# Shortest Feasible Paths with Charging Stops for Battery Electric Vehicles[*]

Moritz Baum[1], Julian Dibbelt[2], Andreas Gemsa[1], Dorothea Wagner[1] and
Tobias Zündorf[1]

[1]Karlsruhe Institute of Technology, Germany,
`moritz@ira.uka.de,gemsa@ira.uka.de,dorothea.wagner@kit.edu,tobias.zuendorf@kit.edu`
[2]Sunnyvale, CA, United States, `algo@dibbelt.de`

### Abstract

We study the problem of minimizing *overall* trip time for battery electric vehicles in road networks. As battery capacity is limited, stops at charging stations may be inevitable. Careful route planning is crucial, since charging stations are scarce and recharging is time-consuming. We extend the Constrained Shortest Path problem for electric vehicles with *realistic* models of charging stops, including varying charging power and battery swapping stations. While the resulting problem is $\mathcal{NP}$-hard, we propose a combination of algorithmic techniques to achieve good performance in practice. Extensive experimental evaluation shows that our approach (CHArge) enables computation of *optimal* solutions on realistic inputs, even of continental scale. Finally, we investigate heuristic variants of CHArge that derive high-quality routes in well below a second on sensible instances.

## 1 Introduction

Electromobility promises independence of fossil fuels, zero (local) emissions, and higher energy-efficiency, especially for city traffic. Yet, most of the significant algorithmic progress on route planning in road networks has focused on conventional combustion-engine cars; see Bast et al. [2] for an overview. Battery electric vehicles (EVs), however, are different: Most importantly, for the foreseeable future, they have limited driving range, charging stations are still much rarer than gas stations, and recharging is time-consuming. Thus, routes can become infeasible (the battery runs empty), and fast routes may be less favorable when accounting for longer recharging time. In this work, we aim to find the overall fastest (capacity-constrained) route, considering charging stops when necessary.

---

**Related Work.** Classic route planning approaches apply Dijkstra's algorithm [18] to a graph representation of the transportation network, using fixed scalar arc weights corresponding to, e. g., driving time. For faster queries, *speedup techniques* have been proposed, with different benefits in terms of preprocessing time and space, query speed, and simplicity [2]. A* search [32] uses vertex potentials to guide the search towards the target. A successful variant, ALT (A*, Landmarks, Triangle Inequality) [27], obtains good potentials from precomputed distances to selected landmark vertices. Contraction Hierarchies (CH) [24], on the other hand, iteratively *contract* vertices in increasing order of (heuristic) importance during preprocessing, maintaining distances between all remaining vertices by adding *shortcut* arcs where necessary. The CH query is then bidirectional, starting from source and target, and proceeds only from less important to more important vertices. Combining both techniques, Core-ALT [4] contracts all but the most important vertices (e. g., the top 1 %), performing ALT on the remaining *core* graph. This approach can also be extended to more complex scenarios, such as edge constraints to model, e. g., maximum allowed vehicle height or weight [23]. More recently, techniques (including variants of CH and ALT) were introduced that allow an additional *customization* after preprocessing, to account for dynamic or user-dependent metrics [11, 17, 19]. Also, approaches towards extended scenarios exist, such as batched shortest paths [12] or shortest via paths [1, 14]. Here, a common approach is to make use of a relatively fast *target selection* phase, precomputing distances to relevant points of interest to enable faster queries. However, these techniques were only evaluated for single-criterion search, where the distance between two vertices is always a unique scalar value. For multi-criteria scenarios, on the other hand, problem complexity and solution sizes increase significantly, and practical approaches are only known for basic problem variants [21, 23, 56]. For a more complete overview of techniques and combinations, see Bast et al. [2].

Regarding route planning for EVs, the battery must not fully deplete during a ride, but energy is recuperated when braking (e. g., going downhill), up to the maximum capacity. These *battery constraints* must be checked during route computation. This can be done implicitly, by weighting each road segment with a *profile*, mapping current *state of charge (SoC)* to actual consumption wrt. battery constraints [10, 20, 54]. If battery constraints have to be obeyed, both preprocessing and queries of speedup techniques are slower compared to the standard shortest path problem [7, 20, 57]. Several natural problem formulations specific to EV routing exist. Some optimize energy consumption [7, 10, 20, 49, 54], often, however, resulting in routes resorting to minor (i. e., slow) roads to save energy. *Constrained Shortest Path (CSP)* formulations [30] ask to find the most energy-efficient route without exceeding a certain driving time—or finding the fastest route that does not violate battery constraints [55]. More recent works additionally take speed planning into account to trade driving speed for energy consumption [9, 33]. These works derive exact and heuristic algorithms, which achieve good query times after preprocessing the input network, but do not include charging stops in route optimization. Another approach separates queries into two phases, optimizing driving time and energy consumption, respectively [28]. This allows for reasonably fast queries (even without preprocessing), but results are not optimal in general, i. e., there might be paths with lower trip time that respect battery constraints.

Without recharging, large parts of the road network are simply not reachable by an EV, rendering long-distance trips impossible. For conventional cars, broad availability of gas stations and short refuel duration allow to neglect this issue in route optimization. Strehler et al. [58] give theoretical insights into the CSP problem for EVs including charging stops. Most importantly, they develop a fully polynomial-time approximation scheme (FPTAS) for this problem. Unfortunately, the algorithm is slow in practice.

Often, previous works dealing with the CSP problem for EVs have considered charging stations under the simplifying assumptions that the charging process takes *constant* time (independent of

the initial SoC) and always results in a *fully* recharged battery [28, 55, 57]. Then, feasible paths between charging stations are independent of source and target, hence easily precomputed. Routes with a minimum number of intermediate charging stops can then be computed in less than a second on subcountry-scale graphs [55, 57]. Nevertheless, the simple model used in these works only applies to battery swapping stations, which are still an unproven technology and business model.

For regular charging stations, charging time depends on both the current SoC and the desired SoC, as well as the charging power provided by the station. Kobayashi et al. [38] distinguish slow and fast charging stations, but assume that the battery is always fully recharged. They propose a heuristic search based on preselected candidates of charging stations for a given pair of source and target. Their approach takes several seconds for suboptimal results on metropolitan-scale routes.

In reality, while nearly linear for low SoC, the charging rate decreases when approaching the battery's limit. Thus, it can be reasonable to only charge up to a *fraction* of the limit. Sweda et al. [59] reflect this behavior by combining a linear with an exponential function for high SoC. Liu et al. [41] model charging between 0 % and 80 % SoC with a linear function, but recharging above 80 % SoC is suppressed altogether. Neither approach was shown to scale to road networks of realistic size. Also, even though omitting the possibility of charging beyond 80% might be appropriate for regions well covered with charging stations, it drastically deteriorates reachability in regions with few stations, where recharging to a full battery can be inevitable [44].

Other works discretize possible charging durations to model different options [35, 58, 63]. This enables search algorithms that closely resemble the well-known bicriteria shortest path algorithm [31].

Some recent works only optimize energy consumption along routes with charging stops, while ignoring travel time entirely to simplify the problem setting [10, 57]. A technique based on dynamic programming to optimize a certain cost function on a route with charging stops is given by Sweda and Klabjan [60], extending previous approaches on refueling strategies for conventional cars [37, 40]. Similarly, Liao et al. [39] use dynamic programming to solve different routing problems for EVs, including the possibility of battery swaps along the way.

The publications mentioned so far deal with shortest path problems for a *single* EV. In addition to that, charging stops are also considered in several works dealing with *vehicle (fleet) routing problems* for EVs [15, 26, 39, 44, 45, 47, 51, 52, 64]. These problems are handled by means of mathematical programming or heuristics. Aiming at complex scenarios (often involving optimization for *multiple* vehicles and variants of the Traveling Salesman Problem), the approaches do not scale to large input instances. Furthermore, most of them rely on the simplified charging models discussed above, e. g., allowing only battery swaps or (uniform) linear charging. A notable exception is the work of Montoya et al. [44], which explicitly models nonlinear charging processes to achieve more realistic results when computing charging stops for a vehicle along a given route. In fact, it builds upon the model we propose (and generalize) in this work, available from previous publications [6, 65].

**Contribution and Outline.**  In this work, we extend the CSP problem to planning routes that, while respecting battery constraints, minimize overall trip time, including time spent at charging stations. Unlike previous works, our solution handles *all* types of charging stations accurately: battery swapping stations, regular charging stations with various charging powers, as well as superchargers. In particular, charging times are *not* independent of the current SoC when arriving at a charging station. Additionally, the charging process can be interrupted as soon as further charging would delay the arrival at the target. This results in a challenging, $\mathcal{NP}$-hard problem, for which even the construction of basic exponential-time algorithms is nontrivial. We propose a label-setting search, which is capable of propagating *continuous* tradeoffs induced by charging stops using labels of *constant* size. Since the problem is $\mathcal{NP}$-hard, we do not guarantee polynomial running times. However, carefully incorporating recharging models in speedup techniques, our approach is the first to

solve a realistic setting *optimally* and within seconds, for road networks of country-scale and beyond. Key ingredients of our speedup techniques are a bicriteria extension of CH and generalizations of A* search, where we incorporate SoC and required charging time into the derived bounds. For even faster queries, we propose heuristic approaches that offer high (empirical) quality. Extensive experiments on detailed and realistic data show that our approach produces sensible results and, even though it is designed to solve a more complex and realistic problem, clearly outperforms and broadens the state-of-the-art.

The remainder of this work is organized as follows. Section 2 sets necessary notation and specifies our model and the problem. Section 3 presents our basic label-propagating approach. It turns out that as its most crucial ingredient, a limited number of new labels must be generated at charging stations to ensure termination and correctness of the approach. We derive methods to construct such labels in Section 4. To improve practical performance of our exponential-time algorithm, we propose tuning based on A* search in Section 5 and the speedup technique CH in Section 6. We also discuss how both techniques can be combined to further reduce practical running times in Section 7. Section 8 introduces heuristics, which drop correctness for faster queries. Section 9 experimentally evaluates all approaches on large, realistic input. We close with final remarks in Section 10.

## 2 Model and Problem Statement

We consider *directed, weighted* graphs $G = (V, A)$ with vertices $V$ and arcs $A \subseteq V \times V$, together with two arc weight functions $d \colon A \to \mathbb{R}_{\geq 0}$ and $c \colon A \to \mathbb{R}$, representing driving time and energy consumption on an arc, respectively. Vertices are *neighbors* if they are connected by an arc. We simplify notation by defining $d(u, v) := d((u, v))$ and $c(u, v) := c((u, v))$ for an arc $(u, v) \in A$.

An *s–t path* in a graph $G$ is defined as a sequence $P = [s = v_1, v_2 \ldots, v_k = t]$ of vertices, such that $(v_i, v_{i+1}) \in A$ for $1 \leq i < k$. If $s = t$, we call $P$ a *cycle*. Given two paths $P = [v_1, \ldots, v_i]$ and $Q = [v_i, \ldots, v_k]$, we denote by $P \circ Q := [v_1, \ldots, v_i, \ldots, v_k]$ their concatenation. The *driving time* on a path $P$ is $d(P) = \sum_{i=1}^{k-1} d(v_i, v_{i+1})$. For energy consumption, this is more involved: First of all, note that consumption can be negative due to recuperation, though cycles with negative consumption are physically ruled out. Moreover, the battery has a limited capacity $M \in \mathbb{R}_{\geq 0}$, and the SoC can neither exceed this limit nor drop below 0. Therefore, we must take *battery constraints* into account. For an *s–t* path $P$ to be *feasible*, the battery's SoC must be within the interval $[0, M]$ at every vertex of $P$. To reflect battery constraints on an *s–t* path $P$, we define its *SoC profile* $b_P \colon [0, M] \to [0, M] \cup \{-\infty\}$ that maps SoC $\beta_s$ at the source $s$ to the resulting SoC $\beta_t$ at the target $t$. Hence, the SoC at $t$ is $\beta_t = b_P(\beta_s)$; see Figure 1 for an example. We use the value $-\infty$ to indicate that the SoC at $s$ is not sufficient to reach $t$, i.e., $P$ is not feasible for the corresponding SoC at $s$. As shown by Eisner et al. [20], $b_P$ can be represented using only three values, namely, the minimum SoC $\mathrm{in}_P \in [0, M]$ required to traverse $P$, its energy consumption $\mathrm{cost}_P \in [-M, M]$ (which can be less than $\mathrm{in}_P$ due to recuperation), and the maximum possible SoC after traversing $P$, denoted $\mathrm{out}_P \in [0, M]$. We then have

$$
b_P(\beta) := \begin{cases} -\infty & \text{if } \beta < \mathrm{in}_P, \\ \mathrm{out}_P & \text{if } \beta - \mathrm{cost}_P > \mathrm{out}_P, \\ \beta - \mathrm{cost}_P & \text{otherwise.} \end{cases}
$$

The SoC profile $b_{[v]}$ of a path consisting of a single vertex $v \in V$ is given by $\mathrm{in}_{[v]} := 0$, $\mathrm{cost}_{[v]} := 0$, and $\mathrm{out}_{[v]} := M$, which yields $b_{[v]}(\beta) = \beta$ for all $\beta \in [0, M]$. For an arc $a = (u, v) \in A$, the profile $b_a = b_{[u,v]}$ is given by $\mathrm{cost}_a := c(a)$, $\mathrm{in}_a := \max\{0, c(a)\}$, and $\mathrm{out}_a := \min\{M, M - c(a)\}$. For two
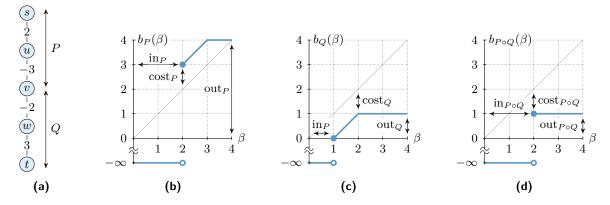
Figure 1: Two SoC profiles and the result after linking them. The battery capacity is $M = 4$. (a) The underlying paths $P = [s, u, v]$ and $Q = [v, w, t]$, with indicated arc costs. (b) The SoC profile $b_P$ is represented by $\text{in}_P = 2$, $\text{cost}_P = -1$, and $\text{out}_P = 4$. (c) The SoC profile $b_Q$ is given by $\text{in}_Q = 1$, $\text{cost}_Q = 1$, and $\text{out}_Q = 1$. (d) Linking the profiles $b_P$ and $b_Q$ yields the profile $b_{P \circ Q}$ with $\text{in}_{P \circ Q} = 2$, $\text{cost}_{P \circ Q} = 1$, and $\text{out}_{P \circ Q} = 1$. Observe that due to a subpath of length $-5 < -M$, the value $\text{cost}_{P \circ Q} = 1$ is greater than the sum $\text{cost}_P + \text{cost}_Q = c(P \circ Q) = 0$.

SoC profiles $b_P$ and $b_Q$ of paths $P$ and $Q$, we obtain the *linked* profile $b_{P \circ Q}$ by setting

$$\text{in}_{P \circ Q} := \max\{\text{in}_P, \text{cost}_P + \text{in}_Q\}$$
$$\text{out}_{P \circ Q} := \min\{\text{out}_Q, \text{out}_P - \text{cost}_Q\}$$
$$\text{cost}_{P \circ Q} := \max\{\text{cost}_P + \text{cost}_Q, \text{in}_P - \text{out}_Q\},$$

provided that $\text{out}_P \geq \text{in}_Q$. Otherwise, the path is infeasible for arbitrary SoC, and thus $b_{P \circ Q} \equiv -\infty$. An example of two SoC profiles as well as the result of linking them is shown in Figure 1. Finally, given two SoC profiles $b_1$ and $b_2$, we say that $b_1$ *dominates* $b_2$ if $b_1(\beta) \geq b_2(\beta)$ holds for all $\beta \in [0, M]$.

**Constrained Shortest Paths.** Given a graph, two (nonnegative) functions representing length (driving time in our case) and consumption on its arcs, and vertices $s \in V$ and $t \in V$, the CSP Problem asks for a path of minimum length such that its consumption does not exceed a certain bound $\beta_s \in \mathbb{R}_{\geq 0}$. Being $\mathcal{NP}$-hard [22, 30], CSP can be solved using an exponential-time bicriteria variant of Dijkstra's algorithm [43], which we refer to as *bicriteria shortest path (BSP)* algorithm. It maintains a *label set $L(\cdot)$* of *labels* for each vertex $v \in V$. In our scenario, each label is a tuple of driving time and SoC. A label *(Pareto) dominates* another label of the same vertex if it is better in one criterion and not worse in the other. Initially, all label sets are empty, except for the label $(0, \beta_s)$ at the source $s$, which is also inserted into a priority queue. In each step, the algorithm *settles* the minimum label $\ell$ of the queue. This is done by extracting the label $\ell = (\tau, \beta)$ assigned to some vertex $u \in V$ and scanning all arcs $(u, v) \in A$ outgoing from $u$. If the new label $\ell' := (\tau + d(u, v), \beta - c(u, v))$ is not dominated by any label in $L(v)$, it is added to $L(v)$ and the queue, removing labels dominated by $\ell'$ from $L(v)$ and the queue. Using driving time as key of tuples in the priority queue (breaking ties by SoC, i.e., giving preference to the label with highest SoC if two or more labels have the same driving time), the algorithm is *label setting*, i.e., extracted labels are never dominated. An optimal (constrained) path is then found once a label with nonnegative SoC is extracted at $t$. Battery constraints can be incorporated on-the-fly by additional checks during the algorithm: When scanning an arc $(u, v) \in A$, we set the SoC of the new label $\ell'$ to $\min\{M, \beta - c(u, v)\}$. If this SoC is negative, we discard the label $\ell'$.
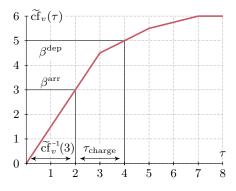
Figure 2: A univariate charging function $\widetilde{\mathrm{cf}}_v$ of a charging station $v \in S$ with minimum SoC value $\beta_v^{\min} = 0$ and maximum SoC value $\beta_v^{\max} = 6$. Reaching $v$ with an arrival SoC of $\beta^{\mathrm{arr}} = 3$ and spending a charging time of $\tau_{\mathrm{charge}} = 2$ yields a departure SoC $\beta^{\mathrm{dep}} = \mathrm{cf}_v(\beta^{\mathrm{arr}}, \tau_{\mathrm{charge}})$, which we obtain by evaluating $\widetilde{\mathrm{cf}}_v(\widetilde{\mathrm{cf}}_v^{-1}(\beta^{\mathrm{arr}}) + \tau_{\mathrm{charge}}) = 5$.

**Speedup Techniques.** A *potential function* $\pi \colon V \to \mathbb{R}_{\geq 0}$ on the vertices is *consistent* (wrt. driving time) if the *reduced arc costs* $\underline{d}(u, v) := d(u, v) - \pi(u) + \pi(v)$ are nonnegative for all $(u, v) \in A$. The A* algorithm [32] adds the potential of a vertex to the keys of its labels, changing the order in which they are extracted from the queue. The A* algorithm was extended to the multicriteria case by Mandow and Pérez-de-la-Cruz [42].

In CH [24], extended to multicriteria scenarios by Geisberger et al. [23] and Funke and Storandt [21], vertices are iteratively contracted during preprocessing according to a given vertex order, while introducing *shortcuts* between their neighbors to maintain distances (wrt. $d$ and $c$). To avoid unnecessary shortcuts, *witness searches* are run between neighbors to identify existing paths that dominate a shortcut candidate. Adding shortcuts may lead to nondominated multi-arcs. To answer $s$–$t$-queries, the original graph is enriched with all shortcuts added during the contraction phase. The CH query then runs a bidirectional variant of the BSP algorithm, starting from both $s$ and $t$, but scanning only arcs leading to vertices of higher importance (wrt. the vertex order).

**Modeling Charging Stops.** In this work, we consider stops at charging stations to recharge the battery (while spending charging time). In our model, a subset $S \subseteq V$ of the vertices represents charging stations. Each vertex $v \in S$ has a designated *charging function* $\mathrm{cf}_v \colon [0, M] \times \mathbb{R}_{\geq 0} \to [0, M]$, which maps *arrival SoC* and the spent charging time to the resulting *departure SoC*. We presume that charging functions are continuous and monotonically increasing wrt. charging time (i.e., charging for a longer time never decreases the SoC). Further, we assume that for arbitrary charging times $\tau_1 \in \mathbb{R}_{\geq 0}$, $\tau_2 \in \mathbb{R}_{\geq 0}$, and SoC values $\beta \in [0, M]$, the *shifting property* $\mathrm{cf}_v(\mathrm{cf}_v(\beta, \tau_1), \tau_2) = \mathrm{cf}_v(\beta, \tau_1 + \tau_2)$ holds. Hence, charging speed only depends on the current SoC, but not on the arrival SoC. These conditions are met by realistic physical models of charging stations [44, 46, 62]. Moreover, exploiting the shifting property, it is possible to represent the (bivariate) charging function $\mathrm{cf}_v$ using a univariate function $\widetilde{\mathrm{cf}}_v \colon \mathbb{R}_{\geq 0} \to [0, M]$ with $\widetilde{\mathrm{cf}}_v(\tau) := \mathrm{cf}_v(0, \tau)$; see Figure 2 and our explanation further below.

Given a vertex $v \in S$ with a charging function $\mathrm{cf}_v$ that has the above properties, we further presume that there is a finite value $\tau_v^{\max} \in \mathbb{R}_{\geq 0}$, such that $\mathrm{cf}_v(\beta, \tau) = \mathrm{cf}_v(\beta, \tau_v^{\max})$ for all $\beta \in [0, M]$ and $\tau \geq \tau_v^{\max}$. In other words, some maximum SoC is reached after a finite charging time (the charging function does not converge to some SoC without reaching it eventually). Then, the minimum SoC value $\beta_v^{\min} := \min_{\tau \in \mathbb{R}_{\geq 0}} \mathrm{cf}_v(0, \tau) = \mathrm{cf}_v(0, 0)$ and the maximum SoC value $\beta_v^{\max} := \max_{\tau \in \mathbb{R}_{\geq 0}} \mathrm{cf}_v(0, \tau) = \mathrm{cf}_v(0, \tau_v^{\max})$ of a charging function induce a range $[\beta_v^{\min}, \beta_v^{\max}]$ of possible SoC values after charging at $v$. We allow the cases $\beta_v^{\min} > 0$ and $\beta_v^{\max} < M$ to model certain

restrictions of charging stations. Thereby, our notion of charging functions is flexible enough to capture features of realistic charging stations. For example, we include swapping stations by setting $\mathrm{cf}_v(\tau, \beta) = M$ for all values $\tau \in \mathbb{R}_{\geq 0}$ and $\beta \in [0, M]$. Hence, we obtain $\beta_v^{\min} = \beta_v^{\max} = M$. Finally, we assign to every charging station $v \in S$ a constant *initialization time* $\tau_{\mathrm{init}}(v)$ that is spent when charging at $v$. This enables us to model time overhead at a charging station for, e. g., parking the car or swapping the battery.

As mentioned above, we represent the bivariate charging function $\mathrm{cf}_v$ of a vertex $v \in S$ with a univariate function $\widetilde{\mathrm{cf}}_v$, as follows. Consider the *inverse function* $\mathrm{cf}_v^{-1}$ mapping a desired departure SoC $\beta \in [\beta_v^{\min}, \beta_v^{\max})$ to the required charging time $\tau \in \mathbb{R}_{\geq 0}$ when the arrival SoC is 0, i. e., $\mathrm{cf}_v^{-1}(\beta) = \tau$ implies that $\widetilde{\mathrm{cf}}_v(\tau) = \mathrm{cf}_v(0, \tau) = \beta$. Since $\widetilde{\mathrm{cf}}_v$ is strictly increasing on the interval $[\beta_v^{\min}, \beta_v^{\max})$ by definition, the function $\mathrm{cf}_v^{-1}$ is well-defined on the domain $[\beta_v^{\min}, \beta_v^{\max})$. Given the minimum charging time $\tau_v^{\max} \in \mathbb{R}_{\geq 0}$ required to charge to an SoC $\beta_v^{\max}$ at $v$ from an arrival SoC of 0, we define the *expanded inverse function* $\widetilde{\mathrm{cf}}_v^{-1} \colon [0, M] \to \mathbb{R}_{\geq 0}$ by setting

$$\widetilde{\mathrm{cf}}_v^{-1}(\beta) := \begin{cases} 0 & \text{if } \beta < \beta_v^{\min}, \\ \tau_v^{\max} & \text{if } \beta \geq \beta_v^{\max}, \\ \mathrm{cf}_v^{-1}(\beta) & \text{otherwise.} \end{cases}$$

This yields the equivalence $\mathrm{cf}_v(\beta, \tau) = \widetilde{\mathrm{cf}}_v(\widetilde{\mathrm{cf}}_v^{-1}(\beta) + \tau)$ for arbitrary values $\beta \in [0, M]$ with $\beta \leq \beta_v^{\max}$ and $\tau \in \mathbb{R}_{\geq 0}$; see Figure 2. Further, we denote by $\mathrm{cf}_v^{-1}(\beta_1, \beta_2) := \widetilde{\mathrm{cf}}_v^{-1}(\beta_2) - \widetilde{\mathrm{cf}}_v^{-1}(\beta_1)$ the time to charge the battery from some arrival SoC $\beta_1 \in [0, M]$ to a desired departure SoC $\beta_2 \in [0, M]$ with $\beta_1 \leq \beta_2$.

Existing models of charging functions use linear, polynomial, and exponential functions, or piecewise combinations thereof [44, 46, 59]. Typically, these functions are also *concave* wrt. charging time (i. e., charging speed only decreases as the battery's SoC increases). However, charging functions in our model are not limited to such functions per se. Section 4 discusses necessary conditions for charging functions besides those mentioned above (continuity, monotonicity, and the shifting property) to ensure that our algorithms terminate. For the sake of simplicity and motivated by data input in our experimental evaluation (see Section 9), examples in subsequent sections use piecewise linear, concave charging functions.

**Problem Statement.** We consider the following objective: For a given source $s \in V$, a target $t \in V$, and an initial SoC $\beta_s \in [0, M]$, we want to find a feasible $s$–$t$ path that minimizes overall *trip time*, i. e., the sum of driving time and total time spent at charging stations. Note that if the input graph contains no charging stations ($S = \emptyset$) and consumption values are nonnegative for all arcs, we have an instance of CSP, hence the considered problem is $\mathcal{NP}$-hard, too.

## 3 Basic Approach

Adapting the bicriteria algorithm described in Section 2 to our setting is difficult, for several reasons: When reaching a charging station, we do not know how much energy should be recharged, since it depends on the remaining route to the target and the charging stations available on this route. To overcome this issue, a key idea of our algorithm is to *delay* this decision until we reach the target or the next charging station. However, since charging functions are *continuous*, there is no straightforward way to apply the bicriteria algorithm in this case, as it might require an infinite number of nondominated labels after settling a charging station with a continuous charging function. In this section, we show how the algorithm can be to generalized to our setting. The *charging*
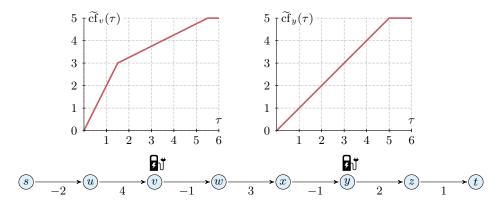
Figure 3: Example graph with indicated charging stations $v$ and $y$ and their respective charging functions $\widetilde{\mathrm{cf}}_v$ and $\widetilde{\mathrm{cf}}_y$. The initialization time is 0 for both stations. Arc labels correspond to their energy consumption, while we presume that the driving time on every arc is 1. On the (unique) path from $s$ to $t$, we seek charging times at $v$ and $y$ that minimize the overall travel time (assuming $\beta_s = 4$ and $M = 5$.)

*function propagating (CFP)* algorithm extends labels to maintain infinite, continuous sets of solutions. The core idea is that a label represents all possible tradeoffs between charging time and resulting SoC induced by the last visited charging station (if it exists), but still has constant size.

Before we describe the CFP algorithm in detail, we illustrate its behavior in a simple example based on charging functions that are piecewise linear and concave. Afterwards, we show how to represent paths containing charging stations with labels of *constant* size and describe the CFP algorithm more formally. In Section 4, we discuss the implementation of its crucial part, namely, generating new labels at charging stations (for general charging functions).

**Running CFP on an Example Graph.** To get a basic understanding of the challenges described above and how they are handled by our algorithm, consider the simple example graph depicted in Figure 3. (We note that neither consumption values nor charging functions in our examples were chosen to be very realistic, but rather to illustrate different cases that may occur both in theory and in practice.) There is only one path from the source vertex $s$ to the target vertex $t$, but it contains two charging stations. Hence, we essentially have to determine the charging time at each station that minimizes the overall trip time. Let the initial SoC at $s$ be $\beta_s = 4$ and assume that the battery capacity is $M = 5$.

Our algorithm propagates labels which represent continuous functions mapping the trip time to the resulting SoC at some vertex, depending on the amount of energy charged at the previous station. Figure 4 shows these functions as they are generated by our algorithm when running on the instance from Figure 3. The algorithm is initialized with a label $\ell_1$ representing a function denoted $f\langle\ell_1\rangle$ at the source vertex $s$; see Figure 4a. The function $f\langle\ell_1\rangle$ evaluates to the initial SoC 4 for arbitrary trip times (since $s$ is not a charging station, spending time at $s$ cannot increase the SoC).

To propagate this function through the graph, the (single) outgoing arc $(s, u)$ with driving time 1 and energy consumption -2 is scanned. Thus, the vertex $u$ is reached after a trip time of 1, which we model by setting the function of the (propagated) label to $-\infty$ for trip times below 1. Moreover, the SoC increases to the maximum SoC 5 (note battery constraints prevent the SoC from exceeding this limit). Next, this label is used to scan the arc $(u, v)$ with trip time 1 and energy consumption 4, which results in a label $\ell_1'$ at $v$ containing the function $f\langle\ell_1'\rangle$ shown in Figure 4a. It evaluates to an SoC of 1 if the trip time is at least 2.
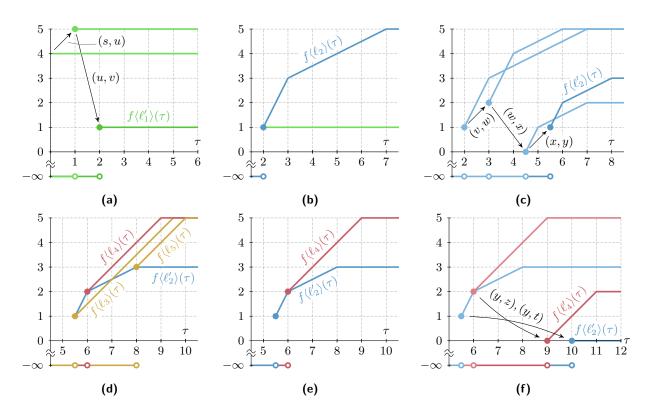
8

Figure 4: Labels generated by the CFP algorithm on the graph shown in Figure 3. (a) The functions (green) created at the vertices $s$, $u$, and $v$ (upon arrival). (b) A new function (blue) corresponding to a label spawned at $v$ to model departure SoC after charging at the station. (c) The label is propagated to $w$, $x$, and $y$. Note that certain function values are "cropped", due to battery constraints. (d) New functions (red, yellow) spawned at the charging station $y$, two of which are dominated (yellow). (e) The resulting nondominated functions upon departure at $y$. Note that in one of them, no energy is charged at $y$. (f) The functions are propagated to the target $t$. The minimum feasible trip time is 9 and requires charging at both stations: two units (departure SoC 3) at $v$ and one unit (departure SoC 3) at $y$.

9

According to the label $\ell_1'$, we reach the first charging station $v$ with an SoC of 1 and after a trip time of 2. Since $v$ is a charging station, we want to incorporate the possibility of charging into the label. A function that reflects tradeoffs between trip time after charging and the resulting SoC is obtained after shifting the charging function $\widetilde{\mathrm{cf}}_v$ along the x-axis such that it intersects the point with x-coordinate (i.e., trip time) 2 and y-coordinate (i.e., arrival SoC) 1. This results in a new label $\ell_2$ which *dominates* the old one, i.e., for each trip time it yields greater or equal SoC compared to $\ell_1$; see Figure 4b (however, this is only the case because we presume that the initialization time is 0, whereas in general we might obtain labels that do not dominate each other). The nondominated label $\ell_2$ is then propagated to the vertices $w$, $x$, and $y$. Again, we apply battery constraints, which in hindsight renders some charging times at $v$ infeasible (not enough energy is charged to traverse the following path) or unprofitable (charging too much energy wastes energy gained through recuperation). We "crop" the image of the resulting function $f\langle\ell_2'\rangle$ by setting corresponding function values to $-\infty$ or $M = 5$, respectively; see Figure 4c.

The search reaches the next charging station at the vertex $y$. Again, we may shift the function $\widetilde{\mathrm{cf}}_y$ until its plot intersects a (nondominated) pair of trip time and SoC of the label $\ell_2'$. However, unlike before, there no longer is a unique *nondominated* pair. As before, we may pick the least feasible trip time (5.5 in this case), which results in the label $\ell_3$. Observe that this yields a function $f\langle\ell_3\rangle$ with lower SoC compared to the function $f\langle\ell_2'\rangle$ for some values of trip time; see Figure 4d. This is due to the fact that, for low SoC values, the charging station at $v$ allows faster charging. Therefore, it pays off to charge more energy at $v$, up to the next breakpoint of the function $f\langle\ell_2'\rangle$ (where its the slope falls below that of the charging function at $y$). Spawning the label $\ell_4$ after picking this breakpoint results in a nondominated function (see Figure 4d). Apparently, breakpoints of functions are good candidates for spawning new labels, since the charging rate changes in these points. Consequently, our algorithm spawns one new label for each breakpoint of the function $f\langle\ell_2'\rangle$. In Section 4, we show that this is actually sufficient to find an optimal solution. Two of these labels ($\ell_3$ and $\ell_5$) are dominated, though, so they can be discarded. (Our algorithm actually only performs *pairwise* dominance checks to reduce overhead, which would spare the label $\ell_3$ as it is only dominated by the upper envelope of $f\langle\ell_2'\rangle$ and $f\langle\ell_4\rangle$; this increases the number of propagated labels, but does not affect correctness of the approach).

The nondominated labels $f\langle\ell_2\rangle$ and $f\langle\ell_4\rangle$ are shown in Figure 4e. They are both propagated to $z$ and finally, to the target $t$. The resulting labels are shown in Figure 4f. Note that both contain feasible pairs of trip time and SoC. The least feasible trip time among both labels is 9. Note that this trip time requires charging at both stations, despite the fact that it is possible to reach the target with a single charging stop at $v$. This is due to the fact that the charging speed at both station differs. At the same time, both stations lie on the (unique) route to the target and there is no initialization overhead, so switching between different charging stations comes at a low cost in this example. In what follows, we formally define labels and their corresponding functions, before we describe the CFP algorithm in more detail.

**Labels and SoC Functions.** More generally, assume we are given a path $P$ from the source $s \in V$ to some vertex $v \in V$, such that $P$ contains a charging station $u \in S$ and the arrival SoC at $u$ is $\beta_u$. Every possible charging time $\tau_{\mathrm{charge}} \in [0, \mathrm{cf}_u^{-1}(\beta_u, \beta_u^{\max})]$ at $u$ results in a certain trip time and an SoC at $v$. In general, this yields an infinite amount of feasible, nondominated pairs of trip time and corresponding SoC for the path. We implicitly represent these pairs in one label by storing the charging station $u$ in the label. However, this no longer allows us to apply battery constraints on-the-fly: For vertices visited after $u$, labels have no fixed SoC, as it depends on how much energy is charged at $u$. Hence, we compute the SoC profile $b_{[u,\dots,v]}$ of the subpath from $u$ to $v$; see Section 2. The label $\ell = (\tau_{\mathrm{trip}}, \beta_u, u, b_{[u,\dots,v]})$ at the vertex $v$ then consists of the trip time $\tau_{\mathrm{trip}}$ of the path
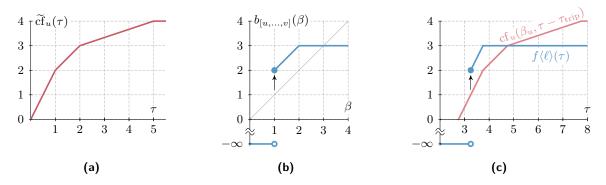
Figure 5: Constructing the SoC function $f\langle\ell\rangle$ of a label $\ell = (\tau_{\text{trip}}, \beta_u, u, b_{[u,\dots,v]})$, with $\tau_{\text{trip}} = 3$ and $\beta_u = 0.5$. (a) The function $\widetilde{\text{cf}}_u$. Assume that the initialization time at $u$ is $\tau_{\text{init}}(u) = 0$. (b) The SoC profile $b_{[u,\dots,v]}$ of the $u$–$v$ subpath. Note that the path has negative consumption (the SoC increases as indicated by the arrow). (c) The SoC function $f\langle\ell\rangle$. The function $\text{cf}_u(\beta_u, \tau - \tau_{\text{trip}})$ (red) reflects pairs of trip time and SoC when charging at $u$, but ignores consumption on the $u$–$v$ subpath. It is equivalent to the function obtained after shifting $\widetilde{\text{cf}}_u$ to the right by $\tau_{\text{trip}} = 3$ minus $\text{cf}_u^{-1}(0, \beta_u) = 0.25$. We apply battery constraints wrt. the $u$–$v$ path to this function and obtain the depicted SoC function $f\langle\ell\rangle$ (blue). Its minimum feasible trip time is $\tau_{\text{min}}(\ell) = 3.25$, because we must spend a charging time of at least $0.25$ at $u$. Moreover, we obtain $f\langle\ell\rangle(\tau) = 3$ for $\tau \geq 3.75$ (charging beyond an SoC of 2 at $u$ never pays off, as it wastes energy gains from recuperation).

from $s$ to $v$ (including charging time on every previous charging station except $u$ on the path from $s$ to $u$), the SoC $\beta_u$ when reaching $u$, the last visited charging station $u$, and the SoC profile $b_{[u,\dots,v]}$ of the subpath from $u$ to $v$. Recall that this SoC profile can be represented by three values; see Section 2. Consequently, even though charging functions can have arbitrary descriptive complexity, we propagate them using labels of constant size. The trip time $\tau_{\text{trip}}$ excludes charging at $u$, but includes its initialization time $\tau_{\text{init}}(u)$. Thus, we can think of $\tau_{\text{trip}}$ as the least trip time to reach $v$ if we stop at $u$ (and ignore battery constraints on the $u$–$v$ path).

Accordingly, we define the *SoC function* $f\langle\ell\rangle$ of a label $\ell$ to represent all feasible pairs of trip time and SoC associated with the label $\ell = (\tau_{\text{trip}}, \beta_u, u, b_{[u,\dots,v]})$. The SoC function $f\langle\ell\rangle \colon \mathbb{R}_{\geq 0} \to [0, M] \cup \{-\infty\}$ mapping trip time to SoC is given as

$$f\langle\ell\rangle(\tau) := \begin{cases} b_{[u,\dots,v]}(\text{cf}_u(\beta_u, \tau - \tau_{\text{trip}})) & \text{if } \tau \geq \tau_{\text{trip}}, \\ -\infty & \text{otherwise.} \end{cases} \tag{1}$$

To obtain the value $f\langle\ell\rangle(\tau)$, i.e., the (arrival) SoC at $v$ when allowing a trip time $\tau \geq \tau_{\text{trip}}$, Equation 1 first evaluates the SoC $\text{cf}_u(\beta_u, \tau - \tau_{\text{trip}})$ after charging at $u$ for a total time of $\tau - \tau_{\text{trip}}$ with an arrival SoC of $\beta_u$. Afterwards, the SoC profile $b_{[u,\dots,v]}$ is applied, which takes account of energy consumption (respecting battery constraints) on the path from the charging station $u$ to the current vertex $v$. This yields the desired SoC at $v$. Note that $f\langle\ell\rangle$ can evaluate to $-\infty$ for values greater than $\tau_{\text{trip}}$, due to battery constraints applied by the SoC profile $b_{[u,\dots,v]}$. We denote by

$$\tau_{\text{min}}(f\langle\ell\rangle) := \min\{\tau \in \mathbb{R}_{\geq 0} \mid f\langle\ell\rangle(\tau) \neq -\infty\}$$

the smallest value for which $f\langle\ell\rangle$ is greater than $-\infty$, i.e., the minimum trip time required for the path represented by $\ell$ to be feasible. Figure 5 shows an example of an SoC function of a label. The definition of SoC functions reflects the interpretation of our labels, which represent all tradeoffs between charging time and resulting SoC induced by the charging function $\text{cf}_u$ of the last station $u$.

11

**Algorithm 1:** Pseudocode of the CFP algorithm.

```
// initialize label sets
```
**1 foreach** $v \in V$ **do**

**2**     $L_{\text{set}}(v) \longleftarrow \emptyset$

**3**     $L_{\text{uns}}(v) \longleftarrow \emptyset$

**4** $v^* \longleftarrow$ dummy vertex without incident edges that is (temporarily) added to $V$

**5** $S \longleftarrow S \cup \{v^*\}$

**6** $\widetilde{\text{cf}}_{v^*} \longleftarrow [(0, \beta_s)]$

**7** $L_{\text{uns}}(s) \longleftarrow \{(0, \beta_s, v^*, b_{[s]})\}$

**8** $Q.\texttt{insert}(s, 0)$

```
// run main loop
```
**9 while** $Q.\texttt{isNotEmpty()}$ **do**

**10**     $v \longleftarrow Q.\texttt{minElement()}$

**11**     $\ell = (\tau_{\text{trip}}, \beta_u, u, b_{[u,...,v]}) \longleftarrow L_{\text{uns}}(v).\texttt{deleteMin()}$

**12**     $L_{\text{set}}(v).\texttt{insert}(\ell)$

**13**     **if** $v = t$ **then**

**14**        **return** $\tau_{\min}(f\langle \ell \rangle)$

 

```
        // handle charging stations; see Section 4
```
**15**     **if** $v \in S \setminus \{u\}$ **then**

**16**        **foreach** $\tau_{\text{charge}} \in T(\angle_{old}^{\ell}, \angle_{new}^{\ell}) \setminus \{\infty\}$ **do**

**17**           $L_{\text{uns}}(v).\texttt{insert}((\tau_{\text{trip}} + \tau_{\text{charge}} + \tau_{\text{init}}(v), f\langle \ell \rangle(\tau_{\text{trip}} + \tau_{\text{charge}}), v, b_{[v]}))$

 

```
        // update priority queue
```
**18**     **if** $L_{\text{uns}}(v).\texttt{isNotEmpty()}$ **then**

**19**        $\ell' \longleftarrow L_{\text{uns}}(v).\texttt{minElement()}$

**20**        $Q.\texttt{update}(v, k(\ell'))$

**21**     **else**

**22**        $Q.\texttt{deleteMin()}$

 

```
        // scan outgoing arcs
```
**23**     **foreach** $(v, w) \in A$ **do**

**24**        $b_{[u,...,w]} \longleftarrow \text{link}(b_{[u,...,v]}, b_{(v,w)})$

**25**        **if** $b_{[u,...,w]}(\beta_u^{max}) \neq -\infty$ **then**

**26**           $\ell' \longleftarrow (\tau_{\text{trip}} + d(v, w), \beta_u, u, b_{[u,...,w]})$

**27**           $L_{\text{uns}}(w).\texttt{insert}(\ell')$

**28**           **if** $\ell' = L_{\text{uns}}(w).\texttt{minElement()}$ **then**

**29**              $Q.\texttt{update}(w, k(\ell'))$

**Algorithm Description.** We are now ready to describe the actual CFP algorithm; see Algorithm 1 for pseudocode. It propagates labels that are quadruples as defined above. Given two labels $\ell$ and $\ell'$, we say that $\ell$ *dominates* $\ell'$ if $f\langle\ell\rangle(\tau) \geq f\langle\ell'\rangle(\tau)$ holds for all $\tau \in \mathbb{R}_{\geq 0}$. The *key* of a label $\ell$, denoted $k(\ell) := \tau_{\min}(f\langle\ell\rangle)$, is defined as its minimum feasible trip time. This value is not stored explicitly in the label, but computed on-the-fly by evaluating the inverse of the charging function of the previous station at the minimum SoC for which the subpath from this station to the current vertex becomes feasible (one could also store the key with the label, but this did not improve performance in our experiments).

The algorithm stores two sets $L_{\mathrm{set}}(v)$ and $L_{\mathrm{uns}}(v)$ for each vertex $v \in V$, containing *settled* (i.e., extracted) and *unsettled* labels, respectively. Sets $L_{\mathrm{uns}}(\cdot)$ are organized as priority queues (implemented as binary heaps), allowing efficient extraction of the unsettled label with minimum key (breaking ties by the corresponding minimum SoC of a label). We maintain the invariant that for each $v \in V$, $L_{\mathrm{uns}}(v)$ is empty or the minimum label $\ell$ (wrt. its key) in $L_{\mathrm{uns}}(v)$ is not dominated by any label in $L_{\mathrm{set}}(v)$. Every time the minimum element of the heap changes, because an element is removed or added, we check whether the new minimum is dominated by a label in $L_{\mathrm{set}}(v)$ and remove it in this case (as we know that it cannot lead to an optimal solution). For piecewise-defined SoC functions, a dominance test requires a linear scan over the subfunctions of both SoC functions. By using heaps for unsettled labels, we avoid unnecessary dominance checks for labels that are never settled. (A more straightforward variant could use a single set per vertex and follow the basic algorithm outlined in Section 2 to identify dominated labels, however, this led to slower running times in preliminary tests.)

Given a source $s \in V$, a target $t \in V$, and the initial SoC $\beta_s \in [0, M]$, the algorithm is initialized in Lines 1–8 with a single label $(0, \beta_s, v^*, b_{[s]})$ at the source $s$, while all other label sets are empty. Note that $v^*$ is a special vertex that is (temporarily) added to the graph as a charging station with the charging function $\mathrm{cf}_{v^*} \equiv \beta_s$. Thereby, we avoid explicit handling of special cases when reaching the first actual charging station. The SoC profile stored in the label is initialized with the identity function $b_{[s]}$ (i.e., the SoC is not affected when applying this function). The source vertex is also inserted into a priority queue (denoted $Q$ in Algorithm 1). The key of a vertex $v \in V$ in the priority queue is the key of the minimum element in $L_{\mathrm{uns}}(v)$, i.e., the minimum feasible trip time among the SoC functions of all unsettled labels.

The algorithm then proceeds along the lines of the BSP algorithm. In each step of the main loop, it first extracts a vertex $v \in V$ with minimum key (breaking ties by SoC) from the priority queue and settles it; see Lines 10–12 of Algorithm 1. Note that at this point, the key of the corresponding label $\ell = (\tau_{\mathrm{trip}}, \beta_u, u, b_{[u,\dots,v]})$ extracted from $L_{\mathrm{uns}}(v)$ is not greater than that of any label that has not been settled yet.

Next, we check whether $v$ is a charging station that differs from the one stored in the current label $\ell$, i.e., $v \in S \setminus \{u\}$. If this is the case, we create new labels to incorporate possible recharging at $v$; see Lines 15–17. This means that we have to spawn new labels $\ell'$ that replace the previous charging station $u$ by $v$. We can do so by fixing a charging time $\tau_{\mathrm{charge}} \in \mathbb{R}_{\geq 0}$ at $u$. For the resulting SoC at $u$, we evaluate the SoC profile $b_{[u,\dots,v]}$ of the $u$–$v$ path to determine the SoC at $v$. We update the trip time accordingly by adding the charging time $\tau_{\mathrm{charge}}$ at $u$ and the initialization time $\tau_{\mathrm{init}}(v)$ at the new charging station $v$. We obtain the new label

$$\ell' := (\tau_{\mathrm{trip}} + \tau_{\mathrm{charge}} + \tau_{\mathrm{init}}(v), b_{[u,\dots,v]}(\mathrm{cf}_u(\beta_u, \tau_{\mathrm{charge}})), v, b_{[v]})$$
$$= (\tau_{\mathrm{trip}} + \tau_{\mathrm{charge}} + \tau_{\mathrm{init}}(v), f\langle\ell\rangle(\tau_{\mathrm{trip}} + \tau_{\mathrm{charge}}), v, b_{[v]}). \tag{2}$$

However, we still face the problem that in general, there are infinitely many possible charging times $\tau_{\mathrm{charge}}$ at the previous charging station $u$ held in $\ell$. In Section 4, we show that for realistic

13

models of charging stations, we only have to consider a small (finite) number of relevant charging times at $u$ when charging at $v$. Thus, spawning a limited number of new labels, each fixing a certain charging time at $u$ and setting the last charging station to $v$, is sufficient to represent all nondominated solutions. Note that the original label $\ell$ is not discarded, to reflect the possibility of not stopping at the charging station $v$.

In Lines 18–22 of Algorithm 1, the key of $v$ in the priority queue is updated. Since the label $\ell$ was settled and new labels may have spawned in case $v$ is a charging station, we update the key of $v$ to the new smallest key of an unsettled label, if it exists. Otherwise, $v$ is removed from the queue.

Afterwards, we scan all outgoing arcs $(v, w) \in A$; see Lines 23–29. Given the current label $\ell = (\tau_{\mathrm{trip}}, \beta_u, u, b_{[u,\ldots,v]})$, traversing the arc $(v, w)$ means to increase trip time by $d(v, w)$ and apply the (constant-time) link operation to the SoC profile $b_{[u,\ldots,v]}$ of $\ell$ and the SoC profile $b_{[v,w]}$ induced by the energy consumption $c(v, w)$; see Section 2. We compute $b_{[u,\ldots,w]} := \mathrm{link}(b_{[u,\ldots,v]}, b_{[v,w]})$ and construct the label

$$\ell' := (\tau_{\mathrm{trip}} + d(v, w), \beta_u, u, b_{[u,\ldots,w]}).$$

Unless the SoC profile $b_{[u,\ldots,w]}$ of $\ell'$ indicates that the $u$–$w$ subpath is infeasible for arbitrary SoC, the new label $\ell'$ is added to the label set at the vertex $w$. Note that we perform no dominance checks at this point (unless the minimum element in the label set $L_{\mathrm{uns}}(w)$ changes), exploiting the fact that labels are organized in two sets per vertex.

When extracting a label $\ell$ at the target vertex $t$ for the first time, we pick the least charging time at the last station such that $t$ can be reached, i. e., the minimum feasible trip time $\tau_{\min}(f\langle\ell\rangle)$ of $f\langle\ell\rangle$, and the algorithm terminates; see Line 14 in Algorithm 1. Correctness of the CFP algorithm follows from Lemma 1 shown in Section 4 below and the fact that the first extracted label $\ell$ at $t$ minimizes the feasible trip time (recall that the algorithm is label setting and minimum feasible trip time is used as key in the priority queue). Theorem 1 at the end of Section 4 summarizes these insights. The asymptotic running time of the algorithm is exponential in the input graph in the worst case (for reasonable charging models; see Section 4).

For path retrieval, we add two pointers to each label, storing its parent vertex and parent label. For a charging station $v \in S$, the vertex $v$ can be its own parent. Two consecutive identical parents then imply the use of a (previous) charging station $u \in S$, which is stored in the former label. The according charging time is the difference between the trip times of both labels.

## 4 Spawning Labels at Charging Stations

As described in Section 3, the CFP algorithm constructs new labels at charging stations to represent all nondominated solutions. We now prove that for reasonable models of charging functions, it suffices to spawn a small number of labels that replace the previous charging station with the new one. The key idea is that we only require labels that correspond to charging at the station that offers the better charging speed at a certain (relative) point in time. We define *switching sequences* for pairs of functions, containing points at which the charging speed of the new function surpasses the old one. Lemma 1 proves that spawning one label per element of the switching sequence suffices. Moreover, switching sequences are finite (and linear in the descriptive complexity) for typical models of charging functions, which implies that the CFP algorithm terminates. Before proving Lemma 1, we introduce helpful tools. We also formalize switching sequences and the slope of an SoC function.

Consider a label $\ell = (\tau_{\mathrm{trip}}, \beta_u, u, b_{[u,\ldots,v]})$ extracted at some charging station $v \in S$. We want to create new labels that reflect charging at $v$. This requires us to fix a charging time $\tau_{\mathrm{charge}} \in \mathbb{R}_{\geq 0}$ at the previous station $u$, so that we can set $v$ as the last visited charging station of a new label $\ell'$;
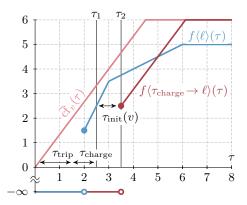
Figure 6: Spawning a label at a charging station. Given the SoC function $f\langle\ell\rangle$ of a label $\ell = (\tau_{\text{trip}}, \beta_u, u, b_{[u,\ldots,v]})$ at a charging station $v \in S$, we can spawn a new label $\tau_{\text{charge}} \to \ell$ by picking a charging time $\tau_{\text{charge}}$ at the station $u$. We compare the slopes of $f\langle\ell\rangle$ at $\tau_1 = \tau_{\text{trip}} + \tau_{\text{charge}}$ and $f\langle\tau_{\text{charge}} \to \ell\rangle$ at $\tau_2 = \tau_{\text{trip}} + \tau_{\text{charge}} + \tau_{\text{init}}(v)$ to determine which one represents the better choice. Note that $\tau_{\text{trip}}$ is smaller than $\tau_{\min}(f\langle\ell\rangle)$, due to battery constraints.

see Equation 2 and Figure 6. We denote the resulting label by $(\tau_{\text{charge}} \to \ell) := \ell'$, as it is obtained after setting the charging time in $\ell$ to $\tau_{\text{charge}}$. Recall that in the label $\tau_{\text{charge}} \to \ell$, we replace the old charging station $u$ by the new station $v$. Moreover, we set $\tau_{\text{trip}} + \tau_{\text{charge}} + \tau_{\text{init}}(v)$ as its overall trip time and $f\langle\ell\rangle(\tau_{\text{trip}} + \tau_{\text{charge}})$ as the corresponding arrival SoC at $v$. The SoC function $f\langle\tau_{\text{charge}} \to \ell\rangle$ represents all tradeoffs between charging time at the new charging station $v$ and resulting SoC. If the label $\tau_{\text{charge}} \to \ell$ is not feasible, i.e., $\tau_{\text{trip}} + \tau_{\text{charge}} < \tau_{\min}(f\langle\ell\rangle)$, we obtain $f\langle\tau_{\text{charge}} \to \ell\rangle \equiv -\infty$.

Not every charging time $\tau_{\text{charge}} \in \mathbb{R}_{\geq 0}$ at $u$ yields a reasonable solution. In particular, if we can find a charging time $\tau'_{\text{charge}} \in \mathbb{R}_{\geq 0}$ such that $f\langle\tau'_{\text{charge}} \to \ell\rangle$ dominates $f\langle\tau_{\text{charge}} \to \ell\rangle$, we know that a charging time of $\tau_{\text{charge}}$ is never beneficial. Intuitively, if the new charging station $v$ allows fast charging, it could pay off to charge less energy at the previous station $u$ and spend more time at $v$ instead, so $f\langle\tau_{\text{charge}} - \varepsilon \to \ell\rangle$ dominates $f\langle\tau_{\text{charge}} \to \ell\rangle$ for some $\varepsilon > 0$. Similarly, if the charging station $u$ offers better charging speed, a charging time $\tau_{\text{charge}} + \varepsilon$ might be the better choice. In other words, the best choice of the value $\tau_{\text{charge}}$ depends on the slopes of the two SoC functions $f\langle\ell\rangle$ and $f\langle\tau_{\text{charge}} \to \ell\rangle$.

We define the *slope* of a given function $f$ at some $\tau' \in \mathbb{R}_{\geq 0}$ as the corresponding *right derivative* $(\partial f(\tau)/\partial\tau)(\tau')$, to ensure that the slope is well-defined also for piecewise-defined SoC functions and at the minimum feasible trip time of an SoC function. As before, let $\ell = (\tau_{\text{trip}}, \beta_u, u, b_{[u,\ldots,v]})$ be a label at a charging station $v \in S$. We introduce a function $\angle^\ell_{\text{old}} \colon \mathbb{R}_{\geq 0} \to \mathbb{R}_{\geq 0} \cup \{\infty\}$ that describes the slope of the SoC function $f\langle\ell\rangle$ at $\tau_{\text{trip}} + \tau_{\text{charge}}$ as a function of the charging time $\tau_{\text{charge}} \in \mathbb{R}_{\geq 0}$. Formally, we define

$$\angle^\ell_{\text{old}}(\tau_{\text{charge}}) := \begin{cases} \frac{\partial f\langle\ell\rangle(x)}{\partial x}(\tau_{\text{trip}} + \tau_{\text{charge}}) & \text{if } \tau_{\text{trip}} + \tau_{\text{charge}} \geq \tau_{\min}(f\langle\ell\rangle), \\ \infty & \text{otherwise.} \end{cases}$$

Note that the slope $\angle^\ell_{\text{old}}(\tau_{\text{charge}})$ of the SoC function $f\langle\ell\rangle$ at $\tau_{\text{trip}} + \tau_{\text{charge}}$ is equivalent to the slope of the charging function $\text{cf}_u$ of the vertex $u$ for the arrival SoC $\beta_u$ and the charging time $\tau_{\text{charge}}$. Hence, we obtain $\angle^\ell_{\text{old}}(\tau_{\text{charge}}) = (\partial\,\text{cf}_u(\beta_u, x)/\partial x)(\tau_{\text{charge}})$ for all $\tau_{\text{charge}} \in \mathbb{R}_{\geq 0}$, unless battery constraints on the $u$–$v$ path render a charging time of $\tau_{\text{charge}}$ infeasible or unprofitable (in which case the slope $\angle^\ell_{\text{old}}$ is either $\infty$ or 0). Thus, $\angle^\ell_{\text{old}}(\tau_{\text{charge}})$ can be interpreted as the charging speed when continuing to charge the battery at $u$ after a charging time of $\tau_{\text{charge}}$.

Alternatively, one could interrupt charging at $u$ after a charging time of $\tau_{\text{charge}}$, continue the journey, and switch to the new charging station upon arrival at $v$. The charging speed that can be achieved in this case is given by the slope of the SoC function $f\langle\tau_{\text{charge}}\to\ell\rangle(\tau)$ at $\tau = \tau_{\text{trip}} + \tau_{\text{charge}} + \tau_{\text{init}}(v)$; see Figure 6. Similar to $\angle_{\text{old}}^{\ell}$, we define the function $\angle_{\text{new}}^{\ell}\colon \mathbb{R}_{\geq 0} \to \mathbb{R}_{\geq 0} \cup \{\infty\}$ as

$$\angle_{\text{new}}^{\ell}(\tau_{\text{charge}}) := \begin{cases} \frac{\partial f\langle\tau_{\text{charge}}\to\ell\rangle(\tau)}{\partial\tau}(\tau_{\text{trip}}+\tau_{\text{charge}}+\tau_{\text{init}}(v)) & \text{if } \tau_{\text{trip}}+\tau_{\text{charge}} \geq \tau_{\min}(f\langle\ell\rangle), \\ \infty & \text{otherwise.} \end{cases}$$

It maps total charging time $\tau_{\text{charge}}$ at $u$ to the slope of the SoC function $f\langle\tau_{\text{charge}}\to\ell\rangle$ at the time $\tau_{\text{trip}} + \tau_{\text{charge}} + \tau_{\text{init}}(v)$ that corresponds to arrival at $v$ and starting to charge the battery. Hence, the value of $\angle_{\text{new}}^{\ell}(\tau_{\text{charge}})$ is equivalent to the slope of the new charging function $\text{cf}_v$ of the vertex $v$ for the arrival SoC $f\langle\ell\rangle(\tau_{\text{trip}} + \tau_{\text{charge}})$ and the charging time 0. More formally, we have $\angle_{\text{new}}^{\ell}(\tau_{\text{charge}}) = (\partial\,\text{cf}_v(f\langle\ell\rangle(\tau_{\text{trip}} + \tau_{\text{charge}}), \tau)/\partial\tau)(0)$ if $\tau_{\text{trip}} + \tau_{\text{charge}}$ is at least $\tau_{\min}(f\langle\ell\rangle)$ and $\angle_{\text{new}}^{\ell}(\tau_{\text{charge}}) = \infty$ otherwise.

Given the two functions $\angle_{\text{old}}^{\ell}$ and $\angle_{\text{new}}^{\ell}$ defined above, we are interested in points in time where $\angle_{\text{new}}^{\ell}$ surpasses $\angle_{\text{old}}^{\ell}$, because at such points it may pay off to interrupt charging at $u$ to benefit from a better charging rate at $v$ later on. Additionally, the minimum charging time $\tau_{\min}(f\langle\ell\rangle) - \tau_{\text{trip}}$ that is necessary to reach the new charging station $v$ may be relevant in cases where the charging function of $v$ has low slope but a large minimum SoC value (e. g., if $v$ is a swapping station). We define the *switching sequence* of $\angle_{\text{old}}^{\ell}$ and $\angle_{\text{new}}^{\ell}$, denoted

$$T(\angle_{\text{old}}^{\ell}, \angle_{\text{new}}^{\ell}) := [\tau_{\min}(f\langle\ell\rangle) - \tau_{\text{trip}} = \tau_1, \tau_2, \ldots, \tau_{k-1}, \tau_k = \infty],$$

as the sequence of all candidate points in time to interrupt charging at the old station $u$, arranged in ascending order. (We only include the value $\tau_k = \infty$ to avoid additional case distinctions in the proof of Lemma 1 below.) Formally, we demand for $T(\angle_{\text{old}}^{\ell}, \angle_{\text{new}}^{\ell})$ that $\tau_1 = \tau_{\min}(f\langle\ell\rangle) - \tau_{\text{trip}}$, $\tau_k = \infty$, $\tau_i < \tau_{i+1}$ for $i \in \{1, \ldots, k-1\}$, and for all $i \in \{2, \ldots, k-1\}$ there exists a value $\varepsilon > 0$ such that for all $0 < \delta < \varepsilon$ it holds that $\angle_{\text{old}}^{\ell}(\tau_i - \delta) \geq \angle_{\text{new}}^{\ell}(\tau_i - \delta)$ and $\angle_{\text{old}}^{\ell}(\tau_i + \delta) < \angle_{\text{new}}^{\ell}(\tau_i + \delta)$. Moreover, we assume $T(\angle_{\text{old}}^{\ell}, \angle_{\text{new}}^{\ell})$ to be *maximal*, i. e., it contains all values with the above property. In general, pairs of functions do not necessarily have a switching sequence of finite length $k \in \mathbb{N}$. At the end of this section we argue that the length of a switching sequence is linear in the descriptive complexity (and thus, finite) for reasonable models of charging functions.

We now prove Lemma 1, stating that we only have to spawn a bounded number of new labels at charging stations. In particular, it suffices to add at most one label per element in the switching sequence induced by the last visited charging station in the current label and the new station.

**Lemma 1.** *For a vertex $v \in S$ and a label $\ell = (\tau_{\text{trip}}, \beta_u, u, b_{[u,\ldots,v]})$ at $v$, let the (finite) switching sequence of $\angle_{old}^{\ell}$ and $\angle_{new}^{\ell}$ be given as $T(\angle_{old}^{\ell}, \angle_{new}^{\ell}) = [\tau_1, \tau_2, \ldots, \tau_{k-1}, \tau_k]$. For every charging time $\tau_{\text{charge}} \in \mathbb{R}_{\geq 0}$, there exists an $i \in \{1, \ldots, k-1\}$ such that the SoC functions $f\langle\ell\rangle$ and $f\langle\tau_i \to \ell\rangle$ together dominate the SoC function $f\langle\tau_{\text{charge}}\to\ell\rangle$, i. e., for all $\tau \in \mathbb{R}_{\geq 0}$ we have the inequality $\max\{f\langle\ell\rangle(\tau), f\langle\tau_i \to \ell\rangle(\tau)\} \geq f\langle\tau_{\text{charge}}\to\ell\rangle(\tau)$.*

*Proof.* Consider an arbitrary charging time $\tau_{\text{charge}} \in \mathbb{R}_{\geq 0}$ at the previous charging station $u$ and the induced label $\tau_{\text{charge}} \to \ell$. If $\tau_{\text{trip}} + \tau_{\text{charge}} < \tau_{\min}(f\langle\ell\rangle)$, the SoC function $f\langle\tau_{\text{charge}}\to\ell\rangle \equiv -\infty$ is dominated by $f\langle\ell\rangle$. Hence, we assume in what follows that $\tau_{\text{charge}} \geq \tau_1 = \tau_{\min}(f\langle\ell\rangle) - \tau_{\text{trip}}$. Then there exists a unique index $i \in \{1, \ldots, k-1\}$, such that $\tau_i \leq \tau_{\text{charge}} < \tau_{i+1}$ holds. We distinguish two cases, depending on the slopes $\angle_{\text{old}}^{\ell}$ and $\angle_{\text{new}}^{\ell}$ at $\tau_{\text{charge}}$ and show that, together with $f\langle\ell\rangle$, the function $f\langle\tau_i \to \ell\rangle$ or the function $f\langle\tau_{i+1} \to \ell\rangle$ dominates $f\langle\tau_{\text{charge}}\to\ell\rangle$.

*Case 1:* $\angle_{\text{old}}^{\ell}(\tau_{\text{charge}}) < \angle_{\text{new}}^{\ell}(\tau_{\text{charge}})$. Intuitively, this means that the new charging station $v$ provides a better charging speed than the old station $u$ for the considered charging time $\tau_{\text{charge}}$. Hence,
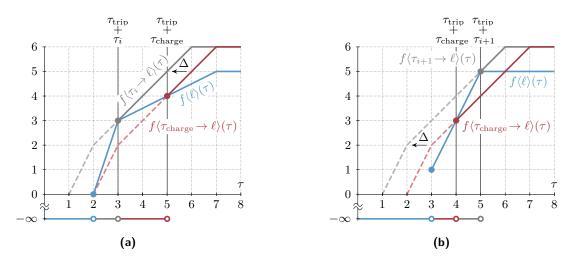
Figure 7: Illustration of dominated SoC functions at a charging station $v \in S$. For simplicity, we assume that $\tau_{\mathrm{init}}(v) = 0$. Dashed segments indicate the (shifted) charging function $\widehat{\mathrm{cf}}_v$ of $v$. Note that a function dominates the area beneath it. (a) The slope of the SoC function $f\langle\ell\rangle$ is lower than the slope of $f\langle\tau_{\mathrm{charge}} \to \ell\rangle$ at $\tau_{\mathrm{trip}} + \tau_{\mathrm{charge}} = 5$. Hence, it pays off to decrease charging time at the previous station to $\tau_i = 3$ and charge more energy at $v$ instead. (b) The slope of $f\langle\ell\rangle$ is greater than the slope of $f\langle\tau_{\mathrm{charge}} \to \ell\rangle$ for $\tau_{\mathrm{trip}} + \tau_{\mathrm{charge}} = 4$. Therefore, charging more energy at the previous charging station pays off and the SoC functions $f\langle\ell\rangle$ and $f\langle\tau_{i+1} \to \ell\rangle$ together dominate the function $f\langle\tau_{\mathrm{charge}} \to \ell\rangle$.

leaving $u$ earlier to charge more energy at $v$ and benefit from faster charging pays off. Consequently, $f\langle\tau_i \to \ell\rangle$ dominates $f\langle\tau_{\mathrm{charge}} \to \ell\rangle$, since the charging speed of $v$ is better (or equally good) for all $\tau \in [\tau_i, \tau_{\mathrm{charge}}]$; see Figure 7a for an example. Since $\tau_i \leq \tau_{\mathrm{charge}}$ and a charging time of $\tau_i$ is sufficient to reach $v$, we have $\tau_{\mathrm{min}}(f\langle\tau_i \to \ell\rangle) \leq \tau_{\mathrm{min}}(f\langle\tau_{\mathrm{charge}} \to \ell\rangle)$. In other words, if the SoC function induced by $\tau_{\mathrm{charge}}$ is finite for some $\tau \in \mathbb{R}_{\geq 0}$, so is the function induced by $\tau_i$. Further, recall that the vertex $v$ is reached with an arrival SoC of $f\langle\ell\rangle(\tau_{\mathrm{trip}} + \tau_i)$ or $f\langle\ell\rangle(\tau_{\mathrm{trip}} + \tau_{\mathrm{charge}})$ when the charging time at $u$ is set to $\tau_i$ or $\tau_{\mathrm{charge}}$, respectively. By assumption, the slope $\angle_{\mathrm{new}}^{\ell}(\tau_{\mathrm{charge}})$ is strictly positive, which means that energy can still be charged at $v$ after the SoC has reached $f\langle\ell\rangle(\tau_{\mathrm{trip}} + \tau_{\mathrm{charge}})$, so it must hold that $f\langle\ell\rangle(\tau_{\mathrm{trip}} + \tau_{\mathrm{charge}}) < \beta_v^{\mathrm{max}}$. Hence, we can define the value

$$\Delta_{\mathrm{charge}} := \mathrm{cf}_v^{-1}(f\langle\ell\rangle(\tau_{\mathrm{trip}} + \tau_i), f\langle\ell\rangle(\tau_{\mathrm{trip}} + \tau_{\mathrm{charge}}))$$

that corresponds to the time to recharge the gap in arrival SoC between $f\langle\tau_i \to \ell\rangle$ and $f\langle\tau_{\mathrm{charge}} \to \ell\rangle$ at the new station $v$. For arbitrary values $\tau \geq \tau_{\mathrm{min}}(f\langle\tau_{\mathrm{charge}} \to \ell\rangle)$, we exploit the shifting property to obtain

$$
\begin{aligned}
f\langle\tau_i \to \ell\rangle(\tau) &= \mathrm{cf}_v(f\langle\ell\rangle(\tau_{\mathrm{trip}} + \tau_i), \tau - \tau_i - \tau_{\mathrm{trip}} - \tau_{\mathrm{init}}(v)) \\
&= \mathrm{cf}_v(f\langle\ell\rangle(\tau_{\mathrm{trip}} + \tau_{\mathrm{charge}}), \tau - \tau_i - \tau_{\mathrm{trip}} - \tau_{\mathrm{init}}(v) - \Delta_{\mathrm{charge}}) \\
&= f\langle\tau_{\mathrm{charge}} \to \ell\rangle(\tau + \tau_{\mathrm{charge}} - \tau_i - \Delta_{\mathrm{charge}}).
\end{aligned}
$$

We claim that $\Delta_{\mathrm{trip}} := \tau_{\mathrm{charge}} - \tau_i - \Delta_{\mathrm{charge}} \geq 0$ holds. This follows immediately from the fact that $\angle_{\mathrm{new}}^{\ell}(\tau) \geq \angle_{\mathrm{old}}^{\ell}(\tau)$ holds for all $\tau \in [\tau_i, \tau_{\mathrm{charge}}]$. Thus, the time $\Delta_{\mathrm{charge}}$ spent at $v$ cannot take longer than $\tau_{\mathrm{charge}} - \tau_i$, the time to charge the same amount of energy at $u$. Consequently, we obtain $\Delta_{\mathrm{trip}} \geq 0$. This, in turn, implies that for all $\tau \geq \tau_{\mathrm{min}}(f\langle\tau_{\mathrm{charge}} \to \ell\rangle)$, we have

$$
\begin{aligned}
f\langle\tau_i \to \ell\rangle(\tau) &= f\langle\tau_{\mathrm{charge}} \to \ell\rangle(\tau + \Delta_{\mathrm{trip}}) \\
&\geq f\langle\tau_{\mathrm{charge}} \to \ell\rangle(\tau).
\end{aligned}
$$

17

*Case 2:* $\angle_{\text{old}}^{\ell}(\tau_{\text{charge}}) \geq \angle_{\text{new}}^{\ell}(\tau_{\text{charge}})$. In this case, the charging station $u$ offers a more (or equally) favorable charging speed, so it pays off to spend more time charging at $u$. Recall that $f\langle\ell\rangle(\tau_{\text{trip}} + \tau_{\text{charge}}) = f\langle\tau_{\text{charge}} \to \ell\rangle(\tau_{\text{trip}} + \tau_{\text{charge}} + \tau_{\text{init}}(v))$ holds by definition. Given that the slope $\angle_{\text{old}}^{\ell}$ of $f\langle\ell\rangle$ is greater or equal for all values $\tau \in [\tau_{\text{charge}}, \tau_{i+1}]$, it follows that the inequality $f\langle\ell\rangle(\tau_{\text{trip}} + \tau) \geq f\langle\tau_{\text{charge}} \to \ell\rangle(\tau_{\text{trip}} + \tau + \tau_{\text{init}}(v))$ holds for arbitrary $\tau \in [\tau_{\text{charge}}, \tau_{i+1}]$. For real-valued $\tau \geq \tau_{i+1}$ (which only exist if $i + 1 < k$), we proceed along the lines of the first case to obtain a nonnegative value $\Delta_{\text{trip}} \geq 0$ that equals the difference between the time to charge from $f\langle\ell\rangle(\tau_{\text{trip}} + \tau_{\text{charge}})$ to $f\langle\ell\rangle(\tau_{\text{trip}} + \tau_{i+1})$ at $u$ and $v$, respectively. Since $u$ offers a charging speed at least as high as the one at $v$ in the whole interval $[\tau_{\text{charge}}, \tau_{i+1}]$, this difference, and therefore $\Delta_{\text{trip}}$, is again nonnegative. Thus, we can show (similar to Case 1) that

$$
\begin{aligned}
f\langle\tau_{i+1} \to \ell\rangle(\tau) &= f\langle\tau_{\text{charge}} \to \ell\rangle(\tau + \Delta_{\text{trip}}) \\
&\geq f\langle\tau_{\text{charge}} \to \ell\rangle(\tau)
\end{aligned}
$$

holds for arbitrary real-valued $\tau \geq \tau_{i+1}$; see Figure 7b for an illustration. Altogether, we obtain that $\max\{f\langle\ell\rangle(\tau), f\langle\tau_{i+1} \to \ell\rangle(\tau)\} \geq f\langle\tau_{\text{charge}} \to \ell\rangle(\tau)$ holds for all $\tau \in \mathbb{R}_{\geq 0}$, which completes the proof of the second case. $\qquad\square$

Given the label $\ell = (\tau_{\text{trip}}, \beta_u, u, b_{[u,\dots,v]})$ at the charging station $v \in S$, we spawn one new label $\tau \to \ell$ for each (finite) element in the switching sequence $T(\angle_{\text{old}}^{\ell}, \angle_{\text{new}}^{\ell})$. If the label is not dominated, it is added to the corresponding label set at $v$; see Lines 16–17 in Algorithm 1. For instance, in the special case that $u$ or $v$ is a swapping station, the switching sequence $[\tau_{\min}(f\langle\ell\rangle) - \tau_{\text{trip}}, \infty]$ consists of two values and exactly one new label is spawned at $v$ (which corresponds to spending the minimum amount of charging time at $u$ such that $v$ can be reached). Recall that the original label $\ell$ is not thrown away, reflecting the possibility of not using the charging station $v$. Furthermore, Lemma 1 implies that all nondominated solutions that extend $\ell$ by charging at $v$ are computed by the algorithm. By induction, correctness is also maintained for routes with two or more charging stops. Therefore, we obtain Theorem 1 given below.

**Theorem 1.** *If the switching sequences induced by arbitrary pairs of charging functions have finite length, the CFP algorithm terminates and finds the shortest feasible path between a given pair of vertices $s \in V$ and $t \in V$ for a given initial SoC $\beta_s \in [0, M]$.*

**Computing Switching Sequences.** In a practical implementation of the CFP algorithm, we need to be able to efficiently compute the switching sequences for given labels and charging functions; cf. Line 16 of Algorithm 1. If certain properties of the available charging functions are known, a reasonable approach is to derive the switching sequences for such functions analytically beforehand and provide a specialized implementation for them.

To give an example, we discuss the case where all charging functions in the network (and hence, all SoC functions) are *piecewise linear* and *concave*, as is the case in our experimental study (see Section 9). We show how a superset of the switching sequence is easily determined in this case. Given a label $\ell = (\tau_{\text{trip}}, \beta_u, u, b_{[u,\dots,v]})$ at a charging station $v \in S$, let its piecewise linear SoC function $f\langle\ell\rangle$ be given as a sequence $F = [(\tau_1, \beta_1), \dots, (\tau_k, \beta_k)]$ of breakpoints. In other words, we have $f\langle\ell\rangle(\tau) = -\infty$ for $\tau < \tau_1$, $f\langle\ell\rangle(\tau) = \beta_k$ for $\tau \geq \tau_k$, and for arbitrary values $\tau_i \leq \tau < \tau_{i+1}$, with $i \in \{1, \dots, k-1\}$, we evaluate the function by linear interpolation between the breakpoints $(\tau_i, \beta_i)$ and $(\tau_{i+1}, \beta_{i+1})$. Observe that we have $\tau_1 = \tau_{\min}(f\langle\ell\rangle)$. The following Lemma 2 shows that the switching sequence of the slopes induced by $f\langle\ell\rangle$ and the charging function $\text{cf}_v$ of $v$ is a subsequence of $[\tau_1 - \tau_{\text{trip}}, \dots, \tau_k - \tau_{\text{trip}}, \infty]$. Note that in particular, the switching sequence does not depend on the charging function $\text{cf}_v$ at $v$.

18

**Lemma 2.** *Given a label $\ell$ at a vertex $v \in S$, let its piecewise linear and concave SoC function $f\langle\ell\rangle$ be defined by the sequence $F = [(\tau_1, \beta_1), \ldots, (\tau_k, \beta_k)]$ of breakpoints. Similarly, let the charging function $\widetilde{cf}_v$ of $v$ be piecewise linear and concave. The switching sequence $T(\angle_{old}^{\ell}, \angle_{new}^{\ell})$ is a subsequence of $\bar{T}(\angle_{old}^{\ell}, \angle_{new}^{\ell}) := [\tau_1 - \tau_{\mathrm{trip}}, \ldots, \tau_k - \tau_{\mathrm{trip}}, \infty]$, i.e., $T(\angle_{old}^{\ell}, \angle_{new}^{\ell})$ contains only values that are also contained in $\bar{T}(\angle_{old}^{\ell}, \angle_{new}^{\ell})$.*

*Proof.* Since both $f\langle\ell\rangle$ and $\widetilde{cf}_v$ are piecewise linear, their corresponding slope functions $\angle_{\mathrm{old}}^{\ell}$ and $\angle_{\mathrm{new}}^{\ell}$ are *piecewise constant* functions, namely, $\angle_{\mathrm{old}}^{\ell} \equiv (\beta_{i+1} - \beta_i)/(\tau_{i+1} - \tau_i)$ holds in the interval $[\tau_i - \tau_{\mathrm{trip}}, \tau_{i+1} - \tau_{\mathrm{trip}})$ for arbitrary $i \in \{1, \ldots, k-1\}$, and an analogous statement holds for $\angle_{\mathrm{new}}^{\ell}$. This implies that each value of the switching sequence must correspond to a breakpoint of $\angle_{\mathrm{old}}^{\ell}$ or $\angle_{\mathrm{new}}^{\ell}$, because both functions are constant between these breakpoints. Moreover, since $f\langle\ell\rangle$ and $\widetilde{cf}_v$ are both concave on their subdomain with finite image, the functions $\angle_{\mathrm{old}}^{\ell}$ and $\angle_{\mathrm{new}}^{\ell}$ are *decreasing*. Consequently, the slope $\angle_{\mathrm{new}}^{\ell}$ can surpass $\angle_{\mathrm{old}}^{\ell}$ only at the breakpoints of $\angle_{\mathrm{old}}^{\ell}$, i.e., at points where the latter function decreases. The x-coordinates of these breakpoints are exactly the finite values of $\bar{T}(\angle_{\mathrm{old}}^{\ell}, \angle_{\mathrm{new}}^{\ell})$. $\qquad\square$

Lemma 2 implies that simply spawning one new label for each breakpoint of $f\langle\ell\rangle$ is sufficient to maintain correctness. Unnecessary labels are detected implicitly during dominance checks. Considering the more general case where charging functions are piecewise linear (but not necessarily concave), it is not hard to see that the length of the switching sequence must be linear in the number of breakpoints of both considered functions. Similar observations can be made for other realistic models of charging based on, e.g., exponential functions or piecewise combinations of linear and exponential functions.

## 5 A* Search

To accelerate basic CFP, we present techniques that extend A* search [32]. The basic idea of A* search is to use vertex potentials that guide the search towards the target. The potential of a vertex is added to the key of a label when updating the priority queue in Line 20 or Line 29 of Algorithm 1. Thereby, vertices that are closer to the target get smaller keys. Below, we first generalize the notion of potential consistency to our setting, before we introduce different consistent potential functions.

**Consistency of Potentials.** We aim at potential functions that are based on backward searches from the target vertex $t$, providing lower bounds on the trip time from any vertex to $t$. A consistent potential function is easily obtained from a single-criterion backward search as follows. It runs Dijkstra's algorithm from $t$, traversing arcs in backward direction and using the cost function $d$, which represents driving time along arcs. This yields, for each vertex $v \in V$, its minimum (unconstrained) driving time to reach $t$. These lower bounds on the remaining trip time induce a consistent potential function on the vertices [61]. However, we can do better, exploiting that the trip time from $v$ to $t$ depends on the SoC of a label. (Observe that both the charging time as well as the route and hence, the driving time, of an optimal solution can change for different SoC values.) We propose a potential function $\pi\colon V \times [0, M] \cup \{-\infty\} \to \mathbb{R}_{\geq 0} \cup \{\infty\}$ taking the current SoC into account, such that $\pi(v, \beta)$ yields a lower bound on the trip time from $v$ to $t$ if the SoC at $v$ is $\beta$, for $\beta \in [0, M] \cup \{-\infty\}$. We define $\pi(v, -\infty) := \infty$. We now discuss how correctness of our approach is maintained in the presence of a potential function that incorporates SoC.

First, we generalize the notion of consistency of a potential function. We say that a potential function $\pi\colon V \times [0, M] \cup \{-\infty\} \to \mathbb{R}_{\geq 0} \cup \{\infty\}$ is *consistent* if it results in nonnegative reduced

19

driving times, i.e., we get

$$\underline{d}((u,v),\beta) := d(u,v) - \pi(u,\beta) + \pi(v, b_{[u,v]}(\beta)) \geq 0$$

for all arcs $(u,v) \in A$ and values $\beta \in [0, M]$, where $b_{[u,v]}$ is the SoC profile of $(u,v)$. Second, for the SoC function $f\langle\ell\rangle$ representing a label $\ell = (\tau_{\mathrm{trip}}, \beta_u, u, b_{[u,...,v]})$ at a vertex $v \in V$, we define its *consistent key* (to be used in the priority queue of CFP) as $k^*(\ell) := \min_{\tau \in \mathbb{R}_{\geq 0}} \tau + \pi(v, f\langle\ell\rangle(\tau))$. We claim that consecutive consistent keys of labels generated after arc scans are increasing if the potential function $\pi$ is consistent. To see this, consider a label $\ell$ at a vertex $u \in V$ and the label $\ell'$ that is created after scanning an arc $(u,v) \in A$. Recall that $f\langle\ell'\rangle(\tau + d(u,v)) = b_{[u,v]}(f\langle\ell\rangle(\tau))$ holds by construction of $\ell'$, so we can substitute $\tau' := \tau + d(u,v)$ below to get

$$
\begin{aligned}
k^*(\ell) &= \min_{\tau \in \mathbb{R}_{\geq 0}} \tau + \pi(u, f\langle\ell\rangle(\tau)) \\
&\leq \min_{\tau \in \mathbb{R}_{\geq 0}} \tau + d(u,v) + \pi(v, b_{[u,v]}(f\langle\ell\rangle(\tau))) \\
&= \min_{\tau' \in \mathbb{R}_{\geq 0}} \tau' + \pi(v, f\langle\ell'\rangle(\tau')) \\
&= k^*(\ell').
\end{aligned}
$$

Similarly, we have to ensure that labels spawned at charging stations never have a smaller consistent key than the original label. Consider the charging function $\mathrm{cf}_v$ of a vertex $v \in S$. We denote by $\mathrm{cf}_{\max}(v)$ the *maximum slope* of the charging function $\mathrm{cf}_v$, which in most cases equals the value $\mathrm{cf}^*_{\max}(v) := \max_{\tau \in \mathbb{R}_{>0}} \{\partial \, \widetilde{\mathrm{cf}}_v(\tau)/\partial\tau\}$. (As before, we use the *right* derivative to ensure that slope is well-defined for piecewise-defined functions.) However, in the special case that $\widetilde{\mathrm{cf}}_v(0) \neq 0$, we incorporate initialization time to obtain a finite slope $\beta_v^{\min}/\tau_{\mathrm{init}}(v)$ for the initial SoC gain, presuming that $\tau_{\mathrm{init}}(v) \neq 0$. In total, we obtain the maximum slope $\mathrm{cf}_{\max}(v) := \max\{\mathrm{cf}^*_{\max}(v), \beta_v^{\min}/\tau_{\mathrm{init}}(v)\}$ of $v$. For instance, we get $\mathrm{cf}_{\max}(v) = M/\tau_{\mathrm{init}}(v)$ if $v$ is a swapping station. To ensure that labels spawned at $v$ do not result in decreasing keys, we demand for the slope of the potential $\pi(v, \beta)$ at $v$ that

$$\frac{\partial \pi(v, \beta)}{\partial \beta} \geq -\frac{1}{\mathrm{cf}_{\max}(v)}.$$

We say that the potential $\pi$ *overestimates charging speed* at $v$ in this case. Observe that overestimation of charging speed implies that the term $\tau + \pi(v, \widetilde{\mathrm{cf}}_v(\tau))$ is *increasing* with $\tau \in \mathbb{R}_{\geq 0}$, so charging at $v$ does not decrease the potential. Finally, we demand that the potential at the target is $\pi(t, \beta) = 0$ for arbitrary SoC $\beta \in [0, M]$. In summary, when using consistent keys, there are the following three requirements to a potential function $\pi$.

1. The potential function $\pi$ is consistent.

2. Potentials at charging stations overestimate charging speed.

3. The target vertex has a potential of 0 (for any finite SoC).

Then, the algorithm is label setting and the correct result is obtained as soon as $t$ is reached (since the key at $t$ equals trip time, so any label extracted at a later point must have a higher trip time). As a simple example, consider the plain CFP algorithm described in Section 3, which uses no potential function. This is equivalent to a potential function that evaluates to 0 at all vertices for arbitrary SoC. Clearly, the smallest feasible trip time of $f\langle\ell\rangle$ is in fact the consistent key of a label $\ell$. Moreover, observe that the potential function $\pi \equiv 0$ is consistent, overestimates charging speed, and equals 0 at the target. In what follows, we derive more sophisticated potential functions that make the search goal directed. For each potential function, we show that the requirements listed above are fulfilled.

**Potentials Based on Single-Criterion Search.** We introduce our first consistent potential function. To obtain a lower bound on the necessary charging time on the path from some vertex $v \in V$ to the target, let $\mathrm{cf}_{\max} := \max_{v \in S} \mathrm{cf}_{\max}(v)$ denote the maximum slope of *any* charging function in $S$, i.e., the maximum charging speed available in the network. We define a new cost function $\omega \colon A \to \mathbb{R}$, given as $\omega(a) := d(a) + (c(a)/\mathrm{cf}_{\max})$ for all $a \in A$. This function adds to the driving time of every arc a lower bound on the time required for charging the energy consumed on the arc. Note that the bound can become negative, due to negative consumption values.

Given the target vertex $t$, prior to running CFP, we perform three (single-criterion) runs of Dijkstra's algorithm from $t$, traversing arcs in backward direction and each using one of the cost functions $d$, $c$, and $\omega$. Thereby, we obtain, for every vertex $v \in V$, the distances $\mathrm{dist}_d(v, t)$, $\mathrm{dist}_c(v, t)$, and $\mathrm{dist}_\omega(v, t)$ from $v$ to $t$ wrt. the cost functions $d$, $c$, and $\omega$, respectively. Note that computation of $\mathrm{dist}_c(v, t)$ and $\mathrm{dist}_\omega(v, t)$ is label correcting, due to negative weights. We can apply potential shifting [36] to remedy this issue, but the effect on overall running time is negligible in practice. Using the obtained distances, the potential function $\pi_\omega \colon V \times [0, M] \cup \{-\infty\} \to \mathbb{R}_{\geq 0} \cup \{\infty\}$ is defined by

$$\pi_\omega(v, \beta) := \begin{cases} \mathrm{dist}_d(v, t) & \text{if } \beta \geq \mathrm{dist}_c(v, t), \\ \mathrm{dist}_\omega(v, t) - \frac{\beta}{\mathrm{cf}_{\max}} & \text{otherwise,} \end{cases} \tag{3}$$

for all $v \in V$ and $\beta \in [0, M]$. It uses the minimum (unrestricted) driving time as lower bound on the remaining trip time in case that the current SoC is greater or equal to the minimum energy required to reach the target without recharging. Otherwise, we know that we have to spend some additional time for charging on the way to the target $t$, so we use a bound induced by the weight function $\omega$. Note that we can discard labels whose consumption exceeds $M$ in the search that computes $\mathrm{dist}_c(v, t)$ to save some time in practice. Lemma 3 formally proves that the potential function $\pi_\omega$ defined above is consistent.

**Lemma 3.** *The potential function $\pi_\omega$ is consistent.*

*Proof.* To prove the claim, we show that reduced costs $\underline{d}(\cdot, \cdot)$ are nonnegative, i.e., the inequality $\underline{d}((u, v), \beta) = d(u, v) - \pi_\omega(u, \beta) + \pi_\omega(v, b_{[u,v]}(\beta)) \geq 0$ holds for all $(u, v) \in A$ and $\beta \in [0, M]$. Consider an arbitrary arc $(u, v) \in A$ and assume that the SoC at $u$ is $\beta \in [0, M]$. We distinguish four cases.

*Case 1:* $\beta < \mathrm{dist}_c(u, t)$ and $b_{[u,v]}(\beta) < \mathrm{dist}_c(v, t)$. The claim follows after a few simple substitutions. We can make use of the fact that $\beta - c(u, v) \geq b_{[u,v]}(\beta)$ holds for all $\beta \in [0, M]$. Consistency then follows directly from the triangle inequality, after the following steps

$$\begin{aligned} \underline{d}((u, v), \beta) &= d(u, v) - \pi_\omega(u, \beta) + \pi_\omega(v, b_{[u,v]}(\beta)) \\ &= d(u, v) - \mathrm{dist}_\omega(u, t) + \frac{\beta}{\mathrm{cf}_{\max}} + \mathrm{dist}_\omega(v, t) - \frac{b_{[u,v]}(\beta)}{\mathrm{cf}_{\max}} \\ &\geq d(u, v) + \frac{c(u, v)}{\mathrm{cf}_{\max}} - \mathrm{dist}_\omega(u, t) + \mathrm{dist}_\omega(v, t) \\ &= \omega(u, v) - \mathrm{dist}_\omega(u, t) + \mathrm{dist}_\omega(v, t) \\ &\geq 0. \end{aligned}$$

*Case 2:* $\beta < \mathrm{dist}_c(u, t)$ and $b_{[u,v]}(\beta) \geq \mathrm{dist}_c(v, t)$. We make use of both preconditions together with the fact that $\beta - c(u, v) \geq b_{[u,v]}(\beta)$ holds for all $\beta \in [0, M]$ to obtain the inequalities

$$\beta \geq b_{[u,v]}(\beta) + c(u, v) \geq \mathrm{dist}_c(v, t) + c(u, v) \geq \mathrm{dist}_c(u, t) > \beta,$$

21

which yield a contradiction. Hence, this case cannot occur.

*Case 3:* $\beta \geq \mathrm{dist}_c(u,t)$ and $b_{[u,v]}(\beta) < \mathrm{dist}_c(v,t)$. We know that, due to the triangle inequality, $d(u,v) + \mathrm{dist}_d(v,t) \geq \mathrm{dist}_d(u,t)$ holds. Moreover, $\mathrm{dist}_c(v,t) - b_{[u,v]}(\beta) > 0$ holds by assumption. We can further exploit that $\mathrm{dist}_\omega(v,t)$ is at least the sum of $\mathrm{dist}_d(v,t)$ and $\mathrm{dist}_c(v,t)/\mathrm{cf}_{\max}$ (note that it is possibly greater if the shortest $u$–$t$ paths in the graph wrt. the cost functions $d$ and $c$ differ). After some substitutions, we thus get

$$
\begin{aligned}
\underline{d}((u,v),\beta) &= d(u,v) - \pi_\omega(u,\beta) + \pi_\omega(v,b_{[u,v]}(\beta)) \\
&= d(u,v) - \mathrm{dist}_d(u,t) + \mathrm{dist}_\omega(v,t) - \frac{b_{[u,v]}(\beta)}{\mathrm{cf}_{\max}} \\
&\geq d(u,v) - \mathrm{dist}_d(u,t) + \mathrm{dist}_d(v,t) + \frac{\mathrm{dist}_c(v,t)}{\mathrm{cf}_{\max}} - \frac{b_{[u,v]}(\beta)}{\mathrm{cf}_{\max}} \\
&\geq 0.
\end{aligned}
$$

*Case 4:* $\beta \geq \mathrm{dist}_c(u,t)$ and $b_{[u,v]}(\beta) \geq \mathrm{dist}_c(v,t)$. This case is trivial; feasibility follows directly from the triangle inequality. $\square$

The potential function $\pi_\omega$ always evaluates to 0 at the target and overestimates charging speed by construction. As for consistent keys, defined for a label $\ell$ as $k^*(\ell) = \min_{\tau \in \mathbb{R}_{\geq 0}} \tau + \pi(v, f\langle\ell\rangle(\tau))$, observe that the corresponding terms $\tau + \mathrm{dist}_d(v,t)$ and $\tau + \mathrm{dist}_\omega(v,t) - (f\langle\ell\rangle(\tau)/\mathrm{cf}_{\max})$ in Equation 3 increase with $\tau$ (the term $f\langle\ell\rangle(\tau)/\mathrm{cf}_{\max}$ increases with a slope of at most 1). Thus, we obtain the consistent key for the label $\ell$ by computing the value of $\tau + \pi_\omega(v, f\langle\ell\rangle(\tau))$ at the minimum feasible trip time $\tau_1 := \tau_{\min}(f\langle\ell\rangle)$ of $f\langle\ell\rangle$ and at the minimum trip time $\tau_2$ with $f\langle\ell\rangle(\tau_2) \geq \mathrm{dist}_c(v,t)$, if it exists. The minimum of both values yields a consistent key, given as $k^*(\ell) = \min\{\tau_1, \tau_2\}$. Together with Lemma 3, this implies correctness of CFP when applying potential shifting with the function $\pi_\omega$.

**Potentials Based on Bound Function Propagation.** Even though the potential function $\pi_\omega$ incorporates SoC, lower bounds may be too conservative in that they presume recharging is possible at any time and with the best charging rate. We attempt to be more precise, while keeping computational effort limited, by explicitly constructing lower bound functions. Again, we run (at query time) a label-correcting backward search from the given target $t$, but this time computing for each vertex $v \in V$ a *piecewise linear function* $\underline{f}\colon \mathbb{R} \to \mathbb{R}_{\geq 0} \cup \{\infty\}$ mapping SoC to a lower bound on the trip time from $v$ to $t$. These piecewise linear functions are represented by sequences $F = [(\beta_1, \tau_1), \ldots, (\beta_k, \tau_k)]$ of breakpoints such that $\underline{f}(\beta) = \infty$ for $\beta < \beta_1$, $\underline{f}(\beta) = \tau_k$ for $\beta \geq \beta_k$, and the function is evaluated by linear interpolation between breakpoints. If the sequence $F$ of breakpoints is empty, denoted $F = \emptyset$, we obtain $\underline{f} \equiv \infty$. During the backward search, each vertex stores a *single* label consisting of such a piecewise linear function. To simplify the search, we ignore battery constraints. Hence, domains of bounds are not restricted to $[0, M]$ and we compute (possibly negative) lower bounds on the SoC necessary to reach $t$. However, we maintain the invariant that all functions are *decreasing* and *convex* on the interval $[\beta_1, \infty)$. This greatly simplifies label updates and the computation of functions, which we describe in detail below.

The algorithm resembles (label-correcting) profile search [13, 53] and its pseudocode is outlined in Algorithm 2. It is initialized in Lines 1–4 with a function $\underline{f}_t$, represented by a single breakpoint $F_t = [(0,0)]$ at the target vertex $t$, i.e., $\underline{f}_t(\beta)$ evaluates to 0 for arbitrary values $\beta \in \mathbb{R}_{\geq 0}$. All other labels are empty, so their functions always evaluate to $\infty$. Each step of the algorithm's main loop (Lines 5–13) scans a vertex with minimum key (the key of a vertex $v \in V$ is the minimum function value $\min_{\beta \in \mathbb{R}} \underline{f}_v(\beta)$ of its label $\underline{f}_v$).

**Algorithm 2:** Pseudocode of the function-propagating potential search for CFP.

```
   // initialize labels
 1 foreach v ∈ V do
 2 │   f_v ⟵ ∅
 3 f_t ⟵ [(0,0)]
 4 Q.insert(t, 0)

   // run main loop
 5 while Q.isNotEmpty() do
 6 │   v ⟵ Q.deleteMin()

   │   // handle charging stations
 7 │   if v ∈ S then
 8 │   │   f_v ⟵ extend(f_v, c̃f_v)

   │   // scan outgoing arcs
 9 │   foreach (u, v) ∈ A do
10 │   │   f ⟵ shift(f_v, [(c(u,v), d(u,v))])
11 │   │   if ∃τ ∈ ℝ: f(τ) < f_u(τ) then
12 │   │   │   f_u ⟵ merge(f_u, f)
13 │   │   │   Q.update(u, k(f_u))
```

Whenever the search reaches a charging station $v \in S$, we have to ensure that the possibility of recharging is reflected in the label of $v$ and that the function overestimates charging speed. To this end, we *extend* the (tentative) lower bound function $\underline{f}_v$ with the charging function $\widetilde{\mathrm{cf}}_v$ (Line 8 of Algorithm 2). The result of this operation is a new (tentative) lower bound on the trip time that incorporates recharging at $v$ and has a slope of at least $-1/\mathrm{cf}_{\max}(v)$ (on its subdomain with finite image). Assume we are given a piecewise linear, convex function $\underline{f}_v$ at $v$ with breakpoints $F_v = [(\beta_1, \tau_1), \ldots, (\beta_k, \tau_k)]$ that maps SoC to trip time without recharging at $v$. We obtain the result $\underline{f}$ of extending $\underline{f}_v$ with $\widetilde{\mathrm{cf}}_v$ as follows; see Figure 8 for an example. In accordance with our requirements for the potential at $v$, we approximate $\widetilde{\mathrm{cf}}_v$ with a lower bound given by a single segment with slope $1/\mathrm{cf}_{\max}(v)$; see Figure 8a. Distributing (lower bounds on) charging time among $\widetilde{\mathrm{cf}}_v$ and stations represented by $\underline{f}_v$ then corresponds to shifting this segment along the y-axis such that it intersects the graph of $\underline{f}_v$; see Figure 8b. To find a lower bound on the best possible distribution, let $i \in \{2, \ldots, k-1\}$ be the unique index (if it exists) such that

$$\frac{\tau_i - \tau_{i-1}}{\beta_i - \beta_{i-1}} \le -\frac{1}{\mathrm{cf}_{\max}(v)} < \frac{\tau_{i+1} - \tau_i}{\beta_{i+1} - \beta_i},$$

i.e., the (negative, inverse) maximum slope of $\widetilde{\mathrm{cf}}_v$ is at least the slope of the segment from $\beta_{i-1}$ to $\beta_i$, but lower than the slopes of all subsequent segments. We set $i := 1$ if the maximum slope of $\widetilde{\mathrm{cf}}_v$ is below the slope of the first segment (from $\beta_1$ to $\beta_2$), and $i := k$ if the maximum slope of $\widetilde{\mathrm{cf}}_v$ is at least the slope of the last segment (from $\beta_{k-1}$ to $\beta_k$). Then, if $\beta_i \le 0$, the function $\underline{f}_v$ remains unchanged, i.e., $\underline{f} = \underline{f}_v$, as charging at $v$ cannot decrease the lower bound in this case. Otherwise, $\underline{f}$ is defined by the sequence $F := [(0, \tau_i + \beta_i/\mathrm{cf}_{\max}(v)), (\beta_i, \tau_i), \ldots, (\beta_k, \tau_k)]$; see Figure 8b. The first segment of this function corresponds to an estimate of the time to charge to an SoC of $\beta_i$ plus the
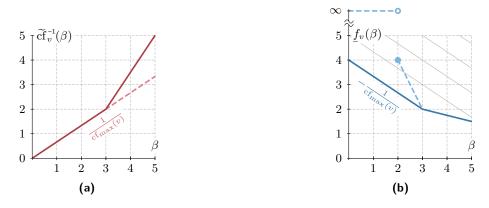
Figure 8: Extending functions in the function propagating backward search. (a) The (expanded) inverse charging function $\widetilde{\mathrm{cf}}_v^{-1}$ of a charging station $v \in S$ maps SoC to charging time. The dashed line indicates its (inverse) maximum slope, which we use to approximate the (inverse) charging function with a single segment. (b) The lower bound function $\underline{f}_v$ (dashed) and the result $\underline{f}$ of extending it with the charging function $\widetilde{\mathrm{cf}}_v$ (dark blue)

remaining trip time to $t$. By construction, $\underline{f}$ is convex and overestimates charging speed (provided that the same holds for $\underline{f}_v$).

Since we ignore battery constraints, scanning an incoming arc $(u, v) \in A$ of $v$ boils down to shifting all breakpoints of the current function $\underline{f}_v$ by $c(u, v)$ and $d(u, v)$ on the x- and y-axis, respectively. More formally, given the breakpoints $F_v = [(\beta_1, \tau_1), \ldots, (\beta_k, \tau_k)]$ of $\underline{f}_v$, we compute a function $\underline{f}$ with

$$F := [(\beta_1 + c(u, v), \tau_1 + d(u, v)), \ldots, (\beta_k + c(u, v), \tau_k + d(u, v))].$$

Then, we check whether the function $\underline{f}$ is smaller than the function $\underline{f}_u$ in the label of $u$ for at least one $\beta \in \mathbb{R}$. If this is the case, we *merge* $\underline{f}$ and $\underline{f}_u$, i.e., we compute the function defined as their pointwise minimum $\underline{f}_u(\beta) := \min\{\underline{f}(\beta), \underline{f}_u(\beta)\}$ for all $\beta \in \mathbb{R}$. This operation requires a linear-time scan over the breakpoints of both involved functions, similar to label-correcting profile searches in time-dependent route planning [3, 8, 13]. The resulting function $\underline{f}_v$ is again piecewise linear, but may no longer be convex. Therefore, we compute, during each merge operation, the *convex lower hull* of the result using Graham's scan [29]. While (slightly) deteriorating the quality of the bound, this reduces the number of breakpoints and simplifies handling of charging stations. We obtain a convex function, which is stored in the label of $v$. The vertex $v$ is also updated in the priority queue.

Again, it is easy to see that the resulting potential function $\pi_f \colon V \times [0, M] \cup \{-\infty\} \to \mathbb{R}_{\geq 0} \cup \{\infty\}$ is consistent when using the computed bounds by setting $\pi_f(v, \beta) := \underline{f}_v(\beta)$ for $v \in V$ and $\beta \in [0, M]$. Lemma 4 proves this formally. The consistent key $k^*(\ell) = \min_{\tau \geq 0} \tau + \pi_f(v, f\langle \ell \rangle(\tau))$ of a label $\ell$ at some vertex $v \in V$ is computed in a linear scan over $f\langle \ell \rangle$ and the breakpoints of the piecewise linear function $\underline{f}_v$.

**Lemma 4.** *The potential function $\pi_f$ is consistent.*

*Proof.* For an arbitrary arc $(u, v) \in A$, consider the piecewise linear functions $\underline{f}_u$ and $\underline{f}_v$ at $u$ and $v$, respectively, after the backward search has terminated. We show that the reduced costs $\underline{d}((u, v), \beta)$ are nonnegative for all $\beta \in [0, M]$. We know that $\underline{f}_u$ is upper bounded by the result of shifting $\underline{f}_v$ by the costs $c(u, v)$ and $d(u, v)$ of the arc $(u, v)$ traversed in backward direction, since this function was merged with $\underline{f}_u$ during the search. This implies that $\pi_f(v, \beta - c(u, v)) + d(u, v) \geq \pi_f(u, \beta)$ holds for arbitrary $\beta \in [0, M]$. Moreover, $b_{[u,v]}(\beta)$ is a lower bound on $\beta - c(u, v)$, so we have $\pi_f(v, \beta - c(u, v)) \leq \pi_f(v, b_{[u,v]}(\beta))$ because $\pi_f$ decreases with increasing SoC. Plugging this into the

above inequality, we obtain $d(u,v) - \pi_f(u,\beta) + \pi_f(v, b_{[u,v]}(\beta)) \geq 0$ for all $\beta \in [0, M]$, which proves the claim. □

**Potential Search on Demand.** Computing the potential function for every vertex in the graph is wasteful for short-range queries. To speed up such queries, we run the backward searches that compute vertex potentials *on demand*: Whenever the CFP search requires the potential of some vertex $v \in V$ that was not scanned by the backward search yet, the backward search is executed until $v$ is reached. It is then suspended and only resumed if the potential of another vertex is required that the backward search has not visited. This procedure yields consistent potentials if the backward search is label setting (otherwise, there is no guarantee that a lower bound was computed when a vertex is scanned for the first time). In case of the potential function $\pi_\omega$, this can be ensured by applying Johnson's shifting technique [36] to its backward searches. For the potential function $\pi_f$, however, the function-propagating search is label correcting, so computing $\pi_f$ on demand becomes more involved. We describe modifications to the search and the lower bounds to ensure that $\pi_f$ is indeed a consistent potential, even if the search is suspended before it terminates.

First, we can ensure that the minimum key in the priority queue of the backward search is nondecreasing during the course of the algorithm: After scanning an arc $(u,v) \in A$, consider two functions $\underline{f}$ and $\underline{f}_u$ corresponding to the result of scanning the arc $(u,v)$ in backward direction and the current label at $u$, respectively, before merging these two functions (see Line 12 in Algorithm 2). Let $\underline{f}_u^*$ denote the result after merging. Since $\underline{f}_u^*$ is the convex lower hull of the minimum of $\underline{f}$ and $\underline{f}_u$, every breakpoint in $\underline{f}_u^*$ must also be contained in $\underline{f}$ or $\underline{f}_u$. Let $(\beta, \tau)$ be the breakpoint with minimum trip time $\tau \in \mathbb{R}_{\geq 0}$ contained in the corresponding sequence $\underline{F}_u^*$ of $\underline{f}_u^*$, but not in the sequence $\underline{F}_u$ of $\underline{f}_u$, i.e., $(\beta, \tau) \in \underline{F}_u^*$ and $(\beta, \tau) \notin \underline{F}_u$. If the result of merging improves the label at $u$, such a point must exist. We set the key of $u$ in the priority queue to the minimum of $\tau$ and its current key. Since driving time is nonnegative, scanning an arc may only increase the driving time of newly added breakpoints. As a result, propagating breakpoints never decreases the minimum key in the priority queue.

Assume that the backward search is suspended at some point and let $\tau^* \in \mathbb{R}_{\geq 0}$ be current minimum key of the priority queue. For each vertex $v \in V$, consider its current label $\underline{f}_v$. Let $\underline{f}_v^*$ denote the function obtained after applying Graham's scan to the result of merging $\underline{f}_v$ and the function induced by the single breakpoint $[(0, \tau^*)]$. We claim that the potential function $\pi_f^*$ with $\pi_f^*(v, \beta) := \underline{f}_v^*(\beta)$ for all $v \in V$ and $\beta \in [0, M]$ is consistent. To see this, consider a (multi-)graph $G' = (V, A')$ constructed from the input graph $G$ by adding, for every $v \in V$, an arc $(v, t)$ with driving time $d(v, t) := \tau^*$ and consumption $c(v, t) := 0$. It is easy to verify that the potential function $\pi_f^*$ on $G$ is equivalent to the potential function $\pi_f$ on $G'$. Hence, $\pi_f^*$ is a consistent potential function for $G'$. Observe that this immediately implies that $\pi_f^*$ is also consistent on $G$, since reduced arc costs must be nonnegative for the subset $A \subseteq A'$. Lemma 5 follows immediately from this observation and Lemma 4.

**Lemma 5.** *The potential function $\pi_f^*$ is consistent.*

Note that we have to keep the original function $\underline{f}_v$ after suspending the backward search, in case it is resumed later. Hence, we do not store $\pi_f^*$ explicitly, but perform the necessary merge operation and Graham's scan on demand when the potential is requested. Given that keys of labels are consistent and by Lemma 3, Lemma 4, and Lemma 5, we obtain Theorem 2, which summarizes our findings on A* search.

**Theorem 2.** *The CFP algorithm computes the correct output when using either of the potential functions $\pi_\omega$, $\pi_f$, or $\pi_f^*$.*

**Implementation Details.** When using potentials on demand, we can suspend the backward search at any time and derive lower bounds for CFP. However, bound quality of $\pi_f^*$ may deteriorate if the search is suspended too early, as it depends on the current minimum key in the priority queue of the backward search. Therefore, we do not abort the search immediately when a vertex $v \in V$ in question is scanned for the first time, but continue until the minimum key $\tau^*$ in the priority queue is significantly greater than the key induced by the first breakpoint $(\beta_1, \tau_1)$ of $\underline{f}_v$ (in our experiments, we suspend the search if $\tau^* > \min\{2\tau_1, \tau_1 + 3\,600\}$, where time is measured in seconds).

# 6 Contraction Hierarchies

During an offline step, CH preprocessing [24] iteratively contracts the vertices of the input graph and adds *shortcuts* in the remaining graph to retain correct distances. These shortcuts then help reducing the search space in online queries. When adapting CH to our scenario, vertex contraction becomes more expensive, as each shortcut represents a pair consisting of driving time and an SoC profile. Moreover, we need a shortcut for every nondominated path. Hence, the resulting search graph may contain multi-arcs.

We compute a *partial* CH, i.e., we contract only some vertices (the *component*), leaving an uncontracted *core* graph—a common approach in complex scenarios [16, 21, 55]. As in Storandt [55], we keep all charging stations in the graph uncontracted. Thus, complexity induced by charging stations only is contained within the core (simplifying the search in the component). Shortcuts store the driving time and the SoC profile (represented by three values as described in Section 2) of the path they represent.

**Witness Search.** During preprocessing, we perform *witness searches* when contracting a vertex, to test whether all shortcut candidates are necessary to maintain distances in the current overlay. Given a shortcut candidate $(u, v)$ with $u \in V$ and $v \in V$, we run a variant of the BSP algorithm that propagates labels consisting of the driving time and the SoC profile of a path (represented by three values), starting from $u$. A label *dominates* another label in this search if its driving time is smaller or equal to the driving time of the other label and its SoC profile dominates that of the other label. Keys in the priority queue follow a lexicographic order of the labels. The witness search stops if either the shortcut candidate is dominated by a label at $v$ or the minimum key in the priority queue exceeds the key induced by the shortcut candidate.

In order to reduce preprocessing time, we simplify the witness search as follows. First, we only search for single witnesses that dominate a shortcut candidate. In other words, we only perform pairwise comparisons between labels at the head of a shortcut candidate and the candidate itself. Thereby, we might insert an unnecessary shortcut whose SoC profile is dominated by the upper envelope of multiple SoC profiles corresponding to labels with lower driving time. Second, during the witness search, we limit the number of labels per vertex to a small constant (10 in our experiments). Whenever this size is exceeded, we identify (in a linear scan over the sorted labels) the pair of labels that has the minimum difference in terms of driving time. Of these two labels, we remove the one with smaller difference to its next closest label (in order to keep the gap between the remaining labels small). Finally, following Geisberger et al. [24], we prune the search after a fixed hop limit (20 in our experiments). Taking these measures, we may possibly insert unnecessary shortcuts. Thus, queries may slow down slightly, but correctness is not affected.

**Queries.** Since we compute a partial CH, the query algorithm consists of two phases. Given a source $s \in V$, a target $t \in V$, and the initial SoC $\beta \in [0, M]$, the first phase runs a backward

CH search from $t$, scanning only upward arcs wrt. the vertex order. This search operates on the component, so it is pruned at core vertices (i. e., core vertices are not inserted into the priority queue). As the component contains no charging stations, a basic variant of the BSP algorithm suffices. Note, however, that the SoC at $t$ is yet unknown and therefore, the search algorithm computes SoC profiles instead of SoC values (as in the witness search). For every nondominated label at any vertex visited by the search, we add a (temporary) shortcut from this vertex to the target. The second phase runs CFP from $s$ and is restricted to upward arcs, core arcs, and the temporary arcs added by the backward search.

**Implementation Details.** During preprocessing of CH, the next vertex in the contraction order is determined from the measures Edge Difference (ED), Deleted Neighbors (DN), and Cost of Queries (CQ) according to Geisberger et al. [24]. The priority of a vertex (higher priority means higher importance) is then set to $64\,\mathrm{ED} + \mathrm{DN} + \mathrm{CQ}$. To reduce the number of necessary witness searches, we cache shortcuts computed during the computation of ED. This requires a simulated contraction; see Geisberger et al. [24]. Whenever witness searches for multiple shortcuts with the same source are required during contraction of some vertex, we run a single multi-target search instead. Further, to improve query times, we reorder vertices after preprocessing, such that core vertices are in consecutive memory.

# 7 CHArge

Combining CH and A\* search (restricting A\* search to the core), we obtain our fastest exact algorithm, *CHArge (CH, A\*, Charging Stops)*. The query algorithm consists of three phases, namely, a unidirectional (backward) phase from $t$ in the component to add temporary shortcuts, a backward search in the (much smaller) core enriched with temporary shortcuts to compute a potential function (either $\pi_\omega$ or $\pi_f^*$), and a forward phase running CFP (augmented with A\* search) from $s$, which is restricted to upward arcs, core arcs, and temporary arcs. Potentials of component vertices are set to 0 for this search. Observe that consistency of the resulting potential function is not violated, since there are no arcs pointing from the core into the component. As described in Section 5, potentials in the core can be computed on demand, in which case the second and third phase are interweaved and their searches are executed alternately.

**Computing Potentials in the Core.** To decrease running time of the second phase, we precompute lower bounds of core shortcuts for the potential function $\pi_f^*$, i. e., for each (ordered) pair $u \in V$ and $v \in V$ of vertices connected by at least one shortcut, we compute a decreasing and convex piecewise linear function that yields a lower bound on driving time from $u$ to $v$ for a given SoC. To this end, we perform Graham's scan [29] on all pairs $[(\beta, \tau)]$ of minimum required SoC $\beta \in [0, M]$ and driving time $\tau \in \mathbb{R}_{\geq 0}$ corresponding to some shortcut arc between $u$ and $v$ (recall that the component may contain multi-arcs). However, this also requires us to adapt the function-propagating search in the second phase, since scanning an arc no longer consists of simply shifting a function by two constant values (cf. Line 10 of Algorithm 2). Instead, piecewise linear functions of labels have to be *linked* with shortcut arcs, which are represented by piecewise linear functions as well. In what follows, we describe how this is done in time linear in the number of breakpoints of both functions.

Consider two piecewise linear functions $f_1 \colon \mathbb{R} \to \mathbb{R}_{\geq 0} \cup \{\infty\}$ and $f_2 \colon \mathbb{R} \to \mathbb{R}_{\geq 0} \cup \{\infty\}$ mapping SoC to lower bounds on the trip time along two paths in the graph, such that both are decreasing and convex on their subdomains with finite image. Let the functions be given by their respective sequences $F_1 = [(\beta_1^1, \tau_1^1), \ldots, (\beta_1^k, \tau_1^k)]$ and $F_2 = [(\beta_2^1, \tau_2^1), \ldots, (\beta_2^\ell, \tau_2^\ell)]$ of breakpoints. The *link operation* takes
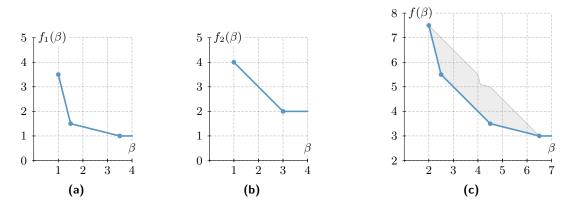
Figure 9: Linking piecewise linear lower bound functions. Linear segments between indicated breakpoints show (finite) values of two functions and the result of linking them. (a) The function $f_1$ is defined by three breakpoints. (b) The function $f_2$ is defined by two breakpoints. (c) Linking $f_1$ and $f_2$ yields the function $f$. It is the lower envelope of the shaded area, which corresponds to (finite) values of $f_1(\beta^*) + f_2(\beta - \beta^*)$ for different choices of $\beta^* \in \mathbb{R}$.

the functions $f_1$ and $f_2$ as input and computes a function $f$ that reflects the concatenation of both paths. Hence, given an arbitrary SoC $\beta \in [0, M]$, the value $f(\beta)$ is a lower bound on the trip time when traversing the paths represented by $f_1$ and $f_2$, such that overall consumption does not exceed the SoC $\beta$. To this end, the link operation identifies values $\beta_1 \in \mathbb{R}$ and $\beta_2 \in \mathbb{R}$, such that $\beta_1 + \beta_2 = \beta$ and $f_1(\beta_1) + f_2(\beta_2)$ is minimized. Hence, we seek to compute the function $f \colon \mathbb{R} \to \mathbb{R}_{\geq 0} \cup \{\infty\}$ with

$$f(\beta) := \min_{\beta^* \in \mathbb{R}} f_1(\beta^*) + f_2(\beta - \beta^*), \tag{4}$$

which yields the desired lower bound on the trip time for traversing $f_1$ and $f_2$. Figure 9 shows an example. Below, we describe an algorithm that computes a sequence $F$ of breakpoints that represents this function $f$. Afterwards, we prove its correctness.

Starting with an empty sequence $F = \emptyset$, the link operation iteratively appends breakpoints to $F$, each of which is the sum of two breakpoints from $F_1$ and $F_2$. For $\beta = \beta_1^1 + \beta_2^1$, there exists exactly one value $\beta^* = \beta_1^1$ in Equation 4 that yields a finite trip time. We obtain $f(\beta_1^1 + \beta_2^1) = \tau_1^1 + \tau_2^1$ and the first breakpoint of $F$ is $(\beta_1^1 + \beta_2^1, \tau_1^1 + \tau_2^1)$. For subsequent breakpoints, the basic idea is to follow the function that offers the better (i.e., lower) slope. Assume that the previous breakpoint added to $F$ is $(\beta_1^i + \beta_2^j, \tau_1^i + \tau_2^j)$ for some $i \in \{1, \ldots, k\}$ and $j \in \{1, \ldots, \ell\}$. Consider the slope

$$\sigma_1^i := \begin{cases} \frac{\beta_1^{i+1} - \beta_1^i}{\tau_1^{i+1} - \tau_1^i} & \text{if } i < k, \\ 0 & \text{otherwise}, \end{cases}$$

of the next segment of the function $f_1$. Let the slope $\sigma_2^j$ be defined symmetrically. Then the next breakpoint is $(\beta_1^{i+1} + \beta_2^j, \tau_1^{i+1} + \tau_2^j)$ if $\sigma_1^i < \sigma_2^j$ and $(\beta_1^i + \beta_2^{j+1}, \tau_1^i + \tau_2^{j+1})$ if $\sigma_1^i > \sigma_2^j$. In other words, we pick the next breakpoint of the currently steeper function (keeping the same point as before for the other function). In the special case $\sigma_1^i = \sigma_2^j$ we obtain three collinear points, so the next breakpoint is $(\beta_1^{i+1} + \beta_2^{j+1}, \tau_1^{i+1} + \tau_2^{j+1})$. The scan is stopped as soon as the last point $(\beta_1^k + \beta_2^\ell, \tau_1^k + \tau_2^\ell)$ is reached and added to $F$.

Clearly, the scan described above runs in linear time in the number $k + \ell$ of breakpoints of $f_1$ and $f_2$. Moreover, as segments are appended in increasing order of original slope, the resulting function $f$ is also decreasing and convex. Lemma 6 formally proves that the link operation in fact computes the correct result.

**Lemma 6.** *Let $f_1\colon \mathbb{R} \to \mathbb{R}_{\geq 0} \cup \{\infty\}$ and $f_2\colon \mathbb{R} \to \mathbb{R}_{\geq 0} \cup \{\infty\}$ be piecewise linear functions defined by sequences $F_1 = [(\beta_1^1, \tau_1^1), \ldots, (\beta_1^k, \tau_1^k)]$ and $F_2 = [(\beta_2^1, \tau_2^1), \ldots, (\beta_2^\ell, \tau_2^\ell)]$ of breakpoints, respectively, such that both functions are decreasing and convex on their subdomains $[\beta_1^1, \infty)$ and $[\beta_2^1, \infty)$ with finite image. The link operation described above computes a sequence $F$ of breakpoints that corresponds to a function $f\colon \mathbb{R} \to \mathbb{R}_{\geq 0} \cup \{\infty\}$ with $f(\beta) = \min_{\beta^* \in \mathbb{R}} f_1(\beta^*) + f_2(\beta - \beta^*)$.*

*Proof.* For $\beta < \beta_1^1 + \beta_2^1$, there exists no value $\beta^* \in \mathbb{R}$ such that $f_1(\beta^*)$ and $f_2(\beta - \beta^*)$ are both finite, so the result of linking equals $\infty$. For $\beta = \beta_1^1 + \beta_2^1$ there is exactly one such value $\beta^* \in \mathbb{R}$ that yields a finite trip time and we obtain $f(\beta_1^1 + \beta_2^1) = \tau_1^1 + \tau_2^1$, which corresponds to the first breakpoint of $F$.

Let $(\beta_1^i + \beta_2^j, \tau_1^i + \tau_2^j)$ with $i \in \{1, \ldots, k\}$ and $j \in \{1, \ldots, \ell\}$ denote the last point that was added to $F$. Without loss of generality, assume that $i < k$ and the next segment of the function $f_1$ is at least as steep as the next segment of $f_2$, i.e., $\sigma_1^i \leq \sigma_2^j$. Hence, our algorithm sets $(\beta_1^{i+1} + \beta_2^j, \tau_1^{i+1} + \tau_2^j)$ as the next breakpoint of $F$ (or adds a segment that contains this point in the special case $\sigma_1^i = \sigma_2^j$). Thus, the slope of $f$ is $\sigma_1^i$ for all $\beta \in [\beta_1^i + \beta_2^j, \beta_1^{i+1} + \beta_2^j]$. To prove the claim, we show that

$$f(\beta) = \tau_1^i + \tau_2^j + \sigma_1^i(\beta - \beta_1^i - \beta_2^j) \leq f_1(\beta^*) + f_2(\beta - \beta^*)$$

holds for all $\beta \in [\beta_1^i + \beta_2^j, \beta_1^{i+1} + \beta_2^j]$ and $\beta^* \in \mathbb{R}$. Since both $f_1$ and $f_2$ are convex, we know that $f_1(\beta) \geq \tau_1^i + \sigma_1^i(\beta - \beta_1^i)$ and $f_2(\beta) \geq \tau_2^j + \sigma_2^j(\beta - \beta_2^j)$ hold for all $\beta \in \mathbb{R}$. This immediately yields

$$
\begin{aligned}
f_1(\beta^*) + f_2(\beta - \beta^*) &\geq \tau_1^i + \sigma_1^i(\beta^* - \beta_1^i) + \tau_2^j + \sigma_2^j(\beta - \beta^* - \beta_2^j) \\
&\geq \tau_1^i + \sigma_1^i(\beta^* - \beta_1^i) + \tau_2^j + \sigma_1^i(\beta - \beta^* - \beta_2^j) \\
&= \tau_1^i + \tau_2^j + \sigma_1^i(\beta - \beta_1^i - \beta_2^j) \\
&= f(\beta).
\end{aligned}
$$

Since $f(\beta) \geq \min_{\beta^* \in \mathbb{R}} f_1(\beta^*) + f_2(\beta - \beta^*)$ must hold for all $\beta \in \mathbb{R}$ by construction, this completes our proof. $\square$

# 8 Heuristic Approaches

With an $\mathcal{NP}$-hard problem at hand, we propose heuristic approaches based on CHArge, which drop optimality to reduce query times. Their basic idea is as follows. During the third phase of CHArge (running the CFP algorithm on the core graph), whenever the search scans multiple shortcuts $(u, v)$ between two vertices $u \in V$ and $v \in V$, at most one new label is added to $L_{\mathrm{uns}}(v)$. This saves time for dominance checks and label insertion in the label set $L_{\mathrm{uns}}(v)$. We use the potential at $v$ to determine a shortcut that minimizes the key of the new label (i.e., the trip time from the source $s$ to $v$ plus a lower bound on the trip time from $v$ to the target $t$, and add only this label to $L_{\mathrm{uns}}(v)$. Recall that the potential depends on the SoC at $v$, hence scanning different shortcuts may result in different potentials at $v$.

Our first heuristic, denoted CHArge-H$_f$, uses the potential function $\pi_f$ to determine the best shortcut. Given a label $\ell$ at some vertex $u \in V$, scanning an outgoing shortcut $(u, v)$ to a vertex $v \in V$ results in some label $\ell'$ at $v$. We compute its key, which minimizes the sum $\tau + \pi_f(v, f\langle \ell' \rangle(\tau))$ for arbitrary $\tau \in \mathbb{R}_{\geq 0}$ (as described in Section 5). However, the label is only added to $L_{\mathrm{uns}}(v)$ afterwards if this key is minimal among all labels at $v$ constructed within the current vertex scan.

The idea of our second heuristic, denoted CHArge-H$_\omega$, is to use the potential function $\pi_\omega$ instead of $\pi_f$. Additionally, when identifying the only label to be inserted into the set $L_{\mathrm{uns}}(v)$ of a vertex $v \in V$, we ignore battery constraints and presume that we are not close to the target, i.e., that we are in the case $\beta < \mathrm{dist}_c(v, t)$ of Equation 3 for arbitrary SoC $\beta \in [0, M]$. Then the best

shortcut $(u, v)$ does not depend on the SoC at $v$, but only on the distance $\text{dist}_\omega(v, t)$. Hence, we can precompute the optimal shortcut for each pair of neighbors $u \in V$ and $v \in V$, namely, the one that minimizes the cost $\omega(u, v)$ of the shortcut. During a query, instead of scanning all shortcuts, we always use the precomputed shortcut for each neighbor $v$ of $u$.

A third, even more aggressive variant, which we denote by CHArge-H$_\omega^\text{A}$, uses the same idea as in CHArge-H$_\omega$ already during vertex contraction for CH, keeping only the *optimal* shortcut wrt. the cost function $\omega$ for each pair of vertices. Thus, we no longer allow the creation of multi-arcs during preprocessing. This significantly reduces the total number of shortcuts in the core graph, allowing the contraction of further vertices. While the resulting search graph can no longer be used for exact queries, CHArge-H$_\omega^\text{A}$ is capable of answering heuristic queries much faster. The query algorithm of CHArge-H$_\omega^\text{A}$ is identical to CHArge-H$_\omega$, however, solutions may differ as it operates on a sparser graph.

**An Optimality Criterion.** Despite their heuristic nature, it is actually possible to formally grasp under which circumstances the heuristics CHArge-H$_\omega$ and CHArge-H$_\omega^\text{A}$ use an optimal shortcut [65]. Basically, we know that if charging is inevitable and charging at a rate of $\text{cf}_\text{max}$ is possible when needed, then $\text{dist}_\omega(\cdot, \cdot)$ yields a *tight* bound on the remaining trip time. The following Lemma 7 formalizes this insight. Recall that when computing shortcut arcs $(u, v)$ for two vertices $u \in V$ and $v \in V$, the only possible charging stations on the underlying $u$–$v$ path are $u$ and $v$ (see Section 6).

**Lemma 7.** *Given two vertices $u \in V$ and $v \in V$, a $u$–$v$ path $P$ that contains no charging stations (except possibly $u$ and $v$) and minimizes the cost $\omega(P)$ wrt. the function $\omega$ among all $u$–$v$ paths is a subpath of a fastest feasible $s$–$t$ path if the following conditions are met.*

1. *The fastest feasible $s$–$t$ path contains $u$ and $v$ in this order, but no charging station on the subpath from $u$ to $v$ (except $u$ and $v$).*

2. *The SoC at $u$ is not sufficient to reach $t$ without recharging.*

3. *The SoC never reaches the capacity $M$ (such that battery constraints need to be applied) on the path from $u$ to the next charging station that is used.*

4. *There is a charging station available on the subpath from $v$ to $t$ before the battery runs out and the uniform charging speed at this station is $\text{cf}_\text{max}$.*

*Proof.* For the sake of simplicity, we assume that paths are unique wrt. the weight function $\omega$. Assume that all four conditions are met and let $P$ be the $u$–$v$ path that minimizes $\omega(P)$. Furthermore, let $\tau_P$ be the minimal trip time from $s$ to $t$ when using $P$ as the subpath from $u$ to $v$. Assume for contradiction that there exists another path $P'$ from $u$ to $v$ (without a charging station except $u$ and $v$) inducing a total trip time from $s$ to $t$ of $\tau_{P'} < \tau_P$.

*Case 1:* We first consider the case of $d(P') \leq d(P)$, i.e., the driving time of $P'$ is less than the driving time of $P$. Since $P$ minimizes the cost $\omega(P)$, it follows that $c(P') > c(P)$ (where $c(P)$ denotes the sum of the consumption values of all arcs in the path $P$). We replace $P'$ in the fastest $s$–$t$ path with $P$. This increases the total trip time of the $s$–$t$ path by $d(P) - d(P') \geq 0$. It remains feasible, since the energy consumption does not increase. By Condition 2, $t$ is not reachable without recharging after reaching $u$. Hence, there has to be a charging station on the $v$–$t$ path that is used. Due to Condition 3, the battery limit is never reached between $v$ and this charging station. Therefore, when replacing $P'$ with $P$, the SoC at the charging station increases by $c(P') - c(P) > 0$. Consequently, we do not have to recharge this amount of energy, which reduces the charging time

by $(c(P') - c(P))/\mathrm{cf}_{\max}$. Thus, the trip time when replacing $P'$ with $P$ is given by

$$\tau_{P'} + d(P) - d(P') - \frac{c(P') - c(P)}{\mathrm{cf}_{\max}} \geq \tau_P.$$

This value cannot be less than $\tau_P$, since $\tau_P$ is the minimal trip time when using $P$. We use that $\tau_P - \tau_{P'} > 0$ holds per assumption, which leads to

$$d(P) - d(P') - \frac{c(P') - c(P)}{\mathrm{cf}_{\max}} \geq \tau_P - \tau_{P'} > 0$$
$$\Leftrightarrow \ d(P) + \frac{c(P)}{\mathrm{cf}_{\max}} > d(P') + \frac{c(P')}{\mathrm{cf}_{\max}}$$
$$\Leftrightarrow \ \omega(P) > \omega(P'),$$

contradicting the fact that $P$ is the $u$–$v$ path that minimizes the cost $\omega$ when going from $u$ to $v$.

*Case 2:* We consider the case of $d(P') > d(P)$. Hence, we have $c(P') \leq c(P)$. Again, we replace $P'$ with $P$ in the $s$–$t$ path, which reduces the driving time by $d(P') - d(P)$ and the SoC at $v$ by $c(P) - c(P')$. Condition 4 ensures that there is a charging station available on the path from $v$ to $t$, which we use to recharge the missing energy. Condition 3 states that the SoC is always below the battery limit between $u$ and this charging station. Therefore, when reaching the charging station, the amount of energy missing is $c(P) - c(P')$ compared to the $s$–$t$ path using $P'$. We use the charging station to recharge this amount of energy, to ensure that the path to $t$ remains feasible. Altogether, the trip time for reaching $t$ when replacing $P'$ with $P$ is given by

$$\tau_{P'} - \big(d(P') - d(P)\big) + \frac{c(P) - c(P')}{\mathrm{cf}_{\max}} \geq \tau_P.$$

As in the first case, this is equivalent to $\omega(P) > \omega(P')$, which is a contradiction to our assumption that $P$ minimizes the cost $\omega$. This completes the proof. $\qquad\square$

## 9 Experiments

All implementations are in C++ using g++ 4.8.3 (-O3) as compiler. Experiments were conducted on a single core of a 4-core Intel Xeon E5-1630v3 clocked at 3.7 GHz, 128 GiB of DDR4-2133 RAM, 10 MiB of L3 cache, and 256 KiB of L2 cache.

**Instances.** Our main instances are based on the road network of Europe (Eur) and the subnetwork of Germany (Ger), kindly provided by PTV AG (`ptvgroup.com`). As in our previous work [7], we extracted average speeds and categories of road segments, augmented by elevation data from the

Table 1: Instances. We report the number of vertices (including charging stations), arcs, negative consumption arcs (as a fraction of total arcs), as well as charging stations alone (CS) obtained from ChargeMap. Note that the Southern Germany instance based on OSM data has many degree-2 vertices, meant for visualization.

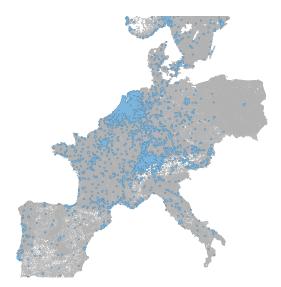| Instance | # Vertices | # Arcs | # Arcs with $c < 0$ | # CS |
|---|---|---|---|---|
| Germany | 4 692 091 | 10 805 429 | 1 119 710 (10.36%) | 1 966 |
| Europe | 22 198 628 | 51 088 095 | 6 060 648 (11.86%) | 13 810 |
| Southern Germany | 5 588 146 | 11 711 088 | 1 142 391 (9.75%) | 643 |

Figure 10: Distribution of charging stations in Europe according to ChargeMap data. We see differences in the distribution of charging stations, which is very dense in the Netherlands and Switzerland. On the other hand, Spain, Italy, and Poland contain relatively few charging stations in our data set. Note that the United Kingdom is not contained in our instance, because we extracted the largest strongly connected component of the input (in which the United Kingdom is only reachable via ferry).

Shuttle Radar Topography Mission (SRTM), v4.1 (`srtm.csi.cgiar.org`). We derived realistic energy consumption from a detailed micro-scale emission model [34], calibrated to a real Peugeot iOn. It has a battery capacity of 16 kWh, but we also evaluate for 85 kWh, as in high-end Tesla models (with a range of 400–500 km). Moreover, we located charging stations on ChargeMap (`chargemap.com`). Figure 10 shows the distribution of charging stations on the Europe instance according to ChargeMap data. We also evaluate the largest instance considered by Storandt [55]. It uses OpenStreetMap (OSM) data of Southern Germany enriched with SRTM, a basic physical consumption model, and 100–1 000 randomly chosen charging stations, but we also test on ChargeMap stations. See Table 1 for an overview of instances and their sizes. We constructed different charging functions to model certain types of stations, namely, battery swapping stations (BSS), superchargers (charging an empty battery to 80 % in 34 min for both vehicle models), as well as regular stations with fast (44 kW), medium (22 kW), or slow (11 kW) charging power. For the three latter types of functions, we used the physical model of Uhrig et al. [62] and approximated the corresponding charging functions with a piecewise linear function (using six breakpoints at 0 %, 80 %, 85 %, 90 %, 95 %, and 100 % SoC). See Figure 11 for a plot showing the resulting charging functions for a battery capacity of 85 kWh. We set initialization time to 3 min for BSS and 1 min for all other types of charging stations. Both the road network and the energy consumption data mentioned above are based on proprietary data sources, but this allowed us to perform experiments on highly detailed and realistic input data.

Table 2: Types of charging stations used in the different scenarios.

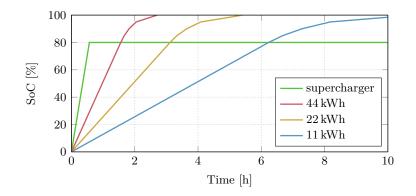| Scenario | 11 kWh | 22 kWh | 44 kWh | Sup. Ch. | BSS |
|---|---|---|---|---|---|
| BSS | — | — | — | — | 100 % |
| Mixed | 30 % | 30 % | 20 % | 10 % | 10 % |
| Realistic | 50 % | 40 % | 10 % | — | — |

Figure 11: Charging functions we use in our experiments for a battery capacity of 85 kWh. Swapping stations (which yield a departure SoC of 100% for any charging time) are not shown in the plot.

If not stated otherwise, queries are always generated by picking source and target vertices uniformly at random and an initial SoC of $\beta_s = M$. In our evaluation, we consider three different distributions of charging stations; see Table 2 for an overview. In the first, all charging stations are BSS. This is the simplest scenario we evaluate, which allows us to compare our algorithm to previous approaches based on simpler models of charging functions [55]. The second uses a mixed composition of chargers, for which we randomly pick 10 % of all stations as BSS, 10 % as superchargers, 20 % as fast chargers, and each 30 % as medium and slow chargers. Containing various types of charging stations, this scenario is much more difficult to solve for our approach. Finally, the third scenario reflects a more realistic distribution of charging stations, where 10 % of all stations are fast, 40 % are medium, and 50 % are slow chargers. This roughly corresponds to the distribution of charging stations in real-world road networks. In each scenario, the composition is fixed, i. e., every charging station retains its type for all queries.

**Evaluating Queries.** We discuss preprocessing and query performance of our algorithms. We only report the fastest exact method (CHArge with the potential $\pi_f$) and our heuristic approaches—for results on plain CFP see below.

Table 3 shows details on preprocessing effort and query performance for different core sizes on Germany (for a battery capacity of 16 kWh). In this experiment, vertex contraction during preprocessing was stopped as soon as the average vertex degree in the core reached the threshold

Table 3: Impact of core size on performance (CHArge, Germany, 85 kWh). We stopped contraction if the average degree in the core graph exceeded a certain threshold (Ø Deg). We report the core size (# Vertices), preprocessing time, and average query times of 1 000 queries answered by CHArge.

| Core size | | Prepr. | Query [ms] | | |
|---|---|---|---|---|---|
| Ø Deg. | # Vertices | [h:m:s] | BSS | Mixed | Realistic |
| 8 | 344 066 (7.33%) | 2:58 | 4 461.9 | 203 576.7 | 179 430.6 |
| 16 | 116 917 (2.49%) | 4:01 | 1 773.4 | 6 739.0 | 4 423.1 |
| 32 | 65 375 (1.39%) | 5:03 | 1 020.1 | 3 037.3 | 2 425.5 |
| 64 | 43 036 (0.91%) | 7:07 | 1 062.2 | 3 184.5 | 2 533.9 |
| 128 | 30 526 (0.65%) | 11:16 | 1 217.2 | 3 728.3 | 2 901.2 |
| 256 | 22 592 (0.48%) | 20:22 | 1 566.0 | 4 987.5 | 4 144.4 |
| 512 | 17 431 (0.37%) | 37:11 | 2 285.8 | 6 936.1 | 5 498.0 |

Table 4: Preprocessing and query performance for different instances, charging station types, and battery capacities. We report regular preprocessing times for CHArge (which also applies to the heuristics CHArge-$H_f$ and CHArge-$H_\omega$) and preprocessing times for CHArge-$H_\omega^A$, the percentage of feasible queries, as well as average exact and heuristic query times of 1 000 queries.

| Ins. | $M$ | CS | Pr. [m:s] | | Feas. | Query [ms] | | | |
|------|-----|-----|------|------|-------|--------|--------|--------|--------|
| | | | reg. | $H_\omega^A$ | | Exact | $H_f$ | $H_\omega$ | $H_\omega^A$ |
| Germany | 16 kWh | BSS | 5:03 | 4:33 | 100 % | 1 607.6 | 1 278.3 | 548.8 | 23.5 |
| | 16 kWh | Mixed | 5:03 | 4:33 | 100 % | 12 191.0 | 4 217.4 | 1 781.9 | 189.2 |
| | 16 kWh | Realistic | 5:03 | 4:33 | 100 % | 6 177.5 | 2 613.0 | 1 119.3 | 103.6 |
| | 85 kWh | BSS | 4:59 | 5:31 | 100 % | 1 020.1 | 988.1 | 174.3 | 31.6 |
| | 85 kWh | Mixed | 4:59 | 5:31 | 100 % | 3 037.3 | 2 346.4 | 555.3 | 51.7 |
| | 85 kWh | Realistic | 4:59 | 5:31 | 100 % | 2 425.5 | 1 405.7 | 358.7 | 45.6 |
| Europe | 16 kWh | BSS | 30:32 | 28:38 | 68 % | 5 958.7 | 3 050.8 | 4 046.7 | 153.0 |
| | 16 kWh | Mixed | 30:32 | 28:38 | 68 % | 35 532.8 | 16 830.7 | 19 338.9 | 3 436.2 |
| | 16 kWh | Realistic | 30:32 | 28:38 | 68 % | 20 310.0 | 11 650.4 | 13 986.7 | 2 146.5 |
| | 85 kWh | BSS | 30:16 | 29:47 | 100 % | 31 548.3 | 22 602.6 | 1 571.3 | 52.8 |
| | 85 kWh | Mixed | 30:16 | 29:47 | 100 % | 100 448.7 | 54 279.8 | 38 520.0 | 718.3 |
| | 85 kWh | Realistic | 30:16 | 29:47 | 100 % | 74 362.2 | 40 828.3 | 13 596.2 | 514.9 |

shown in the first column of the table. We report resulting core graph sizes, preprocessing times, and average query times of CHArge for 1 000 random queries on the BSS, mixed, and realistic station composition, respectively. Apparently, preprocessing effort increases with decreasing core size. We achieve best query performance at an average core degree of 32. At higher degrees, the rather dense core causes query times to increase. Therefore, we use a core degree of 32 as threshold to stop vertex contraction in all further experiments. This results in relative core sizes of 1.3–1.7 % on Germany and Europe. Regarding query times, we observe that the BSS composition is the easiest to solve, because vertex potentials are very accurate with only one available charging station type.

Table 4 shows detailed timings on performance for 1 000 queries on each considered instance. We evaluate three scenarios (BSS only, mixed, and realistic compositions of charging stations) and use our instances Germany and Europe with typical battery capacities (16 kWh and 85 kWh). Preprocessing times are quite practical, considering the problem at hand, ranging from about 5–30 minutes. As before, the mixed and realistic scenarios are harder to solve. As a general observation, increasing the maximum battery capacity leads to faster queries. This can be explained by the fact that less charging is required, so goal direction provided by the potential functions of A* search is more helpful. A notable exception is the capacity of 16 kWh on Europe. In this case, the number of feasible queries drops significantly, due to a highly non-uniform distribution of charging stations (sparse in parts of Southern and Eastern Europe; see Figure 10). This benefits approaches based on the potential function $\pi_f$, as we can often detect infeasibility already during potential computation (the lower bound on trip time evaluates to $\infty$). On the other hand, infeasible queries often deteriorate the performance when using the potential function $\pi_\omega$: Because the target is never reached, large parts of the graph are explored until the queue runs empty.

All in all, running times of the exact algorithm are below 15 seconds on average on Germany and at most 100 seconds on Europe, which is quite notable given that we could not even run a single (long-distance) CFP query on this instance in several hours. Note that the mixed composition is a rather difficult configuration for our algorithms, due to large differences in the charging speeds of the available charging station types (which results in less accurate vertex potentials). When using

Table 5: Detailed query performance of CHArge (Germany, 85 kWh) for mixed and realistic charging stations (CS). For exact CHArge (Ex.) and the different heuristics ($H_f$, $H_\omega$, and $H_\omega^A$), we report average values of the number of settled labels (# Labels), pairwise dominance checks (# Dom.), and running times for 1 000 queries. For the resulting trips, we report the percentage of feasible and optimal paths, average and maximum increase in trip time compared to an optimal solution, as well as average trip time and average total charging time on the resulting routes.

| CS | Algo. | Query | | | Trip Quality | | | | Trip Times | |
|---|---|---|---|---|---|---|---|---|---|---|
| | | # Labels | # Dom. | Time [ms] | Feas. | Opt. | Avg. | Max. | $\tau_{\text{trip}}$ | $\tau_{\text{charge}}$ |
| Mixed | Ex. | 490 421 | 61 183 835 | 3 037.3 | 100 % | 100 % | 1.0000 | 1.0000 | 3:25:45 | 2:01 |
| | $H_f$ | 468 977 | 161 103 | 2 346.4 | 100 % | 84 % | 1.0012 | 1.0519 | 3:26:00 | 2:01 |
| | $H_\omega$ | 313 018 | 7 864 494 | 555.3 | 100 % | 82 % | 1.0004 | 1.0198 | 3:25:50 | 2:01 |
| | $H_\omega^A$ | 16 110 | 45 471 | 51.7 | 100 % | 54 % | 1.0198 | 1.1833 | 3:29:49 | 2:02 |
| Realistic | Ex. | 406 961 | 49 901 425 | 2 425.5 | 100 % | 100 % | 1.0000 | 1.0000 | 3:43:58 | 23:11 |
| | $H_f$ | 345 397 | 112 669 | 1 405.7 | 100 % | 84 % | 1.0006 | 1.0362 | 3:44:05 | 23:11 |
| | $H_\omega$ | 219 620 | 4 061 188 | 358.7 | 100 % | 69 % | 1.0010 | 1.0524 | 3:44:12 | 23:11 |
| | $H_\omega^A$ | 16 837 | 38 691 | 45.6 | 100 % | 59 % | 1.0153 | 1.1575 | 3:47:24 | 23:29 |

the potential function $\pi_\omega$ for (exact) CHArge (not reported in the table), running times increase by up to an order of magnitude, depending on the scenario. Hence, plugging in the more sophisticated potential function $\pi_f$ pays off. For heuristic approaches, we see that—in contrast to CHArge—those based on the potential function $\pi_\omega$ are faster (except for instances with many infeasible queries).

Table 5 reports detailed figures (for the same set of queries as in Table 4) on Germany, using a battery capacity of 85 kWh. We argue that this results in the most sensible queries: Due to a reasonably dense distribution of charging stations in Germany (see Figure 10), all queries are feasible and the target is always reachable with a small number of charging stops. Consequently, we obtain an average trip time of 3 h 25 min 45 sec and an average charging time of 2 min 1 sec on Germany for the mixed composition of charging stations. In contrast to that, queries across Europe evaluated in Table 4 (which we analyze rather to show scalability of our algorithms) yield an average trip time of over 10 h. We also add figures for the realistic scenario, which yields an average trip time of 3 h 43 min 58 sec and an average charging time of 23 min 11 sec. Independent of the scenario and the utilized algorithm, the number of required charging stops is 0.57 on average and no path has more than four charging stops in total. Note that our algorithm does not exhibit any limit on the number of charging stops. In theory, it is therefore easy to construct examples where an optimal route contains a large number of charging stops (even in cases where the target is reachable with only a few or no stops at all). However, this is unlikely to happen in practice, since a charging stop typically involves a detour and we also take initialization time into account. Our experiments confirm this: The number of charging stops is small despite the fact that we assess a rather moderate initialization time of only one minute.

Regarding the performance of our algorithms, we observe that the number of settled labels is a good indicator of the running time. Moreover, all approaches are quite practical. Computing optimal routes takes only a few seconds. The heuristic CHArge-$H_\omega$ provides a good trade-off between running times (300–600 ms) and resulting errors (0.1 % on average, 5 % in the worst case). Our aggressive approach CHArge-$H_\omega^A$ allows query times of about 50 ms on average, which is fast enough even for interactive applications. However, we see that solution quality is up to 18 % off the optimum in the worst case. Still, the average error of all heuristics is very small, and the optimal solution is found in more than half of the cases.
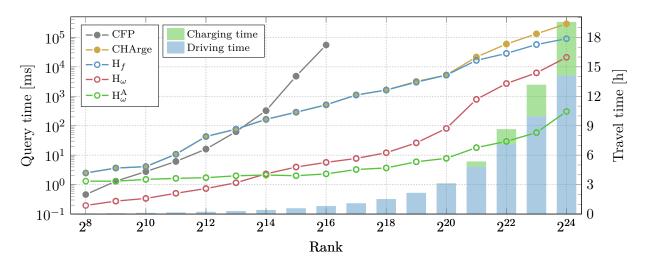
Figure 12: Running times of CHArge subject to Dijkstra rank. Each point in the plot corresponds to the median running time of 100 random queries of the indicated Dijkstra rank on Europe with a 85 kWh battery and the realistic composition of charging stations. Each bar shows the average driving time (blue) and charging time (green) that an optimal solution of the respective rank requires (computed by CHArge).

**Evaluating Scalability.** In our evaluations above, we considered queries with source and target vertices picked uniformly at random. Note that this typically results in long-distance queries (in relation to the graph diameter). To distinguish local and long-distance queries, we consider the *Dijkstra rank* of a query, which equals the number of queue extractions when running Dijkstra's algorithm from the source to the target with unconstrained driving time as arc costs [2, 50]. Thus, higher ranks reflect queries that are more difficult to solve. Moreover, charging en route is only required for queries of high rank.

Figure 12 shows median query running times, as well as average driving and charging times, on our Europe instance with realistic charging stations, subject to their Dijkstra rank. We ran 100 queries per rank. Query times for CFP are only reported up to a rank of $2^{16}$, because for higher ranks at least one query did not terminate within a predefined limit of one hour. The colored bars show that higher Dijkstra ranks correspond to routes with longer travel time. The green part of each bar shows the portion of the overall travel time that is spent at charging stations. Given a relatively large battery capacity of 85 kWh, charging stops are only necessary for queries of the highest Dijkstra ranks, starting from rank $2^{21}$. Consequently, running times of our faster approaches increase significantly at this rank. It also becomes evident that all approaches have exponential running times, with timings of CHArge in the order of minutes for the highest ranks. Nevertheless, performance remains practical for ranks beyond $2^{20}$. Running times of our fastest heuristic (CHArge-$H_\omega^A$) are well below a second for the highest ranks.

**Comparison with Related Work.** We also consider the largest instance that was used to evaluate the fastest approach for routes with charging stops in the literature [55]. It utilizes a simple physical consumption model, where energy consumption of an arc is a linear combination of its horizontal distance and slope. All charging stations in the original work are BSS, picked uniformly at random from all vertices of the input graph. Table 6 shows detailed figures, using 100–1 000 randomly placed charging stations (r100, r1000) and battery capacities that translate to a certain range in the physical model. Query times are average values of 1 000 random queries. For completeness, we also run tests on ChargeMap stations (only BSS), 643 of which are included in the considered road network. As

Table 6: Comparison of CHArge with existing work [55]. For different distributions of charging stations and ranges, we report preprocessing times of the existing approach [55], CHArge, and the heuristic CHArge-$H_\omega^A$. Regarding queries, we show the percentage of feasible paths, as well as average exact and heuristic query times of 1 000 queries. Results from the literature [55] are reported as is from a Core i3-2310M, 2.1 GHz.

| | | Prepr. [m:s] | | | | Query [ms] | | | | |
|---|---|---|---|---|---|---|---|---|---|---|
| CS | $M$ | [55] | CHArge | $H_\omega^A$ | Feas. | [55] | CHArge | $H_f$ | $H_\omega$ | $H_\omega^A$ |
| r1000 | 100 km | 51:41 | 11:37 | 2:30 | 100 % | 539.0 | 122.1 | 93.7 | 38.1 | 2.7 |
| r100 | 150 km | 16:21 | 11:10 | 2:15 | 99 % | 1 150.0 | 206.0 | 116.5 | 26.9 | 1.4 |
| c643 | 100 km | — | 11:21 | 2:32 | 98 % | — | 326.3 | 282.6 | 48.3 | 3.3 |
| c643 | 150 km | — | 11:28 | 2:29 | 99 % | — | 308.3 | 270.0 | 22.7 | 2.9 |

before, query times are average values of 1 000 random queries. We observe that our approach is faster wrt. both preprocessing and queries, even when taking differences in hardware into account. At the same time, CHArge is more general and not inherently restricted to BSS. Furthermore, note that a denser distribution of random charging stations enables faster query times: Although preprocessing effort increases slightly (more charging stations are kept in the core), the potential functions more often rightly assume that a station will be available close to the remaining path. We see that using charging station locations from ChargeMap (c643) gives a slightly harder instance. In conclusion, we are able to solve instances that are harder than those evaluated in the literature.

## 10 Conclusion

We introduced CFP, a novel approach for EV route planning that computes shortest feasible paths with charging stops, minimizing overall trip time. Combining vertex contraction and charging-aware goal-directed search, we developed CHArge, which solves this $\mathcal{NP}$-hard problem optimally and with practical performance even on large realistic inputs. For reasonable distances, our algorithm computes solutions within seconds, making CHArge the first practical exact approach—with running times similar to previous methods that are inexact or use less accurate models [28, 55]. Finally, we proposed heuristic variants that enable even faster queries, while offering high-quality solutions.

Future work includes allowing multiple driving speeds (and consumption values) per road segment, e. g., in order to save energy on the highway by driving at reasonable lower speeds [9, 33]. Moreover, we are interested in more detailed models for energy consumption on turns. In preliminary experiments, we modified our input instances according to a known arc-based approach for integrating turns [25]. Using a simple model that takes consumption overhead for acceleration and deceleration along turns (or when speed limits change) into account, we observed that preprocessing took slightly longer but also resulted in smaller cores, while query times remained nearly the same. This can be explained by the fact that graph size increases when it is enriched with turn costs, but at the same time the number of nondominated solutions decreases (minor detours no longer allow energy savings). This clearly indicates that our techniques can be readily extended to handle turn costs and turn restrictions.

For integration of historic and live traffic, the adaptation of customizable techniques appears to be a natural extension of our approaches. While preliminary experiments showed that our techniques based on CH perform equally fast on vertex orders required by Customizable CH [17], a major challenge is the design of fast customization algorithms that can handle the dynamic data structures required by label sets in our approach. On the other hand, preprocessing time for CHArge is already

in the order of minutes and can even be reduced by increasing core size at the cost of higher query times; see Table 3 in Section 9. This is fast enough for many applications that require frequent updates of the underlying graph.

Furthermore, we are interested in extending our algorithms to more complex settings, e. g., by taking into account that charging stations might currently be in use by other customers. In such a scenario, some charging stations could also be reserved in advance to keep waiting times low [48]. For fast heuristic variants, it might also be useful to precompute potentials for A* search, as in ALT [27]. From a more theoretical perspective, approximability of both problem settings considered in this chapter is an open question. Results by Strehler et al. [58] include an FPTAS for a very similar problem concerning routes with charging stops, which might extend to our setting.

# References

[1] Ittai Abraham, Daniel Delling, Amos Fiat, Andrew V. Goldberg, and Renato F. Werneck. HLDB: Location-Based Services in Databases. In *Proceedings of the 20th ACM SIGSPATIAL International Symposium on Advances in Geographic Information Systems (GIS'12)*, pages 339–348. ACM, 2012.

[2] Hannah Bast, Daniel Delling, Andrew V. Goldberg, Matthias Müller-Hannemann, Thomas Pajor, Peter Sanders, Dorothea Wagner, and Renato F. Werneck. Route Planning in Transportation Networks. In *Algorithm Engineering: Selected Results and Surveys*, volume 9220 of *Lecture Notes in Computer Science*, pages 19–80. Springer, 2016.

[3] Gernot V. Batz, Robert Geisberger, Peter Sanders, and Christian Vetter. Minimum Time-Dependent Travel Times with Contraction Hierarchies. *ACM Journal of Experimental Algorithmics*, 18:1.4:1–1.4:43, 2013.

[4] Reinhard Bauer, Daniel Delling, Peter Sanders, Dennis Schieferdecker, Dominik Schultes, and Dorothea Wagner. Combining Hierarchical and Goal-Directed Speed-up Techniques for Dijkstra's Algorithm. *ACM Journal of Experimental Algorithmics*, 15:2.3:1–2.3:31, 2010.

[5] Moritz Baum. *Engineering Route Planning Algorithms for Battery Electric Vehicles*. Phd thesis, Karlsruhe Institute of Technology, 2018.

[6] Moritz Baum, Julian Dibbelt, Andreas Gemsa, Dorothea Wagner, and Tobias Zündorf. Shortest Feasible Paths with Charging Stops for Battery Electric Vehicles. In *Proceedings of the 23rd ACM SIGSPATIAL International Conference on Advances in Geographic Information Systems (GIS'15)*, pages 44:1–44:10. ACM, 2015.

[7] Moritz Baum, Julian Dibbelt, Thomas Pajor, and Dorothea Wagner. Energy-Optimal Routes for Electric Vehicles. In *Proceedings of the 21st ACM SIGSPATIAL International Conference on Advances in Geographic Information Systems (GIS'13)*, pages 54–63. ACM, 2013.

[8] Moritz Baum, Julian Dibbelt, Thomas Pajor, and Dorothea Wagner. Dynamic Time-Dependent Route Planning in Road Networks with User Preferences. In *Proceedings of the 15th International Symposium on Experimental Algorithms (SEA'16)*, volume 9685 of *Lecture Notes in Computer Science*, pages 33–49. Springer, 2016.

[9] Moritz Baum, Julian Dibbelt, Dorothea Wagner, and Tobias Zündorf. Modeling and Engineering Constrained Shortest Path Algorithms for Battery Electric Vehicles. In *Proceedings of the 25th Annual European Symposium on Algorithms (ESA'17)*, volume 87 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 11:1–11:16. Schloss Dagstuhl–Leibniz-Zentrum für Informatik, 2017.

[10] Moritz Baum, Jonas Sauer, Dorothea Wagner, and Tobias Zündorf. Consumption Profiles in Route Planning for Electric Vehicles: Theory and Applications. In *Proceedings of the 16th International Symposium on Experimental Algorithms (SEA'17)*, volume 75 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 19:1–19:18. Schloss Dagstuhl–Leibniz-Zentrum für Informatik, 2017.

[11] Daniel Delling, Andrew V. Goldberg, Thomas Pajor, and Renato F. Werneck. Customizable Route Planning in Road Networks. *Transportation Science*, 51(2):566–591, 2017.

[12] Daniel Delling, Andrew V. Goldberg, and Renato F. Werneck. Faster Batched Shortest Paths in Road Networks. In *Proceedings of the 11th Workshop on Algorithmic Approaches for Transportation Modeling, Optimization, and Systems (ATMOS'11)*, volume 20 of *OpenAccess Series in Informatics (OASIcs)*, pages 52–63. Schloss Dagstuhl–Leibniz-Zentrum für Informatik, 2011.

[13] Daniel Delling and Dorothea Wagner. Time-Dependent Route Planning. In *Robust and Online Large-Scale Optimization*, volume 5868 of *Lecture Notes in Computer Science*, pages 207–230. Springer, 2009.

[14] Daniel Delling and Renato F. Werneck. Customizable Point-of-Interest Queries in Road Networks. *IEEE Transactions on Knowledge and Data Engineering*, 27(3):686–698, 2015.

[15] Guy Desaulniers, Fausto Errico, Stefan Irnich, and Michael Schneider. Exact Algorithms for Electric Vehicle-Routing Problems with Time Windows. *Operations Research*, 64(6):1388–1405, 2016.

[16] Julian Dibbelt, Thomas Pajor, and Dorothea Wagner. User-Constrained Multi-Modal Route Planning. *ACM Journal of Experimental Algorithmics*, 19:3.2:1–3.2:19, 2015.

[17] Julian Dibbelt, Ben Strasser, and Dorothea Wagner. Customizable Contraction Hierarchies. *ACM Journal of Experimental Algorithmics*, 21:1.5:1–1.5:49, 2016.

[18] Edsger W. Dijkstra. A Note on Two Problems in Connexion with Graphs. *Numerische Mathematik*, 1(1):269–271, 1959.

[19] Alexandros Efentakis and Dieter Pfoser. Optimizing Landmark-Based Routing and Preprocessing. In *Proceedings of the 6th ACM SIGSPATIAL International Workshop on Computational Transportation Science (IWCTS'13)*, pages 25–30. ACM, 2013.

[20] Jochen Eisner, Stefan Funke, and Sabine Storandt. Optimal Route Planning for Electric Vehicles in Large Networks. In *Proceedings of the 25th AAAI Conference on Artificial Intelligence (AAAI'11)*, pages 1108–1113. AAAI Press, 2011.

[21] Stefan Funke and Sabine Storandt. Polynomial-Time Construction of Contraction Hierarchies for Multi-Criteria Objectives. In *Proceedings of the 15th Meeting on Algorithm Engineering & Experiments (ALENEX'13)*, pages 31–54. SIAM, 2013.

[22] Michael R. Garey and David S. Johnson. *Computers and Intractability: A Guide to the Theory of NP-Completeness.* W. H. Freeman & Co., 1979.

[23] Robert Geisberger, Moritz Kobitzsch, and Peter Sanders. Route Planning with Flexible Objective Functions. In *Proceedings of the 12th Workshop on Algorithm Engineering & Experiments (ALENEX'10)*, pages 124–137. SIAM, 2010.

[24] Robert Geisberger, Peter Sanders, Dominik Schultes, and Christian Vetter. Exact Routing in Large Road Networks Using Contraction Hierarchies. *Transportation Science*, 46(3):388–404, 2012.

[25] Robert Geisberger and Christian Vetter. Efficient Routing in Road Networks with Turn Costs. In *Proceedings of the 10th International Symposium on Experimental Algorithms (SEA'11)*, volume 6630 of *Lecture Notes in Computer Science*, pages 100–111. Springer, 2011.

[26] Dominik Goeke and Michael Schneider. Routing a Mixed Fleet of Electric and Conventional Vehicle. *European Journal of Operational Research*, 245(1):81–99, 2015.

[27] Andrew V. Goldberg and Chris Harrelson. Computing the Shortest Path: A* Search Meets Graph Theory. In *Proceedings of the 16th Annual ACM–SIAM Symposium on Discrete Algorithms (SODA'05)*, pages 156–165. SIAM, 2005.

[28] Michael T. Goodrich and Paweł Pszona. Two-Phase Bicriterion Search for Finding Fast and Efficient Electric Vehicle Routes. In *Proceedings of the 22nd ACM SIGSPATIAL International Conference on Advances in Geographic Information Systems (GIS'14)*, pages 193–202. ACM, 2014.

[29] Ronald L. Graham. An Efficient Algorithm for Determining the Convex Hull of a Finite Planar Set. *Information Processing Letters*, 1(4):132–133, 1972.

[30] Gabriel Y. Handler and Israel Zang. A Dual Algorithm for the Constrained Shortest Path Problem. *Networks*, 10(4):293–309, 1980.

[31] Pierre Hansen. Bicriterion Path Problems. In *Multiple Criteria Decision Making – Theory and Application*, volume 177 of *Lecture Notes in Economics and Mathematical Systems*, pages 109–127. Springer, 1980.

[32] Peter E. Hart, Nils J. Nilsson, and Bertram Raphael. A Formal Basis for the Heuristic Determination of Minimum Cost Paths. *IEEE Transactions on Systems Science and Cybernetics*, 4(2):100–107, 1968.

[33] Frederik Hartmann and Stefan Funke. Energy-Efficient Routing: Taking Speed into Account. In *Proceedings of the 37th Annual German Conference on Advances in Artificial Intelligence (KI'14)*, volume 8736 of *Lecture Notes in Computer Science*, pages 86–97. Springer, 2014.

[34] Stefan Hausberger, Martin Rexeis, Michael Zallinger, and Raphael Luz. Emission Factors from the Model PHEM for the HBEFA Version 3. Technical report I-20/2009, University of Technology, Graz, 2009.

[35] Gerhard Huber and Klaus Bogenberger. Long-Trip Optimization of Charging Strategies for Battery Electric Vehicles. *Transportation Research Record: Journal of the Transportation Research Board*, 2497:45–53, 2015.

[36] Donald B. Johnson. Efficient Algorithms for Shortest Paths in Sparse Networks. *Journal of the ACM*, 24(1):1–13, 1977.

[37] Samir Khuller, Azarakhsh Malekian, and Julián Mestre. To Fill or Not to Fill: The Gas Station Problem. *ACM Transactions on Algorithms*, 7(3):36:1–36:16, 2011.

[38] Yuichi Kobayashi, Noboru Kiyama, Hirokazu Aoshima, and Masamori Kashiyama. A Route Search Method for Electric Vehicles in Consideration of Range and Locations of Charging Stations. In *Proceedings of the 7th IEEE Intelligent Vehicles Symposium (IV'11)*, pages 920–925. IEEE, 2011.

[39] Chung-Shou Liao, Shang-Hung Lu, and Zuo-Jun Max Shen. The Electric Vehicle Touring Problem. *Transportation Research Part B: Methodological*, 86:163–180, 2016.

[40] Shieu Hong Lin, Nate Gertsch, and Jennifer R. Russell. A Linear-Time Algorithm for Finding Optimal Vehicle Refueling Policies. *Operations Research Letters*, 35(3):290–296, 2007.

[41] Chensheng Liu, Jing Wu, and Chengnian Long. Joint Charging and Routing Optimization for Electric Vehicle Navigation Systems. In *Proceedings of the 19th International Federation of Automatic Control World Congress (IFAC'14)*, volume 47 of *IFAC Proceedings Volumes*, pages 9611–9616. Elsevier, 2014.

[42] Lawrence Mandow and José-Luis Pérez-de-la-Cruz. Multiobjective A* Search with Consistent Heuristics. *Journal of the ACM*, 57(5):27:1–27:24, 2010.

[43] Ernesto Q. V. Martins. On a Multicriteria Shortest Path Problem. *European Journal of Operational Research*, 16(2):236–245, 1984.

[44] Alejandro Montoya, Christelle Guéret, Jorge E. Mendoza, and Juan G. Villegas. The Electric Vehicle Routing Problem with Nonlinear Charging Function. *Transportation Research Part B: Methodological*, 103:87–110, 2017.

[45] Keisuke Murakami. Formulation and Algorithms for Route Planning Problem of Plug-In Hybrid Electric Vehicles. *Operational Research*, 18(2):497–519, 2018.

[46] Samuel Pelletier, Ola Jabali, Gilbert Laporte, and Marco Veneroni. Battery Degradation and Behaviour for Electric Vehicles: Review and Numerical Analyses of Several Models. *Transportation Research Part B: Methodological*, 103:158–187, 2017.

[47] Sepideh Pourazarm, Christos G. Cassandras, and Tao Wang. Optimal Routing and Charging of Energy-Limited Vehicles in Traffic Networks. *International Journal of Robust and Nonlinear Control*, 26(6):1325–1350, 2016.

[48] Hua Qin and Wensheng Zhang. Charging Scheduling with Minimal Waiting in a Network of Electric Vehicles and Charging Stations. In *Proceedings of the 8th ACM International Workshop on Vehicular Inter-Networking (VANET '11)*, pages 51–60. ACM, 2011.

[49] Martin Sachenbacher, Martin Leucker, Andreas Artmeier, and Julian Haselmayr. Efficient Energy-Optimal Routing for Electric Vehicles. In *Proceedings of the 25th AAAI Conference on Artificial Intelligence (AAAI'11)*, pages 1402–1407. AAAI Press, 2011.

[50] Peter Sanders and Dominik Schultes. Highway Hierarchies Hasten Exact Shortest Path Queries. In *Proceedings of the 13th Annual European Conference on Algorithms (ESA'05)*, volume 3669 of *Lecture Notes in Computer Science*, pages 568–579. Springer, 2005.

[51] Maximilian Schiffer and Grit Walther. An Adaptive Large Neighborhood Search for the Location-Routing Problem with Intra-Route Facilities. *Transportation Science*, 52(2):331–352, 2018.

[52] Michael Schneider, Andreas Stenger, and Dominik Goeke. The Electric Vehicle-Routing Problem with Time Windows and Recharging Stations. *Transportation Science*, 48(4):500–520, 2014.

[53] René Schönfelder and Martin Leucker. Abstract Routing Models and Abstractions in the Context of Vehicle Routing. In *Proceedings of the 24th International Joint Conference on Artificial Intelligence (IJCAI'15)*, pages 2639–2645. AAAI Press, 2015.

[54] René Schönfelder, Martin Leucker, and Sebastian Walther. Efficient Profile Routing for Electric Vehicles. In *Proceedings of the 1st International Conference on Internet of Vehicles (IOV'14)*, volume 8662 of *Lecture Notes in Computer Science*, pages 21–30. Springer, 2014.

[55] Sabine Storandt. Quick and Energy-Efficient Routes: Computing Constrained Shortest Paths for Electric Vehicles. In *Proceedings of the 5th ACM SIGSPATIAL International Workshop on Computational Transportation Science (IWCTS'12)*, pages 20–25. ACM, 2012.

[56] Sabine Storandt. Route Planning for Bicycles - Exact Constrained Shortest Paths Made Practical via Contraction Hierarchy. In *Proceedings of the 22nd International Conference on Automated Planning and Scheduling (ICAPS'12)*, pages 234–242. AAAI Press, 2012.

[57] Sabine Storandt and Stefan Funke. Cruising with a Battery-Powered Vehicle and Not Getting Stranded. In *Proceedings of the 26th AAAI Conference on Artificial Intelligence (AAAI'12)*, pages 1628–1634. AAAI Press, 2012.

[58] Martin Strehler, Sören Merting, and Christian Schwan. Energy-efficient shortest routes for electric and hybrid vehicles. *Transportation Research Part B: Methodological*, 103:111–135, 2017.

[59] Timothy M. Sweda, Irina S. Dolinskaya, and Diego Klabjan. Adaptive Routing and Recharging Policies for Electric Vehicles. *Transportation Science*, 51(4):1326–1348, 2017.

[60] Timothy M. Sweda and Diego Klabjan. Finding Minimum-Cost Paths for Electric Vehicles. In *Proceedings of the 1st International Electric Vehicle Conference (IEVC'12)*. IEEE, 2012.

[61] Chi Tung Tung and Kim Lin Chew. A Multicriteria Pareto-Optimal Path Algorithm. *European Journal of Operational Research*, 62(2):203–209, 1992.

[62] Martin Uhrig, Lennart Weiß, Michael Suriyah, and Thomas Leibfried. E-Mobility in Car Parks – Guidelines for Charging Infrastructure Expansion Planning and Operation Based on Stochastic Simulations. In *Proceedings of the 8th International Electric Vehicle Symposium and Exhibition (EVS28)*, 2015.

[63] Yan Wang, Jianmin Jiang, and Tingting Mu. Context-Aware and Energy-Driven Route Optimization for Fully Electric Vehicles via Crowdsourcing. *IEEE Transactions on Intelligent Transportation Systems*, 14(3):1331–1345, 2013.

[64] Hongming Yang, Songping Yang, Yan Xu, Erbao Cao, Mingyong Lai, and Zhaoyang Dong. Electric Vehicle Route Optimization Considering Time-of-Use Electricity Price by Learnable Partheno-Genetic Algorithm. *IEEE Transactions on Smart Grid*, 6(2):657–666, 2015.

[65] Tobias Zündorf. Electric Vehicle Routing with Realistic Recharging Models. Master's thesis, Karlsruhe Institute of Technology, 2014.