



CIRRELT

Centre interuniversitaire de recherche
sur les réseaux d'entreprise, la logistique et le transport

Interuniversity Research Centre
on Enterprise Networks, Logistics and Transportation

Branch-and-Price-and-Cut for Large-Scale Multicommodity Capacitated Fixed-Charge Network Design

Bernard Gendron
Mathieu Larose

November 2012

CIRRELT-2012-74

Bureaux de Montréal :

Université de Montréal
C.P. 6128, succ. Centre-ville
Montréal (Québec)
Canada H3C 3J7
Téléphone : 514 343-7575
Télécopie : 514 343-7121

Bureaux de Québec :

Université Laval
2325, de la Terrasse, bureau 2642
Québec (Québec)
Canada G1V 0A6
Téléphone : 418 656-2073
Télécopie : 418 656-2624

www.cirrelt.ca

Branch-and-Price-and-Cut for Large-Scale Multicommodity Capacitated Fixed-Charge Network Design

Bernard Gendron^{*}, Mathieu Larose

Interuniversity Research Centre on Enterprise Networks, Logistics and Transportation (CIRRELT) and Department of Computer Science and Operations Research, Université de Montréal, P.O. Box 6128, Station Centre-Ville, Montréal, Canada H3C 3J7

Abstract. We present a branch-and-price-and-cut algorithm for solving large-scale instances of the multicommodity capacitated fixed-charge network design problem. The restricted master problem solved at each column generation iteration is obtained directly from the compact arc-based model by considering only a subset of the commodity flow variables. The pricing subproblem corresponds to a Lagrangian relaxation of the flow conservation and capacity constraints, leaving in the Lagrangian subproblem only the strong inequalities. The column generation procedure is completed by a cut generation step based on strong inequalities. The resulting column-and-row generation procedure is embedded within an enumerative scheme. Computational experiments on a large set of randomly generated instances are presented and analyzed.

Keywords. Multicommodity capacitated fixed-charge network design, column generation, branch-and-price-and-cut.

Acknowledgements. Funding for this project has been provided by the Natural Sciences and Engineering Research Council of Canada (NSERC). We thank the SCIP team of the Zuse-Institut Berlin (ZIB) for their help in using the SCIP library, especially Tobias Achterberg and Stefan Vigerske.

Results and views expressed in this publication are the sole responsibility of the authors and do not necessarily reflect those of CIRRELT.

Les résultats et opinions contenus dans cette publication ne reflètent pas nécessairement la position du CIRRELT et n'engagent pas sa responsabilité.

* Corresponding author: Bernard.Gendron@cirrelt.ca

1 Introduction

In this paper, we present a branch-and-price-and-cut (B&P&C) algorithm for the multicommodity capacitated fixed-charge network design problem (*MCND*), an NP-hard problem [23] defined on a directed graph $G = (N, A)$, where N is the set of nodes and A is the set of arcs. Each commodity $k \in K$ is characterized by a demand $d^k > 0$ to be routed from an origin $O(k)$ to a destination $D(k)$. On each arc (i, j) , there is a capacity $u_{ij} > 0$ on the flow of all commodities circulating on the arc (we assume $u_{ij} \leq \sum_{k \in K} d^k$). The problem is to satisfy the demands at minimum cost, while respecting the capacity constraints. The objective function consists of the sum of transportation costs and fixed design costs, the latter being charged whenever an arc is used. The transportation and fixed design costs on arc (i, j) are denoted $c_{ij} \geq 0$ and $f_{ij} \geq 0$, respectively.

We model the *MCND* as a mixed-integer program (MIP) by using continuous flow variables x_{ij}^k that represent the amount of flow on each arc (i, j) for each commodity k , and 0-1 design variables y_{ij} that indicate if arc (i, j) is used or not:

$$\min \sum_{(i,j) \in A} \sum_{k \in K} c_{ij} x_{ij}^k + \sum_{(i,j) \in A} f_{ij} y_{ij}, \quad (1)$$

$$\sum_{j \in N_i(+)} x_{ij}^k - \sum_{j \in N_i(-)} x_{ji}^k = \begin{cases} d^k, & \text{if } i = O(k), \\ -d^k, & \text{if } i = D(k), \\ 0, & \text{otherwise,} \end{cases} \quad i \in N, k \in K, \quad (2)$$

$$\sum_{k \in K} x_{ij}^k \leq u_{ij} y_{ij}, \quad (i, j) \in A, \quad (3)$$

$$x_{ij}^k \leq d^k y_{ij}, \quad (i, j) \in A, k \in K, \quad (4)$$

$$x_{ij}^k \geq 0, \quad (i, j) \in A, k \in K, \quad (5)$$

$$y_{ij} \in \{0, 1\}, \quad (i, j) \in A, \quad (6)$$

where $N_i(+)$ = $\{j \in N | (i, j) \in A\}$ and $N_i(-)$ = $\{j \in N | (j, i) \in A\}$. Equations (2) are the flow conservation constraints for each node and each commodity. The capacity constraints for each arc, (3), forbid any flow to circulate through an arc that is not chosen as part of the design. The so-called *strong inequalities*, (4), also serve the same purpose and are therefore redundant; however, they significantly improve the linear programming (LP) relaxation bounds [7]. Model (1)-(6) is the strong formulation of the MCND, the weak formulation being obtained by removing constraints (4); their corresponding LP relaxations are, respectively, the strong and the weak relaxations.

LP-based branch-and-bound (B&B) algorithms can be used to solve the MCND, but large-scale instances (with hundreds of commodities) remain computationally elusive. On the one hand, if strong inequalities are removed, the resulting LP relaxation provides too weak lower bounds to be used within an enumerative approach; on the other hand, if all strong inequalities are included *a priori* in the model, the resulting LP relaxation is not only very large, but also highly degenerate. The obvious alternative is to add these valid inequalities in a dynamic way within a cutting-plane algorithm [3] that might also

include other well-known valid inequalities for fixed-charge network flow problems, such as cover and flow cover inequalities. The following key observations can be derived from the computational results obtained with this cutting-plane algorithm [3]:

- It is essential to add the strong inequalities to obtain effective lower bounds that can be computed efficiently; more precisely, cover and flow cover inequalities can be used instead of the strong inequalities to derive equally effective lower bounds, but with a much greater computational effort.
- Obviously, adding cover and flow cover inequalities to the strong relaxation improves the lower bound, but on large-scale instances (with hundreds of commodities), this improvement is so small that it does not pay off when the cutting-plane method is used within B&B.
- Many of these large-scale instances remain difficult to solve.

These findings motivate the development of our B&P&C algorithm: at each node of the enumeration tree, we solve the strong relaxation by adding in a dynamic way not only the strong inequalities, but also the flow variables. The resulting *column-and-row generation* procedure, when embedded within an enumerative scheme, is able to solve large-scale instances more efficiently than the standard cutting-plane method, as we will see in Section 3.

Other decomposition methods have been proposed for solving the MCND, in particular Benders decomposition [4, 5] and Lagrangian-based procedures [7, 8, 14, 15, 19, 22, 26]; in Section 2.6, we examine the relationships between our approach and some of these methods. A large number of heuristic methods have also been proposed for computing high-quality feasible solutions [6, 9, 10, 16, 17, 18, 21, 25]; in practice, near-optimal solutions can be obtained by these heuristics, even for large-scale instances, but the difficulty is in assessing the quality of these solutions with effective lower bounds and in proving optimality, which is the focus of our paper.

The paper is organized as follows. Section 2 presents the B&P&C algorithm, while the results of experiments on a large set of randomly generated instances are reported in Section 3. We conclude this paper with a discussion of future research avenues.

2 Branch-and-Price-and-Cut Algorithm

In this section, we present the different components of our B&P&C algorithm: the restricted master problem, in Section 2.1, the pricing subproblem, in Section 2.2, the cut generation, in Section 2.3, and the variable fixing and branching procedures, in Section 2.4. The overall algorithm is summarized in Section 2.5. We conclude this section by studying the relationships between our approach and other decomposition methods for the MCND; in particular, we show that our algorithm can be seen as a special case of two recently proposed generic B&P&C frameworks [13, 24].

2.1 Restricted Master Problem

To obtain the restricted master problem solved at each column-and-row generation iteration, we enlarge the set of arcs with one *artificial arc* connecting $O(k)$ to $D(k)$ for each commodity k ; this arc is uncapacitated, has no fixed design cost and is given a very large transportation cost. By adding these arcs, we ensure that every restricted master problem is always feasible; in addition, if, at any given node of the B&P&C tree, an optimal solution to the corresponding strong relaxation is obtained that includes at least one artificial arc, then the node is infeasible and therefore fathomed. We denote by A_+ the set containing A along with the artificial arcs.

To define the restricted master problem, we associate with each arc (i, j) in A_+ a commodity subset $\tilde{K}_{ij} \subseteq K$ over which the existing flow variables (x_{ij}^k for $k \in \tilde{K}_{ij}$) are defined (for an artificial arc $(i, j) = (O(l), D(l))$, we always have $\tilde{K}_{ij} = \{l\}$). In addition, we define, for each commodity k , the sets $\tilde{A}^k = \{(i, j) \in A_+ | k \in \tilde{K}_{ij}\}$, and for each node i and commodity k , the sets $\tilde{N}_i^k(+)=\{j \in N | (i, j) \in \tilde{A}^k\}$ and $\tilde{N}_i^k(-)=\{j \in N | (j, i) \in \tilde{A}^k\}$. Only the strong inequalities that are generated as cuts are included in the model; for each arc (i, j) , we represent the generated strong inequalities using a subset \bar{K}_{ij} of \tilde{K}_{ij} . The *restricted master problem* (RMP) has thus the following form, where the dual variables are shown in parentheses on the right-hand side of each constraint:

$$\min \sum_{(i,j) \in A_+} \sum_{k \in \tilde{K}_{ij}} c_{ij} x_{ij}^k + \sum_{(i,j) \in A_+} f_{ij} y_{ij}, \quad (7)$$

$$\sum_{j \in \tilde{N}_i^k(+)} x_{ij}^k - \sum_{j \in \tilde{N}_i^k(-)} x_{ji}^k = \begin{cases} d^k, & \text{if } i = O(k), \\ -d^k, & \text{if } i = D(k), \\ 0, & \text{otherwise,} \end{cases} \quad i \in N, k \in K, \quad (\pi_i^k) \quad (8)$$

$$\sum_{k \in \tilde{K}_{ij}} x_{ij}^k \leq u_{ij} y_{ij}, \quad (i, j) \in A_+, \quad (\alpha_{ij}) \quad (9)$$

$$x_{ij}^k \leq d^k y_{ij}, \quad (i, j) \in A_+, k \in \bar{K}_{ij} \subseteq \tilde{K}_{ij}, \quad (\beta_{ij}^k) \quad (10)$$

$$y_{ij} \leq 1, \quad (i, j) \in A_+, \quad (\gamma_{ij}) \quad (11)$$

$$x_{ij}^k \geq 0, \quad (i, j) \in A_+, k \in \tilde{K}_{ij}, \quad (12)$$

$$y_{ij} \geq 0, \quad (i, j) \in A_+. \quad (13)$$

At the root node of the B&P&C tree, only the variables associated with the artificial arcs are handled by the initial RMP. After solving the RMP at each column generation iteration, flow variables with negative reduced cost are added (see Section 2.2) and the new RMP is solved. This column generation process ends when there are no more flow variables with negative reduced cost. At this point, the cut generation step is performed (see Section 2.3). If cuts are added, the column generation process restarts, with an initial solution derived by the dual simplex method. At subsequent nodes of the B&P&C tree, an initial solution is also derived by the dual simplex method, considering the addition of the branching constraint (see Section 2.4) to the currently available basic solution.

2.2 Pricing Subproblem

To add flow variables with negative reduced cost to the RMP, we need to seek arcs $(i, j) \in A$ and commodities $k \notin \tilde{K}_{ij}$ such that $c_{ij} - \pi_i^k + \pi_j^k + \alpha_{ij} + \beta_{ij}^k < 0$. We note that the values of the dual variables are known after solving the RMP, with the exception of the β_{ij}^k associated with the strong inequalities (10). One simple way of solving this issue is to simply forget about them and generate all flow variables corresponding to arcs $(i, j) \in A$ and commodities $k \notin \tilde{K}_{ij}$ such that $c_{ij} - \pi_i^k + \pi_j^k + \alpha_{ij} < 0$. The resulting column generation method would still be valid, but after every pricing iteration, potentially many flow variables with nonnegative reduced cost could be generated, i.e., those associated with $(i, j) \in A$ and $k \notin \tilde{K}_{ij}$ such that $c_{ij} - \pi_i^k + \pi_j^k + \alpha_{ij} < 0$ and $\beta_{ij}^k \geq -(c_{ij} - \pi_i^k + \pi_j^k + \alpha_{ij})$. A more efficient approach consists in computing the values of the β_{ij}^k variables in order to add to the RMP only flow variables with negative reduced cost.

To achieve this objective, we first write down the dual and the complementary slackness conditions of the LP relaxation of model (1)-(6):

$$\max \sum_{k \in K} d^k (\pi_{O(k)}^k - \pi_{D(k)}^k) - \sum_{(i,j) \in A} \gamma_{ij} \quad (14)$$

$$\pi_i^k - \pi_j^k - \alpha_{ij} - \beta_{ij}^k \leq c_{ij}, \quad (i, j) \in A, k \in K, \quad (15)$$

$$u_{ij} \alpha_{ij} + \sum_{k \in K} d^k \beta_{ij}^k - \gamma_{ij} \leq f_{ij}, \quad (i, j) \in A, \quad (16)$$

$$\alpha_{ij} \geq 0, \quad (i, j) \in A, \quad (17)$$

$$\beta_{ij}^k \geq 0, \quad (i, j) \in A, k \in K, \quad (18)$$

$$\gamma_{ij} \geq 0, \quad (i, j) \in A, \quad (19)$$

$$x_{ij}^k (c_{ij} - \pi_i^k + \pi_j^k + \alpha_{ij} + \beta_{ij}^k) = 0, \quad (i, j) \in A, k \in K \quad (20)$$

$$y_{ij} \left(f_{ij} + \gamma_{ij} - u_{ij} \alpha_{ij} - \sum_{k \in K} d^k \beta_{ij}^k \right) = 0, \quad (i, j) \in A, \quad (21)$$

$$\alpha_{ij} \left(u_{ij} y_{ij} - \sum_{k \in K} x_{ij}^k \right) = 0, \quad (i, j) \in A, \quad (22)$$

$$\beta_{ij}^k (d^k y_{ij} - x_{ij}^k) = 0, \quad (i, j) \in A, k \in K, \quad (23)$$

$$\gamma_{ij} (1 - y_{ij}) = 0, \quad (i, j) \in A. \quad (24)$$

The analysis is based on deriving reduced cost optimality conditions from (14)-(24), for any arc $(i, j) \in A$ and for all commodities $k \in K$. If $k \in \tilde{K}_{ij}$, these conditions are automatically satisfied, since the RMP is solved to optimality. Note that, for $k \in \tilde{K}_{ij} \setminus \bar{K}_{ij}$, the corresponding strong inequality is not in the RMP, which implies that $\beta_{ij}^k \equiv 0$, but the corresponding flow variable x_{ij}^k is already in the RMP, so we do not need to consider this case. The goal is to add to the RMP only the flow variables that correspond to the commodities $k \notin \tilde{K}_{ij}$ that do not satisfy the reduced cost optimality conditions. We

denote by \bar{y}, \bar{x} and $\bar{\pi}, \bar{\alpha}, \bar{\beta}, \bar{\gamma}$, respectively, the optimal solutions to the RMP and to its dual; the optimal solution to the RMP is completed in the obvious way by setting $\bar{x}_{ij}^k = 0$, for any $(i, j) \in A$ and all $k \notin \tilde{K}_{ij}$. For any arc (i, j) , we distinguish two cases:

- $\bar{y}_{ij} > 0$. Let $k \notin \tilde{K}_{ij}$; for the solution to the RMP to be optimal for the LP relaxation of model (1)-(6), we must have, by complementary slackness condition (23):

$$\bar{\beta}_{ij}^k (\underbrace{d^k \bar{y}_{ij}}_{>0} - \underbrace{\bar{x}_{ij}^k}_{=0}) = 0 \implies \bar{\beta}_{ij}^k = 0.$$

The reduced cost optimality condition for $k \notin \tilde{K}_{ij}$ is therefore

$$c_{ij} - \bar{\pi}_i^k + \bar{\pi}_j^k + \bar{\alpha}_{ij} \geq 0,$$

which implies that we add to the RMP the flow variables associated with arc (i, j) and commodity $k \notin \tilde{K}_{ij}$ such that

$$c_{ij} - \bar{\pi}_i^k + \bar{\pi}_j^k + \bar{\alpha}_{ij} < 0.$$

- $\bar{y}_{ij} = 0$. In this case, we have $\bar{x}_{ij}^k = 0$ for all $k \in K$ and

$$\bar{\gamma}_{ij} (\underbrace{1 - \bar{y}_{ij}}_{\neq 0}) = 0 \implies \bar{\gamma}_{ij} = 0.$$

For the solution to the RMP to be optimal for the LP relaxation of model (1)-(6), we therefore must have, by the dual constraints (16), the following reduced cost optimality conditions:

$$f_{ij} - u_{ij} \bar{\alpha}_{ij} \geq \sum_{k \in K} d^k \bar{\beta}_{ij}^k.$$

By the dual constraints (15) and (18), we must have $\bar{\beta}_{ij}^k \geq \max\{0, -(c_{ij} - \bar{\pi}_i^k + \bar{\pi}_j^k + \bar{\alpha}_{ij})\}$ for any $k \in K$. This implies that we add to the RMP the flow variables associated with arc (i, j) such that

$$f_{ij} - u_{ij} \bar{\alpha}_{ij} < \sum_{k \in K} d^k \max\{0, -(c_{ij} - \bar{\pi}_i^k + \bar{\pi}_j^k + \bar{\alpha}_{ij})\},$$

but only for commodities $k \notin \tilde{K}_{ij}$ such that

$$c_{ij} - \bar{\pi}_i^k + \bar{\pi}_j^k + \bar{\alpha}_{ij} < 0.$$

To summarize, the pricing subproblem consists in performing the following tests for all arcs (i, j) :

1. If $\bar{y}_{ij} > 0$, then for any $k \notin \tilde{K}_{ij}$ such that $c_{ij} - \bar{\pi}_i^k + \bar{\pi}_j^k + \bar{\alpha}_{ij} < 0$, we add the flow variables x_{ij}^k to the RMP.
2. If $\bar{y}_{ij} = 0$ and $f_{ij} - u_{ij} \bar{\alpha}_{ij} < \sum_{k \in K} d^k \max\{0, -(c_{ij} - \bar{\pi}_i^k + \bar{\pi}_j^k + \bar{\alpha}_{ij})\}$, then for any $k \notin \tilde{K}_{ij}$ such that $c_{ij} - \bar{\pi}_i^k + \bar{\pi}_j^k + \bar{\alpha}_{ij} < 0$, we add the flow variables x_{ij}^k to the RMP.

2.3 Cut Generation

To solve the RMP at each column generation iteration, one option is to add *a priori* all the strong inequalities (10). However, as mentioned in the Introduction, it is more efficient to add them in a dynamic way. In particular, the separation problem is trivial: for each arc (i, j) , we generate the strong inequalities for all commodities $k \in \tilde{K}_{ij}$ such that $\bar{x}_{ij}^k > d^k \bar{y}_{ij}$. The addition of the corresponding constraints makes the current primal basic solution infeasible, so the dual simplex method is performed until a feasible primal solution is obtained. The column generation procedure is then restarted. Cut generation is performed after the column generation procedure, when no more flow variables with negative reduced cost can be found.

2.4 Variable Fixing and Branching

At any node of the B&P&C algorithm, we apply reduced cost variable fixing at each iteration of the cut generation procedure. The technique is well-known: given the current LP relaxation lower bound Z^l and the best known upper bound Z^* , we test, for each arc (i, j) , if $Z^l + |\bar{f}_{ij}| \geq Z^*$, in which case we can fix y_{ij} to \bar{y}_{ij} , where $\bar{f}_{ij} = f_{ij} - u_{ij}\bar{\alpha}_{ij} - \sum_{k \in K} d^k \bar{\beta}_{ij}^k$.

To show the validity of this test, first note that if $0 < \bar{y}_{ij} < 1$, then $\bar{f}_{ij} = 0$ by complementary slackness conditions (21) and (24), which implies that $Z^l + |\bar{f}_{ij}| = Z^l < Z^*$, otherwise the current node could be fathomed. Then, there are two remaining cases:

- $\bar{y}_{ij} = 0$. By complementary slackness condition (24), we have $\bar{\gamma}_{ij} = 0$. Assume we add the constraint $y_{ij} \geq 1$ with which we associate a dual variable $\lambda_{ij} \geq 0$ that is added to the objective $(+\lambda_{ij})$. The dual constraint (16) is then rewritten as $\lambda_{ij} \leq f_{ij} - u_{ij}\alpha_{ij} - \sum_{k \in K} d^k \beta_{ij}^k$. A dual feasible solution is obtained by setting $\lambda_{ij} = \bar{f}_{ij}$, which implies that $Z^l + \bar{f}_{ij}$ is a lower bound on the problem obtained by adding the constraint $y_{ij} \geq 1$. Therefore, if $Z^l + \bar{f}_{ij} = Z^l + |\bar{f}_{ij}| \geq Z^*$, we necessarily have $y_{ij} < 1$, i.e., $y_{ij} = 0$.
- $\bar{y}_{ij} = 1$. By the objective of the dual (14) and the dual constraints (16) and (19), we have $\bar{\gamma}_{ij} = \max\{0, -\bar{f}_{ij}\}$. Assume we add the constraint $y_{ij} \leq 0$ with which we associate a dual variable $\mu_{ij} \geq 0$. The dual constraint (16) is then rewritten as $-\mu_{ij} - \gamma_{ij} \leq f_{ij} - u_{ij}\alpha_{ij} - \sum_{k \in K} d^k \beta_{ij}^k$. A dual feasible solution is obtained by setting $\mu_{ij} = -\bar{f}_{ij}$ and $\gamma_{ij} = 0$, which implies that $Z^l - \bar{f}_{ij}$ is a lower bound on the problem obtained by adding the constraint $y_{ij} \leq 0$. Therefore, if $Z^l - \bar{f}_{ij} = Z^l + |\bar{f}_{ij}| \geq Z^*$, we necessarily have $y_{ij} > 0$, i.e., $y_{ij} = 1$.

The column-and-row generation procedure is performed until the node is fathomed or no more cuts can be generated. In the latter case, branching is performed. At each node, the LP relaxation of model (1)-(6), with the addition of the variable fixing and branching constraints, is solved by the column-and-row generation procedure. Branching occurs only when $Z^l < Z^*$, a condition that implies that the LP optimal solution is

fractional, i.e., there exists at least one arc (i, j) such that $0 < \bar{y}_{ij} < 1$ (otherwise, the LP optimal solution is feasible and if its value, Z^l , is less than Z^* , it replaces it as the best incumbent value). In this case, we select one such arc and create the two nodes with the added branching constraints $y_{ij} = 0$ and $y_{ij} = 1$. To select the branching variable, we use *reliability branching* (as implemented in the SCIP library [1]), certainly one of the most efficient branching rule available in general-purpose LP-based B&B MIP solvers (our preliminary tests have confirmed this assessment; for more details on LP-based B&B branching rules, see [2]).

2.5 Summary of the Algorithm

The algorithm is outlined as follows (the steps are commented below):

1. Initialize the upper bound Z^* and the incumbent solution.
2. Initialize the node pool \mathcal{L} with the root node.
3. *Selection*: Select the next node to evaluate in \mathcal{L} and remove it from \mathcal{L} .
4. *Lower bound*: Solve the LP relaxation at the current node:
 - (a) Find a pair of primal-dual solutions \bar{y}, \bar{x} and $\bar{\pi}, \bar{\alpha}$ by solving the RMP with the dual simplex method.
Column generation:
 - (b) *Pricing*: For each arc (i, j) , if $\bar{y}_{ij} > 0$ or $(\bar{y}_{ij} = 0$ and $f_{ij} - u_{ij}\bar{\alpha}_{ij} < \sum_{k \in K} d^k \max\{0, -(c_{ij} - \bar{\pi}_i^k + \bar{\pi}_j^k + \bar{\alpha}_{ij})\})$, then for any $k \notin \tilde{K}_{ij}$ such that $c_{ij} - \bar{\pi}_i^k + \bar{\pi}_j^k + \bar{\alpha}_{ij} < 0$, add the flow variables x_{ij}^k to the RMP.
 - (c) If flow variables were added to the RMP, solve the new RMP by the primal simplex method to obtain a pair of primal-dual solutions \bar{y}, \bar{x} and $\bar{\pi}, \bar{\alpha}$; go to step 4b.
 - (d) Let Z^l be the lower bound obtained after the column generation procedure; if \bar{y} is integer and $Z^l < Z^*$, then let $Z^* = Z^l$ and store \bar{y}, \bar{x} as the new incumbent solution.
 - (e) If $Z^l \geq Z^*$, then go to step 6.
Cut generation:
 - (f) *Separation*: For each arc (i, j) and for any $k \in \tilde{K}_{ij}$ such that $\bar{x}_{ij}^k > d^k \bar{y}_{ij}$, add the corresponding strong inequalities to the RMP.
 - (g) If strong inequalities were added to the RMP:
 - i. Solve the new RMP by the dual simplex method to obtain a solution \bar{y}, \bar{x} and a lower bound Z^l .
 - ii. If \bar{y} is integer and $Z^l < Z^*$, then let $Z^* = Z^l$ and store \bar{y}, \bar{x} as the new incumbent solution.

- iii. If $Z^l \geq Z^*$, then go to step 6.
 - iv. *Variable fixing*: For each arc (i, j) , if $Z^l + |\bar{f}_{ij}| \geq Z^*$, then fix y_{ij} to \bar{y}_{ij} ; if some variables were fixed, go to step 4a.
 - v. Go to step 4b.
5. *Branching*: If $Z^l < Z^*$, perform branching to generate two child nodes inserted in \mathcal{L} .
6. If $\mathcal{L} = \emptyset$, stop the algorithm; otherwise, go to step 3.

In step 1, a feasible solution and an upper bound is obtained by any heuristic method to solve the MCND; as mentioned in the Introduction, there are several effective heuristic approaches that identify near-optimal solutions, even for large-scale instances. In step 2, the pool of B&P&C nodes, noted \mathcal{L} , is initialized with the root node. In step 3, the next node to evaluate is selected and removed from \mathcal{L} ; from our computational results, we recommend the best-first selection rule (see Section 3). Step 4 contains the details of the column-and-row generation procedure used at each node. The goal of step 4a is to find an initial pair of primal-dual solutions to start the column generation procedure; at the very beginning, artificial arcs are added and flow is sent along these arcs, while subsequently, the dual simplex method is performed to take into account the effect of the additional constraints derived by variable fixing and branching. Step 4b implements the solution of the pricing subproblem to add flow variables with negative reduced cost to the RMP. In case new flow variables have been added to the RMP, the next column generation iteration proceeds; otherwise, the column generation procedure has converged, which allows to identify a lower bound Z^l . Steps 4d and 4e are the incumbent update and the lower bound test, respectively, to be performed after the column generation procedure. The separation problem for strong inequalities is then solved in step 4f. If strong inequalities have been generated, the RMP with the added cuts is solved by the dual simplex method in step 4(g)i, followed by the incumbent update and the lower bound test in steps 4(g)ii and 4(g)iii, then by variable fixing in step 4(g)iv, after which the control is transferred to the column generation procedure. In case no more strong inequalities have been generated, the column-and-row generation procedure stops. Finally, step 5 performs the branching operation and step 6 verifies the termination of the B&P&C algorithm, i.e., the node pool \mathcal{L} is empty; we can also stop the algorithm when a time limit has been reached.

2.6 Relationships to Other Decomposition Methods

In this section, we discuss the relationships between our B&P&C algorithm for the MCND and other decomposition methods, both specific to the MCND and general ones. We first consider the Lagrangian relaxation with respect to the flow conservation equations (2) and the capacity constraints (3), where we associate with each set of constraints, Lagrange multipliers π and $\alpha \geq 0$, respectively. The corresponding Lagrangian subproblem

decomposes by arc and can be written as follows, for each arc (i, j) :

$$Z_{ij}(\pi, \alpha) = \min \sum_{k \in K} (c_{ij} - \pi_i^k + \pi_j^k + \alpha_{ij}) x_{ij}^k + (f_{ij} - u_{ij} \alpha_{ij}) y_{ij} \quad (25)$$

$$0 \leq x_{ij}^k \leq d^k y_{ij}, \quad k \in K, \quad (26)$$

$$y_{ij} \in \{0, 1\}. \quad (27)$$

The associated Lagrangian dual is:

$$\max_{\pi, \alpha \geq 0} \sum_{k \in K} d^k (\pi_{O(k)}^k - \pi_{D(k)}^k) + \sum_{(i,j) \in A} Z_{ij}(\pi, \alpha). \quad (28)$$

To solve the Lagrangian subproblem, we consider the two possibilities, $y_{ij} = 0$ and $y_{ij} = 1$. If $y_{ij} = 0$, the only feasible solution is to set all flow variables to 0, with an objective function value equal to 0. If $y_{ij} = 1$, an optimal solution is obtained by setting $x_{ij}^k = d^k$, if $c_{ij} - \pi_i^k + \pi_j^k + \alpha_{ij} < 0$, and $x_{ij}^k = 0$, otherwise; the corresponding objective function value is $\sum_{k \in K} -\max\{0, -(c_{ij} - \pi_i^k + \pi_j^k + \alpha_{ij})\} d^k + (f_{ij} - u_{ij} \alpha_{ij})$. The optimal solution is obtained by selecting the alternative with the smallest objective function value; thus, the structure of any optimal solution \tilde{y}, \tilde{x} to the Lagrangian subproblem for arc (i, j) is:

$$\tilde{y}_{ij} = \begin{cases} 1, & \text{if } \sum_{k \in K} -\max\{0, -(c_{ij} - \pi_i^k + \pi_j^k + \alpha_{ij})\} d^k + (f_{ij} - u_{ij} \alpha_{ij}) < 0, \\ 0, & \text{if } \sum_{k \in K} -\max\{0, -(c_{ij} - \pi_i^k + \pi_j^k + \alpha_{ij})\} d^k + (f_{ij} - u_{ij} \alpha_{ij}) > 0, \\ 0 \text{ or } 1, & \text{otherwise,} \end{cases} \quad (29)$$

$$\tilde{x}_{ij}^k = \begin{cases} d^k, & \text{if } \tilde{y}_{ij} = 1 \text{ and } c_{ij} - \pi_i^k + \pi_j^k + \alpha_{ij} < 0, \\ 0, & \text{if } \tilde{y}_{ij} = 0 \text{ or } (\tilde{y}_{ij} = 1 \text{ and } c_{ij} - \pi_i^k + \pi_j^k + \alpha_{ij} > 0), \\ 0 \text{ or } d^k, & \text{otherwise,} \end{cases} \quad k \in K. \quad (30)$$

We claim that the solution to the pricing subproblem presented in Section 2.2 corresponds to one of these optimal solutions for $\pi = \bar{\pi}$ and $\alpha = \bar{\alpha}$, in the sense that a flow variable x_{ij}^k is added to the RMP when this optimal solution satisfies $\tilde{x}_{ij}^k = d^k$.

Indeed, consider first the case $\bar{y}_{ij} > 0$. We then have, for any $k \in K$, $\bar{\beta}_{ij}^k = -(c_{ij} - \bar{\pi}_i^k + \bar{\pi}_j^k + \bar{\alpha}_{ij}) \geq 0$ if $\bar{x}_{ij}^k > 0$, or $\bar{\beta}_{ij}^k = 0$ if $\bar{x}_{ij}^k = 0$, by complementary slackness conditions (20) and (23); this implies that $\bar{\beta}_{ij}^k = \max\{0, -(c_{ij} - \bar{\pi}_i^k + \bar{\pi}_j^k + \bar{\alpha}_{ij})\}$. By complementary slackness condition (21), it follows that

$$\begin{aligned} f_{ij} - u_{ij} \bar{\alpha}_{ij} &= \sum_{k \in K} d^k \bar{\beta}_{ij}^k - \bar{\gamma}_{ij} \\ &= \sum_{k \in K} d^k \max\{0, -(c_{ij} - \bar{\pi}_i^k + \bar{\pi}_j^k + \bar{\alpha}_{ij})\} - \bar{\gamma}_{ij} \\ &\leq \sum_{k \in K} d^k \max\{0, -(c_{ij} - \bar{\pi}_i^k + \bar{\pi}_j^k + \bar{\alpha}_{ij})\}. \end{aligned}$$

Therefore, when $\bar{y}_{ij} > 0$, an optimal solution to the corresponding Lagrangian subproblem with $\pi = \bar{\pi}$ and $\alpha = \bar{\alpha}$ is $\tilde{y}_{ij} = 1$ and, for any $k \in K$, $\tilde{x}_{ij}^k = d^k$, if $c_{ij} - \bar{\pi}_i^k + \bar{\pi}_j^k + \bar{\alpha}_{ij} < 0$, and 0, otherwise.

Now, consider the case $\bar{y}_{ij} = 0$; an optimal solution to the corresponding Lagrangian subproblem with $\pi = \bar{\pi}$ and $\alpha = \bar{\alpha}$ is 1) $\tilde{y}_{ij} = 1$, if $(f_{ij} - u_{ij}\bar{\alpha}_{ij}) < \sum_{k \in K} d^k \max\{0, -(c_{ij} - \bar{\pi}_i^k + \bar{\pi}_j^k + \bar{\alpha}_{ij})\}$ and, for any $k \in K$, $\tilde{x}_{ij}^k = d^k$, if $c_{ij} - \bar{\pi}_i^k + \bar{\pi}_j^k + \bar{\alpha}_{ij} < 0$, and 0, otherwise; or 2) $\tilde{y}_{ij} = 0$, if $(f_{ij} - u_{ij}\bar{\alpha}_{ij}) \geq \sum_{k \in K} d^k \max\{0, -(c_{ij} - \bar{\pi}_i^k + \bar{\pi}_j^k + \bar{\alpha}_{ij})\}$ and $\tilde{x}_{ij}^k = 0$, for all $k \in K$.

In both cases, we conclude that the pricing subproblem is equivalent to solving the Lagrangian subproblem with $\pi = \bar{\pi}$ and $\alpha = \bar{\alpha}$. Thus, the column-and-row generation procedure can be interpreted as a solution method for the Lagrangian dual (28), where the Lagrange multipliers are obtained by solving the restricted master problem (7)-(13). While this approach is similar to Dantzig-Wolfe (DW) decomposition, it is fundamentally different, in that DW decomposition is based on a reformulation of the problem in terms of the (*exponentially many*) extreme points of the convex hull of the feasible solution set to the Lagrangian subproblem. Our column-and-row generation procedure is in fact a special case of the *structured DW (SDW)* decomposition framework presented in [13], which encompasses DW decomposition by allowing other reformulations than the standard DW one. The network design example used to illustrate the SDW decomposition in [13] is based on reformulating general integer variables with binary variables, thus increasing the size of the variable space from a polynomial number to a pseudo-polynomial number. In contrast, our application of SDW to the MCND is working on the same formulation as the original one, i.e., the “reformulation” is the original formulation itself, which has a polynomial number of variables.

Our column-and-row generation procedure can also be seen as a special case of the framework presented in [24] for solving large-scale LPs with so-called *column-dependent-rows*, which are linking constraints (between two types of variables) that are “too many” to be included in the model directly or that can only be known when all variables are explicitly generated. Clearly, the strong inequalities (4) correspond to the first category of column-dependent-rows, since they are linking the flow and the design variables, and they are “too many”, although not exponentially many, as in most applications presented in [24].

The Lagrangian interpretation of our column-and-row generation procedure allows us to clarify its relationships to existing Lagrangian methods for the MCND. These methods are based on two different Lagrangian relaxations: the so-called *shortest path* (or flow) and *knapsack* relaxations. The first one relaxes both the capacity constraints (3) and the strong inequalities (4), leading to a Lagrangian subproblem that decomposes into a set of shortest path problems for each commodity and a problem in y variables solvable by inspection [7, 8, 14, 15]. The second Lagrangian approach relaxes the flow conservation equations (2); the corresponding Lagrangian subproblem decomposes by arc and, for each arc, it consists in solving a continuous knapsack problem [7, 8, 15, 19, 22, 26]. In all these references, the Lagrangian dual is solved either by a subgradient approach or by a bundle method, which is a generalized form of DW decomposition, where the master problem is “stabilized” with the addition of a convex (usually quadratic) term to the objective function [11]. Thus, our column-and-row generation procedure differs from the existing Lagrangian methods in two aspects: first, both the flow conservation equations (2) and the capacity constraints (3) are relaxed, so the resulting Lagrangian

subproblem differs significantly from the shortest path relaxation, while it is similar to the knapsack relaxation, but simpler, since it does not contain the capacity constraints; second, the Lagrangian dual is solved by the column-and-row generation procedure, not by subgradient or bundle methods. It is noteworthy that the master problem used in bundle methods can be seen as a generalized form of the standard DW reformulation, with exponentially many variables, while the master problem in our approach is nothing but the original, compact, formulation.

Column(-and-row) generation methods for multicommodity flow problems are often identified with the classical path-based model, where flow variables are represented by paths between each O-D pair. In fact, for the MCND, the DW master problem associated with the shortest path relaxation corresponds to the path-based model; thus, solving the associated Lagrangian dual by a bundle method implicitly makes use of the path-based formulation. The path-based model has also been used for solving the multicommodity flow subproblems derived from fixing the design variables [22, 26], but also in sophisticated heuristic methods [6, 21]. In particular, the column-and-row generation procedure embedded in the capacity scaling approach of [21] is similar to ours, in that it uses a cutting-plane procedure to generate strong inequalities, but is also different, since the flow variables are represented by paths, not by arcs as in our approach.

3 Computational Results

This section presents computational results obtained by the B&P&C algorithm on a publicly available set of 196 instances (the so-called “Canad” instances, see [12]) used in several papers on the *MCND* (for instance [16, 18, 22]) and described in detail in [8]. These problem instances consist of general transshipment networks with one commodity per origin-destination and no parallel arcs. Associated with each arc are three positive quantities: the capacity, the transportation cost and the fixed design cost. These instances are characterized by various degrees of capacity tightness, with regard to the total demand, and importance of the fixed design cost, with respect to the transportation cost.

The instances are divided into three classes. Class I (the “C” instances in [12]) consists of 31 problem instances with many commodities compared to the number of nodes, while Class II (the “C+” instances in [12]) contains 12 problem instances with few commodities compared to the number of nodes. Class III (the “R” instances in [12]) is divided into two categories, A and B, each containing nine sets of nine problem instances each. Each set is characterized by the numbers of nodes, arcs, and commodities, which are the same for the nine instances, and by instance-specific levels of capacity tightness and importance of the fixed design cost. Class III-A (instances “R01” to “R09”) contains 72 small size problem instances with 10 nodes (nine infeasible instances have been discarded), while Class III-B (instances “R10” to “R18”) contains 81 medium to large size instances with 20 nodes. Table 1 gives the size of the instances in each class.

The B&P&C method is compared against two competitors: a state-of-the-art MIP solver, CPLEX (version 12.3), and the B&C algorithm obtained by performing, at each

Class I (31)	Class II (12)	Class III-A (72)	Class III-B (81)
$ N , A , K $	$ N , A , K $	$ N , A , K $	$ N , A , K $
20,230,40 (3)	25,100,10 (3)	10,35,10 (6)	20,120,40 (9)
20,230,200 (4)	25,100,30 (3)	10,35,25 (6)	20,120,100 (9)
20,300,40 (4)	100,400,10 (3)	10,35,50 (6)	20,120,200 (9)
20,300,200 (4)	100,400,30 (3)	10,60,10 (9)	20,220,40 (9)
30,520,100 (4)		10,60,25 (9)	20,220,100 (9)
30,520,400 (4)		10,60,50 (9)	20,220,200 (9)
30,700,100 (4)		10,85,10 (9)	20,320,40 (9)
30,700,400 (4)		10,85,25 (9)	20,320,100 (9)
		10,85,50 (9)	20,320,200 (9)

Table 1: Classes and problem dimensions (number of instances in parentheses)

node of the B&C tree, a cutting-plane procedure that includes only the strong inequalities as cuts. Each competing method is tested to verify two hypotheses: 1) the B&P&C algorithm is competitive with a state-of-the-art MIP solver; 2) the column-and-row generation procedure is at least as efficient as the cutting-plane procedure, and more efficient on large-scale instances with many commodities (in the order of 100). The initial model given to CPLEX is the weak formulation, i.e., model (1)-(6) without the strong inequalities (4), since including these inequalities in advance would give very poor performance. CPLEX then generates its own cuts, which include, not surprisingly, flow cover inequalities, but also other types of cuts. Both the B&C and the B&P&C algorithms are implemented in C++ using the SCIP library [1]; the LP solver used in both algorithms is also CPLEX (version 12.3). In both B&C and B&P&C, we use reliability branching, as implemented in SCIP, while CPLEX uses its own default branching rule. All experiments are performed on an Intel Xeon X5660 operating at 2.80 GHz under the Linux operating system.

The goal of our experiments is to evaluate the performance of the B&P&C algorithm in providing effective lower bounds; in particular, we wish to evaluate the efficiency of the algorithm for proving the optimality of an already known optimal solution. Thus, for both the B&P&C method and its competitors, we give as initial incumbent the best known feasible solution for each problem instance, which is an optimal one for most of them. When performing CPLEX, we consequently deactivate all features related to the computation of feasible solutions. For all methods, we fix a “reasonable” CPU time limit of 3 hours.

We first look at the results obtained at the root node. We expect CPLEX to deliver better lower bounds than the strong relaxation bound computed by B&C and B&P&C, because CPLEX can generate a large number of cuts of different types. This is indeed the case, but it is interesting to assess the quality of the strong relaxation for various types of instances; thus, for each problem instance, we measure the gap between the strong relaxation bound and CPLEX lower bound. We are also interested in assessing

Class	Method	Gap (%)	CPU (s)	CPU gain (%)
I	CPLEX	0	26.2	0
	B&C	0.2	23.3	11
	B&P&C	0.2	6.3	76
II	CPLEX	0	2.3	0
	B&C	2.5	0.6	74
	B&P&C	2.5	0.6	74
III-A	CPLEX	0	0.09	0
	B&C	0.9	0.02	77
	B&P&C	0.9	0.02	77
III-B	CPLEX	0	6.3	0
	B&C	0.2	5.4	14
	B&P&C	0.2	2.6	59

Table 2: Comparison of the methods at the root node

the efficiency of the three bounding methods, the cutting-plane procedure in CPLEX, the cutting-plane approach that generates only the strong inequalities and our column-and-row generation method; we therefore measure the CPU times obtained by the three methods. We expect the last two approaches to be faster than CPLEX, which is indeed the case. The average results for each class of problem instances are presented in Table 2. Column “Class” corresponds to the class of problem instances; column “Method” relates to the three methods, CPLEX, B&C and B&P&C; column “Gap” provides the average percentage gap between the strong relaxation bound and CPLEX lower bound, which is the best of the two for all problem instances; column “CPU” gives the average CPU time in seconds for each method; finally, column “CPU gain” displays the average relative gain in CPU time obtained by B&C and B&P&C over CPLEX.

These results show that the strong relaxation bound is very close (within 1% on average) to the lower bound computed by CPLEX for problem instances in Classes I, III-A and III-B. For Class II, the average gap between the strong relaxation bound and CPLEX attains 2.5%. This is consistent with the results reported in [3], who show that other inequalities than the strong ones, namely cover and flow cover inequalities, are really useful only for instances in Class II, providing only minor improvements for instances in the other classes. Note that, while Class II instances have few commodities (no more than 30), most instances in Classes I and III-B have many commodities (at least 100). For these two classes, the strong relaxation bound and CPLEX lower bound are almost equal, with an average gap of 0.2%; the CPU times, however, are significantly different: while the average relative gain for B&C is relatively modest (slightly more than 10%), it reaches more than 60% for B&P&C. These results are promising for the B&P&C algorithm, since they suggest significant gains in efficiency for instances in Classes I and III-B, especially for large-scale instances with hundreds of commodities. We verify this hypothesis in the remainder of this section.

To analyze the performance of the different approaches when branching is performed and the enumeration trees are explored, we examine the impact of two classical selection rules, best-first and depth-first. In general, for any B&B algorithm, the best-first search selection rule provides reliable results, independently of the problem instance being tested, since it displays two interesting features: 1) when the algorithm proves the optimality of the solution within the time limit, the total number of nodes is generally smaller than with other selection rules; 2) otherwise, if the algorithm is stopped prematurely without proving the optimality of the solution, the lower bound is generally tighter than with other selection rules. The first advantage vanishes, however, under two conditions: a) the initial incumbent is an optimal solution; b) branching and bounding operations provide the same output, irrespective of the selection rule; indeed, under conditions a) and b), the total number of nodes is the same, independently of the selection rule [20]. Thus, when the initial incumbent is close to an optimal solution, as in our experiments, other selection rules become interesting alternatives to best-first search. In particular, depth-first search has one advantage, when compared to best-first search: it performs less backtracks, which implies faster reoptimization. Therefore, we have tested the two selection rules, best-first and depth-first, for the three approaches, CPLEX, B&C and B&P&C.

The performance analysis depends on the results obtained by the resulting six methods for each problem instance. Table 3 presents the results obtained for the “Easy” instances, which are those that are solved to optimality by the six methods, within the CPU time limit of 3 hours. Column “Class” corresponds to the class of problem instances; column “Method” relates to the six methods that result in combining CPLEX, B&C and B&P&C with depth-first and best-first; columns “CPU”, “Nodes” and “Variables” display for each method, respectively, the average CPU time in seconds, the total number of nodes and the number of variables in the models.

These results show that depth-first search performs better than best-first search when used with CPLEX, since the models generated by CPLEX do not vary with the selection rule, as CPLEX generates cuts only at the root node of the tree. The assumption that branching and bounding operations provide the same output, irrespective of the selection rule, is then verified, which is confirmed by the fact that the number of nodes is almost the same for the two selection rules. Because depth-first search performs less backtracks, reoptimization is faster and the CPU times are better for depth-first search. For B&C and B&C&P, the situation is different, since the formulations vary significantly depending on the selection rule. For B&C, the two selection rules give similar results, while for B&P&C, best-first search emerges as a clear winner.

For Class I instances, B&P&C-best is clearly better than any variant of CPLEX and B&C; even though the total number of nodes is larger for B&P&C-best, the small number of variables generated by the method (four times less than the others) more than compensates, translating into an overall algorithm that is about two times faster than the others, on average. For Class II instances, the situation is completely different: the significantly better lower bounds computed by CPLEX lead to an enumerative approach that is an order of magnitude faster than B&C and B&P&C. In addition, the reduction in the number of variables does not payoff this time, since B&P&C is slower than B&C. For

Class III-A instances, all methods are equally efficient with average CPU times around 2 seconds. Finally, for Class III-B instances, B&P&C-best is competitive with CPLEX: it is slightly slower than CPLEX-depth, but faster than CPLEX-best. As observed for Class I instances, the large number of nodes incurred by B&P&C is more than compensated by the small number of variables generated by the method. Overall, these results point to the conclusion that B&P&C performs well for large-scale problem instances with hundreds of commodities.

Class	Method	CPU (s)	Nodes	Variables
I (12)	CPLEX-depth	845	1487	28942
	CPLEX-best	1204	1487	28942
	B&C-depth	1017	1657	29031
	B&C-best	1013	1756	29031
	B&P&C-depth	474	2352	6692
	B&P&C-best	419	2062	6662
II (8)	CPLEX-depth	26	630	2676
	CPLEX-best	30	627	2676
	B&C-depth	327	6080	2693
	B&C-best	322	5767	2693
	B&P&C-depth	687	7966	1573
	B&P&C-best	496	5870	1557
III-A (72)	CPLEX-depth	2	213	1822
	CPLEX-best	2	215	1822
	B&C-depth	2	285	1850
	B&C-best	3	303	1850
	B&P&C-depth	2	421	745
	B&P&C-best	3	396	743
III-B (60)	CPLEX-depth	766	3516	19409
	CPLEX-best	1051	3507	19409
	B&C-depth	1114	15874	19510
	B&C-best	1213	15926	19510
	B&P&C-depth	901	21131	7739
	B&P&C-best	855	16829	7733

Table 3: Results on “Easy” instances (number of instances in parentheses)

To verify this hypothesis, we focus on the “Easy” instances with at least 100 commodities, which are all taken from Classes I and III-B. Table 4 presents the average results obtained on these instances; the columns have the same meanings as in Table 3. On these 41 instances, we see that the total number of nodes is slightly increased when B&P&C-best is performed, compared to CPLEX and B&C methods, but that the CPU time is better, due to the fact that the number of variables is about three times less.

Next, we look at the results obtained with the “Difficult” instances, which are those

Instances	Method	CPU (s)	Nodes	Variables
$ K \geq 100$ (41)	CPLEX-depth	1246	3150	29983
	CPLEX-best	1733	3020	29983
	B&C-depth	1580	2847	30128
	B&C-best	1669	2826	30128
	B&P&C-depth	1087	3722	11024
	B&P&C-best	1015	3317	11013

Table 4: Results on Large-Scale “Easy” instances (number of instances in parentheses)

that are not solved by any of the six methods, within the 3-hour CPU time limit. Table 5 presents the results obtained for these instances. Only the best-first search variants of the methods are presented, since depth-first search is not competitive in this case. Indeed, because the enumeration is stopped within the time limit without proving optimality, any selection rule that tends to explore the trees in a breadth-first manner, like best-first search does, would deliver final lower bounds that are tighter than those obtained by depth-first search. Thus, in column “Method”, we show only the three methods, CPLEX, B&C and B&P&C with best-first search; columns “Gap”, “Nodes” and “Variables” display for each method, respectively, the average gap in percentage between the lower and the upper bounds, the total number of nodes and the number of variables in the formulations.

For instances in Classes I and III-B, B&P&C-best provides slightly better lower bounds and final gaps, on average, than the other approaches. Indeed, within the 3-hour time limit, the method explores about twice the number of nodes than CPLEX does, which explains why the final gaps are smaller, in spite of the slightly less tight lower bounds computed at the root node. The small number of generated variables explains the efficiency of the B&P&C method: the number of variables is roughly divided by four for Class I instances and by two for Class III-B instances. For Class II instances, CPLEX-best outperforms the other methods because of the quality of its lower bounds, which is reflected in the final gaps. This is in spite of the fact that CPLEX-best explores much less nodes (an order of magnitude less) than the other methods. B&C-best is more efficient than B&P&C-best for these instances, since it is able to explore three times the number of nodes and to reduce the final gaps. Again, these results suggest that B&P&C performs well on instances with hundreds of commodities.

In Table 6, we focus on “Difficult” instances with at least 100 commodities, which are all taken from Classes I and III-B; the columns have the same meanings as in Table 5. On these 31 instances, we see that the final gap is slightly better when B&P&C-best is performed, compared to CPLEX-best and B&C-best, given that the total number of nodes explored within the 3-hour time limit is significantly increased with B&P&C-best: two times more than CPLEX-best and three times more than B&C-best. The number of variables in the models is roughly divided by three when B&P&C is used.

In addition to the “Easy” and “Difficult” instances, there are nine other instances

Class	Method	Gap (%)	Nodes	Variables
I (16)	CPLEX-best	0.9	6792	130563
	B&C-best	0.9	3545	130805
	B&P&C-best	0.7	14640	30364
II (2)	CPLEX-best	3.2	16517	8401
	B&C-best	7.2	429172	8420
	B&P&C-best	7.7	169185	6790
III-B (17)	CPLEX-best	1.3	9900	44966
	B&C-best	1.4	10280	45123
	B&P&C-best	1.3	18550	22353

Table 5: Results on “Difficult” instances (number of instances in parentheses)

Instances	Method	Gap (%)	Nodes	Variables
$ K \geq 100$ (31)	CPLEX-best	1.1	6065	91215
	B&C-best	1.2	4004	91424
	B&P&C-best	1.0	12235	27390

Table 6: Results on Large-Scale “Difficult” instances (number of instances in parentheses)

that are solved by at least one, but not all methods, within the 3-hour CPU time limit. Of these nine instances, CPLEX-depth is able to solve to optimality seven of them, B&C-depth two of them and B&P&C-depth three of them. Among these three instances solved by B&P&C-depth, two of them are large-scale Class I instances that cannot be solved by CPLEX-depth. These results confirm the efficiency of the B&P&C algorithm for solving large-scale instances with hundreds of commodities.

4 Conclusions

In this paper, we have presented a B&P&C algorithm for the multicommodity capacitated fixed-charge network design problem. The restricted master problem solved at each column generation iteration is obtained directly from the compact arc-based model by considering only a subset of the commodity flow variables. The pricing subproblem corresponds to a Lagrangian relaxation of the flow conservation and capacity constraints, leaving in the Lagrangian subproblem only the strong inequalities. A cut generation step based on strong inequalities is also performed. The resulting column-and-row generation procedure is embedded within an enumerative scheme, giving rise to the overall B&P&C algorithm. Our computational experiments show that the B&P&C performs well on large-scale instances with hundreds of commodities. On such instances, the B&P&C algorithm is generally more efficient than a state-of-the-art MIP solver and a B&C algorithm that does not incorporate column generation.

We have explored the relationships between our B&P&C algorithm and other decomposition methods for the problem. Although we have clearly shown the interest of our method by comparing it to a state-of-the-art MIP solver and a B&C algorithm, it would be interesting to perform an extensive analysis of the relative performance of the B&P&C algorithm when compared to other decomposition approaches, in particular Lagrangian-based methods. As mentioned in Section 2.6, our algorithm is a special case of the structured Dantzig-Wolfe framework, which has been applied to another multicommodity network design problem [13]. It would be interesting to extend the approach further to other network design formulations.

Acknowledgments

Funding for this project has been provided by the Natural Sciences and Engineering Council of Canada (NSERC). This support is gratefully acknowledged. We thank the SCIP team of the Zuse-Institut Berlin (ZIB) for their help in using the SCIP library, especially Tobias Achterberg and Stefan Vigerske.

References

- [1] T. Achterberg. Scip: solving constraint integer programs. *Mathematical Programming Computation*, 1:1–41, 2009.
- [2] T. Achterberg, T. Koch, and A. Martin. Branching rules revisited. *Operations Research Letters*, 33(1):42 – 54, 2005.
- [3] M. Chouman, T.G. Crainic, and B. Gendron. Commodity representations and cutset-based inequalities for multicommodity capacitated fixed charge network design. Technical Report CIRRELT-152, CIRRELT, 2011-56.
- [4] A.M. Costa, J.-F. Cordeau, B. Gendron, and G. Laporte. Accelerating benders decomposition with heuristic master problem solutions. *Pesquisa Operacional*, 32:3–19, 2012.
- [5] A.M. Costa, J.F. Cordeau, and B. Gendron. Benders, metric and cutset inequalities for multicommodity capacitated network design. *Computational Optimization and Applications*, 42:371–392, 2009.
- [6] T. G. Crainic, M. Gendreau, and J.M. Farvolden. A simplex-based tabu search method for capacitated network design. *INFORMS Journal on Computing*, 12:223–236, 2000.
- [7] T.G. Crainic, A. Frangioni, and B. Gendron. Multicommodity capacitated network design. In P. Soriano and B. Sanso, editors, *Telecommunications Network Planning*, pages 1–19. Kluwer Academic Publisher, 1999.

- [8] T.G. Crainic, A. Frangioni, and B. Gendron. Bundle-based relaxation methods for multicommodity capacitated fixed charge network design. *Discrete Applied Mathematics*, 112:73–99, 2001.
- [9] T.G. Crainic and M. Gendreau. Cooperative parallel tabu search for capacitated network design. *Journal of Heuristics*, 8:601–627, 2002.
- [10] T.G. Crainic, B. Gendron, and G. Hernu. A slope scaling/Lagrangean perturbation heuristic with long-term memory for multicommodity capacitated fixed-charge network design. *Journal of Heuristics*, 10:525–545, 2004.
- [11] A. Frangioni. About lagrangian methods in integer optimization. *Annals of Operations Research*, 139:163–193, 2005.
- [12] A. Frangioni. <http://www.di.unipi.it/~frangio>, September 2012.
- [13] A. Frangioni and B. Gendron. Satibilized structured dantzig-wolfe. *Mathematical Programming*, 2012.
- [14] A. Frangioni and E. Gorgone. Generalized bundle methods for sum-functions with “easy” components: Applications to multicommodity network design. Technical Report 11-08, Dipartimento di Informatica, Università di Pisa, 2011.
- [15] B. Gendron and T.G. Crainic. Relaxations for multicommodity capacitated network design problems. Technical report, Publication CRT-945, Centre de recherche sur les transports, Université de Montréal, 1994.
- [16] I. Ghamlouche, T.G. Crainic, and M. Gendreau. Cycle-based neighbourhoods for fixed charge capacitated multicommodity network design. *Operations Research*, 51:655–667, 2003.
- [17] I. Ghamlouche, T.G. Crainic, and M. Gendreau. Path relinking, cycle-based neighbourhoods and capacitated multicommodity network design. *Annals of Operations Research*, 131:109–133, 2004.
- [18] M. Hewitt, G.L. Nemhauser, and M.W.P. Savelsbergh. Combining exact and heuristic approaches for the capacitated fixed-charge network flow problem. *INFORMS Journal on Computing*, 22:314–325, 2010.
- [19] K. Holmberg and D. Yuan. A Lagrangian heuristic based branch-and-bound approach for the capacitated network design problem. *Operations Research*, 48:461–481, 2000.
- [20] T. Ibaraki. Enumerative approaches to combinatorial optimization. *Annals of Operations Research*, 10-11:1–602, 1988.
- [21] N. Katayama, M. Chen, and M. Kubo. A capacity scaling heuristic for the multicommodity capacitated network design problem. *Journal of Computational and Applied Mathematics*, 232:90–101, 2009.

- [22] G. Kliewer and L. Timajev. Relax-and-cut for capacitated network design. In *Proceedings of Algorithms-ESA 2005: 13th Annual European Symposium on Algorithms*, pages 47–58. Lecture Notes in Computer Science 3369, 2005.
- [23] T.L. Magnanti and R.T. Wong. Network design and transportation planning: models and algorithms. *Transportation Science*, 18:1–55, 1984.
- [24] I. Muter, S.I. Birbil, and K. Bülbül. Simultaneous column-and-row generation for large-scale linear programs with column-dependent-rows. *Mathematical Programming*, pages 1–36, 2012.
- [25] I. Rodríguez-Martín and J. J. Salazar-González. A local branching heuristic for the capacitated fixed-charge network design problem. *Computers & Operations Research*, 37:575–581, 2010.
- [26] M. Sellmann, G. Kliewer, and A. Koberstein. Lagrangian cardinality cuts and variable fixing for capacitated network design. In *Proceedings of Algorithms-ESA 2002: 10th Annual European Symposium on Algorithms*, pages 845–858. Lecture Notes in Computer Science 2461, 2002.