

Randomized rounding algorithms for large scale unsplittable flow problems

François Lamothe · Emmanuel
Rachelson · Alain Haït · Cedric
Baudoin · Jean-Baptiste Dupe

Received: date / Accepted: date

Abstract Unsplittable flow problems cover a wide range of telecommunication and transportation problems and their efficient resolution is key to a number of applications. In this work, we study algorithms that can scale up to large graphs and important numbers of commodities. We present and analyze in detail a heuristic based on the linear relaxation of the problem and randomized rounding. We provide empirical evidence that this approach is competitive with state-of-the-art resolution methods either by its scaling performance or by the quality of its solutions. We provide a variation of the heuristic which has the same approximation factor as the state-of-the-art approximation algorithm. We also derive a tighter analysis for the approximation factor of both the variation and the state-of-the-art algorithm. We introduce a new objective function for the unsplittable flow problem and discuss its differences with the classical congestion objective function. Finally, we discuss the gap in practical performance and theoretical guarantees between all the aforementioned algorithms.

François Lamothe
ISAE-SUPAERO, Université de Toulouse, France
Corresponding author
ORCID : 0000-0002-1332-9206
E-mail: francois.lamothe@isae-supaeero.fr

Alain Haït
ISAE-SUPAERO, Université de Toulouse, France
ORCID : 0000-0001-5918-0753

Emmanuel Rachelson
ISAE-SUPAERO, Université de Toulouse, France
ORCID : 0000-0002-8559-1617

Cedric Baudoin
Thalès Alenia Space, Toulouse, France

Jean-Baptiste Dupe
Centre national d'études spatiales (CNES), Toulouse, France

Keywords Unsplittable flows · randomized rounding · heuristic · approximation algorithm

Declarations

This work was partially funded by Thales Alenia Space and made in collaboration with several of its members. This work was partially funded by the CNES. Several authors are academically related to ISAE-SUPAERO.

The authors declare that they have no conflict of interest. The datasets and the code used in the experimental section of this work are accessible at https://github.com/SuReLI/randomized_rounding_paper_code.

1 Introduction

The unsplittable flow problem is an extensively studied variant of the classical maximum flow problem. In this problem, one is given a directed or undirected graph, together with capacities on its arcs. A family of commodities, each composed of an origin, a destination, and a demand, is also provided. Each commodity has to route its demand from its origin to its destination through a unique path. The routing must ensure that capacities on the arcs are not exceeded by the flow of the commodities, or at least minimize the violation of the capacities.

This problem is NP-hard as it contains several NP-hard problems as sub-cases. When there are only two nodes linked by one arc, the problem is equivalent to the knapsack problem. When all demands and capacities are 1, the problem is equivalent to the edge-disjoint paths problem.

This problem has various applications especially in telecommunication networks (*e.g.* optical networks, telecommunication satellites), and goods transportation. In these applications, large-scale instances appear with up to 500 nodes, 2000 arcs, and 150 000 commodities. However, only a few methods in the literature can scale to such large instances, such as the approximation algorithm of Raghavan and Tompson (1987) and some meta-heuristics tuned to have small computing times. The algorithm presented by Raghavan and Tompson (1987) uses randomized rounding to compute a solution to the unsplittable flow problem. Even though this $O\left(\frac{\ln m}{\ln \ln m}\right)$ -approximation algorithm has theoretically the best approximation factor achievable, the solution it yields are often far from optimal in terms of solution quality.

That is why, in this work, we focus on an algorithm that can scale to large instances while giving good practical results. This algorithm is a heuristic extension of the randomized rounding algorithm of Raghavan and Tompson (1987). As such, it also uses randomized rounding on the linear relaxation of the unsplittable problem to create an unsplittable solution. This algorithm alternates randomized rounding steps and resolutions of the linear relaxation and will thus be called, in this work, the Sequential Randomized Rounding algorithm (SRR). This heuristic is also an extension of the algorithm proposed

by Coudert and Rivano (2002) for which no complete proof of the approximation factor was given. Compared to the algorithm of Coudert and Rivano (2002), the SRR heuristic offers more flexibility on the number of linear relaxation resolutions and more importantly, takes advantage of the fact that commodities might have different demand levels. We also describe a variation of the SRR heuristic for which we prove the same approximation guarantees as the algorithm of Raghavan and Tompson (1987). This variation will be called the Constrained Sequential Randomized Rounding algorithm (CSR). Moreover, we tighten the analysis of both approximation algorithms and prove that they achieve a $O\left(\frac{\gamma \ln m}{\ln(\gamma \ln m)}\right)$ -approximation factor where γ is a parameter that is small when the commodities demands are small compared to the capacities of the arcs. Finally, we experimentally show that the SRR algorithm scales to large instances. Furthermore, its practical results on large datasets outperform other tested methods.

This paper is structured as follows. In Section 2 we describe the notations used together with several Mixed Integer Linear Programs (MILP) for the unsplittable flow problem. Related work is presented in Section 3. Section 4 presents the SRR heuristic and its complexity analysis. Section 5 describes the CSR algorithm and provides the analysis leading to the $O\left(\frac{\gamma \ln m}{\ln(\gamma \ln m)}\right)$ -approximation factor. In Section 6, we provide experimental results that compare the empirical quality of the various algorithms presented. In Section 7, we discuss different properties of the SRR heuristic and the CSR algorithm. Finally, we conclude and give perspectives in Section 8.

2 The unsplittable flow problem

Throughout this paper, we will use the following notations:

- $G = (V, E)$ is a directed or undirected graph, with V the set of nodes and E the set of arcs
- $L = (o_k, d_k, D_k)_{k \in K}$ is a set of commodities defined by their origin, destination and demand.
- $(c_e)_{e \in E}$ are capacities on the arcs

We also use the Kronecker notation, δ_x^y equals 1 if $x = y$ and 0 otherwise. The sets of arcs incoming and outgoing of node v will be noted $E^-(v)$ and $E^+(v)$ respectively. Finally, $c_{min} = \min_{e \in E} c_e$ and $D_{max} = \max_{k \in K} D_k$ are the smallest capacity and the largest demand.

2.1 Objective functions

Four objective functions for the unsplittable flow problems can be found in the literature: maximizing the served demand (Kolman, 2003), minimizing the cost (Barnhart et al., 2000), minimizing the congestion (Martens and Skutella,

2006), minimizing the number of necessary routing rounds (Aumann and Rabani, 1995). In this work, we focus on minimizing the congestion.

The congestion of a graph is the smallest number Δ by which it is necessary to multiply all the capacities in order to route all the commodities (Martens and Skutella, 2006). The congestion of an arc is the ratio of the flow on this arc to its capacity. This metric emphasizes low capacity arcs. Besides, minimizing the congestion puts no restrictions on the flow going through the arcs that do not have a maximal congestion. In particular, it induces no incentive to minimize the congestion on those arcs. This becomes problematic when an arc is largely more congested in every solution than any other arc because it lifts all restrictions for the other arcs. To prevent this, we introduce a new criterion to minimize the violation of the capacities of the arcs. We use the term *overflow* to describe the quantity of flow Δ_e that exceeds the capacity of an arc e . The overflow of an arc is always non-negative. Our new criterion is to minimize the sum of the overflows $\sum_{e \in E} \Delta_e$. Note that the congestion Δ is not the maximum overflow over all the arcs.

In the following sections, we present two equivalent Mixed Integer Linear Programs (MILP) for the unsplittable flow problem.

2.2 Arc-node formulation

The arc-node formulation is compact as it has a polynomial number of variables and constraints in the number of commodities, nodes, and arcs. It can thus be solved directly in a MILP solver for small instances. This formulation is characterized by the flow conservation constraints which ensure the structural property of flows. The objective function presented is the overflow sum. The meaning of the variables in this formulation is the following:

- f_{ek} indicates whether commodity k pushes flow on arc e ,
- Δ_e represents the overflow on arc e .

The unsplittable flow problem is then:

$$\min_{f_{ek}, \Delta_e} \sum_{e \in E} \Delta_e \quad (1a)$$

such that

$$\sum_{e \in E^+(v)} f_{ek} - \sum_{e \in E^-(v)} f_{ek} = \delta_v^{o_k} - \delta_v^{d_k} \quad \forall k \in K, \forall v \in V, \quad (1b)$$

$$\sum_{k \in K} f_{ek} D_k \leq c_e + \Delta_e \quad \forall e \in E, \quad (1c)$$

$$f_{ek} \in \{0, 1\}, \Delta_e \in \mathbb{R}^+ \quad \forall k \in K, \forall e \in E. \quad (1d)$$

Equation (1b) corresponds to the flow conservation constraints. It ensures that, for each commodity and every node except the origin and the destination of the commodity, the same amount of flow of the commodity goes in and out

of the node. Equation (1c) corresponds to the capacity constraints. It ensures that the capacity of an arc is respected or that the overflow is recorded in Δ_e . The fact that $f_{ek} \in \{0, 1\}$ ensures that the flow is unsplittable.

One can create an arc-node congestion formulation by replacing $c_e + \Delta_e$ by $c_e \Delta$ in Equation (1c) and minimizing over Δ instead of $\sum_e \Delta_e$. Note that the Δ variable is common to all arcs while there was one variable Δ_e for each arc.

2.3 Path formulation

In the path formulation, the flow conservation constraints are unnecessary. The variables represent paths so these constraints are always verified. However, there is an exponential number of variables (in the number of commodities, nodes, and arcs) so the formulation must be solved through a particular MILP technique named Branch and Price (Barnhart et al., 2000). The objective function presented is the overflow sum. The meaning of the variables in this formulation is the following:

- x_{pk} indicates whether commodity k uses path p to push its flow,
- Δ_e represents the overflow on arc e .

The unsplittable flow problem is then:

$$\min_{x_{pk}, \Delta_e} \sum_{e \in E} \Delta_e \quad (2a)$$

such that

$$\sum_{p \in P_k} x_{pk} = 1 \quad \forall k \in K, \quad (2b)$$

$$\sum_{k \in K} \sum_{p \in P_k | e \in p} x_{pk} D_k \leq c_e + \Delta_e \quad \forall e \in E, \quad (2c)$$

$$x_{pk} \in \{0, 1\}, \Delta_e \in \mathbb{R}^+ \quad \forall p \in \bigcup_k P_k, \forall k \in K, \forall e \in E. \quad (2d)$$

Here, P_k denotes the set of all paths usable by commodity k . Equation (2b) ensures that exactly one path is chosen for each commodity. Equation (2c) corresponds to the capacity constraints. It ensures that the capacity of an arc is respected or that the overflow is recorded in Δ_e . The fact that $x_{pk} \in \{0, 1\}$ ensures that the flow is unsplittable.

3 Related work

In this section, we review important solution approaches to the unsplittable flow problem present in the literature. These works are grouped into three subsections: exact methods, approximation algorithms, and meta-heuristics. A

fourth sub-section is dedicated to the linear relaxation of the unsplittable flow problem (the multi-commodity flow problem) whose resolution is paramount to the resolution of the unsplittable flow problem.

3.1 Exact methods

Barnhart et al. (2000) presented a Branch and Price and Cut procedure applied to a path formulation with a cost minimization objective function. Most subsequent works use this baseline as a comparison. A major contribution of their work is their branching strategy. Unlike straightforward branching strategies for this problem, theirs keeps intact the structure of the pricing problem throughout the branching procedure. For a commodity in a non-integer solution, the divergence node is the first node where the flow of the commodity splits. The outgoing arcs of the divergence node are divided into two disjoint subsets E_1 and E_2 . Each set must contain at least one of the arcs used by the commodity. The branching rule is: either forbid the use of E_1 or forbid the use of E_2 . In both cases, the previous non-integer solution is cut from the problem and forbidding sets of arcs keeps the structure of the pricing problem intact. They also included cuts to strengthen the linear relaxation. These cuts are lifted cover inequalities of the capacity constraints.

Park et al. (2003) mixed the path formulation and a knapsack formulation (which is not presented in this work) to derive a new linear formulation of the problem. The linear relaxation of this formulation yields a stronger lower bound, which in turn decreases the time needed to complete the branching procedure. They compared different branching rules and report that the one of Barnhart et al. (2000) produces much better results. They present computational results only for this rule.

Belaidouni and Ben-Ameur (2007) presented a cutting plane method based on super-additive functions to get strong cuts for their Branch and Price method. It appears on small instances that the addition of their cuts derives integer solutions without using a Branch and Bound procedure. Results are compared with those of Barnhart et al. (2000) and large improvements are reported.

Overall, the best results can be found in the articles of Belaidouni and Ben-Ameur (2007) and Park et al. (2003). Belaidouni and Ben-Ameur (2007) compared their results with those of Barnhart et al. (2000) and solved all their instances (at most 30 nodes, 60 arcs, 100 commodities) in less than 10 seconds. Park et al. (2003) did not compare their results with previous works but solved instances of the same magnitude (at most 30 nodes, 80 arcs, 100 commodities) in less than 15 seconds. Note that results were only given for small instances in these approaches.

Other earlier works have been reported in (Parker and Ryan, 1993; Alvelos and De Carvalho, 2003; Park et al., 1996).

3.2 Approximation algorithms and heuristics

As the unsplittable flow problem is NP-hard, a lot of attention has been given to approximation algorithms and heuristics. In particular, the maximum served demand objective has been extensively studied. We refer to the Handbook of approximation algorithms (Group and Gonzalez, 2020) for a detailed survey on approximation algorithms in the context of unsplittable flows. We recall here some works related to the minimum congestion objective.

Approximation algorithms are given with a factor of approximation λ . Let Γ be the objective function of the minimization problem at hand. Solutions generated by a λ -approximation algorithm verify the following inequality:

$$\Gamma(S^*) \leq \Gamma(S) \leq \lambda \Gamma(S^*),$$

where $\Gamma(S)$ and $\Gamma(S^*)$ are respectively the value of the produced solution and the value of the optimal solution. This guarantees that the ratio between the value of the produced solution and the value of the optimal solution is not too high. When considering approximation algorithms, more attention must be paid to the objective function being optimized. The literature uses four different objective functions: maximizing the served demand, minimizing the congestion, minimizing the number of rounds, and minimizing the cost. Approximation hardness results demonstrate that none of these objective functions admits a constant-factor approximation algorithm (Group and Gonzalez, 2020).

In the congestion minimization context, we seek the smallest number by which it is necessary to multiply all the capacities in order to fit all the commodities. The best-known approximation algorithm for congestion is a randomized rounding method introduced by Raghavan and Tompson (1987) which we shall call the RR algorithm in this work. The method proceeds in two steps. First, a solution of the linear relaxation of the problem is computed. Each commodity is allowed to use multiple paths in this solution. The proportion of flow for each commodity on each path is $((x_{pk})_{p \in P_k})_{k \in K}$. Then a path is selected for each commodity. Path p is selected in P_k with probability x_{pk} . Each commodity is then assigned to the selected path to create an unsplittable solution. This procedure produces, with arbitrarily high probability, an unsplittable solution whose congestion is $O\left(\frac{\ln |E|}{\ln \ln |E|}\right)$ larger than the one of the fractional solution. Their algorithm is thus a $O\left(\frac{\ln |E|}{\ln \ln |E|}\right)$ -approximation algorithm that works for directed and undirected graphs. The randomized rounding process can be derandomized using the method of conditional probabilities (Raghavan, 1988). Chuzhoy et al. (2007) showed a tight $\Omega\left(\frac{\ln |V|}{\ln \ln |V|}\right)$ bound on directed graphs, assuming $NP \not\subseteq BPTIME(|V|^{O(\ln \ln |V|)})$. Andrews et al. (2010) showed that minimizing congestion for an unsplittable flow in an undirected graph is hard to approximate within $\Omega(\ln \ln |V| / \ln \ln \ln |V|)$, assuming $NP \not\subseteq ZPTIME(|V|^{polylog(|V|)})$. A $(|K| + 2)$ -approximation algorithm is presented in (Asano, 2000). The reported results show that in practice this algorithm gives results comparable to classical randomized rounding.

In the context of the maximum demand objective function, a parameter $\frac{D_{max}}{c_{min}}$ is introduced. This parameter plays an important role for this objective function. Indeed, when this parameter is upper-bounded by a small constant, several works reported stronger approximation results for their algorithms (Chakrabarti et al., 2007; Shepherd and Vetta, 2015; Azar and Regev, 2006). However, we did not find similar results for the congestion objective function.

A few heuristics have been proposed in previous works. They are either greedy or Linear programming (LP) based heuristics. Coudert and Rivano (2002) introduced an algorithm very similar to the SRR algorithm presented in Section 4, without proving it is an approximation algorithm. Asano (2000) as well as Wang and Wang (1999) proposed greedy algorithms and LP-based algorithms. Reported results show that, except in specific cases, the greedy approaches are usually not competitive with LP-based methods. On the other hand, LP-based heuristics yield results that are similar to the randomized rounding algorithm of Raghavan and Tompson (1987).

3.3 Meta-heuristics

As introduced above, it is NP-hard to find an optimal solution or even to give a constant-factor approximation to the unsplittable flow problem. Thus, the literature investigated various randomized search procedures such as genetic algorithms (Cox, 1991; Masri et al., 2019), tabu search (Anderson et al., 1993; Laguna and Glover, 1993; Xu et al., 1997), local search and GRASP (Santos et al., 2013b,a, 2010; Alvelos and Valério de Carvalho, 2007; Masri et al., 2015, 2019) or ant colony optimization (Li et al., 2010; Masri et al., 2011). One of the major difficulties encountered when solving the unsplittable flow problem with a meta-heuristic is to efficiently create useful paths for the commodities.

Early approaches such as (Cox, 1991; Anderson et al., 1993) encode solutions as permutations of the commodities. The space of permutations is the one searched by the meta-heuristic. The following function is used to create a solution from a permutation and evaluate it. The function goes through the permutation, examining each commodity. The commodity is then allocated to the shortest path where there is still enough capacity to fit the commodity. Once a path is assigned to every commodity, the objective function can be computed.

In (Laguna and Glover, 1993) and (Masri et al., 2015) the k shortest paths are pre-computed for each commodity using the algorithm of Yen (1971). The search space of their meta-heuristics is restrained to the space of solutions using only those paths.

A different idea used in (Santos et al., 2013b,a, 2010; Alvelos and Valério de Carvalho, 2007) is to consider paths extracted from the linear relaxation of the problem. The linear relaxation is solved with a column generation algorithm applied to the path formulation. During the column generation, a set of paths \hat{P}_k is generated for each commodity. A meta-heuristic such as a multi-start

local search is then used to explore the solutions where only paths from \hat{P}_k are used. In (Santos et al., 2013b), after the first linear relaxation is solved, perturbed linear models are solved to create new useful columns and extend the solution space of the meta-heuristic.

Ant colony optimization is also a means to navigate the large solution space of the possible paths and is used in (Li et al., 2010; Masri et al., 2011). In an ant colony optimization approach, at each iteration, each commodity creates a path by taking into account several metrics: path length, path load, and pheromones. Each arc of the graph has a pheromone level for each commodity and the higher the pheromone level the higher the probability of the arc to belong to the path generated. Pheromones are updated through two means. First, the best solutions add pheromones to the arc they use. Second, pheromones decay so that their level does not become excessive thus facilitating the exploration of the solution space.

We refer to the work of Li et al. (2010) and Santos et al. (2013a) for the best performing meta-heuristics. Li et al. (2010) compared their results with the solver CPLEX and were able to solve instances with up to 60 nodes, 400 arcs, and 3500 commodities to optimality in less than 900 seconds. Santos et al. (2013a) show that all their instances (26 nodes, 80 arcs, 500 commodities) are solved in less than 180 seconds with values close to the linear relaxation lower bound.

3.4 Linear multi-commodity flow problem

The multi-commodity flow problem is the linear relaxation of the unsplittable flow problem. The value of its optimal solution is a lower bound for the binary problem and this linear relaxation is used in several exact and approximate methods. As a special case of linear programming, this problem is solvable in polynomial time.

A lot of effort has been invested in exact methods for this problem. Even if a commercial solver can solve the node-arc formulation, this method may take a prohibitive time in large instances. An alternative solution is to use a Lagrangian relaxation of the capacity constraints to decompose the problem into easier sub-problems as in (Retvdri et al., 2004). As reported by Dai et al. (2017), Lagrangian relaxation shows lesser performances in most instances than the competitor method, applying a column generation algorithm to the path formulation. Lagrangian relaxation seems to be the best choice when the number of commodities is very large because its computing time scales only linearly with this parameter. In most other cases, column generation seems to be the best solution.

Several works contributed to increase the performance of column generation algorithms for this problem. First, a primal-dual column generation is presented in (Gondzio et al., 2013; Gondzio and González-Brevis, 2015; Gondzio et al., 2016). An interior-point algorithm is used to solve the master problem and obtain sub-optimal but well-centered solutions. These well-centered solu-

tions are used to compute new columns in the sub-problems which stabilizes the column generation process and reduces the number of iterations needed to achieve convergence. Another approach, which could be combined with the previous one is the use of aggregated variables presented by Bauguion et al. (2013, 2015). In this method, variables do not represent paths but aggregated paths such as trees or more complex structures. The authors report that the sub-problems associated with aggregated variables can be solved efficiently. Aggregated variables reduce the size of the master problem during the algorithm but might induce a larger number of iterations, thus aggregation must be carefully done. Another method presented in (Babonneau et al., 2006) consists of a specialized interior-point method to solve the multi-commodity flow problem. This method has been improved by Castro and Cuesta (2012). Other contributions to linear programming methods can be found in (Moradi et al., 2015; Dai et al., 2016a,b)

For large instances, linear programming methods may take a lot of computing time before finding the optimal solution. Thus the literature focused on combinatorial approximation algorithms and in particular on Fully Polynomial-Time Approximation Schemes (FPTAS). The best results were obtained through the use of exponential length functions. This idea was first introduced by Shahrokhi and Matula (1990). In their algorithm, a length exponential in the passing flow is assigned to each arc. The flow is iteratively augmented on the shortest path connecting any of the source-sink pairs. Their algorithm was improved by Fleischer (2000) who showed that only the computation of an ϵ -shortest path was needed. The algorithm of Fleischer (2000) is the fastest FPTAS in practice while not being the one with the smallest complexity. Indeed, Madry (2010) presented an algorithm with a smaller complexity but Emanuelsson (2016) showed in his work that the algorithm of Fleischer (2000) is faster for instances having less than 100 million arcs. For a detailed survey on the multi-commodity flow problem methods previous to 2005, see the article of Wang (2018).

4 The Sequential Randomized Rounding heuristic

The Sequential Randomized Rounding algorithm (SRR) is a polynomial greedy heuristic for the unsplittable flow problem. This algorithm is similar to the one proposed by Coudert and Rivano (2002) for the light-path assignment problem. We add some features that are specific to our problem such as the sorting of the commodities by decreasing demand. The SRR algorithm also gives the possibility to compute less often the linear relaxation of the problem to reduce the overall computing time of the algorithm. Compared to the approximation algorithm of Raghavan (1988), the SRR algorithm has larger running times, no performance guarantees but returns higher quality solutions on the tested instances. As a heuristic, the SRR algorithm has a shorter computing time than exact solutions and meta-heuristics, especially for large instances. A variation of the SRR heuristic with approximation guarantees is presented in Section

5 together with an analysis of these approximation guarantees. Finally, the SRR algorithm is further discussed in Section 7 to explain its behavior on the tested instances.

4.1 Presentation of the algorithm

The SRR algorithm, presented in Algorithm 1 alternates between two different steps: solve the linear relaxation of the problem and fix some commodities to a unique path. In our case, the linear relaxation is a multi-commodity flow problem. Even though for the sake of clarity we use the notations of the path formulation in the following, in the experimentations, an arc-node formulation paired with the commercial solver (Gurobi Optimization, 2020) is used to solve the linear relaxation. More efficient specialized solvers or approximation algorithms for the multi-commodity flow problem can be found in the literature (see Section 3.4). Solving the linear relaxation provides a distribution of flow among the paths for each commodity: $((x_{pk})_{p \in P_k})_{k \in K}$. After solving the linear relaxation, a path is selected for some commodities. These commodities will be forced to use only these paths for the rest of the algorithm. The path selected for each commodity is chosen through the same randomized rounding procedure introduced by Raghavan and Tompson (1987): for commodity k , path p is selected with probability x_{pk} . The probability that commodity k uses arc e is $f_{ek} = \sum_{p \in P_k | e \in p} x_{pk}$. When solving the next linear relaxations, the fixed commodities will also be forced to only use their single allowed path.

The major difference with the RR algorithm of Raghavan and Tompson (1987) is that, in the SRR heuristic, the linear relaxation is actualized several times during the randomized rounding process. More precisely, after deciding to fix some commodities to a single path, the linear relaxation is solved again with the added constraints that the fixed commodities must use their affected path. To decide when the linear relaxation is actualized, the following procedure is used. In the solution of the linear relaxation, some commodities use multiple paths. After θ of these commodities are fixed to a single path, the linear relaxation is actualized. Choosing the threshold θ trades off between computation time and solution quality. If the linear relaxation is actualized often (low threshold), fixing decisions take into account most of the previous fixing decisions but the computation time is high. If the linear relaxation is never actualized, branching decisions do not take into account previous branching decisions but the computation time is low. The threshold value is fixed to $\frac{|V|}{4}$ in our experiments. A sensitivity analysis of this parameter is given in Section 6.

Another difference with the RR algorithm is that solutions are created using of the overflow sum objective function presented in Section 2.1 instead of the classical congestion objective function. As will be explained in Section 7.1, when using the overflow sum objective function, the SRR heuristic returns solutions with a lower overflow but also a lower congestion.

Compared to the algorithm proposed in (Coudert and Rivano, 2002), the SRR heuristic offers more flexibility on the number of actualizations of the linear relaxation thanks to parameter θ . Moreover, the main difference is that, because the unsplittable flow problem in (Coudert and Rivano, 2002) arises from a Light-Path Assignment problem, all commodities have a demand of one and thus the commodities have their path chosen in no particular order. In the SRR heuristic, paths are assigned to the commodities in decreasing order of commodity's demand, *i.e.* the commodities with larger demands have their paths chosen first. This order is classically used in bin packing heuristics such as the Next Fit Decreasing heuristic (Csirik et al., 1986). The rationale behind this ordering is to allocate commodities with a large demand first, while a large amount of capacity is left in the arcs. Commodities with smaller demands are then used to fill the remaining gaps. In Section 6, it is shown that this rounding order of the variables has a large impact on the quality of the solution returned by the heuristic.

Algorithm 1 The SRR heuristic

Require: $G = (V, E, c)$ a capacitated graph, $L = (o_k, d_k, D_k)_{k \in K}$ a list of commodities

- 1: Sort the commodities by decreasing demand
- 2: $K_{fixed} = \emptyset$ \triangleright K_{fixed} is the set of indices of commodities fixed to a single path
- 3: **for** each commodity k^* in decreasing demand order **do**
- 4: **if** an actualization is needed **then**
- 5: $((x_{pk})_{p \in P_k})_{k \in K} = \text{Solve_Linear_Relaxation}(G, L, K_{fixed}, (p_k)_{k \in K_{fixed}})$
- 6: Draw a path p^* from P_{k^*} with probability $x_{p^*k^*}$
- 7: Add index k^* to K_{fixed} .
- 8: $p_{k^*} = p^*$
- 9: **return** $(p_k)_{k \in K}$

4.2 Complexity analysis

There are $|K|$ iterations and during each of them the algorithm might do the following:

- Solve the linear relaxation: $O(LR(|V|, |E|, |K|))$ operations where $LR(|V|, |E|, |K|)$ is the complexity of the linear relaxation resolution (line 5).
- Select p_* : as the flow of each commodity can be decomposed in at most $|E|$ paths, at most $|E|$ of the variables x_{pk} have a non-zero value (Ford Jr, 1956). Choosing one of these variables requires $O(|E|)$ operations (line 6).

Additionally, sorting the commodities requires $O(|K| \log(|K|))$ operations (line 1). The total time complexity is $O(|K|(\log |K| + |E| + LR(|V|, |E|, |K|)))$.

4.3 Grouping commodities by origin

Grouping commodities by origin is the process of considering a set of commodities with the same origin as one single commodity. This new commodity has a demand equal to the sum of the original demands. To ensure that the flow goes to the right destinations, a super-destination node is created and connected to each original destination with an arc of capacity equal to the original demand of the commodity. When solving the *linear relaxation* with grouped commodities, the same solution is computed at the condition that all commodities of a group originate from the same node. Grouping commodities can greatly reduce the computing time of a linear solution computed with an LP solver when the number of different origins is much smaller than the number of commodities. In our test cases, inspired by practical telecommunication instances, commodities are emitted from a small number of different origins thus it is efficient to group the commodities by origin. However, solutions produced when commodities are grouped yield a little less information. It is necessary to compute exactly what paths each commodity uses. This can be done quickly in $O(|V|(|E| + |K|))$ operations with a flow decomposition algorithm (Ford Jr, 1956).

5 A variation of the heuristic with approximation guarantees

In this section, we present the Constrained Sequential Randomized Rounding algorithm (CSRR). It is a variation of the SRR heuristic for which we prove approximation guarantees similar to the one of the RR algorithm of Raghavan and Tompson (1987). Approximation results are obtained considering the classical congestion objective function. Indeed, for the overflow sum objective, the value of the optimal solution may be zero. This happens when all commodities fit in the capacities. In this case, any approximation algorithm must find the optimal solution. Thus the overflow sum objective function is not suited for approximation proofs.

The RR algorithm is able to yield a solution satisfying a $O\left(\frac{\ln(|E|\epsilon^{-1})}{\ln(\ln(|E|\epsilon^{-1}))}\right)$ -approximation factor with probability $1 - \epsilon$. We extend and tighten the analysis of randomized rounding algorithms by giving a new $O\left(\frac{\gamma \ln(|E|\epsilon^{-1})}{\ln(\gamma \ln(|E|\epsilon^{-1}))}\right)$ factor for the CSRR algorithm and for the RR algorithm of Raghavan and Tompson (1987). In this new factor, γ is the granularity parameter of the instance which is equal to $\frac{D_{max}}{c_{min}\Delta^*}$ where Δ^* is the optimal congestion of the linear relaxation. This parameter is small when the commodities are smaller compared to $c_{min}\Delta^*$. The parameter γ can be related to the parameter $\frac{D_{max}}{c_{min}}$ introduced by Chakrabarti et al. (2007); Shepherd and Vetta (2015); Azar and Regev (2006) to tighten their approximation analysis in the case of the maximum demand objective. The number γ is a decisive parameter of unsplittable flow instances. Indeed, it remains constant when the capacities or the demands are scaled uniformly.

To prove an approximation factor for a randomized rounding where the linear relaxation is actualized (as in the line 5 of Algorithm 1), it appeared necessary to add a constraint to the linear relaxation. Thus the CSRR algorithm is the same as the SRR algorithm with the following additional constraint in the linear relaxation:

$$\sum_{k \in K \setminus K_{fixed}} f_{ek} D_k \leq c_e \Delta^* - \sum_{k \in K_{fixed}} \hat{f}_{ek} D_k, \quad \forall e \in E, \quad (3)$$

where Δ^* is the optimal congestion of the first linear relaxation; K_{fixed} is the set of commodities that were fixed to their respective single paths before the current linear relaxation resolution; f_{ek} are the variables used to optimize the flow of the unfixed variables. The commodities in K_{fixed} sent their flow on several paths in the linear relaxation before they were fixed to one path; \hat{f}_{ek} is the corresponding fractional flow on each arc for these commodities. Note that for the first computation of the linear relaxation, this constraint has no impact on the final solution and can be removed. Thus, this constraint disappears in the algorithm of Raghavan and Tompson (1987), since there is only one resolution of the linear relaxation.

For the sake of clarity, we derive the present proof in the case where the CSRR algorithm makes only a single actualization of the linear relaxation's solution which occurs just after the first rounding step. Extension to the case of several actualizations performed at any time can be done by induction.

In the following, let the discrete random variables F_{ek} indicate the flow of commodity k on arc e in the solution returned by the CSRR algorithm. The variables F_{ek} take the value D_k with probability $f_{ek} = \sum_{p \in P_k | e \in p} x_{pk}$ and 0 otherwise. Thus, their expectation $\mathbb{E}[F_{ek}] = f_{ek} D_k = D_k \sum_{p \in P_k | e \in p} x_{pk}$ is also the flow of commodity k on arc e in the solution of the linear relaxation. Let k_1 be the index of the commodity of largest demand (thus the first one to be fixed in Algorithm 1) and let $F_e = \sum_{k \in K \setminus \{k_1\}} F_{ek}$. Once conditioned by F_{ek_1} , the random variables F_{ek} ($k \neq k_1$) are independent of each other because the linear relaxation is not actualized between their rounding step. However, the random variables F_{ek} ($k \neq k_1$) are not independent of F_{ek_1} ; in particular, we have $\mathbb{E}[F_{ek} | F_{ek_1}] \neq \mathbb{E}[F_{ek}]$. Indeed the realization of F_{ek_1} in the unique rounding step conditions the resolution of the subsequent linear relaxation. Thus, it conditions the values f_{ek} which parametrize the distribution of the random variables F_{ek} . To recall this dependency, we write $f_{ek}(F_{ek_1})$ the fractional flow of commodity k on arc e . Note that constraint (3) added in the CSRR algorithm can be re-written in terms of random variables:

$$\mathbb{E}[F_e | F_{ek_1}] \leq c_e \Delta^* - \mathbb{E}[F_{ek_1}].$$

Recall that in the congestion formulation, the objective function aims at minimizing the minimum multiplicative factor on all arc capacities needed to fit the commodities. We note C^* the optimal congestion for the considered unsplittable flow instance. As introduced above, $F_{ek_1} + F_e$ is the flow on arc e in the solution returned by the CSRR algorithm. Thus, proving a probabilistic

$(1 + \alpha)$ -approximation for this algorithm bound boils down to proving that for all arcs, $F_{ek_1} + F_e$ remains below $(1 + \alpha)c_e C^*$ with high probability. Formally, for a small ϵ :

$$\mathbb{P}(\forall e \in E, F_{ek_1} + F_e \leq (1 + \alpha)c_e C^*) \geq 1 - \epsilon.$$

Conversely, this is equivalent to proving that, with at most probability ϵ , there exists an arc e where the congestion exceeds $(1 + \alpha)C^*$:

$$\mathbb{P}(\exists e \in E, F_{ek_1} + F_e \geq (1 + \alpha)c_e C^*) \leq \epsilon.$$

To that end, we prove in Theorem 1 that for every arc e :

$$\mathbb{P}(F_{ek_1} + F_e \geq (1 + \alpha)c_e \Delta^*) \leq \frac{\epsilon}{|E|}.$$

Indeed, in this case, as Δ^* is a lower bound on C^* , we have:

$$\begin{aligned} \mathbb{P}(\exists e \in E, F_{ek_1} + F_e \geq (1 + \alpha)c_e C^*) &= \mathbb{P}\left(\bigvee_{e \in E} F_{ek_1} + F_e \geq (1 + \alpha)c_e C^*\right) \\ &\leq \mathbb{P}\left(\bigvee_{e \in E} F_{ek_1} + F_e \geq (1 + \alpha)c_e \Delta^*\right) \\ &\leq \sum_{e \in E} \mathbb{P}(F_{ek_1} + F_e \geq (1 + \alpha)c_e \Delta^*) \\ &\leq \sum_{e \in E} \frac{\epsilon}{|E|} \\ &= \epsilon \end{aligned}$$

To ensure that $\mathbb{P}(F_{ek_1} + F_e \geq (1 + \alpha)c_e \Delta^*) \leq \frac{\epsilon}{|E|}$, we first prove through Lemma 2 that the probability $\mathbb{P}(F_{ek_1} + F_e \geq (1 + \alpha)c_e \Delta^*)$ is upper bounded by a quantity $g_e(\alpha)$. Lemma 2 is proved by using the Markov inequality and the probabilistic translation of constraint (3). Proving Lemma 2 requires a preliminary result introduced in Lemma 1. Finally, the proof is completed by showing that there exists a value for α satisfying $1 + \alpha = O\left(\frac{\gamma \ln(|E|\epsilon^{-1})}{\ln(\gamma \ln(|E|\epsilon^{-1}))}\right)$ and for every arc $g_e(\alpha) \leq \frac{\epsilon}{|E|}$.

Without loss of generality and to remove D_{max} from the proof, we assume the considered instances are scaled so that $D_{max} = 1$ and thus $\gamma = (c_{min} \Delta^*)^{-1}$. We now present the two Lemmas together with their proof.

Lemma 1 For any positive scalar α , $\mathbb{E}[(1 + \alpha)^{F_e} | F_{ek_1}] \leq e^{\alpha \mathbb{E}[F_e | F_{ek_1}]}$.

Proof

$$\begin{aligned}
\mathbb{E}[(1 + \alpha)^{F_e} | F_{ek_1}] &= \mathbb{E} \left[\prod_{k \in K \setminus \{k_1\}} (1 + \alpha)^{F_{ek}} | F_{ek_1} \right] \\
&= \prod_{k \in K \setminus \{k_1\}} \mathbb{E} [(1 + \alpha)^{F_{ek}} | F_{ek_1}] \quad \text{because the } F_{ek} | F_{ek_1} \text{ are independent} \\
&= \prod_{k \in K \setminus \{k_1\}} (f_{ek}(F_{ek_1})(1 + \alpha)^{D_k} + 1 - f_{ek}(F_{ek_1})) \\
&\leq \prod_{k \in K \setminus \{k_1\}} (f_{ek}(F_{ek_1})(1 + \alpha D_k) + 1 - f_{ek}(F_{ek_1})) \quad \text{because } D_k \leq 1 \\
&= \prod_{k \in K \setminus \{k_1\}} (1 + \alpha f_{ek}(F_{ek_1}) D_k) \\
&\leq \prod_{k \in K \setminus \{k_1\}} e^{\alpha f_{ek}(F_{ek_1}) D_k} \\
&= e^{\alpha \sum_{k \in K \setminus \{k_1\}} f_{ek}(F_{ek_1}) D_k} \\
&= e^{\alpha \mathbb{E}[F_e | F_{ek_1}]}
\end{aligned}$$

□

Lemma 2 *For any positive scalar α and any instance of the unsplittable flow problem, the flow $F_{ek_1} + F_e$ returned by the CSRR algorithm on arc e satisfies:*

$$\mathbb{P}(F_{ek_1} + F_e \geq (1 + \alpha)c_e \Delta^*) \leq \left[\frac{e^\alpha}{(1 + \alpha)^{1 + \alpha}} \right]^{c_e \Delta^*}$$

Proof We note $\delta = (1 + \alpha)^{(1 + \alpha)c_e \Delta^*}$

$$\begin{aligned}
& \mathbb{P}(F_{ek_1} + F_e \geq (1 + \alpha)c_e\Delta^*) \\
&= \mathbb{P}\left((1 + \alpha)^{F_{ek_1} + F_e} \geq \delta\right) \\
&\leq \delta^{-1} \mathbb{E}\left[(1 + \alpha)^{F_{ek_1} + F_e}\right] && \text{Markov inequality} \\
&= \delta^{-1} \mathbb{E}\left[\mathbb{E}\left[(1 + \alpha)^{F_{ek_1} + F_e} \mid F_{ek_1}\right]\right] \\
&= \delta^{-1} \mathbb{E}\left[(1 + \alpha)^{F_{ek_1}} \mathbb{E}\left[(1 + \alpha)^{F_e} \mid F_{ek_1}\right]\right] \\
&\leq \delta^{-1} \mathbb{E}\left[(1 + \alpha)^{F_{ek_1}} e^{\alpha \mathbb{E}[F_e \mid F_{ek_1}]}\right] && \text{Lemma 1} \\
&\leq \delta^{-1} \mathbb{E}\left[(1 + \alpha)^{F_{ek_1}} e^{\alpha(c_e\Delta^* - \mathbb{E}[F_{ek_1}])}\right] && \text{because of constraint (3)} \\
&= \delta^{-1} e^{\alpha(c_e\Delta^* - \mathbb{E}[F_{ek_1}])} \mathbb{E}\left[(1 + \alpha)^{F_{ek_1}}\right] \\
&= \delta^{-1} e^{\alpha(c_e\Delta^* - \mathbb{E}[F_{ek_1}])} (f_{ek_1}(1 + \alpha)^{D_{k_1}} + 1 - f_{ek_1}) \\
&\leq \delta^{-1} e^{\alpha(c_e\Delta^* - \mathbb{E}[F_{ek_1}])} (f_{ek_1}(1 + \alpha D_{k_1}) + 1 - f_{ek_1}) && \text{because } D_{k_1} \leq D_{max} \leq 1 \\
&\leq \delta^{-1} e^{\alpha(c_e\Delta^* - \mathbb{E}[F_{ek_1}])} e^{\alpha f_{ek_1} D_{k_1}} \\
&= \delta^{-1} e^{\alpha(c_e\Delta^* - \mathbb{E}[F_{ek_1}] + \mathbb{E}[F_{ek_1}])} \\
&= \delta^{-1} e^{\alpha c_e \Delta^*} \\
&= \left[\frac{e^\alpha}{(1 + \alpha)^{1 + \alpha}} \right]^{c_e \Delta^*}
\end{aligned}$$

□

We now present the main theorem which upper-bounds for every arc the probability of high congestion in the solution returned by the CSRR algorithm.

Theorem 1 *For any $\epsilon > 0$, there exists an approximation factor $1 + \alpha$ which is $O\left(\frac{\gamma \ln(|E|\epsilon^{-1})}{\ln(\gamma \ln(|E|\epsilon^{-1}))}\right)$ such that, for any instance of the unsplittable flow problem, the flow $F_{ek_1} + F_e$ returned by the CSRR algorithm on arc e satisfies:*

$$\mathbb{P}(F_{ek_1} + F_e \geq (1 + \alpha)c_e\Delta^*) \leq \frac{\epsilon}{|E|}.$$

Proof Lemma 2 gives us $\mathbb{P}(F_{ek_1} + F_e \geq (1 + \alpha)c_e\Delta^*) \leq \left[\frac{e^\alpha}{(1 + \alpha)^{1 + \alpha}} \right]^{c_e \Delta^*}$. To ensure the veracity of the theorem we need to show that there exists a scalar $1 + \alpha$ which is $O\left(\frac{\gamma \ln(|E|\epsilon^{-1})}{\ln(\gamma \ln(|E|\epsilon^{-1}))}\right)$ and for every arc satisfies:

$$\begin{aligned}
\left[\frac{e^\alpha}{(1 + \alpha)^{1 + \alpha}} \right]^{c_e \Delta^*} &\leq \frac{\epsilon}{|E|} \\
\iff & \\
(1 + \alpha) \ln(1 + \alpha) - \alpha &\geq \frac{\ln(|E|\epsilon^{-1})}{c_e \Delta^*}
\end{aligned}$$

For the arc of capacity c_{min} which gives the highest bound, the lower bound is $B = \gamma \ln(|E|\epsilon^{-1})$ (recall that $D_{max} = 1$). Thus we study the solution of the equation $(1 + \alpha) \ln(1 + \alpha) - \alpha = B$ and show that it satisfies $1 + \alpha = O\left(\frac{B}{\ln(B)}\right)$. By replacing $\ln(1 + x)$ with the classical bounds $2\frac{x-1}{x+1} \leq \ln(1 + x) \leq x$, we have:

$$\begin{aligned} \alpha^2 &\geq (1 + \alpha) \ln(1 + \alpha) - \alpha \geq \alpha - 2 \\ \iff \alpha^2 &\geq B \geq \alpha - 2 \\ \iff \sqrt{B} &\leq \alpha \leq B + 2 \end{aligned}$$

Which implies:

$$B = (1 + \alpha) \ln(1 + \alpha) - \alpha \geq (1 + \alpha) \ln(1 + \sqrt{B}) - B - 2$$

and finally,

$$1 + \alpha \leq \frac{2B + 2}{\ln(1 + \sqrt{B})} \sim \frac{4B}{\ln(B)} = O\left(\frac{B}{\ln(B)}\right)$$

Thus, the solution of the equation $(1 + \alpha) \ln(1 + \alpha) - \alpha = \gamma \ln(|E|\epsilon^{-1})$ satisfies $1 + \alpha = O\left(\frac{\gamma \ln(|E|\epsilon^{-1})}{\ln(\gamma \ln(|E|\epsilon^{-1}))}\right)$ and for this value of α , we have

$$\mathbb{P}(F_{ek_1} + F_e \geq (1 + \alpha)c_e \Delta^*) \leq \frac{\epsilon}{|E|}$$

□

Using Theorem 1, we showed that, for any instance of the unsplitable flow problem, the CSRR algorithm returns a solution whose congestion is less than $O\left(\frac{\gamma \ln(|E|\epsilon^{-1})}{\ln(\gamma \ln(|E|\epsilon^{-1}))} \Delta^*\right)$ with probability $1 - \epsilon$. As for the RR algorithm of Raghavan and Tompson (1987), Lemma 2 and Theorem 1 still apply even though the demonstration of Lemma 2 is simpler because commodity k_1 does not need to be treated separately. We thus have the same approximation results for the RR algorithm.

During the resolutions of the linear relaxation of the CSRR algorithm, the constraint (3) is very restrictive. It almost does not leave any flexibility to the variables to move away from the previous optimum.

Example: We call here *first linear relaxation* the linear program solved before any commodity is fixed to a single path and *second linear relaxation* the linear program solved after the first set of commodities is fixed to a single path. For the first and second linear relaxation respectively, we note f_e^1 and f_e^2 the total flow of all the commodities that are not fixed in the first rounding step. Let us suppose that the first linear relaxation has a unique optimal solution in which every arc has the same congestion Δ^* . Then, in the second linear relaxation, constraint (3) ensures that, $\forall e \in E, f_e^1 \geq f_e^2$. In this case, we must have $\forall e \in E, f_e^1 = f_e^2$. Otherwise, the flow of the unfixed commodities in the first linear relaxation could be replaced by the same flow from the second linear relaxation thus creating a new optimal solution of the first linear relaxation.

This would contradict the assumption that the first linear relaxation has a unique optimal solution. Thus, in this example, the variables of the second linear relaxation cannot move away from the previous optimum at all.

This example highlights that the second linear relaxation can only move away from the solution of the first linear relaxation by using the slack left between the congestion of each arc and Δ^* . Thus, when constraint (3) is used, the actualization step only has a small impact on the linear relaxation's solution and the CSRR algorithm does not plainly benefit from the actualization step. We conjecture that it is why the CSRR algorithm does not yield experimentally good results compared to the SRR heuristic (see experimental results of the CSRR algorithm in Section 6). To overcome this problem, we replace constraint (3) by the following constraint:

$$\sum_{k \in K_{free}} f_{ek} D_k \leq \beta c_e \Delta^* - \sum_{k \in K_{fixed}} \hat{f}_{ek} D_k, \quad \forall e \in E. \quad (4)$$

By replacing Δ^* by $\beta \Delta^*$ and γ by $\beta^{-1} \gamma$ for some $\beta \geq 1$ in the proof, one can prove that using constraint (4) the approximation factor of the CSRR algorithm becomes $O\left(\frac{\gamma \ln(|E| \epsilon^{-1})}{\ln(\beta^{-1} \gamma \ln(|E| \epsilon^{-1}))}\right)$ which is still equal to $O\left(\frac{\gamma \ln(|E| \epsilon^{-1})}{\ln(\gamma \ln(|E| \epsilon^{-1}))}\right)$ for any fixed β . Unlike constraint (3), in most practical cases, constraint (4) is not active in the optimal solutions of the linear relaxations even for $\beta = 1.1$. Thus, constraint (4) has no impact on the practical computations and enables the CSRR algorithm to yield the same experimental results as the SRR heuristic.

To summarize, in this section, we added a constraint to the SRR heuristic and switched back to the congestion formulation to create an algorithm that has the same approximation guarantees as the RR algorithm of Raghavan and Tompson (1987). We also tightened the approximation analysis of both algorithms by introducing a granularity parameter γ . This parameter has the property of remaining constant when commodities and arc capacities are uniformly scaled. Finally, we slightly modified the added constraint to alleviate its impact in practice. This modification increased the approximation factor by a negligible value.

6 Experimental results

In this section, we present experiments that support our claims: the SRR heuristic has a lower computing time on large instances than exact methods and yields solutions with a better overflow than the algorithm of Raghavan and Tompson (1987) and meta-heuristics. The impact of sorting the commodities in decreasing order of demand in the randomized rounding algorithms is also investigated. Moreover, the SRR heuristic is compared to the CSRR approximation algorithm. The datasets and the code used in the experimental section of this work are accessible at https://github.com/SuReLI/randomized_rounding_paper_code. All the code for this work was written in Python 3.

The experiments were made on a server with 48 CPU Intel Xeon E5-2670 2.30GHz, 60 Gbit of RAM, and CentOS Linux 7.

In this section, each figure presents the results for a dataset of instances. Each dataset features ten groups of a hundred instances with each group containing instances created using the same parameters. An exception is made in Figure 2 in which each group contains only ten instances. Each point in a figure reports, for an algorithm, the average result of one group of instances. The 95% confidence intervals are also represented as semi-transparent boxes around the main curve.

6.1 Instance datasets

As stated in (Masri et al., 2019), no standard benchmark of instances is present in the literature for the unsplittable flow problem, especially for large instances. Indeed, most works use small graphs (less than 50 nodes) on which they manually generate a set of commodities. The largest instances (100 nodes) can be found in the work of Masri et al. (2015) and Li et al. (2010). Masri et al. (2015) use a grid graph while Li et al. (2010) create their graphs with an adaptation of the graph generator NETGEN. To compensate for this absence of benchmark, we give a detailed explanation of our instance generation procedure. Moreover, all our instances are given in our GitHub repository together with the code used for their generation.

In our tests, we consider two types of graphs: strongly connected random graphs and grid graphs. For strongly connected random graphs, we use the following method to construct a random very sparse strongly connected graph: select a random node u , select a random v such that there is no path from u to v , add an arc (u, v) to the graph, repeat until the graph is strongly connected. Afterward, random edges are added to control the average degree of the graph (it cannot be less than 2). In our tests, the average degree is fixed to 5 and the probability of a node being an origin is $1/10$. A grid graph is an $n \times m$ toric grid with p additional nodes. Each additional node is randomly connected to q nodes on the grid and serves as an origin of the flow. In our tests, we use $m = n = p = \frac{q}{2}$. Unless mentioned otherwise, the arc capacities are 10^4 .

For both types of graphs the demand is created as follows until no more commodity can be added without breaking the capacity constraints:

- choose a destination node d ;
- choose an origin o which can access d within the remaining capacities;
- compute a random simple path p from o to d using a depth-first search where the visit order of newly discovered nodes is random;
- choose a demand level D uniformly between 1 and \hat{D}_{\max} where \hat{D}_{\max} is a parameter defining the maximum possible demand of a commodity; if the chosen demand level is larger than the remaining capacity on p then truncate it to the remaining capacity;
- decrease the capacities on the path p by D ;
- add (o, d, D) to the list of created commodities;

Note that demands created this way can always be routed within the arc capacities and the optimal congestion is one. Thus, we know optimal solutions have zero overflow and a congestion of one. Hence, these optimal values do not need to be computed with exact optimization methods. Moreover, the parameter \hat{D}_{\max} used to parametrize the size of the commodities and thus the number of commodities. Unless mentioned otherwise the value of \hat{D}_{\max} is fixed to 1500 in our tests.

6.2 Benchmarking meta-heuristics

In this section, we present a benchmark of ACO-MC, one of the ant colony optimization algorithms presented in (Li et al., 2010), and the variable neighborhood search of Masri et al. (2015). Both algorithms were reproduced and the code used is given in our GitHub repository. We compare these algorithms with a handcrafted simulated annealing. The goal is to use only the best one as a comparison for the SRR algorithm in the next sections. The algorithm of Li et al. (2010) was coded exactly as presented in their paper.

Implementation of (Masri et al., 2015): due to differences in the considered unsplittable flow problem, the local search part of their algorithm had to be modified. Their local search creates a new path for a commodity by starting from its destination and adding arcs to the path until the origin is reached. At each step, the next arc is chosen by considering the following heuristic information for each out-going arc of the last node of the current path:

$$I_e = \frac{1}{l_e} + \left(1 - \frac{1}{\hat{c}_e}\right)$$

where l_e is the lead time of arc e (*i.e.* its length) and \hat{c}_e is the remaining capacity on arc e . As we do not have lead times for each arc in our problem, l_e was set to 1 for each arc. Moreover, in the problem studied in (Masri et al., 2015), the remaining capacity \hat{c}_e is positive because no overflow is allowed. Because we can have negative remaining capacities, we replace the function $f : x \rightarrow 1 - \frac{1}{x}$ applied to \hat{c}_e by $g : x \rightarrow \frac{1}{2} \left(1 + \frac{x}{1+|x|}\right)$. Function g was chosen to have similar properties to the function f .

$$\forall x \in (1, +\infty), 0 < f(x) < 1$$

$$\forall x \in (-\infty, +\infty), 0 < g(x) < 1$$

$$f(x) - 1 \underset{+\infty}{\sim} g(x) - 1 \underset{+\infty}{\sim} \frac{-1}{x}$$

$$g(x) \underset{-\infty}{\sim} \frac{-1}{x}$$


In our tests, we also present the results obtained with a version of the algorithm of Masri et al. (2015) where the local search part is disabled.

Our simulated annealing: at each iteration, a solution is created in the neighborhood of the current solution; this modification is accepted with a probability depending on the improvement/deterioration of the solution and a temperature parameter. At each iteration, the temperature parameter is multiplied by $1 - \epsilon$ for some small ϵ depending on the number of iterations. At the beginning of the simulated annealing procedure and similarly to the algorithm of Masri et al. (2015), a list of k -shortest paths of length 10 is computed for each commodity using the algorithm of Jiménez and Marzal (1999). To initialize the solution, each commodity takes a random path in its list of k -shortest paths. At each iteration, a neighborhood solution is created by replacing the path of a commodity with a path randomly chosen in the list of k -shortest paths of the commodity. The stopping criterion is the total number of iterations.

Hyper-parameter setting: we benchmarked different hyper-parameters values for ACO-MC but finally settled to use the same values as in (Li et al., 2010). Except for the size of the largest neighborhood which is not mentioned, both versions of the variable neighborhood search use the hyper-parameter values given in (Masri et al., 2015). After testing different values, the size of the largest neighborhood was set to 3. The hyper-parameters of the simulated annealing are the initial and final temperature chosen to be respectively 200 and 1. The simulated annealing (SA) is given $2|K|^{1.5}$ iterations so that it takes a time comparable to the SRR algorithm in the next sections. ACO-MC and the VNS of Masri et al. (2015) were given respectively 50 and 100 iterations. Although these numbers seem very low, one must consider that for each of their iterations, the algorithms of Li et al. (2010) and Masri et al. (2015) respectively generate $|K|$ and $\frac{|K|}{2}$ new paths. With these numbers of iterations, they already require a longer computing time than the simulated annealing procedure. Finally, the variation of the variable neighborhood search where the local search part is disabled (VNS2) was given $|K|^{1.5}$ iterations.

Results: Figure 1 presents the results of each algorithm on a dataset composed of grid graphs of varying sizes. The simulated annealing procedure clearly outperforms the other algorithms on this dataset. We obtained similar results on all other datasets. That is why the simulated annealing algorithm has been chosen in the next sections to be the comparison point for the SRR algorithm.

Discussion: we now try to explain why the meta-heuristics from the literature are outperformed. At each iteration, the ant colony procedure spends most of its computing time creating an entirely new solution. Indeed, a new path is created for each commodity for a total of $|K|$ paths generated. Conversely, at each iteration, the simulated annealing procedure chooses only one path in a list (choosing a path is much faster than creating one). This leads to a much larger number of iterations for the simulated annealing procedure and thus a lot more solutions evaluated. As for the algorithm of Masri et al. (2015), it seems that the local search part performs poorly in our tests. This is partly because the path generation procedure does not know where its target node is until it is encountered. Indeed, the heuristic information used to choose the



(a) Overflow divided by the total demand — (b) Computing time in seconds — grid graphs

Fig. 1 Performance and computing time versus the number of nodes of various meta-heuristics

next arc of the path considers the length l_e of the arcs and not the length of the shortest path to the target node. Thus, when choosing the next arc, the procedure does not know if it goes toward or away from its goal. We tried to replace l_e by the length of the shortest path to the target node in the local search but disabling it completely still yielded better solutions. Finally, it appears that variable neighborhood search is not a good choice of meta-heuristic for our version of the unsplittable flow problem (especially when the number of commodities is more than a thousand). Indeed, by changing the size k_{max} of the largest neighborhood considered we obtained that the best value for this parameter is $k_{max} = 1$ which is the case where only the smallest neighborhood is considered.

6.3 Results

We compare the following algorithms in our tests :

- RR: the randomized rounding algorithm of Raghavan and Tompson (1987);
- RR sorted: a version of the RR algorithm where the commodities are rounded in order of decreasing demand;
- SRR: the Sequential Randomized Rounding heuristic described in Section 4;
- SRR unsorted: a version of the SRR heuristic where the commodities are not rounded in order of decreasing demand but in random order;
- CSRR: the version of the SRR algorithm with approximation properties presented in Section 5;
- SA and SA2: a handcrafted simulated annealing procedure presented in Section 6.2; SA is given $2|K|^{1.5}$ iterations to have a similar computing time to SRR on grid graphs while SA2 is given $6|K|^{1.5}$ iterations;
- MILP solver: arc-node formulation solved with Gurobi 8.11.

All the randomized rounding algorithms in this list use the overflow-sum objective in the linear relaxation to create their solutions.

An important characteristic of the SRR heuristic is its capacity to scale to large instances. Thus we first illustrate how the algorithm outperforms a MILP method. In Figure 2, a MILP method based on the arc-node formulation is compared with the RR and the SRR algorithms on small grid graphs with very few commodities. The MILP algorithm solves optimally all the small instances but gives poor results on large instances within a 20 minutes time limit. In comparison, the other methods retain a reasonable performance on large instances. For these instances, the capacity of the arcs is 3 and the maximum demand of a commodity is 2 to keep the number of commodities and variables of the MILP formulation low. Due to the large amount of memory required, the MILP algorithm could not be tested on the other larger datasets.

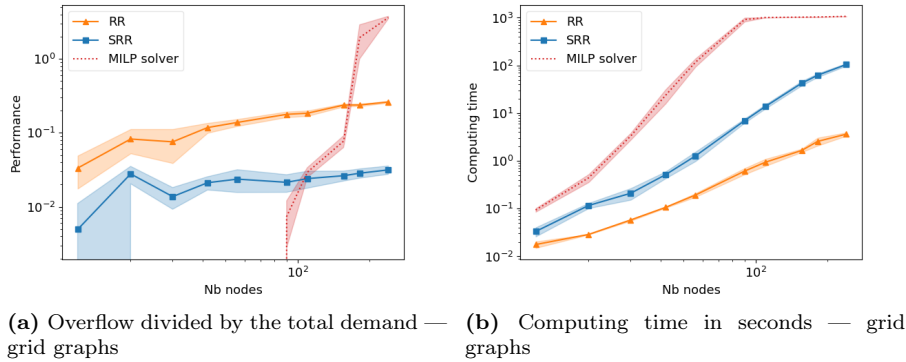


Fig. 2 Performance and computing time versus the number of nodes on small instances

Figure 3 compares different randomized rounding algorithms and the simulated annealing algorithm on grid graphs and random connected graphs. In both cases, the SRR algorithm requires a larger computing time than pure randomized rounding but returns solutions of much higher quality. For small instances of grid graphs, the best results are given by the simulated annealing procedure. However, as the number of nodes increases the SRR heuristics outperforms the simulated annealing procedure with the same computing time (SA). Furthermore, on the largest instances, the simulated annealing procedure is outperformed by the SRR heuristics even when given three times more computing time (SA2). As for strongly connected random graphs, with the described settings, the SRR heuristic returns the highest quality solutions but requires a larger computing time than the simulated annealing. It appears that solving the linear relaxation takes a longer time on random graphs than on grid graphs. The CSRR algorithm returns worse solutions than the SRR heuristic

in a similar computing time. As explained at the end of Section 5, this is due to constraint (3) added in the resolution of the linear relaxation. One can relax this constraint into constraint (4) while keeping a similar approximation factor. In practice, constraint (4) is never active in the optimal solution of the linear relaxation. Thus, with this adaptation, the CSRR algorithm returns exactly the same solutions as the SRR heuristic.

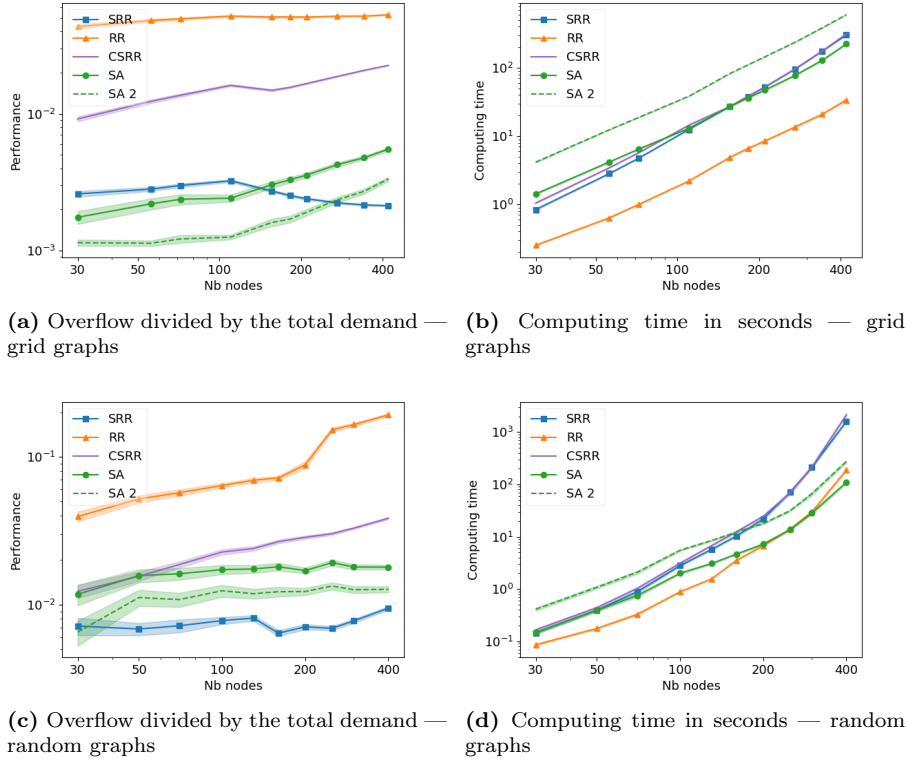
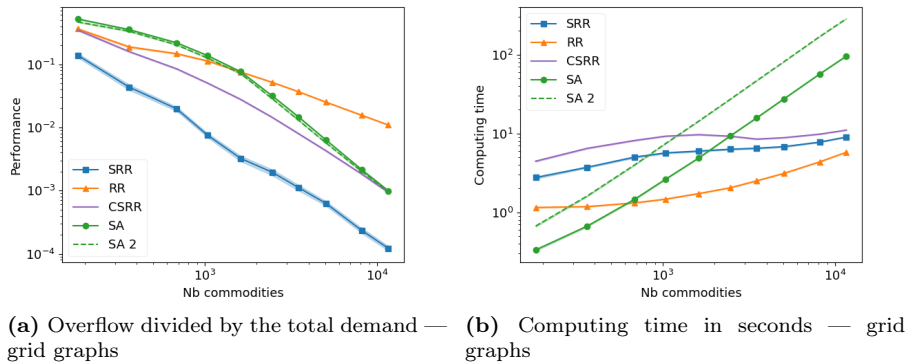


Fig. 3 Performance and computing time versus the number of nodes on large instances

Figure 4 reports the solution performance and the computing time for grid graphs with a constant number of nodes and arcs but a various number of commodities. This variation is obtained by changing the ratio of the arc capacities over the maximum demand of a commodity. Table 1 reports the parameters used to create the instances. As can be seen in Figure 4a, all algorithms produce higher quality solutions when there is a large number of commodities (*i.e.* commodities have small demand compared to the arc capacities). As explained in Section 5, for randomized rounding algorithms, this practical result can be related to the presence of the granularity parameter γ in the approximation

factor of the RR and CSRR algorithms. Figure 4b shows that our use of an aggregated arc-node formulation enables the randomized rounding algorithms to have a computing time that scales very well with the number of commodities. However, this is not the case for simulated annealing which requires a large increase in its number of iterations to obtain the solution reported in Figure 4a.




(a) Overflow divided by the total demand — (b) Computing time in seconds — grid graphs

Fig. 4 Performance and computing time versus the number of commodities

We now analyze the impact of rounding the commodities in order of decreasing demand, based on the results presented in Figure 5. Both pure randomized rounding and the SRR heuristic yield solutions of higher quality when the commodities are sorted. The impact is significantly higher on the SRR algorithm than on pure randomized rounding. A possible explanation for such a strong impact is given in Section 7.2. At first glance, a rounding order should have no impact on pure randomized rounding since the commodities are rounded independently. However, results clearly indicate otherwise and this is due to how the linear relaxation is computed in our experiments. Indeed when using the arc-node formulation of Section 2.2 and the commodity aggregation presented in Section 4.3, the linear program does not directly yield a flow distribution for each commodity. The flow distribution returned is for the aggregated commodities. To get the flow of each commodity, a flow decomposition algorithm is used. This algorithm tends to split less the commodities that are decomposed first. Thus when the commodities with the largest demands are rounded first, they are also decomposed first. This means the biggest commodities have a lower chance of being split and thus a lower chance that their rounding creates a large overflow.


Finally, we describe how we chose the value $\frac{|V|}{4}$ for the actualization parameter θ of the SRR algorithm presented in Section 4. Setting this value implies making a trade-off between solution quality and computing time. Figure 6 compares the performance of SRR for different values of this parameter,



(a) Overflow divided by the total demand — (b) Computing time in seconds — grid graphs

Fig. 5 Influence of commodity sorting on performance and computing time

on a grid graph with 110 nodes. The computing time of the SRR heuristic affinely decreases with the θ parameter. However, the overflow of the returned solution decreases less and less when θ decreases. We decided to choose θ such that the returned solution is close to the best obtainable solution while a lot of computing time is saved compared to choosing θ close to zero. When repeating this experiment on graphs of different sizes and types, it appeared that the chosen trade-off value was close to $\frac{|V|}{4}$. This value is represented with a dashed line in Figure 6.



(a) Overflow divided by the total demand — (b) Computing time in seconds — grid graphs

Fig. 6 Performance and computing time versus the value of the θ parameter of SRR

Arc capacities	1	2	5	10	20
Maximum demand of a commodity	1	1	2	3	4
Average number of commodities	182	362	685	1038	1615
Arc capacities	50	100	200	500	1000
Maximum demand of a commodity	7	10	14	22	31
Average number of commodities	2462	3512	5048	8138	11644

Table 1 Parameters of the instances of the commodity scaling dataset : the graphs have 110 nodes and 580 arcs

7 Discussion

In this section, we discuss different properties of the randomized rounding algorithms. We first show that using the overflow sum objective function presented in Section 2.1 instead of the classical congestion objective function has a positive impact on practical results. Then we highlight a positive interplay between sorting the commodities and actualizing the linear relaxation.

7.1 Impact of the objective function

Two types of objective function coexist in the SRR heuristic. The first one is the objective function of the unsplittable flow problem which will be called *evaluation metric* in this section. The second one is the objective function used in the resolutions of the linear relaxation inside the SRR heuristic which we will call *generation objective*. For both of the previous types, we study the two functions presented in Section 2.1: the sum of the overflow and the congestion. Usually, the generation objective and the evaluation metric are chosen to be the same. However, the overflow sum and the congestion are very close metrics and we want to study the impact of using one for creating solutions for the other. We will show below that using the overflow sum instead of the congestion as generation objective yields solutions of higher quality for both evaluation metrics.

When the evaluation metric is the overflow sum, in our tests, solutions generated using the overflow sum as generation objective have a ten times better solution quality than the one created using the congestion. More surprisingly, the overflow sum used as generation objective also yields the best solutions when the evaluation metric is the congestion. To test this, the SRR heuristic has been applied with the two generation objectives on 100 grid graphs with 120 nodes, 700 arcs, and 4500 commodities. The mean congestion of the returned solutions is reported in the first two columns of Table 2 together with the standard deviation of the results. The mean congestions for the two generation objectives are separated by more than one standard deviation in favor of the overflow sum. Using the values given in the first two columns of Table 2, an unpaired two-sample Student t-test was made. The test showed that the overflow sum used as generation objective performs significantly better than the congestion. The p -value associated with the test is 10^{-18} .

	Overflow sum	Congestion	Mixed
Mean	1.025	1.049	1.027
Standard Deviation	0.016	0.019	0.012

Table 2 Comparison of the congestion of the solutions returned by the SRR heuristic using different generation objectives.

We conjecture that the difference observed above between the two generation objectives is explained by the following reasoning. In the linear relaxation, the congestion objective function does not differentiate solutions having one arc with high congestion from solutions with several arcs with high congestion. However, a larger number of highly congested arcs in the linear relaxation implies, on average, a larger congestion in the solution returned by the randomized rounding process. Furthermore, if the randomized rounding process creates arcs with a larger congestion than the optimal congestion of the linear relaxation, then, in subsequent actualizations of the linear relaxation, the capacity constraints are lifted for all the other arcs. This might lead to even more congested unsplitable solutions.

With this conjecture in mind, we created an algorithm that uses the congestion as its main generation objective but performs as well as the algorithm using the overflow sum as generation objective. To that end, a two-level objective is set in the linear relaxation: find the solution of least overflow sum among the solutions with minimal congestion. The results obtained with this alternative generation objective are presented in the third column of Table 2 under the name “Mixed” objective. These results are comparable to those obtained with the overflow sum objective. Indeed, an unpaired two-sample Student t-test made on the values given in the first and third columns of Table 2 does not show that the difference is statistically significant. The p -value associated with the test is 0.32.

7.2 Interaction between sorting and actualization

In this section, we discuss the positive interplay observed in Section 6 between sorting the commodities by decreasing demand and actualizing the linear relaxation. To that end, we study a toy example where the graph is composed of two nodes linked by two parallel arcs. We also assume the sum of the demands is equal to the sum of the capacities. We compare the solutions with the overflow sum objective and assume that no binary solution has an overflow of zero.

When computing the linear relaxation, optimal extrema of the polyhedron are of the following form: every commodity is assigned to one arc except for one commodity which is split between the two arcs. We assume that each commodity has the same probability of being the split commodity. After applying randomized rounding to a linear solution of this type, the binary solution obtained has an overflow between zero and the demand of the split commodity.

On the other hand, if we actualize the linear relaxation after fixing the split commodity we should often get a better solution. Indeed, after actualizing the linear solution, three situations can occur. Firstly, the remaining free commodities cannot switch arc to compensate any of the overflow generated by rounding the split commodity. In this case, no commodity changes its flow, the new linear solution is a binary solution and this solution is the same with and without actualization. Secondly, the remaining free commodities compensate only a part of the overflow created. In this case, even though some commodities change their flow, the new linear solution is a binary solution whose overflow is lower than the one before the actualization step. Lastly, the remaining free commodities compensate completely the overflow created and a new split commodity is created. Applying randomized rounding from there yields a solution whose overflow is between zero and the demand of the new split commodity. Since the commodities are fixed in order of decreasing demand, the new split commodity has a smaller demand than the old split commodity. The range in which the overflow of the final solution can vary is thus smaller which should on average lead to a better solution. Furthermore, new split commodities are created until the free commodities cannot completely compensate for the rounding of the last split commodity. Most of the time, this happens when only a few commodities remain and thus when only small commodities remain. In this case, the last split commodity has a small demand, and the algorithm yields a small final overflow.

Finally, we can see that in this toy example, actualizing the linear solution and sorting the commodities by decreasing demands yields, on average, solutions with a smaller overflow than pure randomized rounding. We conjecture that this phenomenon generalizes to any graph. Indeed, the fact that the split commodities get smaller and smaller as the algorithm progresses should lead to smaller and smaller final overflows.

8 Conclusion

In this paper, we have presented a heuristic based on randomized rounding for large-scale instances of the unsplittable flow problem which extends the algorithm of Raghavan and Tompson (1987). We experimentally showed that on large-scale instances this heuristic produces solutions of smaller overflow than any other method used for comparison.

We also derived an approximation algorithm from the heuristic by restraining the possible actualization of the linear relaxation. The approximation factor of both this algorithm and the algorithm of Raghavan and Tompson (1987) was then tightened to $O\left(\frac{\gamma \ln(|E|\epsilon^{-1})}{\ln(\gamma \ln(|E|\epsilon^{-1}))}\right)$. This new approximation factor depends on the granularity parameter γ and enables to understand the behavior of randomized rounding when the commodities are small compared to the capacities (*i.e.* $\gamma \ll 1$).

Furthermore, the behavior of the presented heuristic has been analyzed to highlight two of its key particularities. First, the new objective function used

in the algorithm for practical computations yielded solutions of higher quality. Secondly, the actualization of the solution of the linear relaxation enhanced the performances of the heuristic when the commodities are sorted in order of decreasing demand.

Finally, even though the techniques discussed in this paper were presented in the context of unsplittable flows, they apply to other contexts where the randomized rounding method of Raghavan and Tompson (1987) is used (packing problems, covering problems...). Their performances and variations in these contexts could be investigated. Moreover, the impact of backtracking of the decisions made during randomized rounding algorithms seems a promising research direction.

References

- Alvelos F, Valério de Carvalho J (2007) A local search heuristic based on column generation applied to the binary multicommodity flow problem. In: Proceedings of International Network Optimization Conference, INOC, Citeseer, p 6
- Alvelos F, De Carvalho JV (2003) Comparing branch-and-price algorithms for the unsplittable multicommodity flow problem. In: International Network Optimization Conference, pp 7–12
- Anderson CA, Fraughnaugh K, Parker M, Ryan J (1993) Path assignment for call routing: An application of tabu search. *Annals of Operations Research* 41(4):299–312
- Andrews M, Chuzhoy J, Guruswami V, Khanna S, Talwar K, Zhang L (2010) Inapproximability of edge-disjoint paths and low congestion routing on undirected graphs. *Combinatorica* 30(5):485–520
- Asano Y (2000) Experimental evaluation of approximation algorithms for the minimum cost multiple-source unsplittable flow problem. In: ICALP Satellite Workshops, pp 111–122
- Aumann Y, Rabani Y (1995) Improved bounds for all optical routing. In: Proceedings of the sixth annual ACM-SIAM symposium on Discrete algorithms, Citeseer, pp 567–576
- Azar Y, Regev O (2006) Combinatorial algorithms for the unsplittable flow problem. *Algorithmica* 44(1):49–66
- Babonneau F, Du Merle O, Vial JP (2006) Solving large-scale linear multicommodity flow problems with an active set strategy and proximal-accpm. *Operations Research* 54(1):184–197
- Barnhart C, Hane CA, Vance PH (2000) Using branch-and-price-and-cut to solve origin-destination integer multicommodity flow problems. *Operations Research* 48(2):318–326
- Bauguion PO, Ben-Ameur W, Gourdin E (2013) A new model for multicommodity flow problems, and a strongly polynomial algorithm for single-source maximum concurrent flow. *Electronic Notes in Discrete Mathematics* 41:311–318

- Bauguion PO, Ben-Ameur W, Gourdin E (2015) Efficient algorithms for the maximum concurrent flow problem. *Networks* 65(1):56–67
- Belaidouni M, Ben-Ameur W (2007) On the minimum cost multiple-source unsplittable flow problem. *RAIRO-Operations Research* 41(3):253–273
- Castro J, Cuesta J (2012) Improving an interior-point algorithm for multicommodity flows by quadratic regularizations. *Networks* 59(1):117–131
- Chakrabarti A, Chekuri C, Gupta A, Kumar A (2007) Approximation algorithms for the unsplittable flow problem. *Algorithmica* 47(1):53–78
- Chuzhoy J, Guruswami V, Khanna S, Talwar K (2007) Hardness of routing with congestion in directed graphs. In: *Proceedings of the thirty-ninth annual ACM symposium on Theory of computing*, ACM, pp 165–178
- Coudert D, Rivano H (2002) Lightpath assignment for multifibers wdm networks with wavelength translators. In: *Global Telecommunications Conference, 2002. GLOBECOM '02. IEEE*, vol 3, pp 2686–2690 vol.3
- Cox LA (1991) Dynamic anticipatory routing of circuit-switched telecommunications networks. *Hand-book of Genetic Algorithms*
- Csirik J, Galambos G, Frenk HJ, Frieze A, Kan R, Alexander A (1986) A probabilistic analysis of the next fit decreasing bin packing heuristic. *Operations Research Letters* 5(5):233–236
- Dai W, Sun X, Wandelt S (2016a) Finding feasible solutions for multi-commodity flow problems. In: *2016 35th Chinese Control Conference (CCC)*, IEEE, pp 2878–2883
- Dai W, Zhang J, Sun X, Wandelt S (2016b) Node dependency in multi-commodity flow problem with applications to transportation networks. In: *CICTP 2016*, pp 1989–2001
- Dai W, Zhang J, Sun X (2017) On solving multi-commodity flow problems: An experimental evaluation. *Chinese Journal of Aeronautics* 30(4):1481–1492
- Emanuelsson K (2016) Approximating multi-commodity max-flow in practice
- Fleischer LK (2000) Approximating fractional multicommodity flow independent of the number of commodities. *SIAM Journal on Discrete Mathematics* 13(4):505–520
- Ford Jr LR (1956) *Network flow theory*. Tech. rep., Rand Corp Santa Monica Ca
- Gondzio J, González-Brevis P (2015) A new warmstarting strategy for the primal-dual column generation method. *Mathematical Programming* 152(1-2):113–146
- Gondzio J, González-Brevis P, Munari P (2013) New developments in the primal-dual column generation technique. *European Journal of Operational Research* 224(1):41–51
- Gondzio J, González-Brevis P, Munari P (2016) Large-scale optimization with the primal-dual column generation method. *Mathematical Programming Computation* 8(1):47–82
- Group TF, Gonzalez T (2020) *Handbook of Approximation Algorithms and Metaheuristics*, Second Edition: Two-Volume Set. Taylor & Francis Group, URL <https://books.google.fr/books?id=iD2UzQEACAAJ>

- Gurobi Optimization L (2020) Gurobi optimizer reference manual. URL <http://www.gurobi.com>
- Jiménez VM, Marzal A (1999) Computing the k shortest paths: A new algorithm and an experimental comparison. In: International Workshop on Algorithm Engineering, Springer, pp 15–29
- Kolman P (2003) A note on the greedy algorithm for the unsplittable flow problem. *Information Processing Letters* 88(3):101–105
- Laguna M, Glover F (1993) Bandwidth packing: a tabu search approach. *Management science* 39(4):492–500
- Li X, Aneja YP, Baki F (2010) An ant colony optimization metaheuristic for single-path multicommodity network flow problems. *Journal of the Operational Research Society* 61(9):1340–1355
- Madry A (2010) Faster approximation schemes for fractional multicommodity flow problems via dynamic graph algorithms. In: Proceedings of the forty-second ACM symposium on Theory of computing, ACM, pp 121–130
- Martens M, Skutella M (2006) Flows on few paths: Algorithms and lower bounds. *Networks: An International Journal* 48(2):68–76
- Masri H, Krichen S, Guitouni A (2011) An ant colony optimization metaheuristic for solving bi-objective multi-sources multicommodity communication flow problem. In: 2011 4th Joint IFIP Wireless and Mobile Networking Conference (WMNC 2011), IEEE, pp 1–8
- Masri H, Krichen S, Guitouni A (2015) A multi-start variable neighborhood search for solving the single path multicommodity flow problem. *Applied Mathematics and Computation* 251:132–142
- Masri H, Krichen S, Guitouni A (2019) Metaheuristics for solving the biobjective single-path multicommodity communication flow problem. *International Transactions in Operational Research* 26(2):589–614
- Moradi S, Raith A, Ehrgott M (2015) A bi-objective column generation algorithm for the multi-commodity minimum cost flow problem. *European Journal of Operational Research* 244(2):369–378
- Park K, Kang S, Park S (1996) An integer programming approach to the bandwidth packing problem. *Management science* 42(9):1277–1291
- Park S, Kim D, Lee K (2003) An integer programming approach to the path selection problems. In: Proceedings of the International Network Optimization Conference INOC, Evry-Paris, France, pp 448–453
- Parker M, Ryan J (1993) A column generation algorithm for bandwidth packing. *Telecommunication Systems* 2(1):185–195
- Raghavan P (1988) Probabilistic construction of deterministic algorithms: approximating packing integer programs. *Journal of Computer and System Sciences* 37(2):130–143
- Raghavan P, Tompson CD (1987) Randomized rounding: a technique for provably good algorithms and algorithmic proofs. *Combinatorica* 7(4):365–374
- Retvdri G, Bíró JJ, Cinkler T (2004) A novel lagrangian-relaxation to the minimum cost multicommodity flow problem and its application to ospf traffic engineering. In: Proceedings. ISCC 2004. Ninth International Symposium on Computers And Communications (IEEE Cat. No. 04TH8769), IEEE, vol 2,

- pp 957–962
- Santos D, de Sousa A, Alvelos F, Pioro M (2010) Link load balancing optimization of telecommunication networks: A column generation based heuristic approach. In: 2010 14th International Telecommunications Network Strategy and Planning Symposium (NETWORKS), IEEE, pp 1–6
- Santos D, De Sousa A, Alvelos F (2013a) A hybrid column generation with grasp and path relinking for the network load balancing problem. *Computers & Operations Research* 40(12):3147–3158
- Santos D, de Sousa A, Alvelos F, Pióro M (2013b) Optimizing network load balancing: an hybridization approach of metaheuristics with column generation. *Telecommunication Systems* 52(2):959–968
- Shahrokhi F, Matula DW (1990) The maximum concurrent flow problem. *Journal of the ACM (JACM)* 37(2):318–334
- Shepherd FB, Vetta A (2015) The inapproximability of maximum single-sink unsplittable, priority and confluent flow problems. arXiv preprint arXiv:150400627
- Wang IL (2018) Multicommodity network flows: A survey, part ii: Solution methods. *International Journal of Operations Research* 15(4):155–173
- Wang Y, Wang Z (1999) Explicit routing algorithms for internet traffic engineering. In: *Proceedings Eight International Conference on Computer Communications and Networks* (Cat. No. 99EX370), IEEE, pp 582–588
- Xu J, Chiu SY, Glover F (1997) Tabu search for dynamic routing communications network design. *Telecommunication Systems* 8(1):55–77
- Yen JY (1971) Finding the k shortest loopless paths in a network. *management Science* 17(11):712–716