas and P. Tseng [12, 14] and is available at [10]. We used a C++ translation of this code, which was made by A. Frangioni and C. Gentile and is available as part of the MCFClass project [9]. Similarly to CS2, the RelaxIV solver and its previous versions have been used in experimental studies for a long time (see, e.g., [12, 14, 74, 46, 40]).

### 3.2.7 PDNET

PDNET implements the truncated primal-infeasible dual-feasible interior-point algorithm for solving linear network flow problems. It was written in Fortran and C language by L.F. Portugal, M.G.C. Resende, G. Veiga, J. Patrício, and J.J. Júdice [86, 88]. This solver can be used free of charge, its source code is available at [87]. For more information about interior-point network flow algorithms, see, for example, [90, 89, 91].

## 4 Experimental study

This section presents an extensive computational analysis of all considered MCF solvers. First, we describe the problem families used for the experiments. In Section 4.2, we present a comparison of the algorithms implemented by the author in the LEMON library (see Section 3.1). After that, we focus on the most efficient ones among these codes and compare them to other publicly available solvers presented in Section 3.2.

This study improves on [65] by providing a more comprehensive evaluation of the fastest implementations, including three additional solvers: CPLEX, MCFSimplex, and PDNET. Although empirical analysis of MCF algorithms has always been of high interest (see [45, 19, 58, 5, 16, 17, 56, 13, 14, 15, 90, 48, 74, 46, 21, 40]), the author is not aware of any recent study that is comparably extensive and considers networks up to millions of nodes and arcs. Benchmark results on such large graphs, however, turned out to be essential to draw appropriate conclusions.

### 4.1 Test setup

Our test suite comprises numerous networks of various size and characteristics. It can be accessed at <http://lemon.cs.elte.hu/trac/lemon/wiki/MinCostFlowData>.

Most of the problem instances were generated with standard random generators NETGEN, GRIDGEN, GOTO, and GRIDGRAPH, which are available as source codes at the FTP site of the First DIMACS Implementation Challenge [33]. We used these generators with similar parameter settings as previous works (e.g., [48, 15, 14, 74, 46, 21, 40]), but we created larger networks, as well. Other problem families were also generated based on either real-life road networks or maximum flow problems arising in computer vision applications. All instances involve solely integer data.

We used LEMON 1.3 for this study, which is the latest release of the library. The versions of the other solvers we considered can be found in Section 3.2.

The presented experiments were conducted on a machine with AMD Opteron Dual Core 2.2 GHz CPU and 16 GB RAM (1 MB cache), running openSUSE 11.4 operating system. All codes were compiled with GCC 4.5.3 using `-O3` optimization flag.

#### 4.1.1 NETGEN instances

NETGEN is a classic and popular generator developed by D. Klingman, A. Napier, and J. Stutz [67, 33]. It produces random instances of the MCF problem and other network optimization problems. In fact, NETGEN is known to create relatively easy MCF instances.

We generated NETGEN networks as follows. Arc capacities and costs are selected uniformly at random from the ranges  $[1..1000]$  and  $[1..10000]$ , respectively. The number of supply nodes and the number of demand nodes are both set to  $\sqrt{n}$  and the total supply is set to  $1000\sqrt{n}$  (rounded to integer). The three problem families we used differ only in the density of the networks.

- **NETGEN-8.** Sparse networks;  $m = 8n$ .

- **NETGEN-SR.** Dense networks;  $m \approx n\sqrt{n}$ .
- **NETGEN-DEG.** Networks of increasing density;  $n = 4096$  and  $m$  ranges from  $2n$  to  $n^2$ .

#### 4.1.2 GRIDGEN instances

GRIDGEN is a random generator that produces grid-like networks. It was written by Y. Lee and J. Orlin (see [72, 33]).

We used the same parameters for GRIDGEN as for the NETGEN families with the additional setting that the width of the underlying grid is set to  $\sqrt{n}$  for each GRIDGEN network. Actually, the shape of the grid did not turn out to be a significant parameter for GRIDGEN, since this generator selects the supply and demand nodes uniformly at random.

- **GRIDGEN-8.** Sparse networks;  $m = 8n$ .
- **GRIDGEN-SR.** Dense networks;  $m \approx n\sqrt{n}$ .
- **GRIDGEN-DEG.** Networks of increasing density;  $n = 4096$  and  $m$  ranges from  $4n$  to  $n^2$ .

#### 4.1.3 GOTO instances

GOTO is another well-known generator for the MCF problem, which was developed by A.V. Goldberg [48, 33] and intended to produce hard instances. The name of the generator stands for *Grid On Torus*, which reflects to the basic structure of the generated networks. Each GOTO instance has one supply node and one demand node and the supply value is adjusted according to the arc capacities.

As for the previous problem families, the maximum arc capacity is set to 1000 and the maximum arc cost is set to 10000.

- **GOTO-8.** Sparse networks;  $m = 8n$ .
- **GOTO-SR.** Dense networks;  $m \approx n\sqrt{n}$ .

#### 4.1.4 GRIDGRAPH instances

GRIDGRAPH generator was written by M.G.C. Resende and G. Veiga [90, 33]. It produces grid networks similarly to GRIDGEN and GOTO, but applying a stricter scheme. A GRIDGRAPH network consists of transshipment nodes forming a grid of  $W$  rows and  $L$  columns as well as a single source node  $s$  and a single sink node  $t$ . Arcs go from  $s$  to the nodes of the first column; from the nodes of the last column to  $t$ ; and from each transshipment node  $(w, l)$  to nodes  $(w+1, l)$  and  $(w, l+1)$ , except for the last row and column, respectively. The arc capacities and costs are set uniformly at random within a specified range.

We used three GRIDGRAPH families based on grids of different shape as it turned out to be a crucial parameter. The maximum arc capacity and cost are set to 1000 and 10000, respectively, just like for the previous generators.

- **GRID-WIDE.**  $L = 16$  and  $W$  increases.
- **GRID-LONG.**  $W = 16$  and  $L$  increases.
- **GRID-SQUARE.**  $W = L$ .

#### 4.1.5 ROAD instances

Special MCF problems based on real-world road networks were also considered. To generate such instances, we used the TIGER/Line road network files of several states of the USA. These data files are available at the web site of the Ninth DIMACS Implementation Challenge [34].

We selected seven states having road networks of increasing size (namely, DC, DE, NH, NV, WI, FL, and TX) and generated MCF problem instances as follows. The original undirected graphs are converted to directed graphs by replacing each edge with two oppositely directed arcs. The cost of an arc is set to the travel time on the corresponding road section. The number of supply and demand nodes are both set to  $\lfloor \sqrt{n}/10 \rfloor$ . These nodes are selected randomly, and the supply-demand values are determined by a maximum flow computation that maximizes the total supply with respect to the fixed set of supply and demand nodes and fixed arc capacities.

We generated two problem families with different arc capacity settings.

- **ROAD-PATHS.** The arc capacities are uniformly set to one. This means that we are actually looking for a specified number of arc-disjoint directed paths from supply nodes to demand nodes having minimum total cost.
- **ROAD-FLOW.** The capacity of an arc is set 40, 60, 80, or 100 according to the category of the corresponding road section.

#### 4.1.6 VISION instances

Our test suite also includes MCF instances based on large-scale maximum flow problems arising in computer vision applications. The corresponding data files were made available by the Computer Vision Research Group at the University of Western Ontario [76] for benchmarking maximum flow algorithms (see, e.g., [47]).

We used some of the segmentation instances related to medical image analysis, which are defined on three-dimensional grid networks. Those variants were selected in which the underlying graphs are 6-connected and the maximum arc capacity is 100 (the `bone_sub*_n6c100` files). These networks were converted to minimum-cost maximum flow problem instances applying different arc cost functions. The original networks also contain arcs of zero capacity, but we skipped these arcs during the transformation and thereby did not preserve the exact 6-connectivity.

- **VISION-RND.** The arc costs are selected uniformly at random from the range  $[1..100]$ .

- **VISION-PROP.** The cost of an arc is approximately proportional to its capacity:  $c_{ij} = \lfloor \alpha_{ij} u_{ij} \rfloor$ , where  $\alpha_{ij}$  is a random factor selected uniformly from the range  $[0.9, 1.1)$ .
- **VISION-INV.** The cost of an arc is approximately inversely proportional to its capacity:  $c_{ij} = \lfloor \alpha_{ij} K / u_{ij} \rfloor$ , where  $K = 1000$  and  $\alpha_{ij}$  is a random factor selected uniformly from the range  $[0.9, 1.1)$ .

We also experimented with other parameter settings both in case of the generators and in case of the ROAD and VISION instances, but the presented collection turned out to be a representative benchmark suite of reasonable size.

For all problem families, we generated three instances of each network size using different random seeds. To achieve reliable comparison, we report the average running time over such three instances in all cases.

## 4.2 Comparison of MCF algorithms in the LEMON library

In what follows, we present benchmark results obtained for the MCF algorithm implementations of the LEMON library, which are discussed in Section 3.1.

Table 2 and Figure 1 show the results on NETGEN problem families, while Table 3 and Figure 2 show the results on GOTO instances. The first part of each table presents running time in seconds, while the second part reports normalized time results. The best running time is highlighted for each problem size, and a “—” sign denotes the cases when the explicit time limit of one hour was exceeded. The charts display running time in seconds as a function of the number of nodes in the network (logarithmic scale is used for both axes).

In accordance with previous studies, we found that GOTO instances are substantially harder than NETGEN problems, and the relative performance of the various solution methods also turned out to be rather different on them. However, on GRIDGEN networks, we obtained quite similar results as for the NETGEN families, thus we omit GRIDGEN results here.

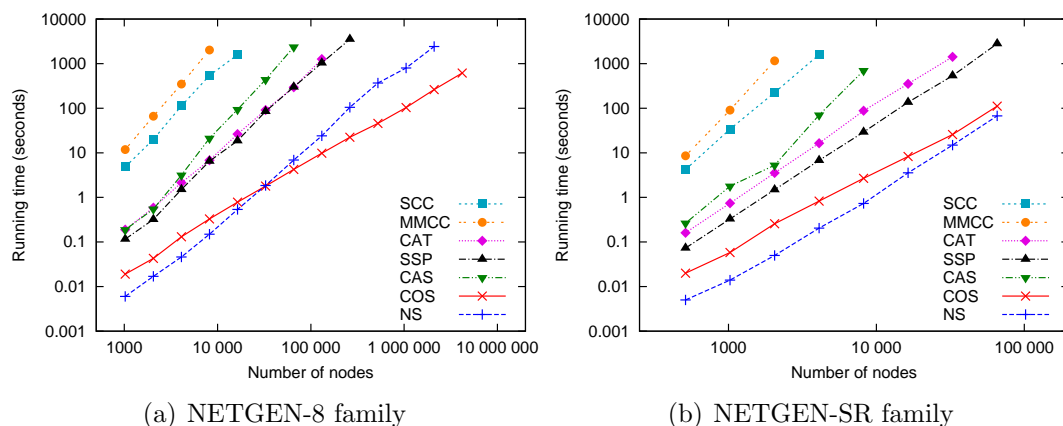


Figure 1: Comparison of LEMON implementations on NETGEN networks

According to these experiments, the simple cycle-canceling methods, SCC and MMCC, are orders of magnitude slower than all other algorithms. CAT, which is an advanced cycle-canceling algorithm, is much faster and usually performs similarly to the dual algorithms SSP and CAS. The relative performance of these three methods greatly depends on the characteristics of the problem instance. The COS and NS algorithms are generally the most efficient. COS shows better asymptotic behavior than NS, thus it is typically faster on the largest networks, while NS clearly outperforms COS on the smaller ones.

Table 4 compares the efficiency of LEMON implementations on GRIDGRAPH families. The shape of the underlying grid is an essential parameter for these networks as it determines the length of augmenting paths or cycles an algorithm should find. This phenomenon has consistently been observed before, for example, see [14]. On GRID-WIDE networks, NS is by far the fastest and COS is the second, while on GRID-LONG instances, the augmenting path algorithms, SSP and CAS, greatly outperform all other solution methods. The GRID-SQUARE results are just similar to the experiments on NETGEN and GOTO families.

Finally, the benchmark results on ROAD networks are presented in Table 5. As one would expect, the SSP algorithm is the fastest on these special instances. On the ROAD-PATHS family, the CAS algorithm works exactly the same as SSP since arc

Running time (seconds)									
Problem family	$n$	$m/n$	SCC	MMCC	CAT	SSP	CAS	COS	NS
NETGEN-8	$2^{10}$	8	4.81	11.85	0.19	0.12	0.19	0.02	<b>0.01</b>
	$2^{13}$	8	540.80	2018.61	6.81	6.48	21.17	0.33	<b>0.15</b>
	$2^{16}$	8	—	—	295.05	298.95	2360.21	<b>4.24</b>	6.88
	$2^{19}$	8	—	—	—	—	—	<b>45.56</b>	368.52
	$2^{22}$	8	—	—	—	—	—	<b>615.42</b>	—
NETGEN-SR	$2^{10}$	32	34.03	90.82	0.74	0.33	1.79	0.06	<b>0.01</b>
	$2^{12}$	64	1592.62	—	16.36	6.77	70.35	0.83	<b>0.21</b>
	$2^{14}$	128	—	—	353.23	136.17	—	8.28	<b>3.55</b>
	$2^{16}$	256	—	—	—	2799.34	—	111.55	<b>67.29</b>
NETGEN-DEG	$2^{12}$	32	657.29	2655.80	7.41	4.40	35.54	0.43	<b>0.13</b>
	$2^{12}$	128	3523.55	—	32.03	12.73	273.76	1.98	<b>0.39</b>
	$2^{12}$	512	—	—	147.58	35.58	2473.16	7.72	<b>1.31</b>
	$2^{12}$	2048	—	—	624.44	99.49	—	47.08	<b>6.68</b>
Normalized time									
Problem family	$n$	$m/n$	SCC	MMCC	CAT	SSP	CAS	COS	NS
NETGEN-8	$2^{10}$	8	481.00	1185.00	19.00	12.00	19.00	2.00	<b>1.00</b>
	$2^{13}$	8	3605.33	13457.40	45.40	43.20	141.13	2.20	<b>1.00</b>
	$2^{16}$	8	—	—	69.59	70.51	556.65	<b>1.00</b>	1.62
	$2^{19}$	8	—	—	—	—	—	<b>1.00</b>	8.09
	$2^{22}$	8	—	—	—	—	—	<b>1.00</b>	—
NETGEN-SR	$2^{10}$	32	3403.00	9082.00	74.00	33.00	179.00	6.00	<b>1.00</b>
	$2^{12}$	64	7583.90	—	77.90	32.24	335.00	3.95	<b>1.00</b>
	$2^{14}$	128	—	—	99.50	38.36	—	2.33	<b>1.00</b>
	$2^{16}$	256	—	—	—	41.60	—	1.66	<b>1.00</b>
NETGEN-DEG	$2^{12}$	32	5056.08	20429.23	57.00	33.85	273.38	3.31	<b>1.00</b>
	$2^{12}$	128	9034.74	—	82.13	32.64	701.95	5.08	<b>1.00</b>
	$2^{12}$	512	—	—	112.66	27.16	1887.91	5.89	<b>1.00</b>
	$2^{12}$	2048	—	—	93.48	14.89	—	7.05	<b>1.00</b>

Table 2: Comparison of LEMON implementations on NETGEN networks

Running time (seconds)									
Problem family	$n$	$m/n$	SCC	MMCC	CAT	SSP	CAS	COS	NS
GOTO-8	$2^{10}$	8	88.85	30.39	0.63	2.27	0.13	0.05	<b>0.01</b>
	$2^{13}$	8	—	—	44.82	241.27	4.41	1.45	<b>0.79</b>
	$2^{16}$	8	—	—	—	—	143.92	<b>41.27</b>	202.47
	$2^{19}$	8	—	—	—	—	—	<b>1563.90</b>	—
GOTO-SR	$2^{10}$	32	2474.46	584.35	2.30	35.66	2.81	0.30	<b>0.08</b>
	$2^{12}$	64	—	—	89.15	—	143.35	5.27	<b>3.16</b>
	$2^{14}$	128	—	—	2650.95	—	—	<b>69.37</b>	163.37
	$2^{16}$	256	—	—	—	—	—	<b>1279.72</b>	—

Normalized time									
Problem family	$n$	$m/n$	SCC	MMCC	CAT	SSP	CAS	COS	NS
GOTO-8	$2^{10}$	8	8885.00	3039.00	63.00	227.00	13.00	5.00	<b>1.00</b>
	$2^{13}$	8	—	—	56.73	305.41	5.58	1.84	<b>1.00</b>
	$2^{16}$	8	—	—	—	—	3.49	<b>1.00</b>	4.91
	$2^{19}$	8	—	—	—	—	—	<b>1.00</b>	—
GOTO-SR	$2^{10}$	32	30930.75	7304.38	28.75	445.75	35.13	3.75	<b>1.00</b>
	$2^{12}$	64	—	—	28.21	—	45.36	1.67	<b>1.00</b>
	$2^{14}$	128	—	—	38.21	—	—	<b>1.00</b>	2.36
	$2^{16}$	256	—	—	—	—	—	<b>1.00</b>	—

Table 3: Comparison of LEMON implementations on GOTO networks

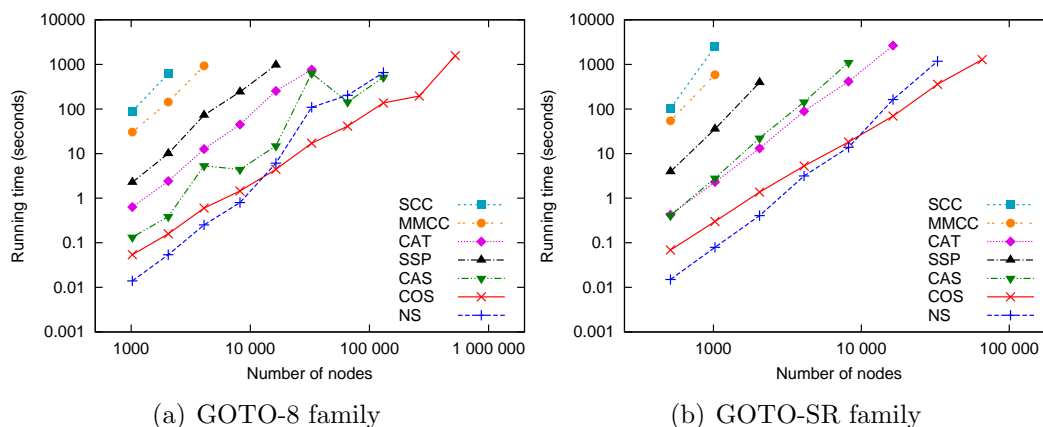


Figure 2: Comparison of LEMON implementations on GOTO networks

capacities are uniformly set to one, but it is slower than SSP on ROAD-FLOW instances. The COS and NS algorithms perform significantly worse than SSP and CAS, especially in case of the ROAD-PATHS family, while the cycle-canceling algorithms are drastically slower than the other methods.

According to these results and many additional experiments (e.g., on GRIDGEN and VISION networks), we can conclude that COS and NS are generally the most efficient algorithms provided by the LEMON library. On particular problem instances, however, the dual-ascent augmenting path algorithms, SSP and CAS, could be significantly faster (e.g., GRID-LONG and ROAD networks).

Running time (seconds)									
Problem family	$n$	$m/n$	SCC	MMCC	CAT	SSP	CAS	COS	NS
GRID-WIDE	$2^{10}$	2.04	0.75	1.32	0.11	0.05	<b>0.01</b>	0.02	<b>0.01</b>
	$2^{13}$	2.06	76.05	144.55	2.55	4.31	1.00	0.32	<b>0.03</b>
	$2^{16}$	2.06	–	–	55.58	395.44	100.19	8.95	<b>0.65</b>
	$2^{19}$	2.06	–	–	1155.56	–	–	522.82	<b>12.68</b>
GRID-LONG	$2^{10}$	1.95	0.61	0.91	0.13	<b>0.01</b>	<b>0.01</b>	0.02	<b>0.01</b>
	$2^{13}$	1.94	58.86	36.21	–	0.11	<b>0.05</b>	0.19	0.08
	$2^{16}$	1.94	2871.00	–	–	0.46	<b>0.35</b>	2.38	5.45
	$2^{19}$	1.94	–	–	–	<b>3.03</b>	3.18	24.05	438.96
GRID-SQUARE	$2^{10}$	2	0.71	1.41	0.13	0.04	<b>0.01</b>	0.02	<b>0.01</b>
	$2^{13}$	2	116.16	194.48	5.21	2.40	0.30	0.34	<b>0.07</b>
	$2^{16}$	2	–	–	301.64	204.89	10.16	6.65	<b>4.11</b>
	$2^{19}$	2	–	–	–	–	1800.86	<b>156.58</b>	487.19
Normalized time									
Problem family	$n$	$m/n$	SCC	MMCC	CAT	SSP	CAS	COS	NS
GRID-WIDE	$2^{10}$	2.04	75.00	132.00	11.00	5.00	<b>1.00</b>	2.00	<b>1.00</b>
	$2^{13}$	2.06	2535.00	4818.33	85.00	143.67	33.33	10.67	<b>1.00</b>
	$2^{16}$	2.06	–	–	85.51	608.37	154.14	13.77	<b>1.00</b>
	$2^{19}$	2.06	–	–	91.13	–	–	41.23	<b>1.00</b>
GRID-LONG	$2^{10}$	1.95	61.00	91.00	13.00	<b>1.00</b>	<b>1.00</b>	2.00	<b>1.00</b>
	$2^{13}$	1.94	1177.20	724.20	–	2.20	<b>1.00</b>	3.80	1.60
	$2^{16}$	1.94	8202.86	–	–	1.31	<b>1.00</b>	6.80	15.57
	$2^{19}$	1.94	–	–	–	<b>1.00</b>	1.05	7.94	144.87
GRID-SQUARE	$2^{10}$	2	71.00	141.00	13.00	4.00	<b>1.00</b>	2.00	<b>1.00</b>
	$2^{13}$	2	1659.43	2778.29	74.43	34.29	4.29	4.86	<b>1.00</b>
	$2^{16}$	2	–	–	73.39	49.85	2.47	1.62	<b>1.00</b>
	$2^{19}$	2	–	–	–	–	11.50	<b>1.00</b>	3.11

Table 4: Comparison of LEMON implementations on GRIDGRAPH networks

Running time (seconds)									
Problem family	$n$	$m/n$	SCC	MMCC	CAT	SSP	CAS	COS	NS
ROAD-PATHS	9 559	3.11	2.29	372.91	2.17	<b>0.01</b>	<b>0.01</b>	0.14	0.04
	116 920	2.27	150.84	–	264.52	<b>0.26</b>	<b>0.26</b>	4.61	2.97
	519 157	2.44	–	–	–	<b>3.29</b>	<b>3.29</b>	35.87	41.38
	2 073 870	2.49	–	–	–	<b>21.99</b>	<b>21.99</b>	238.57	744.36
ROAD-FLOW	9 559	3.11	6.81	625.60	2.94	<b>0.03</b>	<b>0.03</b>	0.20	0.04
	116 920	2.27	415.01	–	289.30	<b>0.71</b>	1.30	6.64	3.45
	519 157	2.44	–	–	–	<b>10.89</b>	17.72	54.54	47.96
	2 073 870	2.49	–	–	–	<b>77.68</b>	230.24	398.62	759.80
Normalized time									
Problem family	$n$	$m/n$	SCC	MMCC	CAT	SSP	CAS	COS	NS
ROAD-PATHS	9 559	3.11	229.00	37291.00	217.00	<b>1.00</b>	<b>1.00</b>	14.00	4.00
	116 920	2.27	580.15	–	1017.38	<b>1.00</b>	<b>1.00</b>	17.73	11.42
	519 157	2.44	–	–	–	<b>1.00</b>	<b>1.00</b>	10.90	12.58
	2 073 870	2.49	–	–	–	<b>1.00</b>	<b>1.00</b>	10.85	33.85
ROAD-FLOW	9 559	3.11	227.00	20853.33	98.00	<b>1.00</b>	<b>1.00</b>	6.67	1.33
	116 920	2.27	584.52	–	407.46	<b>1.00</b>	1.83	9.35	4.86
	519 157	2.44	–	–	–	<b>1.00</b>	1.63	5.01	4.40
	2 073 870	2.49	–	–	–	<b>1.00</b>	2.96	5.13	9.78

Table 5: Comparison of LEMON implementations on ROAD networks

### 4.3 Comparison of various MCF solvers

In this subsection, we compare the best MCF solution methods of LEMON, mainly COS and NS, to other publicly available MCF solvers described in Section 3.2. CS2,



MCFZIB, CPLEX, and RelaxIV are widely known and have served as benchmarks for a long time (see, e.g., [14, 74, 46, 40]), while the LEDA library and PDNET are rarely included in experimental studies, and MCFSimplex is a recent implementation.

All of these codes were compiled with the same compiler and optimization settings as we used for the LEMON implementations (GCC 4.5.3 with `-O3` optimization), and they were also executed applying a time limit of one hour. We used all solvers with their default options as we were interested in evaluating their robustness without exploiting the flexibility they provide in parameter settings. However, flexibility may also be important in many applications.

As in the previous subsection, we always report average running time over three different problem instances of the same size. To avoid overwhelmed charts, the results of MCFSimplex and LEDA are not depicted in them. In general, LEDA performed worse than or similarly to the other two cost-scaling methods, COS in LEMON and CS2; while MCFSimplex has much in common with MCFZIB and thereby their performance is also similar. However, the tables report the results for these two solvers, as well, but note that MCFSimplex is abbreviated as MCFSim.

Table 6 and Figure 3 compare the implementations on NETGEN instances, while Table 7 compares them on GRIDGEN families. The NS code of the LEMON library turned out to be the most efficient on most of these networks, but it is consistently outperformed by the cost-scaling algorithms and RelaxIV on the largest sparse graphs. MCFZIB and MCFSimplex are typically slower than NS by a factor between 2 and 5, while CPLEX turned out to be even less efficient. Our COS code performs similarly to or slightly slower than CS2 on these networks. However, the third cost-scaling implementation, LEDA, is at least 1.5-2 times slower than them. Furthermore, it failed to solve the largest instances due to number overflow errors, which is denoted as “*error*” in the tables. Since the LEDA library has closed source, we could not replace the internally used number types to eliminate this issue. RelaxIV turned out to be highly efficient on these families, it is competitive with or even faster than the other codes. In contrast, PDNET is much slower than all other solvers.

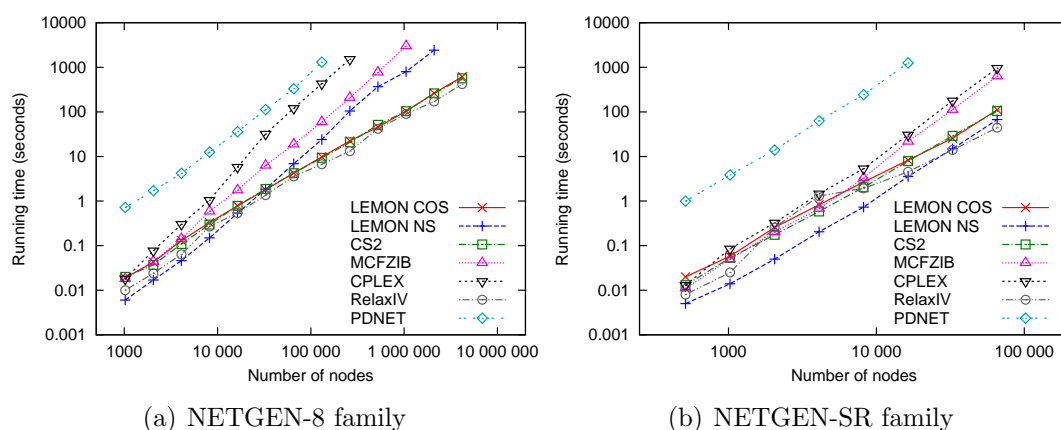


Figure 3: Comparison of MCF solvers on NETGEN networks

The performance results for the GOTO families are presented in Table 8 and Fig-

Running time (seconds)											
Problem family	$n$	$m/n$	LEMON		Other solvers						
			COS	NS	CS2	LEDA	MCFZIB	CPLEX	MCFSim	RelaxIV	PDNET
NETGEN-8	$2^{10}$	8	0.02	<b>0.01</b>	0.02	0.03	0.02	0.02	<b>0.01</b>	<b>0.01</b>	0.72
	$2^{13}$	8	0.33	<b>0.15</b>	0.30	0.46	0.59	1.05	0.43	0.27	12.55
	$2^{16}$	8	4.24	6.88	4.22	8.34	18.78	120.57	24.70	<b>3.58</b>	330.27
	$2^{19}$	8	45.56	368.52	51.15	<i>error</i>	770.84	—	1072.96	<b>41.82</b>	—
	$2^{22}$	8	615.42	—	566.32	<i>error</i>	—	—	—	<b>419.38</b>	—
NETGEN-SR	$2^{10}$	32	0.06	<b>0.01</b>	0.05	0.08	0.05	0.09	0.04	0.03	3.84
	$2^{12}$	64	0.83	<b>0.21</b>	0.59	1.41	0.68	1.46	0.69	1.30	62.88
	$2^{14}$	128	8.28	<b>3.55</b>	7.86	19.61	21.68	30.65	15.18	4.53	1251.55
	$2^{16}$	256	111.55	67.29	104.25	274.06	634.38	958.17	809.51	<b>44.37</b>	—
NETGEN-DEG	$2^{12}$	32	0.43	<b>0.13</b>	0.31	0.70	0.42	1.21	0.41	0.37	24.63
	$2^{12}$	128	1.98	<b>0.39</b>	1.12	2.91	1.32	2.31	1.43	1.40	171.44
	$2^{12}$	512	7.72	<b>1.31</b>	5.15	12.56	4.74	8.37	5.19	3.20	1286.69
	$2^{12}$	2048	47.08	<b>6.68</b>	32.72	69.55	18.52	41.78	15.34	12.81	—

Normalized time											
Problem family	$n$	$m/n$	LEMON		Other solvers						
			COS	NS	CS2	LEDA	MCFZIB	CPLEX	MCFSim	RelaxIV	PDNET
NETGEN-8	$2^{10}$	8	2.00	<b>1.00</b>	2.00	3.00	2.00	2.00	<b>1.00</b>	<b>1.00</b>	72.00
	$2^{13}$	8	2.20	<b>1.00</b>	2.00	3.07	3.93	7.00	2.87	1.80	83.67
	$2^{16}$	8	1.18	1.92	1.18	2.33	5.25	33.68	6.90	<b>1.00</b>	92.25
	$2^{19}$	8	1.09	8.81	1.22	<i>error</i>	18.43	—	25.66	<b>1.00</b>	—
	$2^{22}$	8	1.47	—	1.35	<i>error</i>	—	—	—	<b>1.00</b>	—
NETGEN-SR	$2^{10}$	32	6.00	<b>1.00</b>	5.00	8.00	5.00	9.00	4.00	3.00	384.00
	$2^{12}$	64	3.95	<b>1.00</b>	2.81	6.71	3.24	6.95	3.29	6.19	299.43
	$2^{14}$	128	2.33	<b>1.00</b>	2.21	5.52	6.11	8.63	4.28	1.28	352.55
	$2^{16}$	256	2.51	1.52	2.35	6.18	14.30	21.59	18.24	<b>1.00</b>	—
NETGEN-DEG	$2^{12}$	32	3.31	<b>1.00</b>	2.38	5.38	3.23	9.31	3.15	2.85	189.46
	$2^{12}$	128	5.08	<b>1.00</b>	2.87	7.46	3.38	5.92	3.67	3.59	439.59
	$2^{12}$	512	5.89	<b>1.00</b>	3.93	9.59	3.62	6.39	3.96	2.44	982.21
	$2^{12}$	2048	7.05	<b>1.00</b>	4.90	10.41	2.77	6.25	2.30	1.92	—

Table 6: Comparison of MCF solvers on NETGEN networks

ure 4. In these tests, COS and CS2 also perform similarly, and they are the most efficient on the largest instances of both families. LEDA is also similarly efficient on GOTO-8 networks, but it is 2-3 times slower than CS2 and COS on the GOTO-SR family. NS turned out to be an order of magnitude faster than the other network simplex codes, MCFZIB, CPLEX, and MCFSimplex. Similarly to the earlier results, NS is the most efficient on relatively small graphs, but it is substantially slower than the cost-scaling codes on the large networks, especially on the sparse ones. RelaxIV turned out to be very slow on these hard instances, which is in sharp contrast to its efficiency on the NETGEN and GRIDGEN families. On the other hand, the relative performance of PDNET compared to the other solvers is much better in case of GOTO networks than on easier instances. It outperforms the network simplex codes on the large GOTO-8 instances, with the only exception of our implementation.

Table 9 and Figure 5 show the results for the GRIDGRAPH families, while Table 10 and Figure 6 present the results for the ROAD families. PDNET failed on most of these special MCF instances with an error message: “STOP disconnected network.” Therefore, its running time is not reported for these problem families. However, the CAS algorithm of LEMON is included instead, because the augmenting path methods turned out to be interesting in case of these particular networks.

Running time (seconds)											
Problem family	$n$	$m/n$	LEMON		Other solvers						
			COS	NS	CS2	LEDA	MCFZIB	CPLEX	MCFSim	RelaxIV	PDNET
GRIDGEN-8	$2^{10}$	8	0.02	<b>0.01</b>	0.02	0.03	<b>0.01</b>	0.02	<b>0.01</b>	<b>0.01</b>	0.81
	$2^{13}$	8	0.40	<b>0.14</b>	0.34	0.50	0.52	1.26	0.42	0.34	13.49
	$2^{16}$	8	4.46	4.12	4.27	8.20	11.06	110.56	11.93	<b>3.60</b>	394.16
	$2^{19}$	8	<b>53.78</b>	166.64	66.65	<i>error</i>	644.57	—	586.77	82.06	—
	$2^{22}$	8	513.61	—	440.64	<i>error</i>	—	—	—	<b>352.64</b>	—
GRIDGEN-SR	$2^{10}$	32	0.07	<b>0.01</b>	0.06	0.10	0.05	0.15	0.07	0.04	3.27
	$2^{12}$	64	0.97	<b>0.19</b>	0.72	1.67	0.58	3.04	1.02	1.25	41.41
	$2^{14}$	128	9.59	<b>2.73</b>	8.90	21.25	14.57	70.38	16.04	13.19	840.75
	$2^{16}$	256	147.44	<b>62.29</b>	128.71	339.53	371.56	2522.05	782.92	106.67	—
GRIDGEN-DEG	$2^{12}$	32	0.47	<b>0.10</b>	0.37	0.81	0.34	1.22	0.32	0.55	20.08
	$2^{12}$	128	2.04	<b>0.33</b>	1.32	3.59	1.03	3.02	1.03	3.77	96.64
	$2^{12}$	512	9.25	<b>1.19</b>	5.52	14.40	3.83	9.47	3.85	5.88	381.22
	$2^{12}$	2048	50.15	<b>4.91</b>	31.62	83.16	14.69	61.62	17.87	18.80	1742.13

Normalized time											
Problem family	$n$	$m/n$	LEMON		Other solvers						
			COS	NS	CS2	LEDA	MCFZIB	CPLEX	MCFSim	RelaxIV	PDNET
GRIDGEN-8	$2^{10}$	8	2.00	<b>1.00</b>	2.00	3.00	<b>1.00</b>	2.00	<b>1.00</b>	<b>1.00</b>	81.00
	$2^{13}$	8	2.86	<b>1.00</b>	2.43	3.57	3.71	9.00	3.00	2.43	96.36
	$2^{16}$	8	1.24	1.14	1.19	2.28	3.07	30.71	3.31	<b>1.00</b>	109.49
	$2^{19}$	8	<b>1.00</b>	3.10	1.24	<i>error</i>	11.99	—	10.91	1.53	—
	$2^{22}$	8	1.46	—	1.25	<i>error</i>	—	—	—	<b>1.00</b>	—
GRIDGEN-SR	$2^{10}$	32	7.00	<b>1.00</b>	6.00	10.00	5.00	15.00	7.00	4.00	327.00
	$2^{12}$	64	5.11	<b>1.00</b>	3.79	8.79	3.05	16.00	5.37	6.58	217.95
	$2^{14}$	128	3.51	<b>1.00</b>	3.26	7.78	5.34	25.78	5.88	4.83	307.97
	$2^{16}$	256	2.37	<b>1.00</b>	2.07	5.45	5.97	40.49	12.57	1.71	—
GRIDGEN-DEG	$2^{12}$	32	4.70	<b>1.00</b>	3.70	8.10	3.40	12.20	3.20	5.50	200.80
	$2^{12}$	128	6.18	<b>1.00</b>	4.00	10.88	3.12	9.15	3.12	11.42	292.85
	$2^{12}$	512	7.77	<b>1.00</b>	4.64	12.10	3.22	7.96	3.24	4.94	320.35
	$2^{12}$	2048	10.21	<b>1.00</b>	6.44	16.94	2.99	12.55	3.64	3.83	354.81

Table 7: Comparison of MCF solvers on GRIDGEN networks

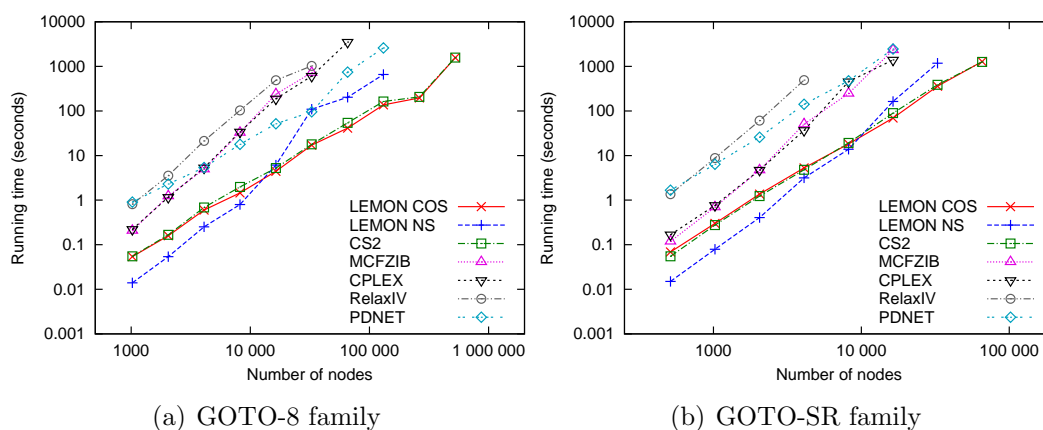


Figure 4: Comparison of MCF solvers on GOTO networks

On the GRID-WIDE instances (see Figure 5(a)), the primal network simplex algorithms are the fastest, probably because they are capable of canceling a lot of short negative cycles efficiently due to the limited depth of the underlying spanning tree data structure. The NS code of LEMON turned out to be the most efficient im-

Running time (seconds)											
Problem family	$n$	$m/n$	LEMON			Other solvers					
			COS	NS	CS2	LEDA	MCFZIB	CPLEX	MCFSim	RelaxIV	PDNET
GOTO-8	$2^{10}$	8	0.05	<b>0.01</b>	0.06	0.08	0.21	0.22	0.17	0.81	0.91
	$2^{13}$	8	1.45	<b>0.79</b>	1.98	1.74	32.76	33.90	32.85	103.14	17.81
	$2^{16}$	8	<b>41.27</b>	202.47	54.05	58.50	–	3501.53	–	–	743.62
	$2^{19}$	8	<b>1563.90</b>	–	1576.96	1839.08	–	–	–	–	–
GOTO-SR	$2^{10}$	32	0.30	<b>0.08</b>	0.28	0.38	0.70	0.76	0.54	8.81	6.38
	$2^{12}$	64	5.27	<b>3.16</b>	4.78	9.57	50.65	36.76	41.75	492.38	140.78
	$2^{14}$	128	<b>69.37</b>	163.37	89.32	184.97	2339.57	1391.05	1702.07	–	2447.43
	$2^{16}$	256	1279.72	–	<b>1259.63</b>	–	–	–	–	–	–

Normalized time											
Problem family	$n$	$m/n$	LEMON			Other solvers					
			COS	NS	CS2	LEDA	MCFZIB	CPLEX	MCFSim	RelaxIV	PDNET
GOTO-8	$2^{10}$	8	5.00	<b>1.00</b>	6.00	8.00	21.00	22.00	17.00	81.00	91.00
	$2^{13}$	8	1.84	<b>1.00</b>	2.51	2.20	41.47	42.91	41.58	130.56	22.54
	$2^{16}$	8	<b>1.00</b>	4.91	1.31	1.42	–	84.84	–	–	18.02
	$2^{19}$	8	<b>1.00</b>	–	1.01	1.18	–	–	–	–	–
GOTO-SR	$2^{10}$	32	3.75	<b>1.00</b>	3.50	4.75	8.75	9.50	6.75	110.13	79.75
	$2^{12}$	64	1.67	<b>1.00</b>	1.51	3.03	16.03	11.63	13.21	155.82	44.55
	$2^{14}$	128	<b>1.00</b>	2.36	1.29	2.67	33.73	20.05	24.54	–	35.28
	$2^{16}$	256	1.02	–	<b>1.00</b>	–	–	–	–	–	–

Table 8: Comparison of MCF solvers on GOTO networks

Running time (seconds)											
Problem family	$n$	$m/n$	LEMON			Other solvers					
			CAS	COS	NS	CS2	LEDA	MCFZIB	CPLEX	MCFSim	RelaxIV
GRID-WIDE	$2^{10}$	2.04	<b>0.01</b>	0.02	<b>0.01</b>	<b>0.01</b>	0.02	<b>0.01</b>	<b>0.01</b>	<b>0.01</b>	0.03
	$2^{13}$	2.06	1.00	0.32	<b>0.03</b>	0.29	0.55	0.10	0.12	0.08	2.89
	$2^{16}$	2.06	100.19	8.95	<b>0.65</b>	6.33	51.88	2.10	2.05	1.78	761.30
	$2^{19}$	2.06	–	522.82	<b>12.68</b>	121.89	<i>error</i>	20.11	24.69	16.84	–
GRID-LONG	$2^{10}$	1.95	<b>0.01</b>	0.02	<b>0.01</b>	<b>0.01</b>	<b>0.01</b>	<b>0.01</b>	<b>0.01</b>	<b>0.01</b>	0.02
	$2^{13}$	1.94	<b>0.05</b>	0.19	0.08	0.12	0.15	0.42	0.42	0.49	0.80
	$2^{16}$	1.94	<b>0.35</b>	2.38	5.45	3.18	2.33	85.34	46.52	<i>error</i>	75.25
	$2^{19}$	1.94	<b>3.18</b>	24.05	438.96	28.05	<i>error</i>	–	–	<i>error</i>	–
GRID-SQUARE	$2^{10}$	2	<b>0.01</b>	0.02	<b>0.01</b>	<b>0.01</b>	0.02	<b>0.01</b>	<b>0.01</b>	<b>0.01</b>	0.03
	$2^{13}$	2	0.30	0.34	<b>0.07</b>	0.21	0.25	0.26	0.19	0.18	2.11
	$2^{16}$	2	10.16	6.65	<b>4.11</b>	4.47	4.87	25.35	13.34	16.30	194.98
	$2^{19}$	2	1800.86	156.58	487.19	<b>77.59</b>	<i>error</i>	1857.09	1522.92	2524.58	–

Normalized time											
Problem family	$n$	$m/n$	LEMON			Other solvers					
			CAS	COS	NS	CS2	LEDA	MCFZIB	CPLEX	MCFSim	RelaxIV
GRID-WIDE	$2^{10}$	2.04	<b>1.00</b>	2.00	<b>1.00</b>	<b>1.00</b>	2.00	<b>1.00</b>	<b>1.00</b>	<b>1.00</b>	3.00
	$2^{13}$	2.06	33.33	10.67	<b>1.00</b>	9.67	18.33	3.33	4.00	2.67	96.33
	$2^{16}$	2.06	154.14	13.77	<b>1.00</b>	9.74	79.82	3.23	3.15	2.74	1171.23
	$2^{19}$	2.06	–	41.23	<b>1.00</b>	9.61	<i>error</i>	1.59	1.95	1.33	–
GRID-LONG	$2^{10}$	1.95	<b>1.00</b>	2.00	<b>1.00</b>	<b>1.00</b>	<b>1.00</b>	<b>1.00</b>	<b>1.00</b>	<b>1.00</b>	2.00
	$2^{13}$	1.94	<b>1.00</b>	3.80	1.60	2.40	3.00	8.40	8.40	9.80	16.00
	$2^{16}$	1.94	<b>1.00</b>	6.80	15.57	9.09	6.66	243.83	132.91	<i>error</i>	215.00
	$2^{19}$	1.94	<b>1.00</b>	7.56	138.04	8.82	<i>error</i>	–	–	<i>error</i>	–
GRID-SQUARE	$2^{10}$	2	<b>1.00</b>	2.00	<b>1.00</b>	<b>1.00</b>	2.00	<b>1.00</b>	<b>1.00</b>	<b>1.00</b>	3.00
	$2^{13}$	2	4.29	4.86	<b>1.00</b>	3.00	3.57	3.71	2.71	2.57	30.14
	$2^{16}$	2	2.47	1.62	<b>1.00</b>	1.09	1.18	6.17	3.25	3.97	47.44
	$2^{19}$	2	23.21	2.02	6.28	<b>1.00</b>	<i>error</i>	23.93	19.63	32.54	–

Table 9: Comparison of MCF solvers on GRIDGRAPH networks




Figure 5: Comparison of MCF solvers on GRIDGRAPH networks

plementation again. The cost-scaling methods perform clearly worse than network simplex algorithms on these networks (even on the largest ones), while CAS and RelaxIV are even much slower. On the other hand, in case of GRID-LONG networks (see Figure 5(b)), the augmenting paths or cycles are much longer, hence CAS, the capacity-scaling augmenting path algorithm, turned out to be the fastest. The cost-scaling codes are significantly slower, but they clearly outperform the network simplex implementations and RelaxIV. The experiments on the GRID-SQUARE family are more similar to the previous results on NETGEN, GRIDGEN, and GOTO networks: NS is the fastest on the smaller graphs, while the cost-scaling algorithms are superior to it on the largest ones.

On the ROAD-PATHS family (see Figure 6(a)), CAS is much faster than the network simplex and cost-scaling solution methods, while RelaxIV is even orders of magnitude slower. The results are similar on the ROAD-FLOW networks, as well (see Figure 6(b)), with the only exception that CAS has smaller advantage due to the larger capacity values. CS2 is superior to CAS on the largest ROAD-FLOW instances, but note that the SSP algorithm of LEMON is even faster (cf. Tables 5 and 10).

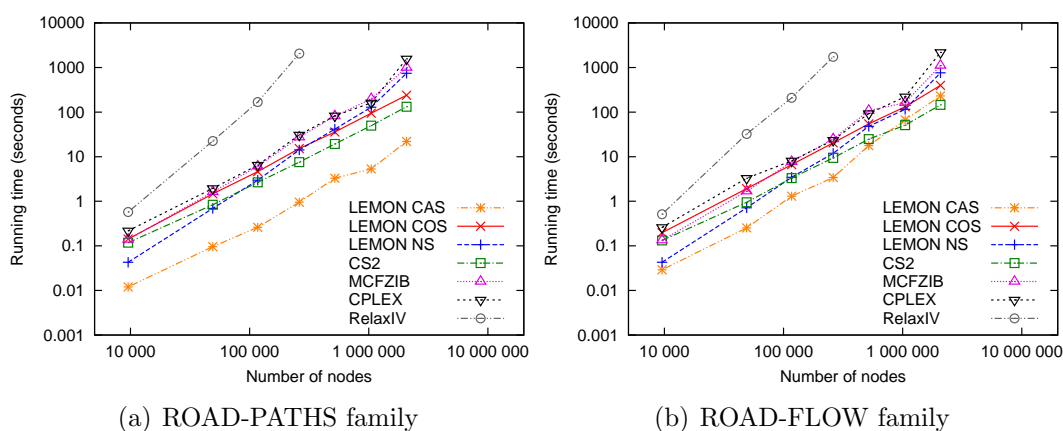


Figure 6: Comparison of MCF solvers on ROAD networks

Running time (seconds)											
Problem family	$n$	$m/n$	LEMON			Other solvers					
			CAS	COS	NS	CS2	LEDA	MCFZIB	CPLEX	MCFSim	RelaxIV
ROAD-PATHS	9 559	3.11	<b>0.01</b>	0.14	0.04	0.12	0.14	0.14	0.22	0.11	0.58
	116 920	2.27	<b>0.26</b>	4.61	2.97	2.67	3.19	6.09	6.48	5.58	168.03
	519 157	2.44	<b>3.29</b>	35.87	41.38	19.26	<i>error</i>	81.07	82.69	76.52	–
	2 073 870	2.49	<b>21.99</b>	238.57	744.36	131.90	<i>error</i>	992.32	1568.30	894.20	–
ROAD-FLOW	9 559	3.11	<b>0.03</b>	0.20	0.04	0.13	0.19	0.13	0.26	0.17	0.51
	116 920	2.27	<b>1.30</b>	6.64	3.45	3.31	5.00	7.69	8.07	7.07	209.21
	519 157	2.44	<b>17.72</b>	54.54	47.96	24.79	<i>error</i>	108.06	92.95	95.39	–
	2 073 870	2.49	230.24	398.62	759.80	<b>145.70</b>	<i>error</i>	1106.14	2181.50	1755.38	–

Normalized time											
Problem family	$n$	$m/n$	LEMON			Other solvers					
			CAS	COS	NS	CS2	LEDA	MCFZIB	CPLEX	MCFSim	RelaxIV
ROAD-PATHS	9 559	3.11	<b>1.00</b>	14.00	4.00	12.00	14.00	14.00	22.00	11.00	58.00
	116 920	2.27	<b>1.00</b>	17.73	11.42	10.27	12.27	23.42	24.92	21.46	646.27
	519 157	2.44	<b>1.00</b>	10.90	12.58	5.85	<i>error</i>	24.64	25.13	23.26	–
	2 073 870	2.49	<b>1.00</b>	10.85	33.85	6.00	<i>error</i>	45.13	71.32	40.66	–
ROAD-FLOW	9 559	3.11	<b>1.00</b>	6.67	1.33	4.33	6.33	4.33	8.67	5.67	17.00
	116 920	2.27	<b>1.00</b>	5.11	2.65	2.55	3.85	5.92	6.21	5.44	160.93
	519 157	2.44	<b>1.00</b>	3.08	2.71	1.40	<i>error</i>	6.10	5.25	5.38	–
	2 073 870	2.49	1.58	2.74	5.21	<b>1.00</b>	<i>error</i>	7.59	14.97	12.05	–

Table 10: Comparison of MCF solvers on ROAD networks

Finally, the benchmark results on the VISION families are presented in Table 11 and Figure 7. CS2 is the fastest to solve these hard instances, while COS is about 1.5-2 times slower than it. All other codes perform much worse, including the third cost-scaling implementation, LEDA. It even failed to solve all VISION-PROP instances due to number overflow errors. NS turned out to be much faster and more robust than the other three network simplex codes. RelaxIV is very slow on these instances, but PDNET performs even much worse, it could not solve any VISION instance within the one-hour time limit.

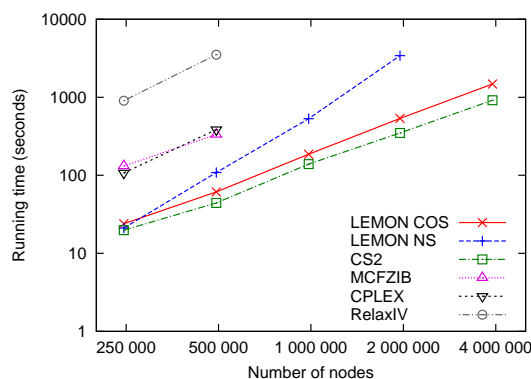


Figure 7: Comparison of MCF solvers on VISION-RND networks

The choice of the cost function has only modest impact on the performance of the algorithms on VISION networks. Surprisingly, the cost-scaling codes tend to be slightly slower on both VISION-PROP and VISION-INV instances than in case of random costs, while other methods are faster on these networks, especially on VISION-INV

Running time (seconds)											
Problem family	$n$	$m/n$	LEMON		Other solvers						
			COS	NS	CS2	LEDA	MCFZIB	CPLEX	MCFSim	RelaxIV	PDNET
VISION-RND	245 762	5.82	23.96	21.06	<b>19.78</b>	54.62	131.66	107.46	180.15	904.74	–
	491 522	5.85	61.25	108.64	<b>44.18</b>	207.51	333.20	385.57	463.05	3518.20	–
	983 042	5.88	186.42	531.59	<b>139.05</b>	1250.56	–	–	–	–	–
	1 949 698	5.91	535.36	3420.76	<b>348.28</b>	–	–	–	–	–	–
	3 899 394	5.92	1475.95	–	<b>916.47</b>	–	–	–	–	–	–
VISION-PROP	245 762	5.82	49.01	<b>20.97</b>	25.37	<i>error</i>	134.48	100.51	143.56	899.03	–
	491 522	5.85	113.35	105.93	<b>61.48</b>	<i>error</i>	227.63	385.75	243.10	1638.86	–
	983 042	5.88	431.25	459.52	<b>206.59</b>	<i>error</i>	3261.30	–	–	–	–
	1 949 698	5.91	937.60	–	<b>477.72</b>	<i>error</i>	–	–	–	–	–
	3 899 394	5.92	2558.97	–	<b>1227.11</b>	<i>error</i>	–	–	–	–	–
VISION-INV	245 762	5.82	40.75	<b>14.05</b>	24.82	48.95	28.76	69.71	35.04	368.44	–
	491 522	5.85	93.94	76.74	58.69	205.60	<b>57.35</b>	311.69	68.76	1465.07	–
	983 042	5.88	301.53	292.31	<b>179.18</b>	616.67	2007.66	–	–	–	–
	1 949 698	5.91	909.44	2637.98	<b>451.09</b>	2355.10	–	–	–	–	–
	3 899 394	5.92	2348.72	–	<b>1144.18</b>	–	–	–	–	–	–

Normalized time											
Problem family	$n$	$m/n$	LEMON		Other solvers						
			COS	NS	CS2	LEDA	MCFZIB	CPLEX	MCFSim	RelaxIV	PDNET
VISION-RND	245 762	5.82	1.21	1.06	<b>1.00</b>	2.76	6.66	5.43	9.11	45.74	–
	491 522	5.85	1.39	2.46	<b>1.00</b>	4.70	7.54	8.73	10.48	79.63	–
	983 042	5.88	1.34	3.82	<b>1.00</b>	8.99	–	–	–	–	–
	1 949 698	5.91	1.54	9.82	<b>1.00</b>	–	–	–	–	–	–
	3 899 394	5.92	1.61	–	<b>1.00</b>	–	–	–	–	–	–
VISION-PROP	245 762	5.82	2.34	<b>1.00</b>	1.21	<i>error</i>	6.41	4.79	6.85	42.87	–
	491 522	5.85	1.84	1.72	<b>1.00</b>	<i>error</i>	3.70	6.27	3.95	26.66	–
	983 042	5.88	2.09	2.22	<b>1.00</b>	<i>error</i>	15.79	–	–	–	–
	1 949 698	5.91	1.96	–	<b>1.00</b>	<i>error</i>	–	–	–	–	–
	3 899 394	5.92	2.09	–	<b>1.00</b>	<i>error</i>	–	–	–	–	–
VISION-INV	245 762	5.82	2.90	<b>1.00</b>	1.77	3.48	2.05	4.96	2.49	26.22	–
	491 522	5.85	1.64	1.34	1.02	3.59	<b>1.00</b>	5.43	1.20	25.55	–
	983 042	5.88	1.68	1.63	<b>1.00</b>	3.44	11.20	–	–	–	–
	1 949 698	5.91	2.02	5.85	<b>1.00</b>	5.22	–	–	–	–	–
	3 899 394	5.92	2.05	–	<b>1.00</b>	–	–	–	–	–	–

Table 11: Comparison of MCF solvers on VISION networks

instances.

Apart from the presented results, many other experiments have also been conducted varying several parameters of the problem instances. A notable observation is that the magnitudes of the supply, capacity, and cost values hardly affect the running time of the algorithms, but it is important whether the capacities incorporate “tight” or “loose” bounds for the feasible solutions with respect to the supply values (see, e.g., [15]). Instances with loose capacity bounds or infinite capacities are easier to solve, especially for network simplex methods, while in case of tight capacities, the cost-scaling algorithms are typically more robust.

We also compared the number of basic operations performed by the cost-scaling and primal network simplex implementations whose source codes we had access to. We found that our COS code tends to perform significantly less push operations than CS2 due to the partial augment method (see Section 3.1.6), but it performs approximately the same number of relabel operations in most cases. On the other hand, CS2 may find admissible arcs faster due to its speculative arc fixing heuristic

(see [48, 46]). The COS implementation could be further improved by applying this technique.

As for the network simplex algorithms, MCFZIB and MCFSimplex perform exactly the same number of iterations as they apply the same pivot strategy. Our NS code, however, tends to perform substantially less iterations than them, although it performs either less or more arc pricing steps depending on the characteristics of the problem instance.

## 5 Conclusions

We have considered several efficient implementations of algorithms for solving the minimum-cost network flow problem, which is essential in a wide range of applications. The performance of these codes has been evaluated on numerous problem instances of various size and characteristics. This work provides a comprehensive experimental survey of currently available minimum-cost flow solvers and thereby also gives guidelines for selecting a solution method for a certain application.

According to the presented results, we can conclude that implementations of the cost-scaling and primal network simplex algorithms are usually the most efficient and robust. In certain cases, if an optimal flow need not be split into many paths, however, the augmenting path algorithms are superior to these methods.

The primal network simplex code developed by the author (NS) turned out to be significantly faster, often by an order of magnitude, than the other implementations of this method: MCF, MCFSimplex, and the NETOPT component of CPLEX. Furthermore, NS is generally the most efficient among all solvers considered in this study on relatively small networks (up to many thousands of nodes). However, the cost-scaling algorithms tend to be superior to simplex-based methods on huge sparse networks due to their better asymptotic behavior in terms of the number of nodes. The cost-scaling code of the author (COS) performs similarly to or slightly slower than CS2, which is a highly efficient authoritative implementation of this algorithm. The LEDA library also implements this method, but it is slower and numerically less stable than COS and CS2. We also considered two other solvers, RelaxIV and PDNET, but they did not turn out to be robust. They tend to be orders of magnitude slower than cost-scaling and primal simplex codes, although RelaxIV is very fast on certain kinds of easy problem instances.

Our implementations are not standalone solvers, but they are included in a versatile C++ network optimization library called LEMON [73, 32]. This makes it easy to combine them with various other algorithms for solving complex network optimization problems. LEMON is an open-source library that can be used in both commercial and non-commercial software development under a permissive license. It is also involved in the COIN-OR project [24]. The author believes that this library with its great variety of efficient algorithms and data structures is a viable alternative to other graph libraries, such as the Boost Graph Library [18, 94] and LEDA [71, 77], as well as to the MCFClass project [9].



## Acknowledgments

The author would like to thank Zoltán Király for his essential contribution to this work, as well as Alpár Jüttner and Balázs Dezső for many useful suggestions related to the implementations. Special thanks are due to the developers of CS2, RelaxIV, MCF, and PDNET for making these codes available for academic use: Andrew V. Goldberg and Boris V. Cherkassky; Dimitri P. Bertsekas and Paul Tseng; Andreas Löbel; Luis F. Portugal, Mauricio G.C. Resende, Geraldo Veiga, João Patrício, and Joaquim J. Júdice; respectively. The author also thanks Antonio Frangioni, Claudio Gentile, and Alessandro Bertolini for developing the MCFClass project including the MCFSimplex solver and the C++ translation of the Fortran code of RelaxIV. Finally, the author is grateful to Antonio Frangioni for his help with the MCFClass project.

## References

- [1] R.K. Ahuja, A.V. Goldberg, J.B. Orlin, and R.E. Tarjan. Finding minimum-cost flows by double scaling. Technical Report Sloan W.P. No. 2047-88, Stanford University, Stanford, CA, 1988.
- [2] R.K. Ahuja, A.V. Goldberg, J.B. Orlin, and R.E. Tarjan. Finding minimum-cost flows by double scaling. *Math. Program.*, 53:243–266, 1992. A preliminary version appeared in [1].
- [3] R.K. Ahuja, T.L. Magnanti, and J.B. Orlin. *Network Flows: Theory, Algorithms, and Applications*. Prentice-Hall, Inc., Englewood Cliffs, NJ, 1993.
- [4] R.D. Armstrong and Z. Jin. A new strongly polynomial dual network simplex algorithm. *Math. Program.*, 78:131–148, 1997.
- [5] R.D. Armstrong, D. Klingman, and D. Whitman. Implementation and analysis of a variant of the dual method for the capacitated transshipment problem. *European J. Oper. Res.*, 4:403–420, 1980.
- [6] F. Barahona and É. Tardos. Note on Weintraub’s minimum-cost circulation algorithm. *SIAM J. Comput.*, 18:579–583, 1989.
- [7] R.S. Barr, F. Glover, and D. Klingman. The alternating basis algorithm for assignment problems. *Math. Program.*, 13:1–13, 1977.
- [8] R.S. Barr, F. Glover, and D. Klingman. Enhancements to spanning tree labelling procedures for network optimization. *INFOR*, 17:16–34, 1979.
- [9] A. Bertolini, A. Frangioni, and C. Gentile. The MCFClass Project. 2011.
- [10] D.P. Bertsekas. *Network Optimization Codes*. 1994.
- [11] D.P. Bertsekas and J. Eckstein. Dual coordinate step methods for linear network flow problems. *Math. Program.*, 42:203–243, 1988.

- 
- [12] D.P. Bertsekas and P. Tseng. The relax codes for linear minimum cost network flow problems. *Ann. Oper. Res.*, 13:125–190, 1988.
- [13] D.P. Bertsekas and P. Tseng. Relaxation methods for minimum cost ordinary and generalized network flow problems. *Oper. Res.*, 36:93–114, 1988.
- [14] D.P. Bertsekas and P. Tseng. RELAX-IV: A faster version of the relax code for solving minimum cost flow problems. Technical Report LIDS-P-2276, Dept. of Electrical Engineering and Computer Science, MIT, Cambridge, MA, 1994.
- [15] R.G. Bland, J. Cheriyan, D. Jensen, and L. Ladányi. An empirical study of min cost flow algorithms. In D.S. Johnson and C.C. McGeoch, editors, *Network Flows and Matching: First DIMACS Implementation Challenge*, DIMACS Series in Discrete Mathematics and Theoretical Computer Science, Vol. 12, pages 119–156, Providence, RI, 1993. American Mathematical Society.
- [16] R.G. Bland and D.L. Jensen. On the computational behavior of a polynomial-time network flow algorithm. Technical Report 661, School of Operations Research and Industrial Engineering, Cornell University, Ithaca, NY, 1985.
- [17] R.G. Bland and D.L. Jensen. On the computational behavior of a polynomial-time network flow algorithm. *Math. Program.*, 54:1–39, 1992. A preliminary version appeared in [16].
- [18] Boost C++ Libraries. <http://www.boost.org/>.
- [19] G. Bradley, G. Brown, and G. Graves. Design and implementation of large scale primal transshipment algorithms. *Manag. Sci.*, 24:1–34, 1977.
- [20] T. Brunsch, K. Cornelissen, B. Manthey, and H. Röglin. Smoothed analysis of the successive shortest path algorithm. In *Proc. 24th ACM-SIAM Symposium on Discrete Algorithms*, SODA '13, pages 1180–1189, Philadelphia, PA, 2013. SIAM.
- [21] U. Bünnagel, B. Korte, and J. Vygen. Efficient implementation of the Goldberg-Tarjan minimum-cost flow algorithm. *Opt. Methods and Software*, 10:157–174, 1998.
- [22] R.G. Busacker and P.J. Gowen. A procedure for determining a family of minimum-cost network flow patterns. Technical Report ORO-TP-15, Operations Research Office, The Johns Hopkins University, Bethesda, MD, 1960.
- [23] B.V. Cherkassky and A.V. Goldberg. On implementing the push–relabel method for the maximum flow problem. *Algorithmica*, 19:390–410, 1997.
- [24] COIN-OR – Computational Infrastructure for Operations Research, 2012.
- [25] IBM ILOG CPLEX Optimization Studio, V12.4, 2011.
- [26] CS2 Software, Version 4.6, 2009. IG Systems, Inc.

- 
- [27] W.H. Cunningham. A network simplex method. *Math. Program.*, 11:105–116, 1976.
- [28] G.B. Dantzig. Application of the simplex method to a transportation problem. In T.C. Koopmans, editor, *Activity analysis of production and allocation*, pages 359–373. John Wiley & Sons, Inc., New York, NY, 1951.
- [29] G.B. Dantzig. *Linear Programming and Extensions*. Princeton University Press, Princeton, NJ, 1963.
- [30] A. Dasdan. Experimental analysis of the fastest optimum cycle ratio and mean algorithms. *ACM Trans. Des. Autom. Electron. Syst.*, 9:385–418, 2004.
- [31] A. Dasdan and R.K. Gupta. Faster maximum and minimum mean cycle algorithms for system performance analysis. *IEEE Trans. Computer-Aided Design*, 17:889–899, 1998.
- [32] B. Dezső, A. Jüttner, and P. Kovács. LEMON – an open source C++ graph template library. *Electron. Notes Theor. Comput. Sci.*, 264:23–45, 2011.
- [33] DIMACS. Network Flows and Matching: First DIMACS Implementation Challenge. 1990-1991.
- [34] DIMACS. Shortest Paths: Ninth DIMACS Implementation Challenge. 2005-2006.
- [35] J. Edmonds and R.M. Karp. Theoretical improvements in algorithmic efficiency for network flow problems. In *Combinatorial Structures and Their Applications*, pages 93–96, New York, NY, 1970. Gordon and Breach.
- [36] J. Edmonds and R.M. Karp. Theoretical improvements in algorithmic efficiency for network flow problems. *J. ACM*, 19:248–264, 1972. A preliminary version appeared in [35].
- [37] EGRES – MTA-ELTE Egerváry Research Group on Combinatorial Optimization, 2012.
- [38] L.R. Ford and D.R. Fulkerson. Constructing maximal dynamic flows from static flows. *Oper. Res.*, 6:419–433, 1958.
- [39] L.R. Ford and D.R. Fulkerson. *Flows in Networks*. Princeton University Press, Princeton, NJ, 1962.
- [40] A. Frangioni and A. Manca. A computational study of cost reoptimization for min-cost flow problems. *INFORMS J. Comput.*, 18:61–70, 2006.
- [41] A. Frank. *Connections in Combinatorial Optimization*. Oxford Lecture Series in Mathematics and Its Applications, Vol. 38. Oxford University Press, Oxford, UK, 2011.

- 
- [42] D.R. Fulkerson. An out-of-kilter method for minimal-cost flow problems. *J. Soc. Indust. Appl. Math.*, 9:18–27, 1961.
- [43] H.N. Gabow and R.E. Tarjan. Faster scaling algorithms for network problems. *SIAM J. Comput.*, 18:1013–1036, 1989.
- [44] L. Georgiadis, A.V. Goldberg, R.E. Tarjan, and R.F. Werneck. An experimental study of minimum mean cycle algorithms. In I. Finocchi and J. Hershberger, editors, *Proc. 6th International Workshop on Algorithm Engineering and Experiments*, ALENEX 2009, pages 1–13, New York, NY, 2009. SIAM.
- [45] F. Glover, D. Karney, D. Klingman, and A. Napier. A computation study on start procedures, basis change criteria, and solution algorithms for transportation problems. *Manag. Sci.*, 20:793–813, 1974.
- [46] A.V. Goldberg. An efficient implementation of a scaling minimum-cost flow algorithm. *J. Algorithms*, 22:1–29, 1997.
- [47] A.V. Goldberg. The partial augment-relabel algorithm for the maximum flow problem. In *Proc. 16th Annual European Symposium on Algorithms*, ESA '08, pages 466–477, Heidelberg, Germany, 2008. Springer-Verlag.
- [48] A.V. Goldberg and M. Kharitonov. On implementing scaling push-relabel algorithms for the minimum-cost flow problem. In D.S. Johnson and C.C. McGeoch, editors, *Network Flows and Matching: First DIMACS Implementation Challenge*, DIMACS Series in Discrete Mathematics and Theoretical Computer Science, Vol. 12, pages 157–198, Providence, RI, 1993. American Mathematical Society.
- [49] A.V. Goldberg, É. Tardos, and R.E. Tarjan. Network flow algorithms. In B. Korte, L. Lovász, H.J. Prömel, and A. Schrijver, editors, *Paths, Flows and VLSI-Layout*, pages 101–164, Berlin, Germany, 1990. Springer-Verlag.
- [50] A.V. Goldberg and R.E. Tarjan. Solving minimum-cost flow problems by successive approximation. In *Proc. 19th ACM Symposium on Theory of Computing*, STOC '87, pages 7–18, New York, NY, 1987. ACM Press.
- [51] A.V. Goldberg and R.E. Tarjan. Finding minimum-cost circulations by canceling negative cycles. In *Proc. 20th ACM Symposium on Theory of Computing*, STOC '88, pages 388–397, New York, NY, 1988. ACM Press.
- [52] A.V. Goldberg and R.E. Tarjan. A new approach to the maximum-flow problem. *J. ACM*, 35:921–940, 1988.
- [53] A.V. Goldberg and R.E. Tarjan. Finding minimum-cost circulations by canceling negative cycles. *J. ACM*, 36:873–886, 1989. A preliminary version appeared in [51].

- 
- [54] A.V. Goldberg and R.E. Tarjan. Finding minimum-cost circulations by successive approximation. *Math. Oper. Res.*, 15:430–466, 1990. A preliminary version appeared in [50].
- [55] D. Goldfarb and J. Hao. Polynomial simplex algorithms for the minimum cost network flow problem. *Algorithmica*, 8:145–160, 1992.
- [56] M.D. Grigoriadis. An efficient implementation of the network simplex method. *Math. Program. Stud.*, 26:83–111, 1986.
- [57] M. Hartmann and J.B. Orlin. Finding minimum cost to time ratio cycles with small integral transit times. *Networks*, 23:567–574, 1993.
- [58] R.V. Helgason and J.L. Kennington. An efficient procedure for implementing a dual simplex network flow algorithm. *AIEE Transactions*, 9:63–68, 1977.
- [59] R.A. Howard. *Dynamic Programming and Markov Processes*. The MIT Press, Cambridge, MA, 1960.
- [60] M. Iri. A new method for solving transportation-network problems. *J. Oper. Res. Soc. Japan*, 3:27–87, 1960.
- [61] W.S. Jewell. Optimal flow through networks. Technical Report No. 8, Operations Research Center, MIT, Cambridge, MA, 1958.
- [62] E.L. Johnson. Networks and basic solutions. *Oper. Res.*, 14:619–623, 1966.
- [63] D.J. Kelly and G.M. O’Neill. The minimum cost flow problem and the network simplex solution method. Master’s thesis, University College, Dublin, Ireland, 1991.
- [64] J.L. Kennington and R.V. Helgason. *Algorithms for Network Programming*. John Wiley & Sons, Inc., New York, NY, 1980.
- [65] Z. Király and P. Kovács. Efficient implementations of minimum-cost flow algorithms. *Acta Universitatis Sapientiae, Informatica*, 4:67–118, 2012.
- [66] M. Klein. A primal method for minimal cost flows with applications to the assignment and transportation problems. *Manag. Sci.*, 14:205–220, 1967.
- [67] D. Klingman, A. Napier, and J. Stutz. NETGEN: A program for generating large scale capacitated assignment, transportation, and minimum cost flow network problems. *Manag. Sci.*, 20:814–821, 1974.
- [68] B. Korte and J. Vygen. *Combinatorial Optimization: Theory and Algorithms*. Algorithms and Combinatorics, Vol. 21. Springer-Verlag, Berlin, Germany, 5th edition, 2012.
- [69] Péter Kovács. Minimum-cost flow algorithms: an experimental evaluation. *Optimization Methods and Software*, 30(1):94–127, 2015.

- 
- [70] H.W. Kuhn. The Hungarian Method for the assignment problem. *Naval Res. Logist. Quart.*, 2:83–97, 1955.
- [71] The LEDA library, Version 5.0, 2004. Algorithmic Solutions Software GmbH.
- [72] Y. Lee. *Computational analysis of network optimization algorithms*. PhD thesis, Dept. of Civil and Environmental Engineering, MIT, Cambridge, MA, 1993.
- [73] LEMON – Library for Efficient Modeling and Optimization in Networks, 2012.
- [74] A. Löbel. Solving large-scale real-world minimum-cost flow problems by a network simplex method. Technical Report SC 96-7, Zuse Institute Berlin (ZIB), Berlin, Germany, 1996.
- [75] A. Löbel. MCF Version 1.3 – A network simplex implementation. 2004.
- [76] Max-flow problem instances in vision, 2008. Computer Vision Research Group, University of Western Ontario.
- [77] K. Mehlhorn and S. Näher. *LEDA: A platform for combinatorial and geometric computing*. Cambridge University Press, New York, NY, USA, 1999.
- [78] G.J. Minty. Monotone networks. *Proc. R. Soc. Lond., Series A*, 257:194–212, 1960.
- [79] J.M. Mulvey. Pivot strategies for primal-simplex network codes. *J. ACM*, 25:266–270, 1978.
- [80] J.B. Orlin. Genuinely polynomial simplex and non-simplex algorithms for the minimum cost flow problem. Technical Report 1615-84, Sloan School of Management, MIT, Cambridge, MA, 1984.
- [81] J.B. Orlin. A faster strongly polynomial minimum cost flow algorithm. In *Proc. 20th ACM Symposium on Theory of Computing, STOC '88*, pages 377–387, New York, NY, 1988. ACM Press.
- [82] J.B. Orlin. A faster strongly polynomial minimum cost flow algorithm. *Oper. Res.*, 41:338–350, 1993. A preliminary version appeared in [81].
- [83] J.B. Orlin. A polynomial time primal network simplex algorithm for minimum cost flows. *Math. Program.*, 78:109–129, 1997.
- [84] J.B. Orlin, S.A. Plotkin, and É. Tardos. Polynomial dual network simplex algorithms. *Math. Program.*, 60:255–276, 1993.
- [85] S.A. Plotkin and É. Tardos. Improved dual network simplex. In *Proc. 1st ACM-SIAM Symposium on Discrete Algorithms, SODA '90*, pages 367–376, Philadelphia, PA, 1990. SIAM.

- 
- [86] L.F. Portugal, M.G.C. Resende, G. Veiga, J. Patrício, and J.J. Júdice. Fortran subroutines for network flow optimization using an interior point algorithm. Technical Report TD-5X2SLN, AT&T Labs Research, 2004.
- [87] L.F. Portugal, M.G.C. Resende, G. Veiga, J. Patrício, and J.J. Júdice. PDNET – Software for network flow optimization using an interior point method. 2004.
- [88] L.F. Portugal, M.G.C. Resende, G. Veiga, J. Patrício, and J.J. Júdice. Fortran subroutines for network flow optimization using an interior point algorithm. *Pesquisa Operacional*, 28:243–261, 2008. A preliminary version appeared in [86].
- [89] M.G.C. Resende and P.M. Pardalos. Interior point algorithms for network flow problems. In J.E. Beasley, editor, *Advances in linear and integer programming*, pages 147–187, Oxford, UK, 1996. Oxford University Press.
- [90] M.G.C. Resende and G. Veiga. An efficient implementation of a network interior point method. In D.S. Johnson and C.C. McGeoch, editors, *Network Flows and Matching: First DIMACS Implementation Challenge*, DIMACS Series in Discrete Mathematics and Theoretical Computer Science, Vol. 12, pages 299–348, Providence, RI, 1993. American Mathematical Society.
- [91] M.G.C. Resende and G. Veiga. An annotated bibliography of network interior point methods. *Networks*, 42:114–121, 2003.
- [92] H. Röck. Scaling techniques for minimal cost network flows. In U. Pape, editor, *Discrete Structures and Algorithms*, pages 181–191, München, Germany, 1980. Carl Hanser.
- [93] A. Schrijver. *Combinatorial Optimization: Polyhedra and Efficiency*. Algorithms and Combinatorics, Vol. 24. Springer-Verlag, Berlin, Germany, 2003.
- [94] J.G. Siek, L.-Q. Lee, and A. Lumsdaine. *The Boost Graph Library: User Guide and Reference Manual*. Addison-Wesley, 2002.
- [95] A. Sifaleras. Minimum cost network flows: Problems, algorithms, and software. *Yugosl. J. Oper. Res.*, 23:3–17, 2013.
- [96] D.D. Sleator and R.E. Tarjan. A data structure for dynamic trees. *J. Comput. System Sci.*, 26:362–391, 1983.
- [97] P.T. Sokkalingam, R.K. Ahuja, and J.B. Orlin. New polynomial-time cycle-canceling algorithms for minimum-cost flows. *Networks*, 36:53–63, 2000.
- [98] V. Srinivasan and G.L. Thompson. Benefit-cost analysis of coding techniques for the primal transportation algorithm. *J. ACM*, 20:194–213, 1973.
- [99] É. Tardos. A strongly polynomial minimum cost circulation algorithm. *Combinatorica*, 5:247–255, 1985.

- 
- [100] R.E. Tarjan. Efficiency of the primal network simplex algorithm for the minimum-cost circulation problem. *Math. Oper. Res.*, 16:272–291, 1991.
- [101] R.E. Tarjan. Dynamic trees as search trees via Euler tours, applied to the network simplex algorithm. *Math. Program.*, 78:169–177, 1997.
- [102] N. Tomizawa. On some techniques useful for solution of transportation network problems. *Networks*, 1:173–194, 1971.
- [103] J. Vygen. On dual minimum cost flow algorithms. In *Proc. 32nd ACM Symposium on Theory of Computing*, STOC '00, pages 117–125, New York, NY, 2000. ACM Press.
- [104] J. Vygen. On dual minimum cost flow algorithms. *Math. Methods Oper. Res.*, 56:101–126, 2002. A preliminary version appeared in [103].
- [105] A. Weintraub. A primal algorithm to solve network flow problems with convex costs. *Manag. Sci.*, 21:87–97, 1974.
- [106] M.A. Yakovleva. A problem on minimum transportation cost. In V.S. Nemchinov, editor, *Applications of Mathematics in Economic Research*, pages 390–399, Izdat. Social'no-Ekon. Lit., Moscow, Russia, 1959.