



FABIO D' ANDREAGIOVANNI^{1,2}, JONATAN KROLIKOWSKI²
AND JONAD PULAJ²

A hybrid primal heuristic for Robust Multiperiod Network Design

¹DFG Research Center MATHEON, Technical University Berlin, Straße des 17. Juni 135,
10623 Berlin, Germany

²Dept. of Optimization, Zuse-Institut Berlin (ZIB), Takustr. 7, 14195 Berlin, Germany,
e-mail: {d.andreagiovanni, krolikowski, pulaj}@zib.de

Herausgegeben vom
Konrad-Zuse-Zentrum für Informationstechnik Berlin
Takustraße 7
D-14195 Berlin-Dahlem

Telefon: 030-84185-0
Telefax: 030-84185-125

e-mail: bibliothek@zib.de
URL: <http://www.zib.de>

ZIB-Report (Print) ISSN 1438-0064
ZIB-Report (Internet) ISSN 2192-7782

A hybrid primal heuristic for Robust Multiperiod Network Design^{*}

Fabio D'Andreagiovanni^{1,2}, Jonatan Krolikowski², and Jonad Pula^j²

¹DFG Research Center MATHEON, Technical University Berlin
Straße des 17. Juni 135, 10623 Berlin

²Department of Optimization, Zuse-Institute Berlin (ZIB)
Takustr. 7, 14195 Berlin, Germany
{d.andreagiovanni,krolikowski,pulaj}@zib.de

Abstract. We investigate the Robust Multiperiod Network Design Problem, a generalization of the classical Capacitated Network Design Problem that additionally considers multiple design periods and provides solutions protected against traffic uncertainty. Given the intrinsic difficulty of the problem, which proves challenging even for state-of-the-art commercial solvers, we propose a hybrid primal heuristic based on the combination of ant colony optimization and an exact large neighborhood search. Computational experiments on a set of realistic instances from the SNDlib show that our heuristic can find solutions of extremely good quality with low optimality gap.

Keywords: Multiperiod Network Design, Traffic Uncertainty, Robust Optimization, Multiband Robustness, Hybrid Heuristics.

1 Introduction

The design of a telecommunication network can be essentially described as the task of establishing the topology of the network and the technological features (e.g., transmission capacity and rate) of its elements, namely nodes and links. The dramatic growth that telecommunications have experienced over the last ten years has greatly increased the complexity and difficulty of the corresponding design problems. The growing need for taking into account data uncertainty, such as that of traffic volumes, has made things even more complicated. In this context, the traditional design approach of professionals, based on a combination of trial-and-error and simulation, may lead to arbitrarily bad design solutions and thus the need for optimization-oriented approaches has arisen.

In this paper, we focus on the development of a new Robust Optimization model to tackle traffic uncertainty in a Multiperiod Network Design Problem

^{*} To appear in *EvoApplications - Proc. of EvoStar 2014*, Esparcia-Alcazar, A. et al. (Eds.), Lecture Notes in Computer Science, Springer, Heidelberg. The final publication is available at link.springer.com.

(MP-NDP). This problem constitutes a natural extension of a classical network design problem, in which we want to decide how to install capacity modules in the network in order to route traffic flows of communications generated by users. The extension implies the design over a time horizon made up of multiple periods. Moreover, traffic uncertainty is taken into account to protect design solutions against deviations of the traffic input data, that may compromise feasibility and optimality of solutions. To the best of our knowledge, the (MP-NDP) has received little attention and just a few works have investigated it (primarily, [1] and [2]). These works point out the difficulty of solving multiperiod problems already for just two periods and (easier) splittable-flow routing [1], and for a pure routing problem in satellite communications [2]. Our direct and more recent computational experience confirmed this behaviour, even for instances of moderate size considering a low number of time periods and solved by a state-of-the-art commercial mixed-integer programming solver.

In this work, our main original contributions are:

1. the first Robust Optimization model for Multiperiod Network Design. The formulation is developed to tackle traffic uncertainty, modeling data uncertainty by Multiband Robustness [3–5], a new model for Robust Optimization recently introduced to refine the classical Bertsimas-Sim model [6];
2. a hybrid solution algorithm, based on the combination of an exact large neighborhood search called RINS [7] with ant colony optimization [8];
3. computational experiments over a set of realistic instances derived from SNDlib, the Survivable Network Design Library [9], showing that our hybrid algorithm is able to produce solutions of extremely high quality associated with very small optimality gap.

The remainder of this paper is organized as follows: in Section 2, we review a canonical network model for joint routing and capacity installation; in Section 3, we introduce the new formulation for Robust Multiperiod Network Design; in Sections 4 and 5, we present our hybrid metaheuristic and computational results.

2 Capacitated Network Design

The *Capacitated Network Design Problem* (CNDP) can be described as follows: given a network and a set of demands whose flows must be routed between vertices of the network, we want to install capacities on network edges and route the flows through the network, so that the capacity constraint of each edge is respected and the total cost of installing capacity is minimized. The CNDP has been a central and highly studied problem in Network Optimization, that appears in a wide variety of real-world applications. For an exhaustive introduction to the topic, we refer the reader to the well-known book [10].

The CNDP is commonly formalized in the following way: we are given 1) a network represented by a graph $G(V, E)$, where V is the set of vertices and E the set of edges; 2) a set of commodities C , each associated with a traffic flow d_c to route from an origin s_c to a destination t_c ; 3) a set of admissible paths P_c

for routing the flow of each commodity c from s_c to t_c ; 4) a cost γ_e for installing one module of capacity $\phi > 0$ on edge $e \in E$. Using this notation, we can model the problem as an *integer linear program*:

$$\begin{aligned}
\min \sum_{e \in E} \gamma_e y_e & && \text{(CNDP-IP)} \\
\sum_{c \in C} \sum_{p \in P_c: e \in p} d_c x_{cp} \leq \phi y_e & && e \in E \quad (1) \\
\sum_{p \in P_c} x_{cp} = 1 & && c \in C \quad (2) \\
x_{cp} \in \{0, 1\} & && c \in C, p \in P_c \\
y_e \in \mathbb{Z}_+ & && e \in E,
\end{aligned}$$

The problem uses two families of variables: the binary variables x_{cp} (*path-assignment variables*) and the non-negative integer variables y_e (*capacity variables*). A path-assignment variable x_{cp} is equal to 1 if the entire flow of a commodity $c \in C$ is routed through path $p \in P_c$ and 0 otherwise. A capacity variable y_e represents instead the number of capacity modules installed on edge $e \in E$. The objective function minimizes the total cost of installation. Capacity constraints (1) impose that the summation of all flows routed through an edge $e \in E$ must not exceed the capacity installed on e (equal to the number of installed modules represented by y_e multiplied by the capacity ϕ granted by a single module). Constraints (2) impose that flow of each commodity $c \in C$ must be routed through a single path.

Remark 1. This is an unsplittable version of the CNDP, namely the traffic flow of a commodity $c \in C$ *cannot be split* over multiple paths going from s_c to t_c , but must be routed on exactly one path. Moreover, the set of feasible routing paths P_c of each commodity is pre-established and constitutes an input of the problem. This is in line with other works based on industrial cooperations (e.g., [11]) and with our experience [12], in which a network operator typically considers just a few paths that meet its own specific business and quality-of-service considerations.

3 Multiband-Robust Multiperiod Network Design

We introduce now a generalization of the CNDP, designing the network over multiple time periods and taking into account traffic uncertainty. The multiperiod design requires the introduction of a time horizon made up of a set of elementary time periods $T = \{1, 2, \dots, |T|\}$. From a modeling point of view, in the optimization problem we simply need to add a new index $t \in T$ to the decision variables, to represent routing and capacity installation decisions taken in each period (we stress however that this greatly increases the size and complexity of the problem).

Concerning traffic uncertainty, we assume that for each commodity $c \in C$ the demand d_c is *uncertain*, i.e. its value is not known exactly, but lies in a known

range. More specifically, we assume to know a nominal value of traffic \bar{d}_c and maximum negative and positive deviations δ_c^-, δ_c^+ from it. The actual value d_c thus belongs to the interval: $d_c \in [\bar{d}_c - \delta_c^-, \bar{d}_c + \delta_c^+]$.

Example 1 (traffic uncertainty). We are given two commodities c_1, c_2 with nominal traffic demands $\bar{d}_{c_1} = 100$ Mb, $\bar{d}_{c_2} = 150$ Mb and we know that these values may deviate up to 10%. So the maximum negative and positive deviations for c_1, c_2 are $\delta_{c_1}^- = \delta_{c_1}^+ = 10$ Mb, $\delta_{c_2}^- = \delta_{c_2}^+ = 15$ Mb, respectively. The actual values of traffic are then $d_{c_1} \in [90, 110]$ Mb, $d_{c_2} \in [135, 165]$ Mb.

The presence of uncertain data in an optimization problem can be very tricky: it is well-known that even small variations in the value of input data may make an optimal solution heavily suboptimal, whereas feasible solutions may reveal to be infeasible and thus completely useless in practice [13]. As a consequence, in our case we cannot optimize just using the nominal demand values \bar{d}_c , but we must take into account the possibility that demands will vary in the ranges $[\bar{d}_c - \delta_c^-, \bar{d}_c + \delta_c^+]$ that we have characterized. We illustrate the bad effects of input data deviations by providing an example.

Example 2 (infeasibility caused by deviations). Consider again the commodities of Example 1 and suppose that in some link we have installed exactly the capacity to handle the sum of their nominal values (i.e. we have installed $100 + 150$ Mb of capacity). This capacity dimensioning neglects that the demands may deviate up to 10%. It is sufficient that one demand increases, while the other remains the same to violate the capacity constraint of the link, making the design solution infeasible in practice.

Over the years, many methods such as *Stochastic Programming* and *Robust Optimization* have been proposed in literature for dealing with data uncertainty in optimization problems. We refer the reader to [13] for a general discussion about data uncertainty and its effects and for an overview of the most studied methodologies to deal with them.

In this paper, we tackle data uncertainty by Robust Optimization (RO), a methodology that has gained a lot of attention over the last decade [13, 6]. RO essentially takes into account data uncertainty by including additional hard constraints in the optimization problem. These constraints eliminate those solutions that are not protected against deviations of the input data from their nominal values. So a robust optimization problem considers only those solutions that are completely protected against specified data deviations. The data deviations that are considered are specified through a so-called *uncertainty set*. More formally, suppose that we are given a generic linear program:

$$v = \max c'x \quad \text{with } x \in \mathcal{F} = \{Ax \leq b, x \geq 0\}$$

and that the coefficient matrix A is uncertain, i.e. we do not know exactly the value of its entries. However, we are able to identify a family \mathcal{A} of coefficient matrices that represent possible valorizations of the uncertain matrix A , i.e. $A \in \mathcal{A}$. This family represents the uncertainty set of the robust problem. Then

we can produce a *robust optimal solution*, i.e. a solution that is protected against data deviations, by considering the *robust counterpart* of the original problem:

$$v^{\mathcal{R}} = \max c'x \text{ with } x \in \mathcal{R} = \{\tilde{A}x \leq b \quad \forall \tilde{A} \in \mathcal{A}, x \geq 0\}.$$

The feasible set \mathcal{R} of the robust counterpart contains only those solutions that are feasible for all the coefficient matrices in the uncertainty set \mathcal{A} . Therefore, \mathcal{R} is a subset of the feasible set of the original problem, i.e. $\mathcal{R} \subseteq \mathcal{F}$. The choice of the coefficient matrices included in \mathcal{A} should reflect the risk aversion of the decision maker.

Providing protection entails the so-called *price of robustness*, namely a deterioration of the optimal value of the robust counterpart w.r.t. the optimal value of the original problems (i.e., $v^{\mathcal{R}} \leq v$). This is a consequence of restricting the feasible set to only robust solutions. The price of robustness reflects the features of the uncertainty set: uncertainty sets expressing higher risk aversion will take into account more severe and unlikely deviations, leading to higher protection but also higher price of robustness; conversely, uncertainty sets expressing risky attitudes will tend to neglect improbable deviations, offering less protection but also a reduced price of robustness.

Example 3 (protection against deviations). Following example 2, a simple way to grant protection would be to install sufficient capacity to deal with the peak deviations of each commodity. So we should install 110+165 Mb of capacity.

3.1 A Robust Optimization model for traffic-uncertain Multiperiod Network Design

If we denote by \mathcal{D} the uncertainty set associated with the demands of the commodities, we can finally state the general form of the robust counterpart of the multiperiod network design problem as follows:

$$\begin{aligned} \min \quad & \sum_{e \in E} \sum_{t \in T} \gamma_{et} y_{et} \\ & \sum_{c \in C} \sum_{p \in P_c: e \in p} \bar{d}_{ct} x_{cpt} + DEV_{et}(x, \mathcal{D}) \leq \phi \sum_{\tau=1}^t y_{e\tau} & e \in E, t \in T & \quad (3) \\ & \sum_{p \in P_c} x_{cpt} = 1 & c \in C, t \in T & \\ & x_{cpt} \in \{0, 1\} & c \in C, p \in P_c, t \in T & \\ & y_{et} \in \mathbb{Z}_+ & e \in E, t \in T, & \end{aligned}$$

Besides the addition of a new index $t \in T$ in the decision variables to represent decisions taken in each time period, the modifications of the model concentrates in the robust capacity constraints (3). Each of these constraints considers: 1) the sum of *nominal* traffic demands \bar{d}_{ct} of commodities using the edge e in period t ; 2) the overall maximum positive deviation $DEV_{et}(x, \mathcal{D})$ that demands may

experience on edge e in period t and are allowed by the uncertainty set \mathcal{D} for a routing vector x ; 3) the overall capacity installed in e since the first period of the horizon (so we sum up the integer variables $y_{e\tau}$ from period 1 to t and multiply them by the basic capacity ϕ of a module).

Structuring the uncertainty set \mathcal{D} . We now have a general definition of the robust counterpart of the multiperiod problem. A question that is still open is how to structure the uncertainty set \mathcal{D} and deciding which deviations from the nominal traffic values \bar{d}_{ct} to take into account to produce robust solutions.

To characterize \mathcal{D} , we use *Multiband Robustness*, a new model for Robust Optimization recently introduced to refine and generalize the classical Γ -robustness model by Bertsimas and Sim [6], while maintaining its accessibility and tractability. For a detailed explanation of Multiband Robustness we refer the reader to [3–5]. Here we directly discuss the adaption of the model to our specific case.

According to the multiband framework, we build a *multiband uncertainty set* as follows:

1. for each commodity $c \in C$ and time period $t \in T$, we know the nominal value \bar{d}_{ct} of the traffic coefficient and maximum negative and positive deviations $\delta_{ct}^-, \delta_{ct}^+$ from it. The actual value d_{ct} is then such that $d_{ct} \in [\bar{d}_{ct} - \delta_{ct}^-, \bar{d}_{ct} + \delta_{ct}^+]$;
2. the overall deviation range $[\bar{d}_{ct} - \delta_{ct}^-, \bar{d}_{ct} + \delta_{ct}^+]$ of each coefficient d_{ct} is partitioned into K bands, defined on the basis of K deviation values:
 $-\infty < \delta_{ct}^- = \delta_{ct}^{K^-} < \dots < \delta_{ct}^{-1} < \delta_{ct}^0 = 0 < \delta_{ct}^1 < \dots < \delta_{ct}^{K^+} = \delta_{ct}^+ < +\infty$;
3. through these deviation values, K deviation bands are defined, namely: a set of positive deviation bands $k \in \{1, \dots, K^+\}$ and a set of negative deviation bands $k \in \{K^- + 1, \dots, -1, 0\}$, such that a band $k \in \{K^- + 1, \dots, K^+\}$ corresponds to the range $(d_{ij}^{k-1}, d_{ij}^k]$, and band $k = K^-$ corresponds to the single value $d_{ij}^{K^-}$;
4. for each capacity constraint (3) defined for an edge $e \in E$, period $t \in T$ and band $k \in K$, a value $\theta_{etk} \geq 0$ is introduced to represent the number of traffic coefficients of the constraint whose value deviates in band k . Of course, $\theta_{etk} \geq 0$ must be less or equal than the number of traffic coefficients that are present in the constraint.

Given the previous characterization of the multiband uncertainty set, the maximum positive deviation of traffic $DEV_{et}(x, \mathcal{D})$ of a constraint (3) can be found by solving a binary linear program (see [3] for details). Since the polytope associated with the binary program is shown to be integral, by considering its relaxation and by exploiting strong duality, it is possible to reformulate the original trivial robust counterpart as the following *linear and compact* robust counterpart (we

refer the reader to [3] for a formal proof of the result):

$$\begin{aligned}
\min \sum_{e \in E} \sum_{t \in T} \gamma_{et} y_{et} & \quad (\text{Rob-MP-CNDP}) \\
& \sum_{c \in C} \sum_{p \in P_c: e \in p} \bar{d}_{ct} x_{cpt} + \\
& + \sum_{k \in K} \theta_{etk} w_{etk} + \sum_{c \in C} \sum_{p \in P_c: e \in p} z_{cpt} \leq \phi \sum_{\tau=1}^t y_{e\tau} \quad e \in E, t \in T \\
w_{etk} + z_{cpt} & \geq \delta_{ctk} x_{cpt} \quad e \in E, c \in C, p \in P_c : e \in p, \\
& \quad t \in T, k \in K \quad (4) \\
w_{etk} & \in \mathbb{R} \quad e \in E, t \in T, k \in K \quad (5) \\
z_{cpt} & \geq 0 \quad e \in E, c \in C, p \in P_c : e \in p, t \in T \quad (6) \\
\sum_{p \in P_c} x_{cpt} & = 1 \quad c \in C, t \in T \\
x_{cpt} & \in \{0, 1\} \quad c \in C, p \in P_c, t \in T \\
y_{et} & \in \mathbb{Z}_+ \quad e \in E, t \in T,
\end{aligned}$$

This formulation includes additional constraints (4) and variables (5),(6) which are derived from the dualization operation that allow to linearly reformulate the original (non-linear) problem including the term $DEV_{et}(x, \mathcal{D})$ in each capacity constraint (see [3] for details).

In principle, we can get a robust optimal solution for (Rob-MP-CNDP) by using any commercial mixed-integer programming software. However, as showed in the computational experiment section, getting feasible solutions to this problem may be a challenge even for a state-of-the-art solver like IBM ILOG CPLEX (<http://www-01.ibm.com>). In the next section, we thus propose a hybrid exact-ant colony primal heuristic that is able to find solutions of very high quality.

4 A hybrid primal heuristic for the Rob-MP-CNDP

Attracted by the effectiveness of MIP-based and bio-inspired heuristics in hard network design problems (see, for example [14, 15, 8, 16]), we present an original hybrid primal heuristic based on the combination of Ant Colony Optimization (ACO) and an exact large neighbourhood search. ACO is a metaheuristic originally proposed by Dorigo and colleagues for combinatorial optimization [17] and later extended to integer and continuous problems (e.g., [8]). Over the years several refinements of the basic algorithm have been proposed (e.g., [18, 19]). ACO was inspired by the behaviour of ants searching for food and is essentially based on the definition of a cycle where a number of feasible solutions are iteratively built in parallel, using information about solutions built in previous executions of the cycle. An ACO algorithm presents the following general structure:

1. UNTIL an arrest condition is reached DO (Gen-ACO)

- (a) Ant-based solution construction
 - (b) Pheromone trail update
2. Daemon actions

We now proceed to detail each phase of the previous sketch for our hybrid ACO-exact algorithm for the (Rob-MP-CNDP). Our approach is hybrid since the canonical ACO construction phase is followed by a daemon-action phase, based on an exact large neighborhood search formulated as a mixed-integer linear program.

Ant-based solution construction. In the step 1 of the cycle, $m \geq 0$ ants are defined and each ant iteratively builds a feasible solution for the optimization problem. At every iteration, the ant is in a *state* corresponding with a *partial solution* of the problem and can further complete the solution by making a *move* and thus fixing the value of a new non-fixed variable. The move is chosen probabilistically, evaluating pheromone trail values. For a more detailed description of the elements and actions of step 1, we refer the reader to the paper [19] by Maniezzo. This paper presents ANTS, an improved ANT algorithm that we have taken as reference for our work. We considered ANTS particularly attractive as it proposes a series of improvements for ACO that allow to better exploit polyhedral information about the problem. Furthermore, ANTS is based on a reduced number of parameters and uses more efficient mathematical operations.

Before describing how our ANTS implementation is structured, we make some preliminary considerations. The formulation (Rob-MP-CNDP) is based on four families of variables: 1) the path assignment variables x_{cpt} ; 2) the capacity variables y_{et} ; 3-4) the auxiliary variables w_{etk} , z_{ept} coming from robust dualization. Though we have to deal with four families, we can notice that routing decisions taken over the time horizon entirely determine the capacity installation of minimum cost. Indeed, once the values of all path assignment variables are fixed, the routing is completely established and the worst traffic deviation term $DEV_{et}(x, \mathcal{D})$ can be efficiently derived without the auxiliary variables w_{etk} , z_{ept} [3,4]. So we can derive the total traffic D_{et} sent over an edge e in period t in the worst case. The minimum cost installation can then be derived through a sequential evaluation from period 1 to period T , keeping in mind that we must have $\left\lceil \frac{D_{et}}{\phi} \right\rceil$ capacity modules on e in t to accommodate the traffic. As a consequence, in the ant-construction phase we can limit our attention to the binary assignment variables and we introduce the concept of *routing state*.

Definition 1. *Routing state (RS):* let $P = \bigcup_{c \in C} P_c$ and let $R \subseteq C \times P \times T$ be the subset of triples (c, p, t) representing the assignment of path $p \in P_c$ to commodity c in period $t \in T$. A routing state is an assignment of paths to a subset of commodities in a subset of time periods which excludes that multiple paths are assigned to a single commodity. Formally:

$$RS \subseteq R: \exists (c_1, p_1, t_1), (c_2, p_2, t_2) \in RS: c_1 = c_2 \wedge p_1, p_2 \in P_{c_1} \wedge t_1 = t_2.$$

We say that a routing state RS is *complete* when it specifies the path used by each commodity in each time period (thus $|RS| = |C||T|$). Otherwise the RS is

called *partial* and we have $|RS| < |C||T|$.

In the ANTS algorithm that we propose, we decided to assign paths considering time periods and commodities in a pre-established order. Specifically, we establish the routing in each time period separately, starting from $t = 1$ and continuing up to $t = |T|$, and in each time period commodities are sorted in descending order w.r.t. their nominal traffic demand. Formally, this can be sketched through the following cycle that builds a *complete routing state*:

```

FOR  $t := 1$  TO  $|T|$  DO
  1. sort  $c \in C$  in descending order of  $\bar{d}_{ct}$ .
  2. FOR (sorted  $c \in C$ ) DO
    (a) assign a single path  $p \in P_c$  to  $c$ ;
  END FOR
END FOR

```

For an iteration (t, c) of the above nested cycles, the assignment of a path to a commodity corresponds with an ant moving from a partial routing state RS_i to a partial routing state RS_j such that: $RS_j = RS_i \cup \{(c, p, t)\}$ with $p \in P_c$.

We note that by the definitions of routing state a sequence of moves is actually a sequence of fixings of decision variables, as done in [19].

The probability that an ant k moves from a routing state i to a more complete routing state j , chosen among a set of feasible routing states, is defined by the improved formula of [19]:

$$p_{ij}^k = \frac{\alpha \tau_{ij} + (1 - \alpha) \eta_{ij}}{\sum_{f \in F} \alpha \tau_{if} + (1 - \alpha) \eta_{if}},$$

where $\alpha \in [0, 1]$ is a parameter assessing the relative importance of trail and attractiveness. As discussed in [19], the trail values τ_{ij} and the attractiveness values η_{ij} should be provided by suitable lower bounds of the considered optimization problem. In our particular case: 1) τ_{ij} is derived from the values of the variables in the solution associated with the linear relaxation of the robust counterpart (Rob-MP-CNDP); 2) η_{ij} is equal to the optimal solution of the linear relaxation of the nominal multiperiod network design problem, i.e. the problem that does not consider the traffic uncertainty. The optimum of this problem can be quickly computed and its computation becomes faster as more variables are fixed.

Daemon actions: Relaxation Induced Neighborhood Search. At the end of the ant-construction phase, we try to improve the quality of the feasible solution found by executing an *exact local search* in a *large neighborhood*. In particular, we adopt a modified *relaxation induced neighborhood search* (RINS) (see [7] for an exhaustive description of the method). Let (\bar{x}, \bar{y}) be a feasible solution of (Rob-MP-CNDP) found by an ant and (x^{LR}, y^{LR}) be an optimal (continuous) solution of the linear relaxation of (Rob-MP-CNDP). Moreover, let $(\bar{x}, \bar{y})_j, (x^{LR}, y^{LR})_j$ denote the j -th component of the vectors. Our modified RINS (*mod-RINS*) solves a sub-problem of (Rob-MP-CNDP) where:

1. we fix the variables whose value in (\bar{x}, \bar{y}) and (x^{LR}, y^{LR}) differs of at most $\epsilon > 0$, i.e.:

$$\begin{aligned}(\bar{x}, \bar{y})_j = 0 \cap (x^{LR}, y^{LR}) \leq \epsilon &\implies (x, y)_j = 0 \\(\bar{x}, \bar{y})_j = 1 \cap (x^{LR}, y^{LR}) \geq 1 - \epsilon &\implies (x, y)_j = 1\end{aligned}$$

2. impose a solution time limit of T .

A time limit is imposed since the subproblem may be difficult to solve, so the exploration of the neighbourhood may need to be truncated. Note that in point 1 we generalize the fixing rule of RINS, in which $\epsilon = 0$.

Pheromone trail update. At the end of each ant-construction phase h , the pheromone trails of a move $\tau_{ij}(h-1)$ are updated according to an improved formula proposed in [19]:

$$\tau_{ij}(h) = \tau_{ij}(h-1) + \sum_{k=1}^m \tau_{ij}^k \quad \text{with} \quad \tau_{ij}^k = \tau_{ij}(0) \cdot \left(1 - \frac{z_{curr}^k - LB}{\bar{z} - LB}\right), \quad (7)$$

where the values $\tau_{ij}(0)$ and LB are set by using the linear relaxation of (Rob-MP-CNDP): $\tau_{ij}(0)$ is set equal to the values of the corresponding optimal decision variables and LB equal to the optimal value of the relaxation. Additionally, z_{curr}^k is the value of the solution built by ant k and \bar{z} is the moving average of the values of the last ψ feasible solutions built. As noticed in [19], adopting formula (7) allows to replace the pheromone evaporation factor, a tricky parameter, with the moving average ψ whose setting has been shown to be much less critical.

Algorithm 1 details the structure of our original hybrid exact-ACO algorithm. The algorithm includes an outer loop repeated until a time limit is reached. At each execution of the loop, an inner loop defines m ants to build the solutions. Pheromone trail updates are done at the end of each execution of the inner loop. Once the ant construction phase is over, mod-RINS is applied so to try to get an improvement.

Algorithm 1. Hybrid ACO-exact algorithm for (Rob-MP-CNDP).

1. Compute the linear relaxation of (Rob-MP-CNDP) and initialize the values of $\tau_{ij}(0)$ by it.
2. UNTIL time limit is reached DO
 - (a) FOR $\mu := 1$ TO m DO
 - i. build a complete routing state;
 - ii. derive a complete feasible solution for (Rob-MP-CNDP);
 - END FOR
 - (b) Update $\tau_{ij}(t)$ according to (7).
3. apply mod-RINS to the best feasible solution.

5 Experimental results

We tested the performance of our hybrid algorithm on a set of 15 instances based on realistic network topologies from the SNDlib [9] defined in collaboration with industrial partners from former and ongoing projects. The experiments were performed on a machine with a 2.40 GHz quad-core processor and 16 GB of RAM and using IBM ILOG CPLEX 12.4. All the instances led to very large and

hard to solve robust multiperiod network design problems. We observed that even a state-of-the-art solver like CPLEX had troubles identifying good feasible solutions and in all the cases the final optimality gap was over 90%. In contrast, as clear from Table 5, in most cases our hybrid primal heuristic was able to find very high quality solutions associated with very low optimality gaps.

After executing preliminary tests, we found that an effective setting of the parameters of the heuristic was: $\alpha = 0.5$ (balancing attractiveness and trail level), $m = 3$ ants, $\psi = m$ (width of the moving average equal to the number of ants), $\epsilon = 0.1$ (tolerance of fixing in mod-RINS), $T = 20$ minutes (time limit in mod-RINS). The overall time limit for the execution of the heuristic was 5 hours. The same time limit was imposed on CPLEX when used to solve the robust counterpart (Rob-MP-CNDP). Each commodity admits 5 feasible paths, i.e. $|P_c| = 5, \forall c \in C$ and 3 positive deviations bands including the null deviation band. For each instance, in Table 5 we report its ID and features ($|V|$ = no. vertices, $|E|$ = no. edges, $|C|$ = no. commodities, $|T|$ = no. time periods). Moreover, we show the performance of the hybrid solution approach, that is denoted by the three measures $c^*(ACO)$, $c^*(ACO+RINS)$, $gapAR\%$, which represent the value of the best solution found by pure ACO, the value of the best solution found by ACO followed by RINS and the corresponding final optimality gap). We also show the performance of CPLEX, which is denoted by measures $c^*(IP)$ and $gapIP\%$ representing the value of the best solution found and the corresponding final optimality gap.

The best solutions found by our hybrid algorithm have in most cases a value that is at least one order of magnitude better than those found by CPLEX (2700% better on average). The results are of very high quality and, given the very low optimality gap, we can suppose that some of these solutions are actually optimal. We notice that in most cases executing RINS after the ant-construction phase can remarkably improve the value of the best solution found by the ants.

6 Conclusion and future work

We studied a Robust Optimization model for the Multiperiod Network Design Problem to tackle uncertainty of traffic demands. Robust solutions are deterministically protected against deviations of input traffic data, that may compromise the quality of produced solutions. The increase in complexity and dimension of the problem caused by considering multiple periods and robustness prevents state-of-the-art commercial solvers from finding good quality solutions, so we have defined a hybrid heuristic based on the combination of ant colony optimization and an exact large neighborhood search. Computational experiments on a set of realistic instances from the SNDlib showed that our heuristic can find solutions of extremely good quality. As future work, we plan to refine the heuristic (for example, by improving the ant-construction phase) and to integrate it with a branch-and-cut algorithm to enhance its computational performance.

Table 1. Experimental results

ID	V	E	C	T	$c^*(ACO)$	$c^*(ACO+RINS)$	gapAR%	$c^*(IP)$	gapIP%
Germany50	50	88	662	5	1.16E07	5.68E06	29.8	1.37E08	97.1
				7	2.12E07	9.02E6	15.5	3.48E08	97.8
				10	6.66E07	5.75E08	96.2	1.25E09	98.2
Pioro40	40	89	780	5	5.89E06	2.34E06	1.3	9.52E07	97.6
				7	1.42E07	5.10E06	3.1	2.40E08	97.9
				10	4.78E07	1.62E07	0.4	8.45E08	98.1
Norway	27	51	702	5	6.41E06	3.04E06	23.0	6.01E07	96.1
				7	1.44E07	5.73E06	12.8	1.47E08	96.6
				10	4.91E07	1.74E07	7.7	5.15E08	96.9
Geant	22	36	462	5	1.55E06	6.04E05	2.2	1.74E07	96.6
				7	3.61E06	1.29E06	1.6	4.32E07	97.1
				10	1.23E07	4.30E06	0.5	1.24E08	96.5
France	25	45	300	5	2.55E05	1.02E05	4.9	1.50E06	93.5
				7	5.97E05	2.18E05	2.2	3.01E06	92.9
				10	2.00E06	6.81E05	1.0	1.62E07	95.8

Acknowledgments

This work was partially supported by the *German Research Foundation* (DFG), project *Multiperiod Network Optimization*, by the DFG Research Center Math-eon (www.matheon.de), Project B3, and by the *German Federal Ministry of Education and Research* (BMBF), project *ROBUKOM* [12], grant 05M10PAA.

References

1. Lardeux, B., Nace, D., Geffard, J.: Multiperiod network design with incremental routing. *Networks* **50** (1) (2007) 109–117
2. Gamvros, I., Raghavan, S.: Multi-period traffic routing in satellite networks. *Europ. J. Oper. Res.* **219** (3) (2012) 738–750
3. Büsing, C., D’Andreagiovanni, F.: New results about multi-band uncertainty in robust optimization. In Klasing, R., ed.: *Experimental Algorithms*. Volume 7276 of LNCS., Springer, Heidelberg (2012) 63–74
4. Büsing, C., D’Andreagiovanni, F.: A new theoretical framework for robust optimization under multi-band uncertainty. In Helber, S., et al., eds.: *Operations Research Proceedings 2012*, Springer, Heidelberg (2014) 115–121
5. Büsing, C., D’Andreagiovanni, F., Raymond, A.: 0-1 multiband robust optimization. In Huisman, D., et al., eds.: *Operations Research Proceedings 2013*, Springer, Heidelberg (2014) To appear.
6. Bertsimas, D., Sim, M.: The price of robustness. *Oper. Res.* **52** (1) (2004) 35–53
7. Danna, E., Rothberg, E., Le Pape, C.: Exploring relaxation induced neighborhoods to improve mip solutions. *Math. Program.* **102** (2005) 71–90
8. Dorigo, M., Di Caro, G., Gambardella, L.: Ant algorithms for discrete optimization. *Artificial Life* **5** (2) (1999) 137–172

9. Orlowski, S., Wessäly, R., Pioro, M., Tomaszewski, A.: Sndlib 1.0 - survivable network design library. *Networks* **55** (3) (2010) 276–286
10. Ahuja, R., Magnanti, T., Orlin, J.: *Network Flows: Theory, Algorithms, and Applications*. Prentice Hall, Upper Saddle River (1993)
11. Bley, A., Grötschel, M., Wessäly, R.: Design of broadband virtual private networks: Model and heuristic for the b-win. Volume 53 of DIMACS SDMTCS. (2000) 1–16
12. Bauschert, T., Büsing, C., D’Andreagiovanni, F., Koster, A.M.C.A., Kutschka, M., Steglich, U.: Network planning under demand uncertainty with robust optimization. *IEEE Communications Magazine* (2014) To appear.
13. Ben-Tal, A., El Ghaoui, L., Nemirovski, A.: *Robust Optimization*. Springer, Heidelberg (2009)
14. D’Andreagiovanni, F.: On improving the capacity of solving large-scale wireless network design problems by genetic algorithms. In Di Chio, C., et al., eds.: *EvoApplications 2011*. Volume 6625 of LNCS., Springer, Heidelberg (2011) 11–20
15. D’Andreagiovanni, F., Mannino, C., Sassano, A.: Negative cycle separation in wireless network design. In Pahl, J., Reiners, T., Voss, T., eds.: *Network Optimization*. Volume 6701 of LNCS., Springer, Heidelberg (2011) 51–56
16. Kambayashi, Y.: A review of routing protocols based on ant-like mobile agents. *Algorithms* **6** (3) (2013) 442–456
17. Dorigo, M., Maniezzo, V., Coloni, A.: Ant system: Optimization by a colony of cooperating agents. *IEEE Trans. Syst. Man Cybern. B* **26** (1) (1996) 29–41
18. Gambardella, L.M., Montemanni, R., Weyland, D.: Coupling ant colony systems with strong local searches. *Europ. J. Oper. Res.* **220** (3) (2012) 831–843
19. Maniezzo, V.: Exact and approximate nondeterministic tree-search procedures for the quadratic assignment problem. *INFORMS J. Comp.* **11** (4) (1999) 358–369