Enabling E-Mobility: Facility Location for Battery Loading Stations

Sabine Storandt
Albert-Ludwigs-Universität Freiburg
Institut für Informatik
79110 Freiburg, Germany
storandt@informatik.uni-freiburg.de

Stefan Funke
Universität Stuttgart
Institut für Formale Methoden der Informatik
70569 Stuttgart, Germany
funke@fmi.uni-stuttgart.de

Abstract

The short cruising range due to the limited battery supply of current Electric Vehicles (EVs) is one of the main obstacles for a complete transition to E-mobility. Until batteries of higher energy storage density have been developed, it is of utmost importance to deliberately plan the locations of new loading stations for best possible coverage. Ideally the network of loading stations should allow driving from anywhere to anywhere (and back) without running out of energy. We show that minimizing the number of necessary loading stations to achieve this goal is NP-hard and even worse, we can rule out polynomial-time constant approximation algorithms. Hence algorithms with better approximation guarantees have to make use of the special structure of road networks (which is not obvious how to do it). On the positive side, we show with instance based lower bounds that our heuristic algorithms achieve provably good solutions on real-world problem instances.

Introduction

Battery-powered Electric Vehicles (EVs) play an important role in reducing the global consumption of fossil fuels. Not only can they be powered using renewable energies, but also have the ability to recuperate energy, e.g. when going downhill or in deceleration phases. Despite their environmental advantages EVs still wait for their great breakthrough. One reason might be their limited cruising range (up to 150km) due to the battery's maximal capacity. Making use of loading stations (LSs) is inevitable on longer tours, unfortunately, LSs are still sparse in most countries. Therefore, deliberate planning of new LSs positions is crucial to extend practical cruising ranges with EVs – an important step towards the transition to E-mobility.

Ideally, the set of LSs in a network should allow for travelling from an arbitrary source to an arbitrary target (and back) when starting fully loaded and choosing an energetically reasonable route. The problem of computing energy-optimal routes for an EV was introduced in (Artmeier et al. 2010) and differs from route planning for conventional cars in two aspects: (a) the underlying graph bears partly negative edge weights due to energy recuperation and (b) battery constraints (no overcharging or running out of energy) have to be taken into account.

In (Storandt and Funke 2012), graph concepts like reachability and connectivity were extended to the scenario of energy-aware routing in the context of EVs: Let $G(V, E, c)$ with $|V| = n, |E| = m$ be a strongly connected (di)graph with $c : E \to \mathbb{R}$ being a cost function representing energy consumption, and $M \in \mathbb{R}^+$ the battery capacity of the EV. W.l.o.g. assume that $c(e) \leq M$ for all $e \in E$ and all cycles in G have positive cost (otherwise we would have a perpetual motion machine).

The maximal battery load that can be achieved at node $w \in V$ when starting fully loaded at $v \in V$ is called $b_v(w)$. A path from v to w that realizes a final battery load ≥ 0 is called ev-feasible. Accordingly the set of ev-reachable nodes of v contains all nodes on ev-feasible paths starting at v, i.e. $R(v) := \{w \in V : b_v(w) \geq 0\}$. The minimal required battery load to reach v when starting at w is denoted by $b'_v(w)$. Similarly a node w is called inverse reachable from v if $b'_v(w) \leq M$ and the respective set containing all these nodes is called $R^{-1}(v)$. The set of strongly ev-connected nodes of v is defined as $C(v) := \{w \in V : b_v(w) \geq b'_v(w)\}$ and represents the set of feasible roundtours $v, \cdots, w, \cdots, v$. A loading station (LS) is a node $l \in V$ that leads to a fully loaded battery whenever it is visited. The set of loading stations is denoted by $L \subseteq V$. Of course, the presence of LSs might augment the sets of ev-reachable or strongly ev-connected nodes for $v \in V$. The respective sets are called $R^L(v)$ and $C^L(v)$ and can be computed efficiently. The topic considered in this paper is how to choose a small set L of LSs such that the augmented ev-reachable and ev-connected node sets are V.

Contribution

In this paper we study two problems related to identifying reasonable sets of LSs. As placing LSs generates costs, we always aim for computing a minimum subset $L \subseteq V$ that fulfils certain connectivity requirements. At first we consider the problem of finding L such that for each pair of vertices $s, t \in V$, there exists an ev-feasible path from s to t. Next we augment the problem by the constraint that returning back to s without running out of energy must be possible as well (which somewhat surprisingly is considerably harder to achieve). For both of these problems we prove NP-hardness and even hardness of approximation, and provide approximation algorithms. With both problems not only being NP-hard but also hard to approximate, we develop practical, heuristic algorithms and prove their good performance on real-world instances by computing strong lower bounds on the concrete problem instances.

Getting Anywhere

For the transition to E-mobility to succeed, a necessary condition is to make sure that any target in the road network can be reached when starting with a fully charged EV (possibly with recharging at LSs). Let us formalize this:

Definition 1 (EV-Reachability Cover) *Given a (di)graph $G(V, E, c), M \in \mathbb{R}^+$, we refer to the problem of converting a subset $L \subseteq V$ of minimal cardinality into loading stations such that $\forall v \in V : R^L(v) = V$ as the EV-Reachability Cover Problem (ERC).*

Note that the following formulation for ERC is equivalent as soon as at least one LS is required (otherwise the problem is trivial, of course): Select a minimum subset $L \subseteq V$ as loading stations such that

1. $\forall v \in V \exists l, l' \in L : v \in R(l), l' \in R(v)$
2. $\forall l \in L : C^L(l) \supseteq L.$

The first condition assures that for any source vertex a LS can be reached with the EV, and for any target there is a LS it can be reached from. The second condition then claims that one can go from any LS to any other, resulting together with the first condition in complete pairwise ev-reachability.

Note though, that this formulation does not exclude the case that one gets stuck at the destination if no loading facility is available there. We will come back to that problem later on.

Approximation via Strongly Connected Dominating Set

The second formulation of ERC implies a close connection to Strongly Connected Dominating Set (SCDS) (Li et al. 2009), which is a variant of the classical Dominating Set (DS)(Kann 1992) problem.

Definition 2 (DS) *Given a graph $G(V, E)$, select a minimum subset $D \subseteq V$ such that $\forall v \in V \setminus D : \exists \{v, w\} \in E$ with $w \in D$.*

Definition 3 (SCDS) *Given a digraph $G(V, E)$, select a minimum subset $D \subseteq V$ such that $\forall v \in V \setminus D : \exists (u, v), (v, w) \in E$ with $u, w \in D$ and D induces a strongly connected subgraph.*

The SCDS description differs from ERC in terms of 'domination distance'; while in SCDS only adjacent nodes can serve as dominators in ERC all ev-connected nodes are candidates for covering or dominating a node. This gap can be closed by creating the reachability graph for an ERC instance:

Definition 4 (Reachability Graph) *Given $G(V, E, c), M \in \mathbb{R}^+$, we define the reachability graph $RG(V, E^+)$ with $\forall v, w \in V : (v, w) \in E^+ \Leftrightarrow w \in R(v)$.*

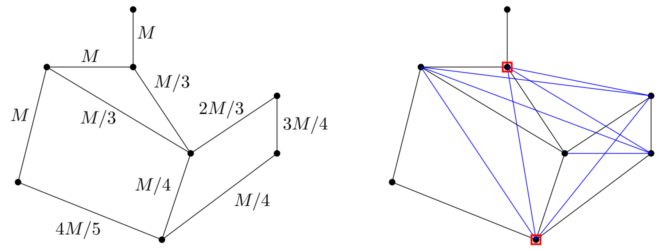

Figure 1: Left: Instance of ERC. Right: Reachability graph upon the ERC instance together with a (SC)DS/ERC solution (red boxes).

So the reachability graph augments the original graph with direct edges between nodes that are connected via an ev-feasible path. On that basis ERC can be rephrased as determining a SCDS in the respective reachability graph, see Figure 1. The two conditions given above are then naturally fulfilled, since choosing $L = D$ for a solution D of SCDS assures that for any non-loading station there is an ev-feasible path towards/from a loading stations and moreover any strongly connected node set in RG implies a strongly ev-connected node set in G, if the contained nodes are converted into LSs.

This reformulation immediately leads to a $3 \ln n$ approximation for ERC via the SCDS algorithm in (Li et al. 2009). Note that this is asymptotically optimal as no polytime algorithm can approximate SCDS within $c \ln n$ for $c < 1$(Kann 1992; Li et al. 2009). To prove that the same inapproximability bound applies for ERC, we observe that for an instance of ERC with $c(e) = M$ for every edge $e \in E$ the reachability graph equals the original one ($RG = G$). Therefore given an instance of SCDS, we can construct an instance of ERC in polytime by augmenting the graph with the cost function $c : E \to M$ for an arbitrary selected $M \in \mathbb{R}^+$. There is a one-to-one correspondence between the set of loading stations L and the strongly dominating set D, with $|L| = |D|$. Therefore any approximation guarantee for ERC is transferable to SCDS, and hence ERC must be at least as hard to approximate:

Theorem 5 *ERC can be $3 \ln n$ approximated in polynomial time; there exists no polynomial time algorithm approximating ERC within $c \ln n$ for $c < 1$.*

Hence, if the goal are algorithms with better approximation guarantees (like a constant factor approximation or even an approximation scheme), one has to take into account the special structure of the underlying graph, e.g. (near) planarity, or other graph properties like treewidth etc. (which is far from obvious how to do it).

Getting Anywhere and Back

While ensuring EV-reachability might allow us to get anywhere we want, we might be stuck there forever if no reloading facility is available at the destination. So what we really want is to get anywhere and back again – without getting stuck. Hence for any source/target pair $s, t \in V$ there must exist an ev-feasible roundtour $s, \cdots, t, \cdots, s$. This requirement implies strong ev-connectivity between any pair of nodes.

Definition 6 (EV-Connectivity Cover) *Given a (di)graph $G(V, E, c)$ and $M \in \mathbb{R}^+$, we refer to the problem of converting a subset $L \subseteq V$ of minimal cardinality into loading stations such that $\forall v \in V : C^L(v) = V$ as EV-Connectivity Cover (ECC).*

Again the problem can be rephrased (if $L = \emptyset$ is not a solution), leading to the following alternative formulation: Select a minimum size subset $L \subseteq V$ as loading stations such that:

1. $\forall v \in V \exists l, l' \in L : p = l, \cdots, v, \cdots, l'$ is ev-feasible
2. $\forall l \in L : C^L(l) \supseteq L$.

The first condition assures that there exists a LS from which the node v can be reached with a sufficient high final battery load such that continuing to another (or the same) loading station is possible for sure. The second condition is identical to the ERC problem formulation.

Let us first define three related problems of the above: in the weak ECC (wECC) problem, we only require condition one (existence of l, l'). In the simple ECC (sECC) problem, we insist on both conditions but additionally require for the first condition $l = l'$. Note that sECC is a rather natural problem variant, as nodes with differing l and l' only seem to be required in mountainous areas or very long one-way streets. Also observe that sECC is equivalent to ECC if the input graph is undirected. Finally, in the simple weak ECC (swECC) problem, only the first condition has to be met with $l = l'$.

Approximation Algorithms

Obviously any solution to ECC is also feasible for ERC. Hence the necessary number of loading stations to solve the ECC problem is lower bounded by the optimal ERC solution for the same input. As we have seen a SCDS in the respective reachability graph solves ERC and in particular ensures the second condition. The goal is now to augment this set of loading stations to extend the ERC solution to an ECC solution.

Lemma 7 *Given an instance of (s)ECC. A solution for the respective (s)wECC problem together with a solution for SCDS in the respective reachability graph provides a feasible solution for (s)ECC.*

Proof. Let L' be the solution for the (s)wECC problem, D the SCDS in RG and $L = L' \cup D$ their union. The first condition is naturally fulfilled by L because of the nodes in L', hence it remains to show that L is strongly ev-connected. Observe that any node $l \in L'$ is either contained in D or has an in-neighbour d_1 and an out-neighbour d_2 in RG with $d_1, d_2 \in D$. Therefore there exists an ev-feasible path to and from any loading station in D, hence the nodes in L are pairwise strongly ev-connected.■

Assuming we can solve (s)wECC with an approximation guarantee of α, we would obtain $(3 \ln n + \alpha)$ approximation guarantee for (s)ECC (since SCDS as well as (s)wECC lower bound (s)ECC).

We first consider swECC. Here the task is placing a minimum number of loading stations such that their strongly ev-connected node sets cover the whole network. This can be solved via an instance of Set Cover (SC) which is known to be approximable within $\ln n$ using a standard greedy algorithm and inapproximable to anything better than $c \ln n$ for $c < 1$ (Feige 1998):

Definition 8 (Set Cover) *Given a universe of elements $U = \{1, \cdots, n\}$ and a collection of subsets $S = \{S_1, \cdots, S_k\}, S_i \subseteq U$, select a minimum subset $S' \subseteq S$ with $\bigcup_{S_i \in S'} S_i = U$.*

For our application the universe is the set of nodes $U = V$ and the collection of subsets is defined as the strongly ev-connected node sets of all nodes, i.e. $S_i = C(v_i)$ for $i = 1, \cdots, n$. Therefore swECC can be solved with an approximation guarantee of $\ln n$ in polynomial time via the standard greedy SC algorithm. Hence sECC can be solved within a factor $4 \ln n$ of the optimal solution in polynomial time.

Again the question arises whether a better approximation guarantee might be possible. We will prove that the bound is asymptotically optimal by a solution preserving reduction from DS (which cannot be approximated better than $c \cdot \ln n$ for $c < 1$ since it is just a variant of SC):

Given an instance $G(V, E)$ of DS, we create an instance of (s)ECC by choosing $M \in \mathbb{R}^+$ arbitrarily and setting $c(e) = M/3$ for all $e \in E$.

Lemma 9 *Every DS $D \subseteq V$ in $G(V, E)$ is a (s)ECC in $(G, V, E, c), M$.*

Proof. We show that setting $L = D$, we can guarantee a battery load of $b(v) \geq 2M/3$ for any $v \in V$ when starting at an arbitrary source node $s \in V$. Let $p = s, \cdots, w$ be an arbitrary path in G and $v \in V$ the first vertex in p that is reached with a battery load of $M/3$. Because D is a DS either v or one of its neighbours must be a LS. Therefore visiting this LS and returning to v leads to battery load of $\geq 2M/3$. Hence the successor of v in p can again be reached with a battery load $\geq M/3$. Repeating the argument, $C^L(v) = V$ is obviously fulfilled for every $v \in V$.■

Lemma 10 *Every non-empty ECC $L \subseteq V$ in $(G, V, E, c), M$ is a DS in $G(V, E)$.*

Proof. We first show that L is empty iff G is a complete graph. If G is complete every energy-optimal roundtour from u to v and back uses $\{u, v\} \in E$ twice and therefore consumes energy equal to $2M/3$, which is of course ev-feasible. If G is not the complete graph there must be nodes u, w which are not connected via a direct edge. Hence the energy-optimal path from u to w and back must visit at least one other node and therefore consumes in total $4M/3$, which is not ev-feasible. Hence if G is incomplete it follows $L \geq 1$. Accordingly if $L = \emptyset$ setting $D = \{v\}$ for an arbitrary node $v \in V$ solves DS in $G(V, E)$. From now on assume G is incomplete. Furthermore assume there exists a star (tree of depth 1) T rooted in $v \in V$ and containing all nodes adjacent to v but no LS. If $V(T) = V$, i.e. the star contains all graph nodes, then there must exist neighbours $u_1, u_2 \in V$ of v which are not connected via a direct edge (because G is incomplete). Therefore one possible energy-optimal path from one node to the other passes v and therefore the roundtour from u_1 over u_2 is not ev-feasible which contradicts L being a ECC. Now let $V(T) \subset V$, therefore there exist $w \in V \setminus V(T), u \in V(T)$ with $\{w, u\} \in E$. Accordingly when starting at w the highest possible battery load in v is $b(v) = {}^{M}\!/3$. Using any edge $\{v, x\} \in E$ to complete the roundtour back to w we receive $b(x) = 0$ and hence the EV is stuck at x. Hence $v \notin C^L(w)$, contradiction to L is (s)ECC. So if G is incomplete no induced star in G is free of loading stations, therefore for any node v either $v \in L$ or it exists $u \in L$ with $\{v, u\} \in E$. Therefore $D = L$ is a DS in G.■

Theorem 11 *There exists no polytime algorithm that approximates (s)ECC within $c \ln n$ for $c < 1$.*

Proof. According to Lemma 9 and 10 an optimal solution for DS can be obtained by solving the respective instance of (s)ECC. Therefore any polytime approximation algorithm for (s)ECC solves DS with the same approximation guarantee. So the inapproximability bound for DS transfers to (s)ECC. ■

Again, this result implies that without making use of the special structure of the underlying road network, there is little hope to efficiently achieve a provable approximation guarantee better than $O(\log n)$.

For the wECC problem, things get even worse and the proposed technique to achieve a $O(\log n)$ approximation for swECC does not carry over since the choice of loading stations is mutually dependent. wECC is closely related to the much harder Set Cover by Pairs (SCP) (Hassin and Segev 2005) problem.

Definition 12 (Set Cover by Pairs) *Given a universe of elements $U = \{u_1, \cdots, u_n\}$, a set of cover elements A and for every $\{a, a'\} \subseteq A$ a subset $S_{a,a'} \subseteq U$, select a minimum subset $A' \subseteq A$ such that $\bigcup_{\{a,a'\} \in A'} S_{a,a'} = U$.*

For wECC the universe of elements and the set of cover objects equal both the set of nodes, so $U = A = V$. For any two nodes, the nodes that can be covered are the ones on ev-feasible paths between them, i.e. $S_{u,w} = \{v \in V | p = u, \cdots, v, \cdots, w \text{ is ev-feasible}\}$ (hence $S_{u,u} = C(u)$). The selection of a subset of cover objects instead of a subcollection of sets allows to take care of the dependencies, so e.g. choosing the nodes a, b, c automatically all of the nodes in $S_{a,a}, S_{a,b}, S_{a,c}, S_{b,b}, S_{b,c}$ and $S_{c,c}$ are covered.
The best known SCP algorithm provides a $\mathcal{O}(\sqrt{n \log n})$ approximation for the general case, hence wECC and according to Lemma 7 also ECC can be solved with the same guarantee. Moreover SCP was proven to be inapproximable within $2^{\log^{(1-\epsilon)} n}$ for any $\epsilon > 0$ under the assumption $NP \nsubseteq$

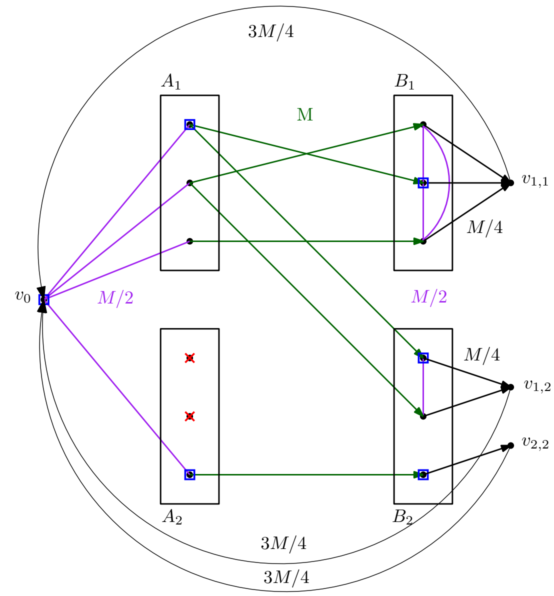

Figure 2: Derived ECC instance for a given instance of MR with $k = 3$ and $q = 2$. Original edges are coloured green. A possible solution for ECC is implied by the blue boxes. The restriction to nodes inside the large black boxes provides the respective solution for MR.

$DTIME(n^{\mathcal{O}(polylog(n))})$ by reduction from Min Rep (MR) (Breslau et al. 2011).

Definition 13 (Min Rep (Kortsarz 2001)) *Let $G(A, B, E)$ be a bipartite graph with A and B being partitioned into q subsets of size k, i.e. $A = \biguplus A_i, B = \biguplus B_i, i = 1, \cdots, k$ with $|A_i| = |B_i| = k$. The graph $SG(A', B', E')$ is called the supergraph of G, with A' and B' containing a node for each of the partitions of A or B, and $E' := \{(A_i, B_j) \in A' \times B' | a \in A, b \in B_j, \{a, b\} \in E\}$ being the set of superedges. The goal is to select a minimum size $W \subseteq V = A \uplus B$ such that the subgraph of G induced by V' has the same supergraph as G.*

It was proven that the NP-hardness of MR remains the same even if G is a collection of stars rooted in A. We will show that MR with the star-property can be reduced to ECC such that the inapproximability bound transfers. To that end we create for a given MR instance an instance $G(V', E', c), M$ of ECC as follows: First we prune all nodes with a degree of zero. Then we direct all edges in E from A to B and augment them with costs $c(e) = M$ for some fixed chosen $M \in \mathbb{R}^+$. Furthermore for every superedge (A_i, B_j) we create a dummy node v_{ij} and augment the graph with edges (b, v_{ij}) for $b \in B_j, (a, b) \in E, a \in A_i$. These edges receive a cost of ${}^{M}\!/4$. Moreover we insert edges with costs of ${}^{M}\!/2$ between any two nodes in b that are adjacent to the same dummy node. Finally we insert a further dummy node v_0. We connect v_0 to A by adding all edges $\{v_0, a\}$ for $a \in A$ with costs ${}^{M}\!/2$. Also we insert directed edges (v_{ij}, v_0) for all previously created dummy node with costs ${}^{3M}\!/4$; see Figure 2 for an example.

Lemma 14 *Every solution W for MR augmented with v_0 is an ECC.*

Proof. We first show that $L = W \cup v_0$ is a wECC, i.e. for all nodes exists an ev-feasible path between two LSs in L that visits this node. For all $a \in A$ this is trivially true because $a \in C(v_0)$. For all dummy nodes v_{ij} at least one of the adjacent nodes $b \in B$ must be in L, otherwise W would not realize the superedge (A_i, B_j), hence b, v_{ij}, v_0 is ev-feasible. Also it follows that for all $b \in B$ either $b \in L$ or $e = \{b, b'\} \in E'$ with $c(e) = M/2$ and $b' \in L$. It remains to show that L is strongly ev-connected. To that end observe that for any superedge realization $a_i, b_j \in W$ the roundtour v_0, a, b, v_{ij}, v_0 is ev-feasible. ∎

Lemma 15 *For every ECC L the set $L \cap V$ is a MR solution.*

Proof. We have to prove that in $W = L \cap V$ every superedge is realized. Assume for contradiction that (A_i, B_j) is not covered by W. Observe that any path from v_0 to v_{ij} has to traverse one of the candidate edges (a_i, b_j) (because of the star property). If $a_i \notin L$ the node b_j can not be reached via this edge, because the summed costs would exceed M. If $a_i \in L$ but $b_j \notin L$ the battery load at b_j equals 0, therefore the EV would be stuck there. Hence if no candidate edge is realized it yields $v_{ij} \notin R(v_0), C(v_0)$, which contradicts L being an ECC. Moreover we observe that v_0 always has to be part of L because all paths from some v_{ij} to A have to pass this node and the costs would otherwise be $5M/4$. Finally if L is an optimal solution it follows $|W| = |L| - 1$ because any LS placed at some dummy node v_{ij} is superfluous as there exists an ev-feasible roundtour $b_j, v_{ij}, v_0, a_i, b_j$ in any case. ∎

Theorem 16 *If $NP \nsubseteq DTIME(n^{\mathcal{O}(polylog(n))})$, no polytime algorithm can be guaranteed to solve ECC within a factor of $2^{(0.5log(n))^{1-\epsilon}-1}$.*

Proof. From Lemma 14 and 15 we obtain for a given (approximate) solution L' of the ECC instance an (approximate) solution W' for MR with $|W'| \leq |L'| - 1$. Moreover for an optimal solution L equality holds, i.e. $|W| = |L| - 1$. Therefore any α approximation for ECC leads to a solution W' for MR with $|W'| \leq (\alpha + 1)|W| + \alpha \leq (2\alpha + 1)|W|$. Assume there would exist a polytime algorithm with approximation factor $\alpha < 2^{(0.5log(n))^{1-\epsilon}-1}$ for ECC. Because the number of nodes for the ECC instance is bounded by $2kq + k^2 + 1 \leq 2(kq)^2$ this would result in a $2\alpha < 2^{log^{1-\epsilon} n}$ approximation for MR, which would contradict the known inapproxmability bound.∎

In particular the results for ECC are rather discouraging. Unless one can make use of the special properties of the road network things seem pretty hopeless. In the following we will see, that real-world instances do not exhibit this hardness, and results with quality guarantees can be obtained efficiently – they can only be given on a per instance basis and not a priori, though.

Heuristics and Instance Based Approximation Guarantees

Not only do the inapproximability results for ERC and ECC seem rather discouraging for solving these problems in practice, but the described approximation algorithms also require the knowledge of $R(v)/R^{-1}(v)/C(v)$ for all $v \in V$. Normally these sets are not given explicitly but have to be computed on the basis of $G(V, E, c)$ and M. For large cruising ranges this task is comparable to solving the all pair shortest path problem, moreover it would require almost quadratic space to store the resulting sets. Therefore we aim for heuristics that need only a (small) subcollection of these sets.

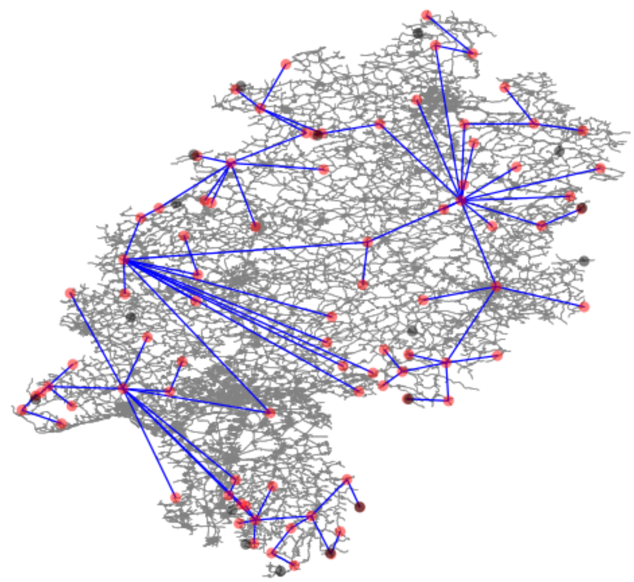

Figure 3: wERC solution (red) along with a PEC (black) providing a lower bound, and the two overlayed MSTs (blue edges) which are the basis for lifting the wERC to an ERC.

To solve ERC/ECC we proceed in two phases: In the first phase we create a weak ERC/ECC L' which is augmented in the second phase to a strongly ev-connected superset $L \supseteq L'$. For the first phase we propose two strategies: *Restricted Random (RR)* and *k-Greedy (kG)*. RR selects in each round a vertex v u.a.r and adds it to the solution. Then the search space is reduced by all the nodes for which the weak cover condition is already fulfilled. kG always maintains a set of k candidates for the next LS, preferring the one which increases the number of nodes satisfying the first condition most. Then the list is refilled by selecting an arbitrary node which is not already weakly covered.

If the resulting set L' is not already strongly ev-connected (which can easily be checked in RG), we compute the most energy-efficient path between each pair of LSs in L'. Then we build a complete digraph on L' with edge costs equal to the costs of the respective energy-optimal paths. On that basis we select an arbitrary LS $s \in L'$ and compute the *minimum spanning tree (MST)* rooted in s in this auxiliary graph and the respective MST in the reversed auxiliary graph (compare (Li et al. 2009)). For all edges contained in one of the trees we make the respective energy-optimal path between the LSs feasible by placing additional LSs as required. Obviously the resulting set of LSs is strongly ev-connected. In Figure 3, we provide an illustration of our approach on a small example graph. The MSTs in the image, which are based on the energy metric, connect sometimes LSs directly which seem to be very far apart. Note, that this is the result of some close nodes being only reachable via steep climbs and therefore being not close in terms of energy costs (the artificially small cruising range chosen for visualisation purposes amplifies this effect).

As heuristics like the ones suggested above cannot guarantee to find a good solution a priori, we will provide instance based lower bounds which can be used to prove the quality of the retrieved solution. These bounds are based on solving the Partial EV-Cover (PEC) problem.

Definition 17 (Partial EV-Cover) *Given a (di)graph $G(V, E, c)$ and a maximal battery load M, convert a minimum subset $L \subseteq V$ into loading stations such that $\forall v \in V \exists l \in L : v \in R(l)$.*

Observe that any solution of ERC or ECC would also be a (non optimal) solution for PEC on the same input. Hence an optimal solution for PEC would be a valid lower bound for the optimal solution of ERC/ECC. But to solve PEC exactly we must again have knowledge about every $R(v)$ or $R^{-1}(v)$ respectively. Luckily we can find a lower bound for the optimal solution of PEC which requires less computational effort: We call two vertices $v, v' \in V$ *distinct*, if $\nexists w \in V : v, v' \in R(w)$. Observe that this is equal to demand that $R^{-1}(v) \cap R^{-1}(v') = \emptyset$, therefore this property can be checked efficiently. Note that any set of distinct vertices provides a lower bound for the solution of PEC. A maximal set of distinct nodes can be retrieved easily by selecting repeatedly vertices $v \in V$ and checking whether they are distinct wrt to the previously selected ones; if this is the case they are added to the solution, otherwise they get excluded from the search space.

We implemented the proposed heuristics and the lower bound algorithm in C++ and evaluated them on the street graph of Germany (GER, based on OpenStreetMap[1]) with $n = 15.015.877, m = 30.771.648$ for varying EV cruising ranges. The edge costs were chosen as proposed in (Eisner, Funke, and Storandt 2011) as a linear combination of euclidean distance and the difference of the incident node's elevations (retrieved from SRTM[2]). The results of our experiments can be found in Table 1. We observed, that using RR or kG with $k = 10.000$ did not produce a very different output for (w)ERC (hence we reported the results of only one strategy for each problem type). But we could save about 10% of the LSs using kG instead of RR for (w)ECC (but paying with a twenty times longer runtime, namely about 9h30min instead of 30 min). Because of this redundancy and due to space restrictions, we did not put all the numbers in the table, but report only the explicit results for solving ERC with RR and ECC with kG. The ratio of obtained solution and lower bound ranges from 2.7 to 3.3 for ERC and from 7.2 to 9.8 for ECC. Noticing that the lower bound does not take connectivity of LSs into account, our resulting sets seem to be very close to the optimal one in size.

[1]http://www.openstreetmap.org

[2]http://www2.jpl.nasa.gov/srtm

ACR	lower bound	wERC RR	ERC +MST	wECC kG	ECC + MST
75 km	105	194	+ 145	808	+ 4
100 km	49	96	+ 51	372	+ 7
125 km	37	61	+ 42	255	+ 13
150 km	22	39	+ 22	170	+ 17
175 km	14	26	+ 20	118	+ 20

Table 1: Experimental results for the GER graph. ACR denotes the average cruising range of the EV. Values are the maximum (lower bound) or minimum of three runs. kGreedy (kG) was performed with $k = 10.000$.

Conclusions

In this paper we investigated the problem of placing as few loading stations as possible for electric vehicles to achieve mobility goals like reachability or connectivity. Unfortunately, it turned out that if the special structure of the road network is not made use of, the respective optimization problems are all very hard to approximate. In particular, even for the easier goal of establishing reachability only, it is not possible to come up with a polynomial-time constant approximation algorithm. On the positive side, we devised simple heuristic algorithms which do not provide any priori approximation guarantee. Still, using instance-based lower bounds we could prove their results to be close-to-optimal for our real-world problem instances derived from the road network of Germany.

Acknowledgement

This work was partially supported by the Google Focused Grant Program on Mathematical Optimization and Combinatorial Optimization in Europe.

References

Artmeier, A.; Haselmayr, J.; Leucker, M.; and Sachenbacher, M. 2010. The shortest path problem revisited: Optimal routing for electric vehicles. In *KI*, 309–316.

Breslau, L.; Diakonikolas, I.; Duffield, N. G.; Gu, Y.; Hajiaghayi, M. T.; Johnson, D. S.; Karloff, H. J.; Resende, M. G. C.; and Sen, S. 2011. Disjoint-path facility location: Theory and practice. In *ALENEX*, 60–74. SIAM.

Eisner, J.; Funke, S.; and Storandt, S. 2011. Optimal route planning for electric vehicles in large networks. In *AAAI*.

Feige, U. 1998. A threshold of ln n for approximating set cover. *J. ACM* 45(4):634–652.

Hassin, R., and Segev, D. 2005. The set cover with pairs problem. In Sarukkai, S., and Sen, S., eds., *FSTTCS 2005:*

Foundations of Software Technology and Theoretical Computer Science, volume 3821 of *Lecture Notes in Computer Science*. Springer Berlin / Heidelberg. 164–176.

Kann, V. 1992. On the approximability of np-complete optimization problems. In *PhD thesis*.

Kortsarz, G. 2001. On the hardness of approximating spanners. *Algorithmica* 30:432–450. 10.1007/s00453-001-0021-y.

Li, D.; Du, H.; Wan, P.-J.; Gao, X.; Zhang, Z.; and Wu, W. 2009. Construction of strongly connected dominating sets in asymmetric multihop wireless networks. *Theoretical Computer Science* 410(810):661 – 669.

Storandt, S., and Funke, S. 2012. Cruising with a battery-powered vehicle and not getting stranded. In *AAAI*.