Transportation Research Part B 79 (2015) 93–120

Contents lists available at ScienceDirect

Transportation Research Part B

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

The recoverable robust facility location problem

Eduardo Álvarez-Miranda [a,*], Elena Fernández [b], Ivana Ljubić [c,d]

[a] Industrial Engineering Department, Universidad de Talca, Curicó, Chile
[b] Departament of Statistics and Operational Research, Universitat Politècnica de Catalunya, Barcelona, Spain
[c] Department of Statistics and Operations Research, Universität Wien, Vienna, Austria
[d] ESSEC Business School, Cergy-Pontoise, France

ARTICLE INFO

Article history:
Received 31 January 2015
Received in revised form 31 May 2015
Accepted 1 June 2015
Available online 17 June 2015

Keywords:
Facility location
Two-stage robust optimization
Branch-and-cut
L-shaped cuts

ABSTRACT

This work deals with a facility location problem in which location and allocation (transportation) policy is defined in two stages such that a first-stage solution should be robust against the possible realizations (scenarios) of the input data that can only be revealed in a second stage. This solution should be robust enough so that it can be *recovered* promptly and at low cost in the second stage. In contrast to some related modeling approaches from the literature, this new *recoverable robust* model is more general in terms of the considered data uncertainty; it can address situations in which uncertainty may be present in any of the following four categories: *provider-side* uncertainty, *receiver-side* uncertainty, uncertainty *in-between*, and uncertainty with respect to the cost parameters.

For this novel problem, a sophisticated branch-and-cut framework based on Benders decomposition is designed and complemented by several non-trivial enhancements, including scenario sorting, dual lifting, branching priorities, matheuristics and zero-half cuts. Two large sets of instances that incorporate spatial and demographic information of countries such as Germany and US (transportation) and Bangladesh and the Philippines (disaster management) are introduced. They are used to analyze in detail the characteristics of the proposed model and the obtained solutions as well as the effectiveness, behavior and limitations of the designed algorithm.

1. Introduction

Nowadays, we are more and more aware of the growing presence of dynamism and uncertainty in decision making. Fortunately, as the decisions become more complex, the availability of modeling, algorithmic and computational tools increases as well. Facility location and allocation decisions are among the most relevant decisions in several private and public transportation contexts and they usually involve strategic and operative policies with mid and long term impacts. Precisely because of the practical relevance of these decisions, it is important that they incorporate the uncertainty that naturally appears during the planning, modeling and operative process. Such uncertainty can be represented by different realizations of the input data: *customers* that actually require a commodity or a service, *locations* where the facilities can be located, the *transportation network* that can be used for connecting customers with facilities, and the corresponding *costs*. The true values of this data usually become available later in the decision process. In such cases a standard deterministic

* Corresponding author at: Industrial Engineering Department, Universidad de Talca, Merced 437, Curicó, Chile. Tel.: +56 075 201706.
E-mail address: ealvarez@utalca.cl (E. Álvarez-Miranda).

http://dx.doi.org/10.1016/j.trb.2015.06.001
0191-2615/

optimization model that considers a single possible outcome of the input data can lead towards solutions that are unlikely to be optimal, or for that matter even feasible, for the final data realization.

Supply chain management is a strategical area in which both *uncertainty* and *facility location* are core elements. For instance, as it is pointed out in Snyder and Daskin (2005), supply chains are particularly vulnerable to disruptions (intentional or accidental), imposing the need of taking into account the possible availability of depots and roads and different structures of the demand. Likewise, short-term phenomena such as fluctuations in commodity prices (such as oil) or long-term public policies (such as new toll road concessions) might lead to operational cost increases that should be considered when deciding the transportation network to be used.

In another context, natural events such as tsunamis, hurricanes or blizzards can produce disastrous effects with unpredictable intensity on populated areas and on the transportation infrastructure. Countries such as Bangladesh and the Philippines are two typical examples; both of them are regularly hit by hydrological disasters such as floods and typhoons. According to the Department of Disaster Management of Bangladesh (DDM, 2014), every year around 18% of the country is flooded, which produces over 5000 causalities and the destruction of more than 7 millions of homes. However, flooded areas my exceed the 75% of the country in case of severe events (as in 1988, 1998 and 2004). In the case of the Philippines, between 6 and 9 typhoons make landfall every year producing thousands of human losses and incalculable urban destruction; in November of 2013, typhoon Haiyan produced 6241 causalities and material damage of over 809 millions USD (see PAGASA, 2014). In these examples, it is crucial to be able to count with a robust system of humanitarian relief facilities that even in the worst possible scenario can provide assistance with the quickest possible response reducing the number of human loses after the occurrence of the event.

The Uncapacitated Facility Location Problem (UFL), also referred as the Simple Plant Location Problem, is one of the fundamental models in the wide spectrum of *Facility Location* problems (see, e.g., recent overviews presented in Eiselt and Marianov (2011), Daskin (2013), and Laporte et al. (2015)). In the classical deterministic version of the UFL one is given the set of customers, the set of locations, the facility set-up costs and the transportation costs. The goal is to define where to open facilities and how to allocate the customers to them so that the sum of set-up plus transportation costs is minimized.

In practice, it is usually the case that from the moment that the information is gathered until the moment in which the solution has to be implemented, some of the data might change with respect to the initial setting. As mentioned above, even if some (rough) idea about customers and locations is known, changes in demographic, socio-economic, or meteorological factors can lead to changes in the structure of the demand during the planning horizon, and/or the availability of a given location to host a facility (even if a facility has been already installed). This means that the solution obtained using a classical method might become infeasible and a new solution might have to be redefined from scratch. In these cases it would be better to recognize the presence of different scenarios for the data and find a solution comprised by first- and second-stage decisions.

Two well-known approaches to deal with uncertainty in optimization are Two-stage Stochastic Optimization (2SSO) and Robust Optimization (RO). In 2SSO (see Birge and Louveaux, 2011) the solutions are built in two stages. In the first stage, a *partial* collection of decisions is defined which are later on completed (in the second stage), when the true data is revealed. Hence, the objective is to minimize the cost of the first-stage decisions plus the *expected* cost of the recourse (second-stage) decisions. The quality of the solutions provided by this model strongly depends on the accuracy of the random representation of the parameter values (such as probability distributions) that allow to estimate the second-stage expected cost. Nonetheless, sometimes such accuracy is not available so the use of RO models dealing with *deterministic uncertainty* arises as a suitable alternative (see Kouvelis and Yu, 1997; Bertsimas and Sim, 2004; Ben-Tal et al., 2010). On the one hand these models do not require assumptions about the distribution of the uncertain input parameters; but on the other hand, they are usually meant for calculating single-stage decisions that are immune (in a certain sense) to all possible realizations of the parameter values.

A novel alternative that combines RO and 2SSO is Two-stage Robust Optimization (2SRO); as in RO, no stochasticity of the parameters is assumed, and as in 2SSO, decisions are taken in two stages. In this case, the cost of the second-stage decision is computed by looking at the worst-case realization of the data. Therefore, the goal of 2SRO is to find a *robust* first-stage solution that minimizes both the first-stage cost plus the worst-case second-stage cost among all possible data outcomes. 2SRO is a rather generic classification of models; for references on different 2SRO settings we refer the reader to Ben-Tal et al. (2004) and Zhao and Zeng (2012).

One of the possibilities in the 2SRO framework is *Recoverable Robustness* (see Liebchen et al., 2009). Recalling our practical motivation, assume that the facility location and allocation policy is defined in two stages such that we are required to find a first-stage solution that should be *robust* against the possible realizations (*scenarios*) of the input data in a second stage. This means that the first-stage solution is expected to perform *reasonably well*, in terms of feasibility and/or optimality, for *any* possible realization of the uncertain parameters. An essential element of this approach is the possibility of *recovering* the solution built in the first stage (i.e., to modify the previously defined location–allocation policy in order to render it feasible and/or cheaper) once the uncertainty is resolved in a second stage. The *recovery plan* is comprised by *recovery actions* which are known in advance and whose costs might also depend on the possible scenario. This recovery plan is *limited*, in the sense that the effort needed to recover a solution may be limited algorithmically (in terms of how a solution may be modified) and economically (in terms of the total cost of recovery actions). Therefore, instead of looking for a static solution that is robust against all possible scenarios without allowing any kind of recovery (which is the case for many RO approaches, see Ben-Tal