Variable Neighborhood and Greedy Randomized Adaptive Search for Capacitated Connected Facility Location

Markus Leitner and Günther R. Raidl

Institute of Computer Graphics and Algorithms
Vienna University of Technology, Vienna, Austria
{leitner,raidl}@ads.tuwien.ac.at
http://www.ads.tuwien.ac.at

Abstract. The Connected Facility Location problem combining facility location and Steiner trees has recently gained stronger scientific interest as it can be used to model the extension of last mile communication networks in so-called fiber-to-the-curb scenarios. We consider a generalization of this problem which considers capacity constraints on potential facilities and aims at maximizing the resulting profit by potentially supplying only a subset of all customers. In this work, we discuss two metaheuristic approaches for this problem based on variable neighborhood search and greedy randomized adaptive search. Computational results show that both approaches allow for computing high quality solutions in relatively short time.

Keywords: connected facility location, network design, variable neighborhood search, greedy randomized adaptive search procedure.

1 Introduction

Nowadays, telecommunication companies are confronted with rising bandwidth demands of customers and thus they need to upgrade existing networks. Among others, *fiber-to-the-curb* is a popular deployment strategy in which parts of the existing connection between some central office and a customer is replaced by new fiber-optic technology. In addition, certain facilities bridging between fiber-optic and the previously existing – usually copper based – technology need to be installed. As long as the distance between a customer and its correspondingly assigned facility is not too high, a noticeable increase of the provided bandwidth can be achieved while avoiding the usually significantly higher costs for realizing the entire network by fiber-optic technology, i.e. fiber-to-the-home.

The resulting optimization problems have been formalized as variants of the *Connected Facility Location Problem* (ConFL) [19] which combines facility location and the Steiner tree problem in graphs. In this work, we consider the *Capacitated Connected Facility Location Problem* (CConFL) [16] which resembles a prize collecting variant of ConFL and additionally considers capacity constraints on potential facility locations.

R. Moreno-Díaz et al. (Eds.): EUROCAST 2011, Part I, LNCS 6927, pp. 295–302, 2011.

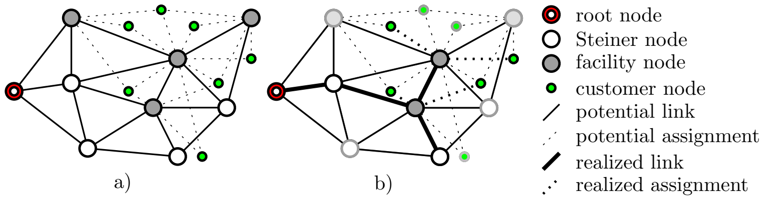

Fig. 1. An exemplary problem instance and a solution to CConFL

Formally, we are given an undirected, weighted graph $G = (V, E)$, with edge costs $c_e \geq 0, \forall e \in E$. The node set $V = \{r\} \cup F \cup T$ is the disjoint union of the root node r, potential facility locations F, and possible Steiner nodes T. Each facility $i \in F$ has associated opening costs $f_i \geq 0$ and a maximum assignable capacity $D_i \in \mathbb{N}$. Furthermore, we are given a set of potential customers C, with individual capacity demands $d_k \in \mathbb{N}$ and prizes $p_k \geq 0$, $\forall k \in C$, the latter corresponding to the expected profit when supplying customer k. Each customer $k \in C$ may be assigned to at most one facility of a subset $F_k \subseteq F$, with assignment costs $a_{ik} \geq 0$, $\forall i \in F_k$. A solution to CConFL $S = (R_S, T_S, F_S, C_S, \alpha_S)$ consists of a Steiner tree (R_S, T_S), $R_S \subseteq V$, $T_S \subseteq E$, connecting the set of opened facilities $F_S \subseteq F$ and the root node r. $C_S \subseteq C$ is the set of customers feasibly (i.e. respecting the capacity constraints) assigned to facilities F_S, whereas the actual mapping between customers and facilities is described by $\alpha_S : C_S \rightarrow F_S$. The objective value of a feasible solution S is given by $c(S) = \sum_{e \in T_S} c_e + \sum_{i \in F_S} f_i + \sum_{k \in C_S} a_{\alpha_S(k),k} + \sum_{k \in C \setminus C_S} p_k$, and we aim at identifying a most profitable solution minimizing this function. See Figure 1a for an exemplary instance and Figure 1b for an exemplary solution to this instance.

Since CConFL combines the (prize collecting) Steiner tree problem on a graph with the single source capacitated facility location problem, which are both strongly $\mathcal{NP}$-hard [14,5], CConFL is strongly $\mathcal{NP}$-hard, too.

2 Previous and Related Work

Karger and Minkoff [13] discussed the so-called maybecast problem which can be modeled as a connected facility location problem. The name connected facility location has been introduced by Gupta et al. [9] in their work on virtual private networks. A number of constant factor approximation algorithms have been proposed for different variants of ConFL, see e.g. [25,12], among which the currently best by Eisenbrand et al. [6] yields an approximation factor of 4.23 in its derandomized variant.

Ljubić [19] proposed a hybrid metaheuristic approach combining variable neighborhood search (VNS) and reactive tabu search, while Tomazic and Ljubić [27] discussed a greedy randomized adaptive search procedure (GRASP) for the unrooted version of ConFL. A more general variant of ConFL has been introduced by Bardossy and Raghavan [24,3] who combined dual ascent with local search to derive lower and upper bounds in their approach.

A large number of different integer linear programming (ILP) based models and solution approaches for ConFL have been described by Gollowitzer and Ljubić [8] and in [20,22,21] for a variant with hop constraints.

We discussed two compact multi-commodity flow based ILP models and a Lagrangian relaxation based approach for CConFL in [16,18]. The latter has been further hybridized with local search and very large scale neighborhood search. Furthermore, we presented additional, theoretically stronger ILP models [17]. The resulting approaches for solving them based on branch-and-cut and branch-cut-and-price, respectively, showed to significantly outperform all previous ones from a computational point of view.

To the best of our knowledge, no pure metaheuristic approaches for CConFL have been proposed so far, besides two VNS variants for a version of CConFL without assignments and opening costs by the current authors [15].

3 Greedy Solution Construction

We use a greedy approach to construct an initial feasible solution S. Initially, S consists of the root node only, i.e. $R_S = \{r\}$, $T_S = F_S = C_S = \emptyset$. In each iteration a single facility $i \in F \setminus F_S$ with a correspondingly assigned set of customers $C'_i \subseteq C \setminus C_S$ is added to S and connected to the current Steiner tree (R_S, T_S). We calculate a score $\delta_i = \frac{-f_i + \sum_{k \in C'_i}(p_k - a_{ik})}{\hat{c}_i}$ for each facility $i \in F \setminus F_S$ to decide which facility to add next. Hereby, the optimal set of customers $C'_i \subseteq C \setminus C_S$ still assignable to facility i is computed by solving a binary knapsack problem with an item with profit $p_k - a_{ik}$ and weight d_k for each customer $k \in C \setminus C_S$ and total knapsack capacity D_i. We apply the Combo algorithm [23] for solving these knapsack problems. Furthermore, we need to compute $\hat{c}_i \geq 0$ denoting the costs for connecting facility $i \in F \setminus F_S$, i.e. the costs of a least-cost path (V_i, E_i), $V_i \subseteq V$, $E_i \subseteq E$, from r to i where already included edges $e \in T_S$ are assigned zero costs. In each step, we add the facility $i \in F \setminus F_S$ with maximal score, i.e. $\text{argmax}_{i \in F \setminus F_S} \delta_i$. It is then connected to the partially constructed Steiner tree and the customers C'_i are assigned to it. Given a current partial solution S', solution S after adding facility i is defined as $F_S = F_{S'} \cup \{i\}$, $C_S = C_{S'} \cup C'_i$, $R_S = R_{S'} \cup V_i$, $T_S = T_{S'} \cup E_i$, and $\alpha_S(k) = i$, $\forall k \in C'_i$. This process is repeated as long as at least one facility $i \in F \setminus F_S$ exists for which the achievable profit exceeds the additional connection costs, i.e. as long as $\exists i \in F \setminus F_S : \delta_i > 1$.

4 Metaheuristic Approaches

In the following, we describe a variable neighborhood descent which is further embedded within a variable neighborhood search (VNS) as well as a GRASP.

4.1 Variable Neighborhood Descent

We use variable neighborhood descent (VND) [11] to improve solutions using four different neighborhood structures. These neighborhood structures, which are detailed in the following, focus on different aspects of a solution and are applied in the given order.

Key-path improvement. Neighborhood structures based on the concept of key-paths have been previously used for the Steiner tree problem on graphs [28] as well as on several related problems including CConFL [18]. The main idea is to replace paths between so-called key-nodes by cheaper ones and thus to reduce the total edge costs of a solution. As introduced in [18], for CConFL the set of key-nodes consists of the root node, all open facilities, and all other nodes of degree greater than two. The set of key-paths of solution S is given by all paths in S connecting two key-nodes that do not contain further key-nodes. The key-path neighborhood iteratively considers all key-paths $\mathcal{K}_S$ of solution S. For each such key-path $(V', E') \in \mathcal{K}_S$, $V' \in R_S$, $E' \in T_S$, connecting key-nodes $u, v \in R_S$ a minimum cost path (V'', E''), $V'' \subseteq V$, $E'' \subseteq E$, connecting u and v with respect to edge costs c' defined as $c'_e = c_e$, $\forall e \in (E \setminus T_S) \cup E'$, and $c'_e = 0$, $\forall e \in T_S \setminus E'$ is determined. In case $\sum_{e \in E'} c_e > \sum_{e \in E''} c'_e$ replacing (V', E') by (V'', E'') yields an improved solution.

Customer Swap. The customer swap neighborhood, which has been previously used by Contreras et al. [4] for the single source capacitated facility location problem as well as by the current authors for CConFL [18], tries to reduce the assignment costs of a solution S. It consists of all solutions S' reachable from S by swapping exactly two assignments. More precisely, given two facilities $i, j \in F_S$, $i \neq j$, and two customers $k, l \in C_S$, $k \neq l$, with $\alpha_S(k) = i$ and $\alpha_S(l) = j$, each move transforms S into a solution S' where $\alpha_{S'}(k) = j$ and $\alpha_{S'}(l) = i$.

Single Customer Cyclic and Path Exchange. This very large scale neighborhood search approach proposed in our previous work [18] generalizes above described customer swap neighborhood by considering changes of multiple assignments simultaneously. Furthermore, currently unassigned customers may be added to the solution, customers may be released, and facilities may be opened and closed. As done for related problems [1] a so-called improvement graph w.r.t. solution S with arc costs corresponding to resulting changes of the objective value is defined. Each feasible and improving single customer cyclic or path exchange then corresponds to a negative cost subset disjoint cycle in this improvement graph. Since deciding whether a graph contains a negative cost subset disjoint cycle is NP-hard [26] we adopted a heuristic approach originally proposed by Ahuja et al. [2] to find improving moves. We refer to our previous work [18] for a complete description.

Single Facility Swap. Our last neighborhood structure focuses on the set of opened facilities. The single facility swap neighborhood of a solution S consists of all solutions S' for which the set of opened facilities of S and S' differs by exactly one facility, i.e. one facility may be opened or closed. It is searched by iteratively considering all facilities $i \in F$ and calculating the corresponding objective value change δ_i due to opening or closing i. If $i \notin F_S$, $\delta_i = f_i + \sum_{k \in C'_i} (a_{ik} - p_k) + \hat{c}_i$ where $C'_i \subseteq C \setminus C_S$ is the optimal set of customers currently assignable to facility i and $\hat{c}_i \geq 0$ are the costs for connecting facility i to the current Steiner tree (R_S, T_S). Both C'_i and $\hat{c}_i$ are computed as described in Section 3. On the contrary, for facilities $i \in F_S$ that may be closed $\delta_i = -f_i + \sum_{k \in C_S : \alpha_S(k)=i}(p_k - a_{ik}) - \hat{c}_i$ holds. Here, a lower bound for the savings due to pruning the Steiner tree after closing facility i is used to estimate $\hat{c}_i$. In case i is a leaf node of (R_S, T_S), $\hat{c}_i$ is set to the costs of the unique key-path of S containing node i, while we set $\hat{c}_i = 0$ for all facilities $i \in F_S$ "inside" the Steiner tree.

4.2 Variable Neighborhood Search

We embed the VND as local improvement procedure in a VNS [10] approach. Shaking to escape from local optima is performed by applying random moves in generalizations of above described single facility swap neighborhood, swapping $l = 2, \ldots, l_{\max}$ randomly chosen facilities simultaneously.

4.3 Greedy Randomized Adaptive Search Procedure

We further embed above described VND into a GRASP [7] utilizing a randomized version of aforementioned constructive heuristic. Let $F' = \{i \in F \setminus F_S : \delta_i > 1\}$ be the actual set of facilities for which the achievable profit exceeds the connection costs, and in case $|F'| \neq \emptyset$, $\delta_{\min} = \text{argmin}_{\delta_i}\{i \in F'\}$ and $\delta_{\max} = \text{argmax}_{\delta_i}\{i \in F'\}$ denote the minimal and maximal scores among all relevant facilities, respectively. Rather than adding the facility with maximal score in each step, the randomized variant of above described constructive heuristic used in the GRASP approach randomly chooses one among the facilities $i \in F'$ for which $\delta_{\max} - \beta(\delta_{\max} - \delta_{\min}) \leq \delta_i \leq \delta_{\max}$ holds.

5 Computational Results

Computational tests have been performed on the benchmark instances from [18] using a single core of an Intel Core 2 Quad with 2.83GHz and 8GB RAM for each experiment. The VND has been configured as follows: We apply the neighborhood structures in the same order as introduced above, but switch back to the first – i.e. the key-path – neighborhood after changing the set of opened facilities only. The single customer cyclic exchange neighborhood is searched using a next improvement strategy, while best improvement is applied for all other neighborhood structures. VNS is terminated after ten consecutive non-improving iterations of the outermost largest shaking move of size $l_{\max} = \min\{|F|, 10\}$. We set $\beta = 0.2$ and generate 100 initial solutions for the GRASP approach. Each experiment has been repeated 30 times.

Table 1 summarizes relative minimum, average, and maximum objective values in percent, corresponding standard deviations, and relative median CPU times of the VNS and GRASP in relation to the branch-and-cut-and-price approach (dBCP) from [17], which performed best among the previously presented methods. We also report the total number of instances (#) of each group as well as the number of instances solved to proven optimality ($\#_{\text{opt}}$) by the branch-and-cut-and-price within the applied CPU-time limit of 7200 seconds.

Table 1. Relative minimum, average, and maximum objective values in % and relative median CPU times in seconds for GRASP and VNS in relation to branch-and-cut-and-price (dBCP) from [17]. Standard deviations for average values are reported in parentheses. Instances have been grouped according to $|F|$ and $|C|$ and each experiment has been repeated 30 times for GRASP and VNS.

				relative objective value								CPU time	
				$\frac{\text{GRASP}-\text{dBCP}}{\text{dBCP}}$ in %				$\frac{\text{VNS}-\text{dBCP}}{\text{dBCP}}$ in %				$\frac{\text{GRASP}}{\text{dBCP}}$	$\frac{\text{VNS}}{\text{dBCP}}$
$\|F\|$	$\|C\|$	#	$\#_{\text{opt}}$	min	avg		max	min	avg		max	median	
75	75	12	10	2.72	5.02	(2.08)	8.86	2.23	3.90	(1.45)	7.31	0.09	0.17
100	100	12	9	2.37	4.09	(1.30)	7.29	2.14	3.45	(1.16)	5.65	0.23	0.27
200	200	12	11	2.07	4.07	(2.03)	7.95	1.95	3.25	(1.14)	4.86	1.79	1.81
75	200	12	7	1.41	1.86	(0.32)	2.51	0.58	0.95	(0.28)	1.34	0.58	0.59
200	75	12	7	-56.78	-6.23	(23.59)	4.50	-56.40	-5.62	(23.71)	5.29	0.01	0.01

From Table 1 we conclude that both GRASP and VNS generally compute solutions only slightly worse than those of the state-of-the-art exact approach based on branch-cut-and-price. Except for the instance set with $|F| = 200$ and $|C| = 200$ both metaheuristic approaches also need considerably less CPU-time than $dBCP$. We further note that the solutions of both VNS and GRASP are significantly better than those obtained by dBCP within the given time limit of two hours for some instances with $|F| = 200$ and $|C| = 75$ that seem to be particularly hard. While needing slightly more computing time, VNS generally outperforms GRASP with respect to solution quality. The solutions obtained by VNS are less than 4% worse than those of dBCP on average and the maximum quality loss never exceeded 7.31%. Since both VNS and GRASP use the same VND as embedded local improvement procedure, we believe that the slight advantages of VNS over GRASP are due to its greater capabilities to explore the search space. More precisely, due to larger and random changes regarding the set of open facilities, VNS may on the contrary to GRASP also consider to open facilities, which do not seem to pay off at a first glance.

6 Summary and Outlook

In this article, we considered a prize collecting variant of the connected facility location problem with capacity constraints on potential facility locations. We proposed the use of metaheuristics to obtain high quality solutions to instances of CConFL within relatively short time when providing optimality gaps is not necessary. After introducing a variable neighborhood descent utilizing four different neighborhood structures we discussed its integration as local search component in VNS and GRASP approaches, respectively. Computational results on previously proposed benchmark instances show that both VNS and GRASP allow for generating high quality solutions in relatively short time and showed slight advantages for VNS.

In future, we might consider approaches combining the individual strengths of state-of-the-art exact methods for CConFL and the metaheuristics proposed in the current paper. On the one hand, one could integrate metaheuristic components into ILP based approaches to avoid huge gaps due to poor primal solutions after terminating the exact method due to a given time limit. On the other hand, restricted variants of existing and quite efficient exact approaches may be used within metaheuristics by means of large neighborhood searches. Finally, we also aim to further analyze the contributions of the different components of the proposed metaheuristics to the overall success in more details as well as to conduct a deeper computational study involving additional, larger instances.

References

1. Ahuja, R.K., Orlin, J.B., Pallottino, S., Scaparra, M.P., Scutella, M.G.: A multi-exchange heuristic for the single-source capacitated facility location problem. Management Science 50(6), 749–760 (2004)
2. Ahuja, R.K., Orlin, J.B., Sharma, D.: Multi-exchange neighborhood structures for the capacitated minimum spanning tree problem. Mathematical Programming 91(1), 71–97 (2001)
3. Bardossy, M.G., Raghavan, S.: Dual-based local search for the connected facility location and related problems. INFORMS Journal on Computing 22(4), 584–602 (2010)
4. Contreras, I.A., Diaz, J.A.: Scatter search for the single source capacitated facility location problem. Annals of Operations Research 157(1), 73–89 (2008)
5. Cornuejols, G., Nemhauser, G.L., Wolsey, L.A.: The uncapacitated facility location problem. In: Mirchandani, P.B., Francis, R.L. (eds.) Discrete Location Theory, pp. 119–171. Wiley, Chichester (1990)
6. Eisenbrand, F., Grandoni, F., Rothvoß, T., Schäfer, G.: Connected facility location via random facility sampling and core detouring. Journal of Computer and System Sciences 76(8), 709–726 (2010)
7. Feo, T., Resende, M.: Greedy randomized adaptive search procedures. Journal of Global Optimization 6(2), 109–133 (1995)
8. Gollowitzer, S., Ljubić, I.: MIP models for connected facility location: A theoretical and computational study. Computers & Operations Research 38(2), 435–449 (2011)
9. Gupta, A., Kleinberg, J., Kumar, A., Rastogi, R., Yener, B.: Provisioning a virtual private network: a network design problem for multicommodity flow. In: Proceedings of the 33rd Annual ACM Symposium on Theory of Computing, pp. 389–398 (2001)
10. Hansen, P., Mladenovic, N.: An introduction to variable neighborhood search. In: Voss, S., Martello, S., Osman, I.H., Roucairol, C. (eds.) Meta-heuristics, Advances and trends in local search paradigms for optimization, pp. 433–458. Kluwer Academic Publishers, Dordrecht (1999)
11. Hansen, P., Mladenović, N.: Variable neighborhood search: Principles and applications. European Journal of Operational Research 130(3), 449–467 (2001)
12. Hasan, M.K., Jung, H., Chwa, K.: Approximation algorithms for connected facility location problems. Journal of Combinatorial Optimization 16(2), 155–172 (2008)
13. Karger, D.R., Minkoff, M.: Building Steiner trees with incomplete global knowledge. In: Proceedings of the 41st Annual Symposium on Foundations of Computer Science, pp. 613–623. IEEE Computer Society, Los Alamitos (2000)

14. Karp, R.M.: Reducibility among combinatorial problems. In: Miller, E., Thatcher, J.W. (eds.) Complexity of Computer Computations, pp. 85–103. Plenum Press, New York (1972)
15. Leitner, M., Raidl, G.R.: Variable neighborhood search for a prize collecting capacity constrained connected facility location problem. In: Proceedings of the 2008 International Symposium on Applications and the Internet, pp. 233–236. IEEE Computer Society, Los Alamitos (2008)
16. Leitner, M., Raidl, G.R.: A Lagrangian decomposition based heuristic for capacitated connected facility location. In: Voß, S., Caserta, M. (eds.) Proceedings of the 8th Metaheuristic International Conference (MIC 2009), Hamburg, Germany (2009)
17. Leitner, M., Raidl, G.R.: Branch-and-cut-and-price for capacitated connected facility location. Journal of Mathematical Modelling and Algorithms (2011), `http://dx.doi.org/10.1007/s10852-011-9153-5`
18. Leitner, M., Raidl, G.R.: Combining Lagrangian decomposition with very large scale neighborhoood search for capacitated connected facility location. In: Post-Conference Book of the Eight Metaheuristics International Conference – MIC (2009) (to appear)
19. Ljubić, I.: A hybrid VNS for connected facility location. In: Bartz-Beielstein, T., Blesa Aguilera, M.J., Blum, C., Naujoks, B., Roli, A., Rudolph, G., Sampels, M. (eds.) HM 2007. LNCS, vol. 4771, pp. 157–169. Springer, Heidelberg (2007)
20. Ljubić, I., Gollowitzer, S.: Hop constrained connected facility location. Tech. Rep. 2009–09, University of Vienna (2009)
21. Ljubić, I., Gollowitzer, S.: Layered graph approaches to the hop constrained connected facility location problem. Tech. Rep. 2010-08, University of Vienna (2010) (submitted)
22. Ljubić, I., Gollowitzer, S.: Modelling the hop constrained connected facility location problem on layered graphs. In: International Symposium on Combinatorial Optimization (ISCO 2010), Hammamet, Tunisia. Electronic Notes in Discrete Mathematics, vol. 36, pp. 207–214 (2010)
23. Martello, S., Pisinger, D., Toth, P.: Dynamic programming and strong bounds for the 0–1 knapsack problem. Management Science 45(3), 414–424 (1999)
24. Raghavan, S., Bardossy, M.G.: Dual based heuristics for the connected facility location problem. In: Scutellà, M.G., et al. (eds.) Proceedings of the International Network Optimization Conference 2009 (2009)
25. Swamy, C., Kumar, A.: Primal-dual algorithms for connected facility location problems. Algorithmica 40(4), 245–269 (2004)
26. Thompson, P.M., Orlin, J.B.: The theory of cyclic transfers. Tech. Rep. OR 200-89, Massachusetts Institute of Technology, Operations Research Center (1989)
27. Tomazic, A., Ljubić, I.: A GRASP algorithm for the connected facility location problem. In: Proceedings of the 2008 International Symposium on Applications and the Internet, pp. 257–260. IEEE Computer Society Press, Los Alamitos (2008)
28. Voß, S.: Steiner's problem in graphs: heuristic methods. Discrete Applied Mathematics 40, 45–72 (1992)