Battery and Energy Management in Fleets of Switchable Battery EVs

Vladimir Zdornov and Yitzhak Birk

***Abstract*—This paper addresses the challenge of managing battery switching and charging in fleets of switchable battery electric vehicles (SBEVs). The goal of efficient management is to optimize resource utilization by the fleet, and thus the operational costs, under restricted power supply during operation hours. The resources include spare batteries, battery switching mechanisms, sophisticated infrastructure, as well as the availaibility of the charging power from the grid. We analyze performance limiting factors and formulate heuristic algorithms to tackle them. Furthermore, we evaluate the algorithms in simulations based on a synthetic travel schedule and energy demand model. The collected results expose interesting trade-offs between different resources that should be taken into account when designing the fleet's depot.**

***Index Terms*—Electric Vehicles, Switchable battery EVs, EV fleets, battery management.**

Fig. 1: Battery switching mechanism of "Better Place" company

I. Introduction

THE aspiration towards using electric vehicles as a greener and more efficient alternative to fossil fuel transportation has resulted in the promotion of plug-in hybrid (PHEV) and pure battery electric vehicle (BEV) technologies. The advantages of BEVs are the absence of a combustion engine and the complete withdrawal from using fossil fuels. However, a typical BEV has a significantly larger, and thus significantly more expensive, battery. Also, unlike a PHEV, it presently cannot be rapidly refueled when its battery becomes empty. In order to address this limitation, the switchable battery electric vehicle (SBEV) concept was adopted by several companies [1]. An SBEV can have its empty battery replaced with a full one in a matter of minutes at a special battery replacement station. An example of a switching mechanism used in such a station is shown in Fig. 1.

Consider a company operating a fleet of SBEVs in an urban environment from a centralized depot. Such a company can provide delivery or logistic services or be a public transportation company. Every day, a certain set of service trips, dubbed *trips*, has to be performed by vehicles in the depot. Each trip begins and ends in the central depot. The energy required for the set of trips assigned to a given SBEV on a given day may exceed the capacity of its battery. This could be due to battery size limitations or to the choice of battery for cost reduction. As a result, battery switching must take place so that long idle periods of the vehicle are avoided. In order for the fleet not to generate excessive load on the grid, or to reduce its energy cost by consuming at off-peak hours, the system must obey load shaping by scheduling massive power consumption at off-peak periods [2]. This means that the bulk of the energy exploited by the daily operation has to be drawn and stored in batteries at night before operation hours.

V. Zdornov and Y. Birk are with the Department of Electrical Engineering, Technion Israel Institute of Technology, Haifa, 32000 Israel. E-mail: zdv@tx.technion.ac.il, birk@ee.technion.ac.il

Our goal is to minimize fleet's operational costs while delivering the service. The following factors have to be considered when attempting to reach this goal:

- The high battery price causes the capital expense of purchasing the batteries to constitute a significant portion of the company's operational costs. Consequently, it is important to reduce the overall battery capacity required by the depot.
- A vehicle's battery can only be replaced by a special switching mechanism. These mechanisms are very expensive, so we wish to require as few of them as possible.
- Even if some power is available from the grid during the operation hours, its price is expected to be extremely high and therefore it should be used with caution.
- Additional depot complexity, such as allowing energy to be transferred between batteries, has its cost too.

Having these in mind, we formulate several heuristic policies aimed at minimizing resource usage under varying assumptions on the depot abilities. Simulations based on synthetically generated workloads are used to assess effects of the policies. The collected results and the exposed trade-offs provide insights that can be used in depot designs.

The remainder of the paper is organized as follows. In Section II we provide a detailed description of the problem and present the our synthetic workload generation approach. In Section III, we concentrate on managing battery replacement. The benefits of applying inter-battery energy transfer and exploiting limited charging are discussed in Section IV. Finally, Section V offers concluding remarks.

II. Problem Statement and Workload Generation

A. Problem Statement

For each vehicle $v_i \in V$ operating from the depot, we define $R^{v_i} = \{r_j^{v_i}\}$ to be the set of trips assigned to v_i. Trip j of vehicle v_i is denoted $r_j^{v_i} = (d_j^{v_i}, a_j^{v_i}, e_j^{v_i})$, where $d_j^{v_i}$, $a_j^{v_i}$ and $e_j^{v_i}$ denote the departure time (from the depot), the arrival time (back to the depot) and the required energy respectively. In the general case, the latter two members of the tuple are prone to some amount of uncertainty.

For a given time t, we define $R^{v_i}(t) = \{r_j^{v_i} | t \leq d_j^{v_i}\}$ to be v_i's trips starting after t. We define two more sets to contain trips of all vehicles: $R = \cup R^{v_i}$ and $R(t) = \cup R^{v_i}(t)$. Finally, we denote by e_{th} the maximum $e_j^{v_i}$ for $r_j^{v_i} \in R(t)$, i.e. the energy of the trip requiring the highest amount of energy among trips starting after t. In a similar manner, e_{tl} is such a trip that requires the lowest amount of energy.

In addition to vehicles, the depot holds a set of batteries B. We denote the energy level of a battery b at time t as $e(b,t)$. We assume all batteries to have the same capacity E and to be fully charged at the beginning of the day. Therefore, $0 \leq e(b,t) \leq E$ and $e(b,0) = E$. At any time b can be in one of the following three states: a) at the depot, not in a vehicle, b) at the depot, in a vehicle, or c) in a vehicle outside the depot. We use $b^{v_i}(t) = b$ to denote that b is in v_i at t. If for some reason v_i has no battery at t, then $b^{v_i}(t) = \phi$.

Let $B^d(t) = \{b_j^d(t)\}$ be the set of batteries present at the depot at t both inside and outside of vehicles. Next, let $B_{th} = \{b \in B^d(t) | 0 < e(b,t) \leq e_{th}\} = \{b_1, ..., b_n\}$ be the sorted set of non-empty batteries having at most e_{th} energy, such that $e(b_j^d(t), t) \leq e(b_{j+1}^d(t), t)$.

Each vehicle v_i starts the day with a full battery and a set of trips. In order for the trip $r_j^{v_i}$ to be *successfully completed*, v_i has to leave the depot at $d_j^{v_i}$ with $e_j^{v_i} \leq e(b^{v_i}(d_j^{v_i}), d_j^{v_i})$. In other words, v_i must leave at the scheduled time and must have a sufficiently full battery at departure as otherwise it won't be able to complete the trip. We say that we have a *feasible run* if during the day of operation all vehicles successfully complete all their trips. Achieving a feasible run is a fundamental requirement from any solution provided by the management policies.

The initial energy of $b^{v_i}(0)$ alone is generally insufficient for successful completion of all v_i's trips. Therefore, v_i has to switch battery several times during the day. Depending on the scenario, v_i is either forced to always have a non-empty battery or is allowed to move without a battery inside the depot. In any case, all switching-related activities are performed at one of the *switching mechanisms* $m_i \in M$. Each switching mechanism has a separate battery pool and its content at t is denoted by $B^{m_i}(t) \subseteq B^d(t)$.

Unless allowed to be charged, a battery's energy is monotonically non-increasing throughout the day. A battery b_1 can be charged only while at the depot, i.e. $b_1 \in B^d$. The charging energy may come from two distinct sources: a) another battery $b_2 \in B^d$, or b) grid connection. In both cases, we assume the depot's local power delivery network not to introduce any constraints. The maximum battery charging and discharging rates are assumed to be equal and are denoted by c. In all scenarios, we only allow batteries to be charged/discharged at this maximum rate (or otherwise their energy remains unchanged). The available grid power p is assumed to be limited and is an integer multiple of c.

B. Synthetic Workload Generation

Let us now describe the method for generating synthetic trip sets assigned to vehicles that we will later use for different benchmarks. Before we start, note that throughout the paper we use distance units (kms) to express both distance and energy. The latter should be understood as the equivalent amount of energy required for a vehicle to travel a given distance. Of course, the energy required by a trip is not uniquely defined by its length. However, we find this approximation to be convenient while sufficiently precise for our needs. Similarly we express battery capacity in kilometers. Note that the use of departure and arrival times as well as energy (distance) for trip specification permits the accommodation of diverse traffic situations, and there is no restriction to a fixed speed.

For all simulations, the trips were assumed to take place during 18 hours of operation between 06:00 and 24:00. The size of the fleet was set to be 160 vehicles. A vehicle was randomly chosen to begin its day at 06:00, 06:30 or 07:00. From then on, an intermittent series of randomly generated trips and breaks was constructed. The process stopped when the last trip was scheduled to complete after 24:00. The last trip was then discarded.

When generating a trip $r_j^{v_i}$, the departure time $d_j^{v_i}$ is given by the end time of the latest break (or the chosen start time). We need to choose its duration to derive the arrival time $a_j^{v_i}$, and to set its energy $e_j^{v_i}$. We fix both by drawing the trip's *distance* from a distribution that is uniform in (20,35) and (55,70) kilometers intervals. The average travel speed for all vehicles is assumed to be 15 km/h. The chosen distance equals $e_j^{v_i}$ while the arrival time is calculated as: $a_j^{v_i} = d_j^{v_i} + e_j^{v_i}/15$. The following departure time $d_{j+1}^{v_i}$ is then calculated by adding a uniform (15,40) minute break time to the arrival time.

Let us make several remarks. First, the relatively low average speed is attributed to operating in an urban environment and anticipated stops during the trip (at bus stops or for unloading the cargo). Second, while most of the techniques presented later operate transparently when noise is introduced to $a_j^{v_i}$ and $e_j^{v_i}$, in the presented results we neglect this aspect for the sake of simplicity. Finally, under the above model the fleet travels on average (over different input generations) $\sim$33,000 kilometers in $\sim$750 trips.

III. Battery Switching

In this section we present some of our policies. The discussion here is restricted to situations wherein batteries cannot be charged during the day. This means that all optimization policies presented here deal only with intelligent battery switching. We first consider a simplified version of the problem, and then introduce more realistic assumptions.

A. Simplified Problem

Here, battery switching is instantaneous. Accordingly, any battery with sufficient energy can be used for $r_j^{v_i}$ trip (but the battery must be at the depot, i.e. in $B^d(d_j^{v_i})$). If we consider all trip parameters (departure and arrival times, and required energy) to be precisely known at the beginning of the day, we can formulate the problem as an instance of *bin packing with conflicts* [3].

The basic bin packing deals with packing items of given weights into bins of equal weight-capacity. The goal is to use as few bins as possible while packing all items without violating capacity constraints. In bin packing with conflicts, some items are defined to have a conflict (in pairs) and thus are forbidden to be packed into the same bin. In our case, bins are the batteries in B. As all are initially fully charged and the assignment is made at the beginning of the day, the bin capacities are all E. The items are the trips in R with weights being their energy requirements $e_j^{v_i}$. A pair of trips $r_{j1}^{v_{i1}}$ and $r_{j2}^{v_{i2}}$ have a conflict if the intervals $(d_{j1}^{v_{i1}}, a_{j1}^{v_{i1}})$ and $(d_{j2}^{v_{i2}}, a_{j2}^{v_{i2}})$ intersect. Intuitively, this means that trips overlapping in time cannot use the same battery.

Both the bin-packing and bin-packing with conflicts are known to be NP-complete problems and therefore lack an efficient optimal algorithm. The commonly used heuristic for finding an approximate solution for the basic bin-packing problem is *first-fit decreasing* (FFD) heuristic:

1) Start with zero used bins
2) Go over the items in the descending weight order
3) Find an earliest added used bin in which current item can be feasibly packed (without violating bin's capacity) if no such bin can be found add a new bin and pack the item in it.

When the number of bins in the optimal solution is high (as in our case) FFD achieves a solution that has at most $\sim \frac{11}{9}$-times more bins than the optimal solution [3].

We use a slightly modified version of FFD in an attempt to find good solutions for instances of bin packing with conflicts that are of interest to us. In the modified FFD, an item can be feasibly packed into a bin only if the bin has sufficient capacity and holds no item conflicting with the current one.

In [3], an alternative heuristic is proposed. It has a guaranteed approximation ratio of 2.5-2.7 (depending on certain assumptions on the input) between the number of bins its solutions require and the optimum number of bins. This heuristic involves constructing a conflict graph in which nodes are the items of the original problem and arcs are added between a pair of nodes that represent conflicting items. In our case, nodes are the trips with arcs being added between trips that occur simultaneously. After the graph is constructed, it is colored in a way that avoids coloring neighboring nodes with the same color. After colors have been assigned to all items, FFD is applied to each color group independently and the final result is achieved through the union of results for individual color groups.

We propose yet another approach, ubbed *bi-modal heuristic* (BMH). BMH is designed to avoid two fundamental issues that may increase the number of the required batteries. The first issue arises if while choosing a battery to be used in a trip $r_j^{v_i}$, the aggregate energy of batteries in B^d is greater than $e_j^{v_i}$, but each individual battery holds less than $e_j^{v_i}$. The second issue arises when the system reaches a state in which energy is over-concentrated in a small number of batteries. Once those batteries are away from the depot, a trip cannot be succesfully completed. In this case, a run could have a feasible completion if the energy were more evenly distributed among batteries. Interestingly, analogies can be found between the above phenomena and *fragmentation* witnessed in storage and memory in computer systems.

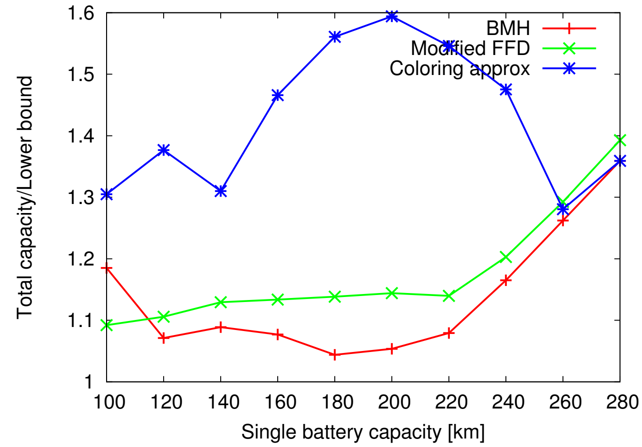

Fig. 2: Performance of packing algorithms compared to lower bound

As implied by its name, BMH has a bi-modal nature and it packs trips into batteries in a chronological order, i.e. in the order of their departure times. In fact, the choice of a battery for $r_j^{v_i}$ is from $B^d(d_j^{v_i})$, i.e. at departure time. Recall, that we denote by $e_{th}(t)$ the longest trip in $R(t)$ (that is the longest trips to depart in the future). At the beginning of the day, when the batteries in $B^d(d_j^{v_i})$ are full and their mean energy level is above $e_{th}(d_j^{v_i})$, $r_j^{v_i}$ is packed into the fullest battery in $B^d(d_j^{v_i})$. As a result, load-balancing is achieved and energy over-concentration is avoided. When batteries become relatively empty, BMH moves to its second mode and $r_j^{v_i}$ is packed into a battery in $B^d(d_j^{v_i})$ that has the lowest energy level which is higher than $e_j^{v_i}$. The purpose of the second mode is to minimize unusable energy leftovers trapped in the batteries towards the end of the day.

The performance of the three heuristics is compared through simulations on random workloads generated as described in Section II. The results are depicted in Fig. 2. The abcissa is the size of a single battery (100-280km). Recall, that the trip-length distribution is uniform in (20,35), (55,70). The ordinate is the ratio between the minimum total battery capacity required for achieving a feasible run according to a given policy (the number of batteries that it required times the battery capacity) and the actual total energy for all trips. The latter is a lower bound on the optimal solution. For each value of the battery capacity, the presented results are averages over 30 randomly chosen input sets. For the BMH we

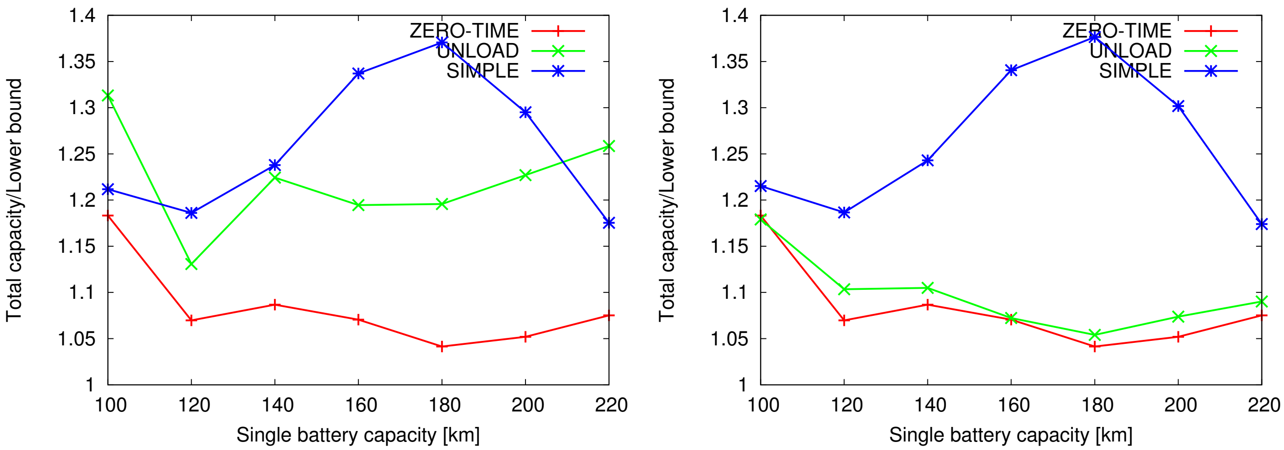

(a) Minimum number of mechanisms

(b) Minimum number of mechanisms+1

Fig. 3: Performance of realistic battery switching schemes

used exhaustive search in order to find the minimum required capacity under each every workload.

TABLE I: Minimum Required Number of Switching Mechanisms

	Battery capacity						
Scheme	100	120	140	160	180	200	220
Simple	3.3	3.0	3.4	3.13	2.83	2.0	2.0
Unload	6.0	5.93	5.56	5.46	5.43	5.43	5.43

The coloring approximation algorithm appears to yield the worst results. Although it indeed guarantees the formal approximation ratio, other heuristics outperform it dramatically by up to 55% of the lower bound. The modified FFD and BMH provide results of comparable quality, yet BMH appears to be better for a wider range of battery capacities.

For 100 km batteries, the modified FFD saves 9.3% of the lower bound compared to BMH. On the other hand, for 180 km batteries BMH is better by 9.4%. Apparently, the modified FFD is good with small batteries because the role of conflicts is reduced and FFD does what it's good at – packing without conflicts. However, for larger batteries it loses its advantage and BMH prevails. For battery sizes above 220 km the number of batteries is principally dictated by the maximum number of intersecting (in time) trips, each requiring a separate battery, which is an obvious lower bound on the number of batteries. (It is closely linked to the total number of vehicles and never larger than it). Once this number of batteries suffices, any further increase in battery capacity does not alter the number of required batteries and thus results in a linear increase in total required capacity regardless of the policy.

After considering the results, we decided to use BMH as the switching policy in the rest of the work. This choice is motivated by two factors: 1) BMH is better over a wider range of battery capacities, 2) unlike FFD, BMH operates incrementally in chronological order which allows for a simple combination between BMH and other policies that alter the state of the batteries during the day. Moreover, the second argument can be re-stated when the tentative arrival time and energy consumption of a trip are only approximate until the trips end.

As a conlcuding remark, let us state that we chose to collect average results because we assume that for a given depot instance the set of the trips is fixed. Therefore, 30 random workloads represent 30 different depots and the average is a good measure of the general fitness of the heuristics. In real-life scenarios when the depot is required to deal with workloads changing on the daily basis, the worst-case among the workloads will dictate the actual number of batteries that has to be held.

B. Realistic Battery Switching

While providing insightful results, the assumption of instantaneous switching is obviously far from realistic. Once we relax this assumption and assume switching times on the order of minutes, queues develop at the depot. The number of switching mechanisms, each holding its own battery pool, becomes important. Moreover, the exact policy of how and when to switch batteries directly affects the minimum feasible number of mechanisms required for a run to complete feasibly. We next examine two realistic switching schemes.

With the first scheme, the only operation allowed at a switching mechanism m_k is replacing a vehicle's battery with another one from m_k's pool B^{m_k}. As a result, since we assume every vehicle to have a battery at the beginning of the day, each vehicle holds a battery at any given moment. Upon returning to the depot at $a_j^{v_i}$, v_i's battery energy $e(b^{v_i}(a_j^{v_i}))$ is compared with $e_{j+1}^{v_i}$, the energy required for the next v_i's trip. If $e_{j+1}^{v_i} \leq e(b^{v_i}(a_j^{v_i}))$, v_i keeps its current battery for $r_{j+1}^{v_i}$. Otherwise, v_i immediately attempts to switch it. Battery pools of all idle switching mechanisms are examined and a suitable battery is chosen from them according to BMH (e_{th} is computed as before). Then v_i proceeds to the relevant mechanism. There it takes t_{sw} until its battery is replaced with the chosen one.

However, if at $a_j^{v_i}$ all mechanisms are busy, or no battery in the idle mechanisms has sufficient energy ($e_{j+1}^{v_i}$), v_i enters a waiting queue. In this queue, vehicles are sorted according to their pending departure times. When some switching mechanism m_k becomes idle, batteries in B^{m_k} are checked against the requirements of the queued vehicles. The queue is searched starting with a vehicle having the earliest departure time. If possible, i.e. if there is a sufficiently charged battery in B^{m_k}, the battery for replacement is chosen from B^{m_k} according to BMH and the examined vehicle is called in for switching. Finally, after completing the last scheduled trip, every vehicle switches its battery for the one with lowest energy level in the depot (for simplicity we assume that even "empty" batteries still have enough energy for moving around the depot).

With the second scheme, vehicles are allowed to move around the depot without a battery. This can be achieved by using a small secondary battery or connecting vehicles to a power supply in the depot. Thus, a vehicle can switch, unload or load a battery at a switching mechanism. We assume all operations to last the same t_{sw} amount of time. We still assume all vehicles to have a battery on-board at the beginning of the day. The scheme is described below.

Upon returning to the depot at $a_j^{v_i}$, v_i's remaining time until next departure $d_{j+1}^{v_i} - a_j^{v_i}$ is compared to a predefined threshold t_{th}. If $d_{j+1}^{v_i} - a_j^{v_i} < t_{th}$, v_i behaves according to the first scheme. Otherwise, v_i proceeds to unload its battery into one of the switching mechanisms. If at least one of the mechanisms is idle, the one with the fewest batteries is chosen. After unloading the battery, v_i moves to parking. Shortly before its next departure, at $d_{j+1}^{v_i} - t_w$, v_i once again turns to the switching mechanisms, this time in order to load a suitable battery. The battery is chosen according to BMH from battery pools of idle mechanisms. Whenever v_i cannot be immediately served, it enters the queue which is processed exactly as before in the simple scheme.

As mentioned above, given an input workload, each scheme has a minimum number of switching mechanisms required for providing a feasible run. We used simulation combined with exhaustive search in order to find that minimum. Then, the corresponding optimal battery capacity was found (for the smallest feasible number of mechanisms). Note that the smallest number of mechanisms does not necessarily enable the achievement of globally minimal value of battery capacity.

We collected minimum feasible number of switching mechanisms and corresponding optimal capacity results for 30 random input sets. The following parameter values were fixed: $t_{th} = 25$ minutes, $t_{sw} = 4$ minutes, $t_w = 15$ minutes. For each input set, the minimum capacity was found under the smallest feasible number of mechanisms and under that number of mechanisms incremented by one.

The averaged results are summarized in Table I and Fig. 3, and lead us to the following conclusions. First, the more elaborate switching scheme that relies on battery unloading requires a larger number of switching mechanisms to be deployed at the depot. This is not surprising since this scheme generates much higher switching activity. Moreover, increasing the number of mechanisms by one in order to better cope with that activity improves the achieved battery capacity. Interestingly, the further addition of mechanism has no impact. With the slightly larger than minimum number of mechanisms, the unloading scheme saves up to 32% of the lower bound in battery capacity relative to the simple switch-only scheme, but requires up to more than twice switching mechanisms (up to 3.43 additional mechanisms). These results expose a trade-off between increased operational complexity (the ability to unload the battery) combined with higher number of switching mechanisms and significant savings in the required battery capacity.

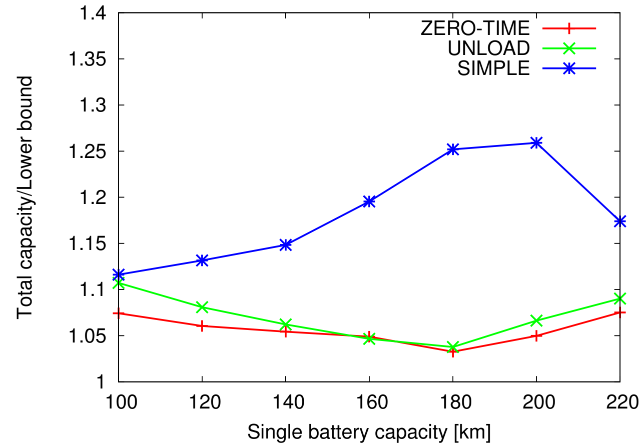

Fig. 4: Overall required battery capacity under inter-battery energy transfer

TABLE II: Normalized Improvement in Battery Capacity due to Inter-Battery Energy Transfer

	Battery capacity						
Scheme	100	120	140	160	180	200	220
Simple	0.10	0.06	0.09	0.15	0.12	0.04	0.00
Unload	0.07	0.02	0.04	0.03	0.02	0.01	0.00

IV. Battery Charging

In this section, we explore the benefit of battery charging during the day. In Section IV-A, we consider the transfer of energy between batteries. In Section IV-B, exploiting limited power from the grid is discussed.

A. Inter-Battery Energy Transfer

As mentioned in Section III-A, one of the factors limiting efficient battery capacity utilization is the fragmentation phenomenon, wherein available energy is spread thinly among batteries and is unusable. For instance, if at t it holds that $e(b,t) < e_{tl}(t)$, i.e. b has less energy than required by any of the trips occurring after t, then b and its energy are effectively wasted until the end of the day. We propose to transfer this otherwise unusable energy to other batteries in order to avoid the need in additional batteries. The idea of using vehicle battery as a power source was previously proposed in [4], [5]. However, these works discuss vehicle-to-grid (V2G) instead of inter-battery energy transfer. The exact heuristic policy that we apply is described below.

Let us assume that energy transfer decisions are made at time t, i.e. charging rates from $-c, 0, c$ are assigned to batteries in a feasible manner. Recall that $B_{th}(t) = \{b \in B^d(t) | 0 < e(b,t) \leq e_{th}\} = \{b_1, ..., b_n\}$, such that $e(b_j) \leq e(b_{j+1})$. Begin with the pair b_1, b_n. If $e(b_1) < e_{tl}$, assign b_1 to be discharged and b_n to be charged both at the same rate c. Proceed to b_2, b_{n-1}. Repeat until all pairs were traversed or both batteries in the current pair have more than e_{tl} energy. The trigger for assigning the rates is any change of the state of batteries in B^d. For example, a battery leaving or returning to the depot, a battery becoming empty or its energy reaching e_{th} from above or below. Note that the energy transfer is orthogonal to battery switching, wherein switching decisions are independent of charging/discharging decisions.

The reasoning behind the above policy is as follows. We don't charge fullest batteries in order to avoid energy over-concentration mentioned in Section III-A. Instead, we choose batteries that may lack a relatively small amount of energy in order to be suitable for a future trip. The charging is assigned first to b_n, then to b_{n-1} etc. because ultimately a run fails when a trip requires more energy than the fullest battery in B^d can provide.

Yet again, we assess the impact of inter-battery energy transfer by collecting and averaging simulation results for 30 random input sets. The charging/discharging rate was fixed $c = 3$ km/min so that a 180 km battery can be completely charged in an hour. The minimum required number of switching mechanisms for all inputs was found to be identical to that in Section III-B. This indicates that the peak rate of switching operations remains unchanged, but fewer batteries are required as energy distribution among the batteries in the depot not in the vehicles is better.

The results presented in Fig. 4 and Tab. II were acquired for the minimum required number of mechanisms incremented by one (as in Fig. 3b). Fig. 4 shows the lowest achievable aggregate battery capacity, while Table II summarizes the improvement in aggregate battery capacity for the different policies (in units normalized to the lower bound value) caused by applying energy transfer.

Apparently, under the simple switching scheme the depot benefits significantly more from inter-battery energy transfer (up to 15% of the lower bound) compared to the improvement under elaborate switching scheme (up to 7% of the lower bound). This behavior is attributed to the fact that the elaborate switching scheme is much closer to the ideal zero-time switching BMH which is designed to mitigate the negative impact of fragmentation. Overall, it seems that when battery unloading is used for switching, further complicating depot's infrastructure by allowing inter-battery energy transfer has a questionable profitability. It thus appear that one should best use either the switching-only plus inter-battery energy transfer or the unload without this transfer. The trade-off between those two is mostly the number of switching mechanisms and the ability to move battery-less vehicles in the depot versus required aggregate battery capacity.

B. Exploiting External Power Supply

High power price and its low availability may render massive consumption from the grid during the depot's operation hours impractical. Nevertheless, we are interested to explore how limited charging power affects the depot's essential resources. The available power p is assumed to be an integer multiple of battery charging rate c. Since limited power is assumed to be sufficient for charging only a few batteries simultaneously, some policy for choosing which batteries to charge is required.

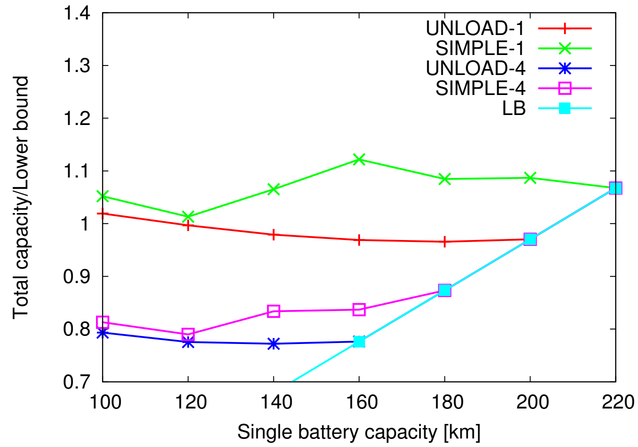

Fig. 5: Overall required battery capacity under external charging

We propose a heuristic policy inherently similar to that described in Section IV-A. Exactly as there, we begin with B_{th}. As long as not all of p is exploited, batteries of B_{th} are picked in a decreasing order of their energy levels and are assigned charging rate of c. By the end of this procedure, if all batteries in B_{th} are scheduled to be charged but some power is still left, we go over the rest of the batteries in an increasing order of their energy levels and assign them charging rate of c as well. On the other hand, if some of the batteries in B_{th} are not assigned charging power, we attempt to improve their state by applying on them the inter-battery energy transfer heuristic from Section IV-A.

As mentioned in Section II-B, the sum distance of all trips for the random inputs in our use is ∼33,000km. Assuming as before that all this energy is charged during 6 night hours, we conclude that the depot has a peak charging power of ∼92 km/min. We collect simulation results for $p = c$ and $p = 4c$. For $c = 3$ km/min, the former constitutes ∼3.2% of peak charging power and the latter constitutes ∼12.8%. The lowest achievable required aggregate battery capacity, averaged over 30 input sets is depicted in Fig. 5. Since some of the energy used by the trips is charged from the grid during the day, the battery capacity can drop below the sum distance of all trips (and thus the ratio value can drop below 1.0). Therefore, we use the number of vehicles multiplied by single battery size as an additional lower bound (curve LB).

The results suggest that having $p = c$ has a mild impact on the required battery capacity, 5%-15% of the total travel distance depending on the switching scheme and battery size. While for $p = 4c$ the savings are 25%-35%. When considering the impact of grid power on the number of switching mechanisms, for the unloading switching scheme we witness a minor reduction of 0.3-0.4 on average. For the simple switching scheme the number of mechanisms is virtually unchanged until sufficiently large battery sizes are reached. However, for external power of $4c$ there is no need for switching mechanisms if 180 km batteries are used, as every vehicle can hold its original battery throughout the work day. In fact, the depot's designer has a choice between paying for power during peak hours and using large batteries versus investing in a more sophisticated battery switching.

V. Conclusions

In this work, we have considered the problem of resource management in a fleet of SBEVs operating from a centralized depot under limited available power. Our goal was minimizing operational costs through efficient resource utilization. To the best of our knowledge, this is the first work to deal with this kind of a problem. As is typical of a new type of problem, both settings and solutions had to be invented, and the main goal was to explore and gain insights and directions.

We began by considering a setup in which no charging was allowed and battery switching took zero time. After analyzing the limiting factors, we proposed applying a novel heuristic algorithm named BMH. Under synthetically generated random inputs, the algorithm was shown to outperform other, commonly known, heuristics as came to within 10% of the lower bound for a wide range of battery capacities.

Next, we abandoned the zero time switching assumptions and defined two battery switching schemes: a) the simple scheme that performs lazy (only when necessary) switching, and b) the more elaborate scheme that relies on battery unloading and more closely implements the ideal BMH. We showed that the elaborate scheme succeeds to save up to 32% of the lower bound capacity compared to the simple one. However, it uses significantly more battery switching mechanisms and requires vehicles to move without a battery inside the depot. This behavior demonstrates a clear trade-off between the depot's complexity and the total battery capacity held at it.

Then, we explored the benefits of inter-battery energy transfer in an attempt to overcome remaining-energy fragmentation. This technique was observed to have no impact on the number of switching mechanisms, but was shown to reduce the required aggregate battery capacity especially for the simple switching scheme (up to 15% of the lower bound). Apparently, using a more elaborate switching scheme in the first place greatly reduces the profitability of allowing energy to be transferred between batteries.

Finally, the effects of charging batteries from the grid during the day were observed. As expected, the total battery capacity used by the depot is significantly reduced when such charging is applied. For larger sized batteries and sufficiently high charging power, we saw that a vehicle can travel all day with the same battery. This fact, represents yet another trade-off that needs to be taken into consideration when designing the fleet's depot.

In the future work, we intend to use insights presented in this paper in conjuncture with real-life data in order to analyze and find optimal working points for various types of depots. Such data includes among other things traces of trips, costs of charging power, switching mechanisms and batteries of different capacities.

References

[1] T. A. Becker, "Electric vehicles in the United States: A new model with forecasts to 2030," Berkley, CA: Centre for Entrepreneurship and Technology, Tech. Rep., 2009.

[2] G. A. Putrus, P. Suwanapingkarl, D. Johnston, E. C. Bentley, and M. Narayana, "Impact of electric vehicles on power distribution networks," in *Proc. IEEE Vehicle Power and Propulsion Conf. VPPC '09*, 2009, pp. 827–831.

[3] K. Jansen and S. Ohring, "Approximation algorithms for time constrained scheduling," in *Proc. IEEE First Int Algorithms and Architectures for Parallel Processing ICAPP 95. IEEE First ICA/sup 3/PP. Conf*, vol. 2, 1995, pp. 670–679.

[4] W. Kempton and J. Tomic, "Vehicle-to-grid power implementation: From stabilizing the grid to supporting large-scale renewable energy," *Journal of Power Sources*, vol. 144, issue 1, pp. 280–294, 2005.

[5] C. Guille and G. Gross, "Design of a conceptual framework for the v2g implementation," in *Proc. IEEE Energy 2030 Conf. ENERGY 2008*, 2008, pp. 1–3.

Vladimir Zdornov received the B.Sc. degree in computer engineering and the M.Sc. degree in electrical engineering from the Technion - Israel Institute of Technology in 2003 and 2009, respectively. Today he is a Ph.D. student at the Electrical Engineering Department of the Technion.

His research interests include computer systems and networks as well as algorithmic optimization approaches in electric vehicle infrastructures.

Yitzhak Birk received the B.Sc. and M.Sc. degrees from the Technion - Israel Institute of Technology in 1975 and 1982, respectively, and the Ph.D. degree from Stanford University in 1987, all in electrical engineering. Today he is an Associate Professor on the faculty of the Electrical Engineering Department at the Technion where he has been since October 1991.

Professor Birk's research interests include computer systems and subsystems, as well as communication networks. Recently, he has been particularly interested in energy efficiency aspects of these and other types of systems as he is attempting "cross fertilization" between his various areas of research.