Computers & Operations Research 77 (2017) 44–57

Contents lists available at ScienceDirect

Computers & Operations Research

journal homepage: www.elsevier.com/locate/caor

A parallel local search framework for the Fixed-Charge Multicommodity Network Flow problem

Lluís-Miquel Munguía [a,*], Shabbir Ahmed [b], David A. Bader [a], George L. Nemhauser [b], Vikas Goel [c], Yufen Shao [c]

[a] College of Computing, Georgia Institute of Technology, Atlanta GA 30332, United States
[b] Industrial and Systems Engineering, Georgia Institute of Technology, Atlanta GA 30332, United States
[c] ExxonMobil Upstream Research Company, Houston, TX 77098, United States

ARTICLE INFO

Article history:
Received 8 October 2015
Received in revised form
5 July 2016
Accepted 20 July 2016
Available online 22 July 2016

Keywords:
FCMNF
Parallel computing
Primal heuristics
Discrete optimization
Multicommodity capacitated network design

ABSTRACT

We present a parallel local search approach for obtaining high quality solutions to the Fixed Charge Multicommodity Network Flow problem (FCMNF). The approach proceeds by improving a given feasible solution by solving restricted instances of the problem where flows of certain commodities are fixed to those in the solution while the other commodities are locally optimized. We derive multiple independent local search neighborhoods from an arc-based mixed integer programming (MIP) formulation of the problem which are explored in parallel. Our scalable parallel implementation takes advantage of the hybrid memory architecture in modern platforms and the effectiveness of MIP solvers in solving small problems instances. Computational experiments on FCMNF instances from the literature demonstrate the competitiveness of our approach against state of the art MIP solvers and other heuristic methods.

1. Introduction

The Fixed-Charge Multicommodity Network Flow (FCMNF) problem is a classic optimization problem arising in numerous applications. Given a directed capacitated network and a set of commodities, the objective is to route every commodity from its origin to destination through the network so as to minimize the total cost. The cost associated with an arc is the sum of a fixed cost derived from its use and a variable cost proportional to the flow going through it. The total cost is derived from the sum of all arc costs.

The FCMNF problem was proven to be NP-Hard [1]. In practice, realistic sized instances of the FCMNF problem are extremely difficult to solve to optimality. Consequently a variety of heuristic approaches and integer programming techniques have been developed and proven to be effective means to achieve high quality solutions quickly. In this paper, we introduce a local search heuristic framework for the FCMNF problem that is explicitly designed for both parallel shared-memory systems and distributed-memory systems. Our method finds competitive solutions by exploring a large number of local search neighborhoods concurrently. Given a feasible solution *s*, the local searches proceed by solving restricted instances of the problem where flows of certain commodities are fixed to those in the solution *s* while that of the other commodities are optimized. We take advantage of a state-of-the-art Mixed Integer Programming (MIP) solver to drive these local searches.

Recent works have introduced successful heuristic methods for obtaining high quality solutions. Most common heuristics consist of embedding a problem-specific mechanism for improving solutions in the context of a metaheuristic search framework. Ghamlouche et al. [2] identify cycles in the network as a heuristic strategy for finding alternative flow routes. The same methodology is used in further works in combination with machine learning techniques in order to improve and guide the local search [3]. Chouman et al. [4] use a similar approach to identify arc-balanced cycles in combination with a Tabu Search. A different heuristic approach is presented by Yaghini et al. [5], where the authors define local search neighborhoods based on simplex pivots in the context of a simulated annealing framework. Other meta-heuristic frameworks for the FCMNF problem are based on Evolutionary algorithms [6] and Scatter Search procedures [7–9]. The latter works were developed more than a decade ago and differ in the generation of the original population, and the mechanisms used for solution improvement and recombination.

Heuristic strategies can also be used in the context of an exhaustive search framework. An example is the local branching technique introduced by Fischetti et al. [10]. Their method uses linear inequalities to branch on smaller subproblems, which are

* Corresponding author.
E-mail address: lluis.munguia@gatech.edu (L.-M. Munguía).

http://dx.doi.org/10.1016/j.cor.2016.07.016
0305-0548/

solved by a black-box MIP solver. Examples of applications of local branching for the FCMNF problem are studied by Rodríguez-Martín et al. [11]. The efficacy of the previously cited work resides in the use of heuristics algorithms in combination with exact mixed-integer programming techniques.

Katayama et al. [12] develop a column generation, path-based formulation enhanced by strong inequalities in conjunction with an arc capacity scaling approach. In [13], the same scheme is improved by using local branching ideas to polish the solutions obtained through arc capacity scaling strategies.

Despite providing high quality solutions quickly, heuristic methods cannot provide optimality certificates because of their exclusive focus on primal solutions. Hewitt et al. [14] introduce an algorithm that provides lower bounds on the optimal solution in addition to primal solution improvements. Such improvements are found by solving strategically restricted MIP subproblems while tighter lower bounds are found with mathematical programming approaches. In further work, their approach combines the use of restricted MIPs in the context of a branch-and-price framework that also provides a performance guarantee upon completion [15]. The authors take advantage of parallelism to solve the pricing problems and restrictions.

Parallelizations of large neighborhood search algorithms have been successfully implemented in other applications such as the LNG inventory routing problem [16]. To our knowledge, parallel computing remains a relatively unexplored field for the FCMNF problem. Crainic et al. [17] propose an asynchronous parallel Tabu Search where every processor communicates with a centralized solution pool. They introduce and test several communication policies as well as strategies for handling the exchanged information. In [18], special emphasis is put on the control of the information diffusion between the different processors. The authors present a multilevel parallel local search algorithm that employs parallel cycle-based Tabu Searches defined by sets of fixed arcs. Their approach differs greatly from ours in many aspects. These include the solution improvement method used, the fact that our method has a solution recombination step, the arrangement and synchronization of parallel resources, the communication protocol, and the information exchanged between processors. Crainic et al. [19] provide a comprehensive literature review on the application of parallelism in meta-heuristics. Our contribution is a highly scalable parallel algorithm specifically designed to find quality primal solutions of large-scale FCMNF problems. Many algorithmic enhancements are combined in order to attain competitive levels of parallel performance: a novel parallel decomposition procedure based on the problem structure, a highly parallelizable local search scheme, and a tiered parallel procedure that is able to combine large numbers of partial solutions quickly. Solution crossover methods such as the one used in our approach have already been introduced and discussed previously [20,21]. In contrast to these works, we introduce a parallelization of the method that enables the recombination of a large number of solutions simultaneously.

We present experimental results that show the effectiveness of our parallel local search approach. For the instances in the C problem set [22], our method identifies primal solutions that are within an average optimality GAP of 0.58% with respect to the best known lower bound in an average time of 152 s per instance. We also test our parallel algorithm against the GT problem set [14], which contains substantially bigger instances. Our method takes less than 200 s on average to obtain a better solution than the best one found by CPLEX running for 5 h. We are able to identify considerably better solutions in more time. In addition, we present parallel scalability and load balancing performance results, which show that our novel implementation is able to take advantage of a large number of parallel processors to effectively reduce computation times in a load balanced execution.

The remainder of the paper is structured as follows. Section 2 presents an arc based MIP formulation of the FCMNF problem. Sections 3 and 4 provide a detailed description of our local search methodology and its parallel implementation on hybrid-memory parallel architectures, respectively. Section 5 presents computational experiments and results on standard instances from the literature. Finally, Section 6 provides some concluding remarks.

2. Problem description

Our local search approach is based on an arc-based MIP formulation of the FCMNF problem, which is described as follows. Let $G = (V, A)$ be a directed network, where V is the set of vertices and A the set of arcs. Let K be a set of commodities to be routed through G. Each commodity $k \in K$ is specified by a source vertex $s_k \in V$, a destination $t_k \in V$ and a quantity q_k of flow to be routed. Each arc $(i, j) \in A$ has an associated fixed cost f_{ij} that is imposed only when commodities are routed through it. Arcs also have a variable cost c_{ij} that is proportional to the flow traversing it and a maximum flow capacity u_{ij}. The problem consists of finding a routing for every commodity in K such that the arc capacities are respected and the costs are minimized. Let the flow variable x_{ij}^k denote the proportion of commodity $k \in K$ that is routed through the arc $(i, j) \in A$. In addition to the flow variables, we also introduce the binary variables y_{ij}, which reflect whether each arc (i, j) is used. The FCMNF problem can then be formulated as the following MIP:

$$\min_{x,y} \sum_{k\in K} \sum_{(i,j)\in A} c_{ij} q_k x_{ij}^k + \sum_{(i,j)\in A} f_{ij} y_{ij} \tag{1}$$

subject to:

$$\sum_{(i,j)\in A} x_{ij}^k - \sum_{(j,i)\in A} x_{ji}^k = d_i^k \quad \forall\, i \in V, \quad \forall\, k \in K \tag{2}$$

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

$$y_{ij} \in \{0, 1\} \quad \forall\, (i, j) \in A \tag{4}$$

$$0 \le x_{ij}^k \le 1 \quad \forall\, (i, j) \in A, \ \forall\, k \in K. \tag{5}$$

Restriction (2) ensures the conservation of flow. The flow differential for a vertex i and a commodity k is expressed by d_i^k, which is defined as:

$$d_i^k = \begin{cases} 1 & \text{if } i = s_k \\ -1 & \text{if } i = d_k \\ 0 & \text{otherwise.} \end{cases}$$

The coupling constraints (3) guarantee that the flow through each arc does not exceed the arc capacity. The capacity restrictions have a two-fold function, as they also ensure that the fixed cost is imposed when an arc is used. All commodity flow variables relative to the same arc are aggregated in the same constraint. A tighter and stronger LP relaxation can be obtained by introducing a set of $|A|\cdot|K|$ independent constraints:

$$x_{ij}^k \le y_{ij} \quad \forall\, (i, j) \in A, \quad \forall\, k \in K. \tag{6}$$

These are redundant with respect to (3). We choose not to include them in our model due to performance issues resulting from their