

Backup Battery Analysis and Allocation against Power Outage for Cellular Base Stations

Fangxin Wang, *Student Member, IEEE*, Xiaoyi Fan, *Student Member, IEEE*, Feng Wang, *Member, IEEE*, and Jiangchuan Liu, *Fellow, IEEE*

Abstract—Base stations have been widely deployed to satisfy the service coverage and explosive demand increase in today's cellular networks. Their reliability and availability heavily depend on the electrical power supply. Battery groups are installed as backup power in most of the base stations in case of power outages due to severe weathers or human-driven accidents, particularly in remote areas. The limited numbers and capacities of batteries, however, can hardly sustain a long power outage without a well designed allocation strategy. As a result, the service interruption occurs along with an increasing maintenance cost. Meanwhile, a deep discharge of a battery in such case can also accelerate the battery degradation and eventually contribute to a higher battery replacement cost. In this paper, we closely examine the base station features and backup battery features from a 1.5-year dataset of a major cellular service provider, including 4,206 base stations distributed across 8,400 square kilometers and more than 1.5 billion records on base stations and battery statuses. Through exploiting the correlations between the battery working conditions and battery statuses, we build up a deep learning based model to estimate the remaining lifetime of backup batteries. We then develop *BatAlloc*, a battery allocation framework to address the mismatch between the battery supporting ability and diverse power outage incidents. We present an effective solution that minimizes both the service interruption time and the overall cost. Our real world trace-driven experiments show that *BatAlloc* cuts down the average service interruption time from 4.7 hours to nearly zero with only 85% of the overall cost compared to the current practical allocation.

Index Terms—Mobile network, Backup power system, Battery feature profiling, Deep learning, Battery allocation.

1 INTRODUCTION

Wireless mobile networks, particularly wide-area cellular networks, have seen deep penetration and broad coverage in the past decades. Base stations play a key role in today's cellular networks. Their reliability and availability heavily depend on the electrical power supply, for such modules as transceivers, air conditioners, monitoring system are all power hungry. The modern power grid is known to be highly reliable in urban areas, but still suffers from outages due to the severe weather (e.g., storm, hurricane, fire, earthquake) or human-driven accidents (e.g., vandalism or theft) [2], [3]. In many rural areas, the outages can be quite frequent, no matter in developing or developed countries.

To avoid service interruptions, most base stations are equipped with energy-storage battery groups as the backup power. These batteries are usually kept in the float charge state. Yet when a power outage happens, they will be activated to maintain cellular services until the electrical grid recovers or diesel generators are launched. The capacity of a backup battery group is limited, which typically lasts 10 to 12 hours during power outage. For remote areas or during extreme weather, however, the power recovery can take a long time (e.g., during the severe windstorm in March 2010, the power outage in southwestern Connecticut as well as parts of Long Island and New Jersey lasted for tens of hours, and in some of the rural communities the outage lasted as

long as 6 days [4]), so for technicians to arrive at the base station with diesel generators, not to mention that many base stations would be affected at the same time. As such, a long power outage without timely rescue will inevitably drain the backup battery, resulting in service interruption during the extended power outage. In this situation, these base stations have to rely on diesel generators whose operating cost is about ten times greater than powering through the electric grid [5]. Besides the possible long time duration, some areas may suffer from frequent power outages due to the bad weathers, e.g., it is reported that there were as many as 5 severe power outages in Okanagan Valley area in Canada in the first half of 2017 with an average duration of 8.4 hours [6]. These situations seriously affect the user experience and undermines the telecom operators' service commitments, particularly considering the clients' high reliance on the network during the incident.

Moreover, different from batteries for phones or electrical vehicles which regularly experience full charge/discharge cycles, a deep discharge of an energy-storage battery group (typically lead-acid) will severely affect its internal structure, reducing its capacity and lifetime. Given the long time interval between regular maintenances (usually three months [2]), the poor working condition of the battery after a deep discharge will further accelerate its degradation. In the worst case, an overdischarge can permanently damage the battery. Considering the transportation and labor costs, an emergent battery replacement and maintenance can be prohibitively expensive, particularly for remote areas.

In this paper, we closely examine the power outage

A preliminary version of this work [1] appeared in ACM e-Energy 2017. Fangxin Wang, Xiaoyi Fan and Jiangchuan Liu are with the School of Computing Science, Simon Fraser University, Burnaby, B.C., Canada. E-mail: {fangxinw,xiaoyif}@sfu.ca, jcliu@cs.sfu.ca. Feng Wang is with the Department of Computer and Information Science, the University of Mississippi, MS, USA. E-mail: fwang@cs.olemiss.edu.

events and the backup battery activities from a 1.5-year dataset of a branch of China Mobile (the biggest cellular service provider in China), including 4206 base stations and more than 1.5 billion records on base stations and batteries. Our analysis of the data reveals the ineffectiveness of existing battery allocation strategies during power outages. In particular, there is a clear mismatch between the battery supporting ability and the diverse power outage events.

Based on the logs of batteries, we further identify the impact of power outages on the conditions of the battery groups, and estimate the battery lifetime and reserve time (indicating the duration a battery group can support) through a deep learning based model. We accordingly develop *BatAlloc*, a battery allocation framework that allocates proper numbers of battery groups to each base stations to address the mismatch between the battery supporting ability and the diverse power outage incidents. We present an effective solution that minimizes both the service interruption time and the overall cost. Our trace-driven experiments show that *BatAlloc* reduces the average service interruption time from 4.7 hours to almost zero (i.e., nearly full service availability) with only 85% of the overall cost, as compared to the current real deployment.

The rest of paper is organized as follows. §2 introduces the related researches in the base station and battery management. §3 introduces the background and analysis of our dataset on base stations and backup battery groups. §4 summarizes the existing problems in current base stations and proposes the *BatAlloc* framework. §5 formulates the multi-objective optimization problem for battery allocation followed by a deep learning based model and an effective solution for optimization. §6 shows our experiments on voltage estimation and battery allocation. We provide some discussions in §7 and conclude our work in §8.

2 RELATED WORK

In this section, we first introduce some recent works related to our research, including energy aware resource allocation and battery feature profiling.

Many researches on energy related resource allocation have been proposed towards better performance and cost effectiveness. Wang et al. [7] proposed a novel resource allocation scheme to improve the performance of D2D communications. They considered battery lifetime as the optimization goal and employed a game-theoretic approach to achieve effective power control and radio resource allocation. Holtkamp et al. [8] focused on minimizing the base station supply power consumption by exploring the trade-offs between three basic power-saving mechanisms. Ramamonjison et al. [9] considered the resource allocation in a two-tier wireless system and proposed new mechanisms to efficiently allocate available energy over time. These two works aimed to increase the base station service availability by reducing the power consumption. Chamola et al. [10] considered both the grid energy cost and the quality of service, and proposed a framework to explore the tradeoff between the two aspects. Most of these works, however, focused on reducing the power cost or improving the quality of service given the fixed batteries [8], [10], [11]. Our work complements those aforementioned by investigating the

TABLE 1
Statistics on number of battery groups for more than 4200 base stations.

number	1	2	3	4	other
percent	65.9%	27.5%	4.7%	1.6%	0.1%

energy related problem in the base station from a different angle, where we propose a battery allocation framework to achieve better service availability and reduce the overall cost in base stations.

Estimation of capacity and reserve time of batteries is an everlasting topic and has attracted many efforts due to its considerable importance for continuity of service and wide use in large or small systems. A lot of approaches have been proposed to estimate the battery state of charge (SOC) and lifetime based on battery features, such as open circuit voltage, ampere-hour characteristics, charge or discharge curves etc. Kutluay et al. [12] proposed an online battery SOC estimation based on the discharge rate versus discharge time and coulometric measurement given in the manufacturer's data sheets. Anbuky et al. [13], [14] built up an estimation model to predict battery SOC and reserve time only based on the battery discharge voltage, which is robust against discharge rate, ambient temperature, battery degradation situation etc. with the error ratio less than 10%. Coleman et al. [15] incorporated the changes occurring due to terminal voltage, current load and internal resistance to predict electromotive force (EMF) of battery, and further estimate SOC based on the EMF. Bhangu et al. [16] utilized a Kalman Filter based approach for real-time SOC estimation, and further predicted the state of battery health. These methods all utilized traditional electrochemical theory to analyze the battery characteristics, while they may fail to achieve a high accuracy. Different from prior works, our work strives to profile the battery features from the aspect of big data analysis. With the advance of deep learning, we propose a learning based approach for battery profiling considering multi-battery deployment, which is based on a large-scale real world dataset of base station batteries.

3 BACKGROUND AND DATA ANALYSIS

In this section, we mainly profile the collected dataset and the related observation on the base stations and backup battery groups. We collaborate with a branch of China Mobile (the biggest cellular service provider in China) and collect a dataset from July-28-2014 to February-17-2016, which covers 4206 base stations distributed across 8400 square kilometers with over 10,616,000 clients. The dataset consists of more than 1.5 billion¹ records on battery activities, including such information as the base station locations, battery voltages and event records (e.g., power outage, low voltage alert, high voltage alert, etc.), which are used to analyze the current situation of base stations.

1. When batteries are in floating charging state, the first-year dataset has higher resolution and the remaining half-year dataset has relatively lower resolution. Yet for all other battery states, the resolution is the same across the whole dataset. Thus, this imbalance does not affect our data driven observation and evaluation.

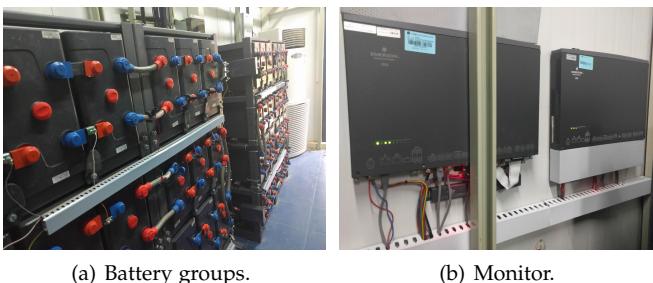


Fig. 1. The backup batteries and the monitor system.

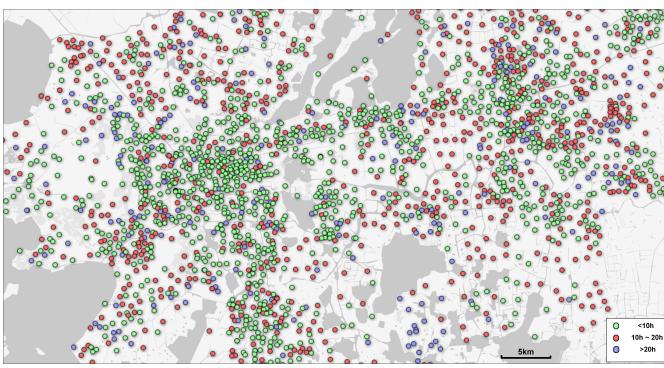


Fig. 2. Base stations distribution and their power outage situations. Each point shows the location of corresponding base station and the color represents the maximum single power outage duration in a year.

3.1 Power Supply in Base Stations

We first introduce a generic backup power system in the base stations of mobile networks. The equipment in base stations is usually supported by the utility grid, where the battery group is installed as the backup power. In case that the utility grid interrupts, the battery discharges to support the communication switching equipment during the period of the power outage. Fig. 1(a) shows two lead-acid battery groups in a mobile network base station and each battery group contains 24 cell batteries (the rated voltage of each battery cell is 2v). The rated capacity of a battery group is usually 500AH and it can support about 10-12 hours (i.e., the *reserve time* of a battery group is 10-12 hours). Compared to other types of batteries (e.g., Li-ion battery), lead-acid battery groups demonstrate some important advantages such as the mature technologies, safe storage, high capacity and low price, which make them widely used in base stations. We observe the number of battery groups from more than 4200 base stations and show it in Tab. 1. We find that about 93.4% of base stations are equipped with one or two battery groups while only very few base stations have more. In Fig. 1(b), the monitoring system connects to each cell of the battery group and periodically records the voltage and status in both normal and abnormal situations.

When the monitoring system reports an alert status, the emergency repairing service is scheduled depending on the accident severity. For instance, grid transmission lines can be cut off in case of extreme weather (e.g., storm, hurricane and heavy snow). Then the monitoring system in base stations will report the power outage to the maintenance center and an emergent maintenance should be scheduled

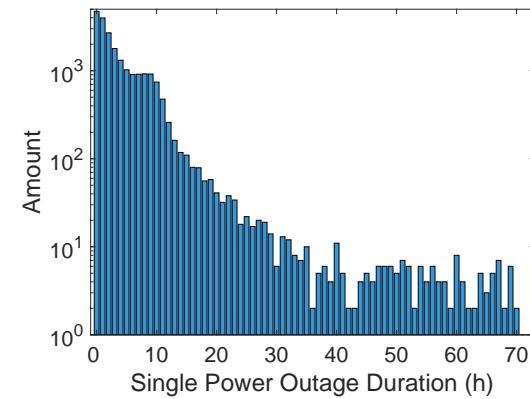


Fig. 3. Statistics of power outage duration each time for all base stations.

historybattery			
equipmentid	recordtime	floatvalue	signalseverity
12705658	2015-02-05 11:28 PM	2.22315	255
12705658	2015-02-06 06:15 AM	1.95235	0
12705658	2015-02-06 12:48 AM	1.62205	0
historystatus			
equipmentid	starttime	endtime	meanings
12705658	2015-02-05 11:26 PM	2015-02-06 06:15 AM	power outage
12705658	2015-02-06 00:45 AM	2015-02-06 04:37 AM	voltage low
12705658	2015-02-06 09:27 AM	2015-02-06 12:53 AM	voltage low

Fig. 4. A part of real logs of batteries.

according to priorities of different base stations. Since few base stations have the diesel generators permanently installed on site, maintenance engineers have to spend a long time to take diesel generators as well as other necessary devices to the corresponding base stations. The power outage can occur frequently and severely in the rural areas and developing countries due to the unstable utility grid. To make it even worse, the construction of infrastructure often makes that the base stations are difficult to reach, e.g., slippery rock trails in the mountains, where the workers have to manually carry the heavy generators to the site. So the power recovery time is quite uncertain and can not be guaranteed.

We extract power outage situations of the base stations and illustrate the practical distribution of the base stations as well as the power outage situations in Fig. 2. It is clear that in the urban regions most base stations have relatively good power supply, while in the remote rural areas base stations can suffer from long-time power outages. Fig. 3 shows the statistics of power outages of all base stations, from which we can find that quite a few power outages last very long time. However, according to the current battery allocation in Tab. 1, base stations with inefficient backup batteries are not able to sustain the long-time power outage without timely emergent maintenance, which can lead to service interruptions and cause serious consequences.

3.2 Backup Battery Features Analysis

In our dataset, we have obtained huge amounts of logs from batteries of 4206 base stations with totally 531 tables and 1,550,032,984 rows. As shown in Fig. 4, the main logs in our

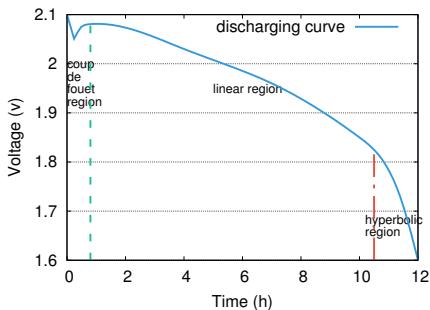


Fig. 5. Typical discharge voltage versus time characteristics.

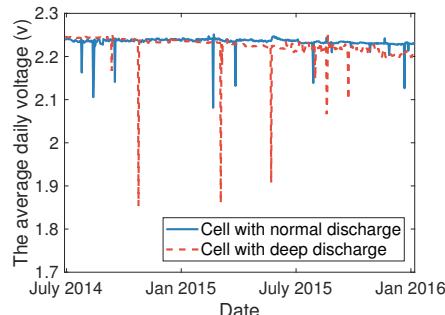


Fig. 6. The comparison of two battery cells under different discharge situations.

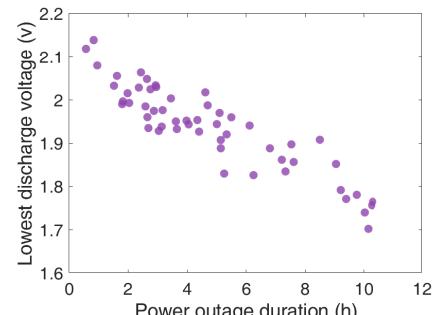


Fig. 7. The relationship between power outage duration and voltage drop.

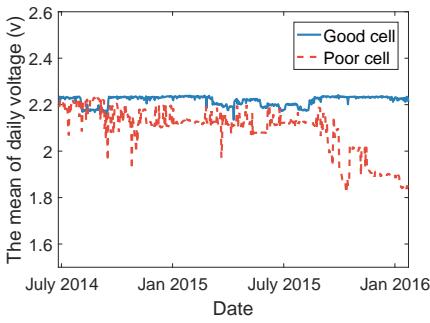


Fig. 8. Mean voltage versus battery status.

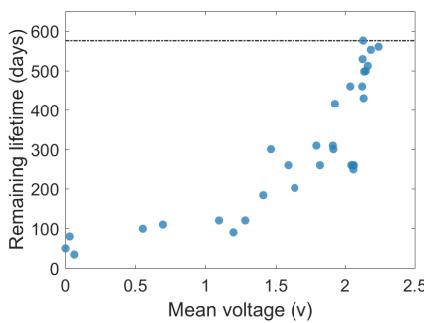


Fig. 9. Correlation between the remaining life and mean voltage.

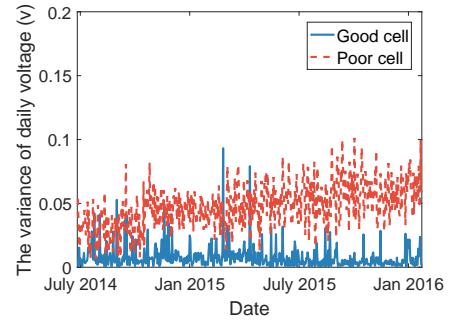


Fig. 10. Voltage variances versus battery status.

collected dataset include two parts, i.e., historybattery and historystatus. The historybattery logs record the collected information of each battery cell such as equipmentid (the unique device number related to a battery cell), recordtime (the timestamp when this log record was generated), floatvoltage (the monitored float voltage for this cell) and signalseverity (the level of emergency which decreases as the value grows). The historystatus logs describes the status of battery or external environment, such as the equipmentid (the unique identification of each battery cell), starttime (the start time of a status), endtime (the end time of a status) and the meanings (the specific status, such as power outage, voltage low, voltage high, etc.).

3.2.1 Battery Discharge Analysis

Base station batteries are connected to the electrical grid and kept in float charging state to compensate the capacity loss due to the slow self-discharging process. When there is a power outage, the backup batteries begin to discharge to support base station services. The battery discharging process can be divided into three regions: the coup-de-fouet region [17], the linear region [14] and the hyperbolic region [18]. Fig. 5 illustrates a typical discharging curve for a lead-acid cell. The coup-de-fouet region appears at the start of battery discharging, where the battery voltage first falls quickly below its open circuit voltage and then rises to a higher plateau voltage in a short time. This kind of voltage change is a special characteristic usually observed from lead-acid batteries. Then the discharging process goes into a long linear region, where the voltage drop has an approximately linear relationship with the discharging time. The discharging characteristic is robust to variations in

operating conditions as well as battery conditions, such as the discharging mode (constant current or constant power), ambient temperature, battery degradation condition [14], etc. A battery will release most of its energy during the linear region. In the last hyperbolic region, the voltage falls very fast while it can only release a very small fraction of power.

During a long power outage, the backup batteries may need to discharge to a deep level (e.g., to the hyperbolic region in Fig. 5), which further exert an impact on the battery conditions. The conditions of lead-acid batteries are largely dependent on the depth of discharge (DoD). If a lead-acid battery frequently discharges to high DoDs, the lead in the negative plate will form large lead sulfate crystals adhered to the negative plate and further accelerate the battery sulfation. This degradation process is accumulative, which as a result greatly reduces the capacity and lifetime of lead-acid batteries. Therefore, it is not desirable to allow a battery group to discharge completely, because the battery group will be permanently damaged and become incapable of being fully recharged to its rated capacity again. According to the industry standard, a battery should be replaced once its capacity falls below the 80% of the rated capacity. So the fast battery degradation contributes to a high battery replacement cost. Fig. 6 presents a comparison of the voltage change between two battery cells, one of which was in good condition and the other suffered from several deep discharges. We can see that the cell suffering from deep discharges degrades quickly with the float voltage showing a clear decreasing trend.

In base station power management, a low voltage disconnect (LVD) strategy is applied for battery protection.

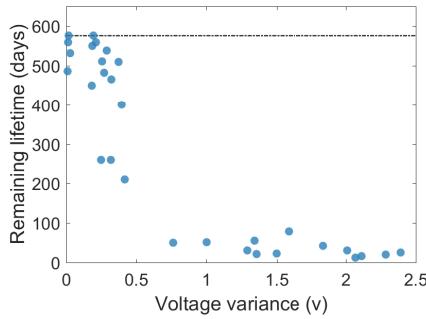


Fig. 11. Correlation between the remaining life and voltage variance.

Rank	Category	Count	Percentage
1	Alert	1479914	28.09%
2	Faulty Cell	1075533	20.42%
3	Discharge	563469	10.70%
4	Low float voltage	457767	8.69%
5	Too high	292966	5.56%
6	Low	263625	5.00%
14	Power outage	48605	0.92%
26	Failure	9100	0.17%
27	Voltage too low	7060	0.13%
30	Low cell voltage	2830	0.05%
34	Generator on charge	960	0.02%

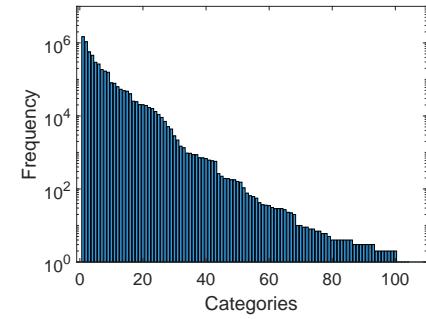
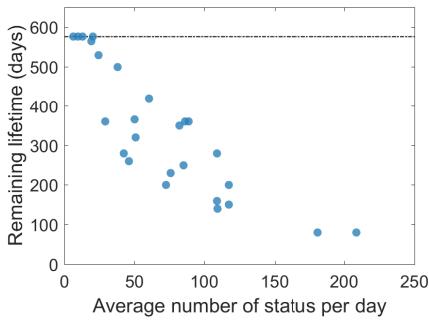
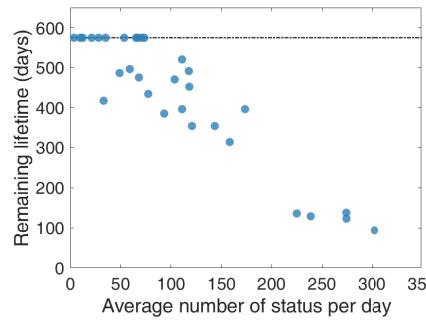


Fig. 12. The list of battery status categories.

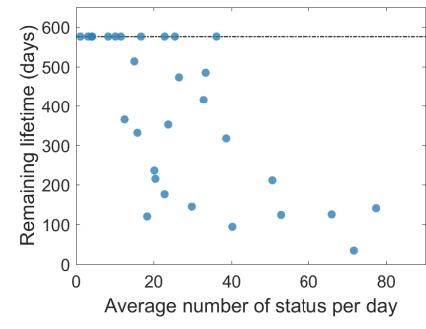
Fig. 13. Distribution of battery status categories.



(a) Low float voltage status.



(b) Discharge status.



(c) Faulty cell status.

Fig. 14. Correlation between the remaining life and the number of different statuses.

When the battery voltage falls below a first pre-defined threshold, the lead-acid battery groups will be disconnected from the secondary devices and only provide backup power to primary communication devices. When the voltage continues to drop below a second predefined threshold, power system cuts off all the loads to avoid the battery groups from being drained. Base stations usually have a low LVD settings to prolong the backup power supply, yet actually the deep discharge before LVD has already exerted an impact on battery degradation process. Fig. 7 plots the relationship between the power outage duration and the voltage drop (to avoid the impact of battery group numbers, we only choose those base stations with one battery group). We observe that the discharge voltage could fall below 1.71v during a long power outage, which in fact will seriously damage the battery condition.

3.2.2 Battery Voltage Analysis

The voltage of each cell battery is the most important feature that we have measured, as it reflects the power output pattern of the battery. In general, we have observed two representative categories of cell batteries, where we manually choose 1578 batteries as the newly-installed group and put 1459 batteries into the nearly-dead group depending on the repair records. The rated voltage of a cell is around 2.23v and the rated voltage of a battery group is 53.5v, where 24 cell batteries are connected in serial as one battery group. Based on this, we further analyze the typical status of the voltage patterns inside the two representative cell battery categories. Fig. 8 shows the significant differences in mean voltage between the newly-installed and nearly-

dead batteries. The blue solid line plots the mean voltage of newly-installed batteries, which judders between 2.14v and 2.24v. The red dotted line shows the decay trend on the mean voltage of the nearly-dead batteries. There is a clear downward trend close to the failure date, where the battery power frequently falls down and becomes quickly exhausted, causing many issues and alerts in the mobile network base station.

Fig. 9 further plots how the mean voltage and the length of remaining lifetime correlate with each other, which indicates that the mean voltage has strong correlations with the battery life. Fig. 10 shows the results on the voltage variances, where the blue solid line represents the newly-installed battery can output a steady power and the variance of the voltage keeps very close to zero. The red dotted line illustrates that the variance of the nearly-dead batteries increases much faster than the newly-installed batteries.

Fig. 11 illustrates that the voltage variance has a correlation with the length of the remaining lifetime, indicating that the variance of the output voltage from the batteries over time also reflects the aging trend of battery quality degradation. These observations motivate us to correlate battery working conditions with the battery historical voltages.

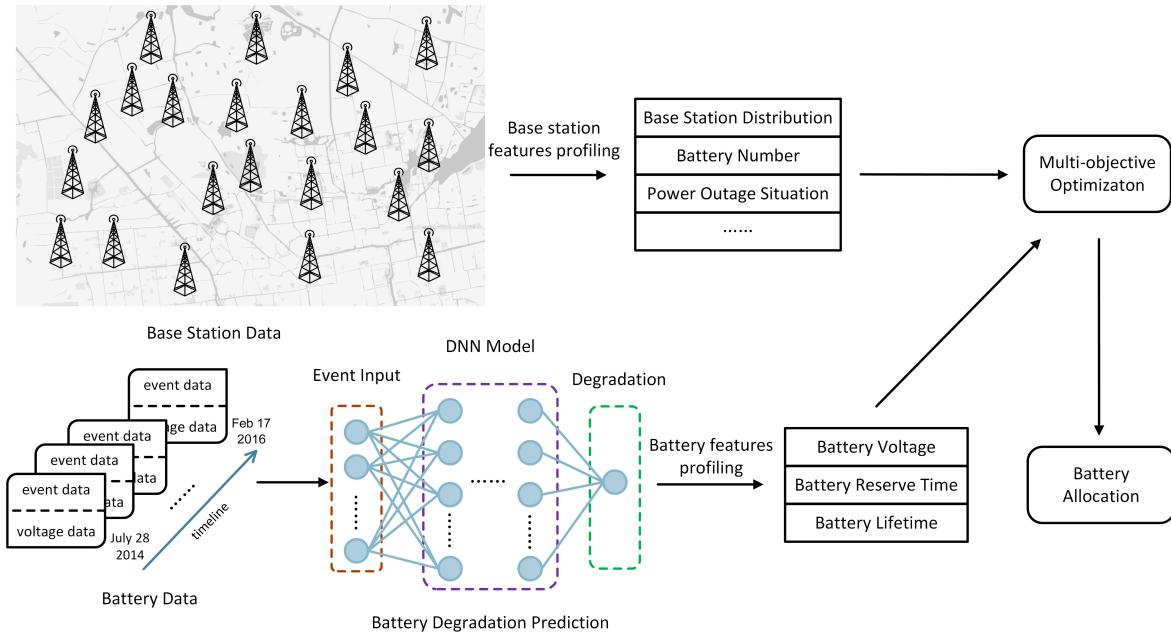


Fig. 15. Systematical design of BatAlloc framework.

3.2.3 Battery Status Analysis

Fig. 12 lists the number and percentage of some selected status categories², and Fig. 13 also shows the frequency distribution among all the 105 categories. We can see that the distribution is highly skewed: the most popular category is Alert (meaning that there are warnings such as irregular voltage change), at about 28.09%; the second is Faulty cell (meaning that the system infers that the corresponding cell may have fault), at about 20.42%; and the third is Discharge (indicating that the cell is discharging), at about 10.70%.

We take three statuses as examples to further investigate the correlations between the status and battery remaining lifetime. These statuses are *Low float voltage* (i.e., the monitored float voltage falls below a threshold), *Discharge* and *Faulty cell* as shown in Fig. 14(a), (b) and (c), respectively. We count the specific status number for each battery until the batteries are replaced, and pick up 30 batteries with different number of statuses for observation. There are 576 days in our dataset, where the remaining lifetime of most batteries in our dataset is longer than 576 days. Therefore dash lines represent that those batteries on it have longer remaining lifetime than 576 days. Fig. 14(a) and (b) plot the correlation between *Low float voltage*, *Discharge* and remaining lifetime. They clearly demonstrate that there exists a strong correlation between battery remaining lifetime and *Low float voltage*, as well as between battery remaining lifetime and *Discharge*. We further plot the remaining lifetime against the number of *Faulty cell* status in the system in Fig. 14(c), which does not show noticeable correlation between them. These results imply that the remaining lifetime is comprehensively affected by some statuses rather than a specific one. The observations suggest that the diverse statuses have different

2. In the figure, the status category of generator on charge indicates that the generator is providing power. Too high means the monitored float voltage of the corresponding cell surpasses a threshold. Failure means that the power system or the communication system goes wrong.

influences on the battery working conditions, thus it is necessary to discriminately differentiate these statuses for the accurate lifetime prediction.

4 BATALLOC FRAMEWORK

Our real trace-driven data analysis clearly reveals that in the battery allocation strategy currently used in practice, there exists a mismatch between the supporting ability of backup batteries and the power outage situations in each base station. The mismatch can lead to serious problems in base stations. First, due to the limited numbers and capacities of backup battery groups, long time power outages can result in service interruptions in many base stations. It is even worse during severe weather in rural areas or remote places, where maintenance engineers are not guaranteed to arrive timely. Besides, as the emergent maintenance is accompanied with service interruptions, more service interruptions also contributes to extra cost on emergent maintenance. What is more, long time power outages can drain the battery capacity, affecting battery structures and accelerating battery degradation. The results further lead to sooner battery replacement and higher overall cost.

One intuitive solution is to allocate as many battery groups as possible for every base station, yet such an overprovision will cause a large waste of resources and dramatically increase the overall cost. To this end, we propose BatAlloc, a battery allocation framework to carefully address this mismatch by allocating an appropriate amount of backup battery groups for each base station. As shown in Fig.15, our framework consists of three major stages, namely, Base Station Feature Profiling, Battery Feature Profiling, and Battery Allocation Optimization, which will be further explained as follows:

Base Station Feature Profiling: In this stage, we mainly extract the features of base stations from massive data, including the practical distribution of base stations, numbers

of battery groups equipped in base stations, power outage situations, etc. The profiling results lay the foundation for later analysis such as the severity of power outage, the impact of service interruption, as well as the cost for emergent maintenance and battery replacement.

Battery Feature Profiling: This stage conducts a solid analysis on the battery features, so that battery capacity, battery lifetime and battery degradation under different levels of discharges can be accurately estimated. Although the lead-acid battery technology is mature, due to the large variations of real world factors, it is still very difficult, if not impossible, to do such estimations directly by the domain knowledge. To this end, we develop a deep learning based approach to well model the complicated relationships between different real world events and various battery conditions, which will serve as a key component for the battery allocation optimization in next stage.

Battery Allocation Optimization: Based on the feature profiling results of previous two stages, the battery allocation can then be formulated as an optimization problem. This problem involves multiple optimization goals, e.g., to minimize the service interruptions and minimize the overall cost. In addition, a number of real world factors can also be considered into the optimization, such as the importance of different base stations, the available budget, and the practical limitations on the number of battery groups that can be installed on a base station. Besides, due to the large number and space span of base stations, the optimization solution should also be very efficient for computation.

It is worth noting that different battery types (e.g., lead-acid batteries and Li-ion batteries) may have quite different degradation characteristics. The battery feature profiling is only one component of our framework and the remained part can still be well applied to other chemical battery scenarios, as long as the battery profiling model is updated as needed. We have analyzed the base station features in the previous section. In next section, we present the deep learning based battery feature profiling model and the solutions for battery allocation.

5 BATTERY ALLOCATION SOLUTIONS

In this section, we first formulate the battery allocation optimization stage in our BatAlloc framework as a multi-objective optimization problem. Then we propose a deep learning based approach integrated with battery discharge features to model the battery reserve time and battery lifetime for a base station equipped with different numbers of batteries. At last, we propose an efficient algorithm to solve the formulated optimization problem. Tab.2 lists the notations to be used in this section.

5.1 Problem Formulation

Current base stations are mostly equipped with one or two battery groups, which are often insufficient to provide uninterrupted backup power during a long power outage. Assume that we assign n_s battery groups for a particular base station $s \in \mathcal{N}$, where \mathcal{N} is the set of all base stations. We then need to calculate how long the n_s battery groups can support this base station during a power outage. Recall

TABLE 2
Notations

n_s	number of battery groups at base station s
r_{s,n_s}^t	the reserve time for station s at t with n_s battery groups
o_s^t	duration from power outage to grid recovery or generator launch for stations s at t
ω_s	the importance factor of station s on service interruption severity
T_{s,n_s}	the expected lifetime of each battery group when stations s is equipped with n_s battery groups
\mathcal{I}	the normalized total service interruption time
\mathcal{T}	the time-based index range
\mathcal{N}	the set of all the base stations
c_b	the replacement cost of a battery group
\mathcal{C}_b	normalized total cost on battery group replacement
x_s^t	a variable indicating whether station s needs an emergent maintenance at t
$c_{m,s}$	emergent maintenance cost for station s
\mathcal{C}_m	the normalized total emergent maintenance cost
\mathcal{C}_{all}	the normalized overall cost
n_L	lower limit of battery group number in a station
n_U	upper limit of battery group number in a station
\mathcal{B}	the budget limit
v_{if}^k	float voltage of battery i in k -th segment
s_i^k	voltage slope of battery i in k -th segment
d_i^k	degradation of battery i in k -th segment
\mathcal{E}_i^k	the event set for battery i in k -th segment
\mathcal{E}	the event set
\mathcal{D}	the degradation set
ϵ	the voltage drop at start of discharging
Φ_s^t	percentage of remaining capacity of a battery group
v_i^t	the plateau discharging voltage of battery i at t
v_E	the end discharging voltage in linear region
v_P	plateau voltage at the beginning of discharging
τ	the rated reserve time before end voltage

that the battery has already severely suffered from deep discharge at the hyperbolic region. To protect the battery, when the battery discharges to the end of linear region (as illustrated in Fig. 5), we disconnect it from the workload. To this end, we denote r_{s,n_s}^t as the total reserve time for station s with n_s battery groups at time t .

We denote the time duration from the beginning of power outage to electrical grid recovery or diesel generator launch in station s as $\mathbf{o}_s = \{o_s^{t_1}, o_s^{t_2} \dots o_s^{t_i}\}$, where t_i is a time-based index. Once the duration exceeds the battery reserve time, there will be a service interruption. We assign importance factor ω_s to represent the service interruption severity (e.g., the service interruptions in core station have more serious consequences). We use T_{s,n_s} to denote the expected lifetime of each battery when station s is equipped with n_s battery groups under its specific working situations. Thus, we have our first optimization objective, which minimizes the total service interruption time:

$$\begin{aligned} \text{Min} : \mathcal{I} &= \sum_{s \in \mathcal{N}} \mathcal{I}_s \\ &= \sum_{s \in \mathcal{N}} \frac{\omega_s \sum_{t \in \mathcal{T}} [\max(0, o_s^t - r_{s,n_s}^t)]}{T_{s,n_s}} \end{aligned} \quad (1)$$

where \mathcal{T} is the time-based index range of the considered period. We use T_{s,n_s} as denominator for normalization(i.e., representing annual service interruption time).

Besides achieving as short service interruption time as possible, telecom operators may also want to reduce the

overall cost, which includes the battery replacement cost and emergent maintenance cost. Then the battery replacement cost \mathcal{C}_b (including purchase and installment) can be represented as follows:

$$\mathcal{C}_b = \sum_{s \in \mathcal{N}} \mathcal{C}_{s,b} = \sum_{s \in \mathcal{N}} \frac{n_s c_b}{T_{s,n_s}} \quad (2)$$

where c_b is the replacement cost of a single battery group of base stations. For simplification, we assume an average for the shipment and labor cost, and combine all these costs including the purchase cost as the replacement cost.

When there is a long power outage that the battery capacity is not sufficient enough, engineers may be scheduled an emergent maintenance to the corresponding base station for power generation.

We use x_s^t as a binary indicator that indicates whether an urgent maintenance is demanded during a power outage (x_s^t is set to 1 if the battery reserve time is not enough when there is a power outage at time t , and 0 otherwise). And $c_{m,s}$ is the emergent maintenance cost of station s . Then we can get the total emergent maintenance cost \mathcal{C}_m for all the base stations as follows:

$$\mathcal{C}_m = \sum_{s \in \mathcal{N}} \mathcal{C}_{m,s} = \sum_{s \in \mathcal{N}} \frac{\sum_{t \in \mathcal{T}} (x_s^t c_{m,s})}{T_{s,n_s}} \quad (3)$$

Based on Equation 2 and Equation 3, we then have our second optimization objective, i.e., minimizing the overall cost \mathcal{C}_{all} for telecom operators:

$$\begin{aligned} \text{Min : } \mathcal{C}_{all} &= \mathcal{C}_b + \mathcal{C}_m \\ &= \sum_{s \in \mathcal{N}} \frac{n_s c_b + \sum_{t \in \mathcal{T}} (x_s^t c_{m,s})}{T_{s,n_s}} \end{aligned} \quad (4)$$

In practice, there may be other requirements that limit the number of battery groups being installed at a base station. We thus use n_L and n_U to denote the lower limit and upper limit on the number of battery groups that can be installed, respectively, and have the following constraint:

$$\forall s, n_L \leq n_s \leq n_U, n_s \in \mathbb{N}^+ \quad (5)$$

Besides, telecom operators usually want to control the overall cost within a give upper budget limit \mathcal{B} . So we also have the following constraint:

$$\mathcal{C}_{all} \leq \mathcal{B} \quad (6)$$

5.2 Deep Learning Based Battery Profiling

In order to solve the optimization problem on battery allocation, we first need to model the lifetime and reserve time of the batteries in a base station. Given that the battery voltage is often used as a criterion for battery working conditions as well as battery capacity, we thus can conduct battery profiling to build up the models based on the historical battery activities under different events recorded in our data logs. Traditional time series estimation models such as ARIMA [19] and linear regression [20] only explore the time series features of battery voltages, while they are not able to capture the impact that external events have on batteries. Although [21] considered the impacts of events on voltage trend, its proposed approach can only be used to model a single battery group. To this end, we develop a

deep learning based approach that utilizes the deep neural network (DNN) to accurately model the voltage trend based on previous events and voltages with the consideration of multiple battery groups.

The degradation process of a battery is relatively a long period impact derived from battery activities. So we focus on the voltage trend rather than every single voltage value at each time point. We first filter out the noise voltage data generated during battery activities (e.g., charging and discharging) and only extract the effective float voltage data. Given a time series of float voltages for battery i , we divide them into a number of time segments where the length of each segment is l . For each segment k , we fit the voltage decreasing trend by linear regression and obtain the voltage change slope s_i^k as well as the initial voltage value v_{if}^k . Then each time segment can be represented as $\{(v_{if}^1, s_i^1), (v_{if}^2, s_i^2) \dots (v_{if}^k, s_i^k)\}$. We define voltage degradation term as the rate of change on voltage slope for a battery. Then we have battery degradation d_i^k as following:

$$d_i^k = s_i^k - s_i^{k-1} \quad (7)$$

For each segment, the battery voltage degradation is ascribed to the battery activities, which are directly reflected by the event logs. We define $\mathbf{e}_i^k = \{e_{i,1}^k, e_{i,2}^k, \dots, e_{i,m}^k\}$ as the input events for battery i in time segment k , where m is the number of event categories. When a base station is equipped with multiple battery groups, the impact of activities is actually shared by all these batteries. Then the impact on each single battery should be proportionally reduced. Thus, we can build up a learning model from events $\frac{\mathbf{e}_i^k}{n_s}$ to the battery degradation d_i^k in segment k , where n_s is the number of battery groups in base station s .

Formally, the inputs are the event sets associated with related segments. Let \mathcal{E} denote the input space of the historical events and we have $\mathcal{E} = \{\frac{\mathbf{e}_1}{n_1}, \frac{\mathbf{e}_2}{n_2}, \dots, \frac{\mathbf{e}_N}{n_N}\}$ with N examples. The outputs are voltage degradations for each segment. Let \mathcal{D} denote the output space of voltage degradation and we have $\mathcal{D} = \{d_1, d_2, \dots, d_N\}$. The modeling process is actually a mapping from \mathcal{E} to \mathcal{D} .

As illustrated in Fig. 15, we build up a DNN to model the battery degradation process. Each node in the input layer is associated with one kind of events and output layer has one node for degradation estimation. Assuming the base station's situations keep statistically consistent every year, we can then obtain the voltage degradation utilizing our deep learning model. Given that the target time t falls in segment $k+1$ and $v_{if}^{t_1}$ is the initial voltage value, the float voltage can be calculated as follows:

$$v_{if}^t = v_{if}^{t_1} + \sum_{j=1}^k (d_i^j + s_i^{j-1})l + (d_i^{k+1} + s_i^k)(t - kl) \quad (8)$$

With the domain knowledge, a battery is judged in poor quality when its float voltage is below a pre-defined threshold θ . Then we can obtain the lifetime of battery i in station s if the float voltage falls below θ at segment $k+1$:

$$T_i = \frac{\theta - v_{if}^{t_1} - \sum_{j=1}^k (d_i^j + s_i^{j-1})l}{d_i^{k+1} + s_i^k} + kl \quad (9)$$

When there is a power outage, the batteries begin to discharge to provide backup power. According to the electrochemistry knowledge of base station battery features [22], there is a voltage drop from the float charge state to the plateau discharging state mostly due to the cell internal resistance and polarization. We denote the voltage drop as ϵ and we can calculate the plateau discharging voltage as $v_i^t = v_{if}^t - \epsilon$.

Recall that the battery degradation will lead to the battery capacity decrease. The battery discharge characteristics can be utilized to estimate the battery state of charge (SOC) and battery reserve time [14], [23] in the linear region. The scaled discharge curves of batteries with different degradation keep highly consistent, and the plateau discharge voltage drops with the degradation level. Thus we can build the mapping from the plateau discharge voltage to the corresponding capacity in the linear region. Let v_E denote the end voltage and v_P is the plateau voltage of discharging phase for a new battery cell. We use Φ_s^t to represent the percentage of remaining capacity of a battery group in the linear region at t . Then we can calculate Φ_s^t based on the discharging voltage v_i^t :

$$\Phi_s^t = \frac{v^t - v_E}{v_P - v_E} \quad (10)$$

Let τ denote the rated battery reserve time of a new battery before the end voltage. The reserve time r_{s,n_s}^t defined in the previous subsection can thus be calculated as follows:

$$r_{s,n_s}^t = \tau n_s \Phi_s^t \quad (11)$$

5.3 Battery Allocation Algorithm

With the profiling results of base station features and battery features, we next solve this battery allocation optimization problem. Recall that our objectives are minimizing both the service interruption time (Equation 1) and the overall cost (Equation 4). Then we have two constraints: the number of battery groups in each station falls within the limit range (Equation 5) and the overall cost does not exceed the budget limit (Equation 6). This multi-objective optimization problem is actually a multi-objective integer programming problem, where the battery group number n_s must be an integer between a lower bound n_L and an upper bound n_U . This makes the problem NP-hard and we thus design a heuristic algorithm to solve it efficiently.

Before jumping to the algorithm design, we first briefly analyze the characteristics of this optimization problem. Intuitively, given the same external incidents happening to a base station, the base station can sustain longer power outages when equipped with more battery groups. The total service interruption time is thus reduced. Meanwhile, since the emergent maintenance is accompanied with service interruptions, fewer service interruptions also cut down the cost of emergent maintenance. Thus in our allocation model when the battery group number keeps increasing, both the service interruption time and the emergent maintenance cost will monotonously decrease until no service interruption occurs.

However, the battery replacement cost is different, where the process can be divided into two stages: In the first stage, when the battery group number of a base station increases,

Algorithm 1: Battery allocation

Input: Results of base station and battery features profiling.
Output: The allocation results n_s for every station s .

- 1 **foreach** s in \mathcal{N} **do**
- 2 Set initial battery assignment as $n_s = n_L$ and calculate \mathcal{I}_s and $\mathcal{C}_{s,all}$;
- 3 **while** $n_s \leq n_U$ **do**
- 4 Increase n_s when both \mathcal{I}_s and $\mathcal{C}_{s,all}$ keep decreasing; Record n_s that results in the smallest \mathcal{I}_s and $\mathcal{C}_{s,all}$;
- 5 When $\mathcal{C}_{s,all}$ begin to rise as n_s increases, prune this branch and switch to next station;
- 6 **while** $\mathcal{C}_{all} \leq \mathcal{B}$ **do**
- 7 Try to pre-allocate one more battery group for each station s and calculate the \mathcal{I}_s and $\mathcal{C}_{s,all}$;
- 8 Choose station s that leads to maximum $Gain$ and still keeps the correspondingly calculated $\mathcal{C}_{all} \leq \mathcal{B}$;
- 9 Do add one battery group for station s and update the \mathcal{C}_{all} and \mathcal{I} ;
- 10 **return** n_s for all the station s ;

the additional backup power helps the base station sustain long power outages and reduce deep discharging of batteries. So the battery lifetime is prolonged, achieving a lower battery replacement cost. In the second stage, if we continue to increase the battery group number, the extra backup power becomes redundant due to enough power supply. Then the service interruption time remains unchanged or decreases very little, while the battery replacement cost increases due to the unavoidable battery degradation process. Note that the lead-acid battery itself has a self-aging process so that too many battery groups will lead to a high average replacement cost (due to the self-aging process). Therefore, there may exist three situations for battery replacement cost according to different conditions of the corresponding base station, i.e., the cost first drops and then rises (both stage 1 and stage 2), the cost keeps decreasing (only stage 1), and the cost keeps increasing (only stage 2). As the sum of battery replacement cost and emergent maintenance cost, the overall cost can also have this characteristic when the battery replacement cost dominates, which is future verified by our real data-driven experiments in §6.

Based on the above analysis, it is easy to see that the two objectives in our model are conflicting and multiple Pareto optimal solutions may exist. Considering the practical situation of telecommunication industry, the most important objective for telecom operators is to provide more reliable cellular communication services. So we utilize a lexicographic method [24] to solve this problem. The lexicographic method assumes that the objectives can be ranked in the order of importance, and repeatedly solve the most important objective by fixing other less important objectives with a bound. We first consider minimizing the service interruption time when the overall cost has an upper limit \mathcal{B} . Then we strive to minimize the overall cost without increasing the service interruption time.

The designed heuristic algorithm is shown in Algorith-

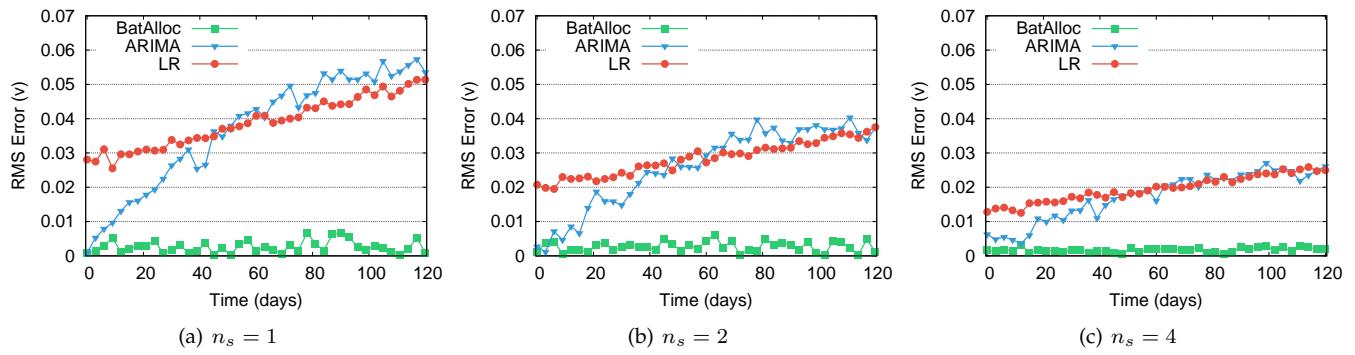


Fig. 16. RMS errors on the voltage estimation by three models under different battery group number.

1, where we divide the solving process into two stages. In the first stage (line 1-5), for each base station we keep increasing the battery group number until the overall cost begins to rise. We thus stop and record the battery allocation results in the first stage. The first stage of our allocation algorithm has the following property as shown in Theorem 1.

Theorem 1: The allocation result of the first stage is optimal with the minimum possible budget constraint.

Proof: Based on our analysis, in the first stage of allocation, both \mathcal{I}_s and $\mathcal{C}_{s,all}$ are monotonously decreasing as we increase the battery group number for each base station. So the two optimization objectives are currently not conflicting. The allocation result of the first stage has the minimum overall cost in any case, because a lower budget is not sufficient for normal management of all the base stations and batteries. The allocation results in the first stage thus must be optimal under the same budget limit. \square

In the second stage (line 6-9), the two objectives are conflicting because the battery replacement cost begins to rise. As aforementioned, we consider reducing the service interruption time when the overall cost does not exceed the budget limit. To better balance the tradeoff between them, we define *Gain* as the ratio of the weighted service interruption decrease and the overall cost increase:

$$Gain = \frac{\mathcal{I}_s(n_s) - \mathcal{I}_s(n_s + 1)}{\mathcal{C}_{s,all}(n_s + 1) - \mathcal{C}_{s,all}(n_s)} \quad (12)$$

We each time select the base station with the maximum *Gain* and add one battery group to it until we reach the budget. By utilizing such a greedy approach we guarantee to reduce the most service interruption time with the least cost increase for each step.

We next analyze the complexity of our heuristic algorithm to show its efficiency. In the first stage, we only access each station once and the complexity is $O(n)$ where n is the total number of the base stations. In the second stage, we calculate $\mathcal{C}_{s,all}$ and iteratively select station with the maximum *Gain*, which contributes the complexity of $O(n \log(n))$. So the total complexity of this heuristic algorithm is $O(n \log(n))$.

6 EVALUATION

In this section, we present the evaluation of our BatAlloc framework based on real trace-driven experiments. We first evaluate our battery feature profiling process and compare

our model with commonly used time series estimation methods, such as ARIMA [19] and Linear Regression (LR) [20]. Based on the base station and battery profiling results, we present the performance evaluation on the overall BatAlloc framework.

6.1 Experiment Setup

We conduct data processing on our dataset from China Mobile and extract useful features on base stations and backup batteries. We process the massive data on our workstation as illustrated in Fig. 17, including dual Gigabyte AORUS GeForce GTX 1080 Ti Xtreme Edition 11GB Video Card, dual Intel i7-6850K BROADWELL-E Processor 6 Core 15M Cache 3.6GHz CPU, Corsair Dominator Platinum 32GB 2x16GB DDR4 3000MHz Memory Kit, Samsung 850 EVO 1TB SATA 3 Solid State Drive, and etc. We construct the deep learning based model using Keras [25], which is a neural network library on top of TensorFlow [26] and Theano [27].

The parameter settings of our experiments are extracted from our dataset as well as adapted from the typical settings based on the domain knowledge. The normal float voltage is 2.25v and the plateau discharging voltage v_P is set as 2.08v for a new battery cell. According to the industry standard, the battery used in cellular communication base station is designed to provide power supply for about 10 to 12 hours and we thus set τ to 10. The second low voltage disconnect (LVD) of base stations is usually set as 1.8v, and we set the end voltage v_E as 1.85v to avoid extreme deep level discharge. In our experiments, we set the importance factor ω_s based on the population that a base station covers, which is normalized to a value between 0 and 1. According to the real world market [28], we set the price of battery cost c_b as \$5000. The emergent maintenance cost c_m mainly consists of two parts: labor cost for a visit and diesel consumption cost for power generation. The labor cost is calculated as $c_e(t_d + t_m)$, where c_e is the average wage for engineers per hour (\$30/hr with at least two engineers [29]), t_d is the time spent on road (calculated by location information), and t_m is the emergent maintenance time at a base station. The diesel cost is computed by $c_d t_m$, where c_d is the diesel cost per hour and we set c_d as \$7.6/hr [2].

6.2 Experiment on Voltage Prediction

We first evaluate the performance of our deep learning based battery profiling model. We use the data of the first



Fig. 17. Our work station for data processing.

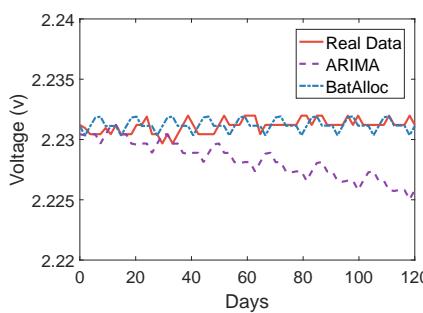


Fig. 18. Voltage prediction of a battery using different predicting methods.

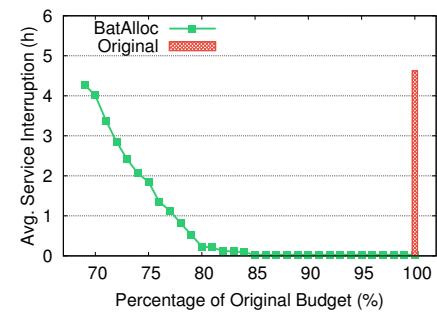


Fig. 19. The average overall service interruption duration on different budget limit. X-axis is based on the percentage of original budget.

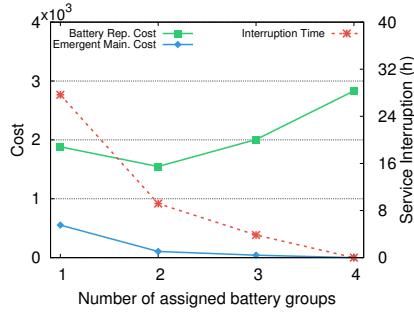


Fig. 20. Various metrics for a typical base station when equipped with different numbers of battery groups.

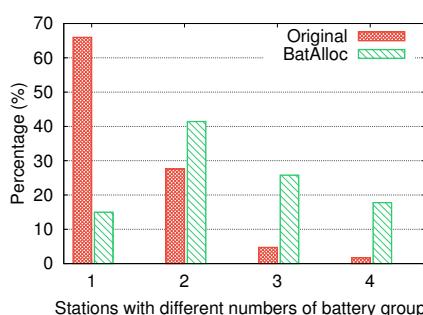


Fig. 21. The percentage of base stations with different numbers of battery groups.

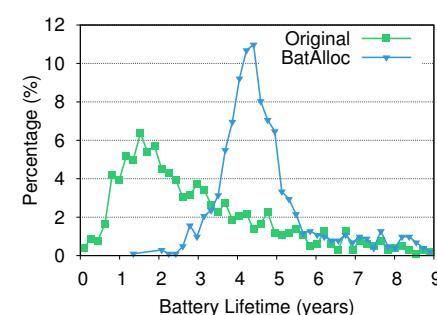


Fig. 22. The comparison of battery lifetime in original allocation and lifetime after re-allocation based on BatAlloc.

365 days as the training set and the data of the next 120 days as the testing set. And we compare our model with ARIMA and LR. In our evaluation, the LR and ARIMA method only capture the time series features of the past voltages, e.g., the voltage trend and the voltage variance, and use the captured features to predict the future change of battery voltages. We use a multiplicative seasonal ARIMA model to learn the voltage features considering the variance of the wave and estimate the parameters based on the Matlab economics toolbox automatically.

Fig. 16 shows the root mean square (RMS) errors between the estimation results by the three models and the actual voltage data. We can see that under various numbers of battery groups, the deep learning based model used in the BatAlloc framework can always achieve better accuracies with the RMS error less than $0.008v$. This means our deep learning based model can effectively capture the influences that different events exert on battery conditions. ARIMA and LR only extract the features of time series from the voltages and make corresponding estimations. Although their RMS errors are relatively small due to the stationary voltage trend when n_s is large, they become worse when n_s is small. Integrated with the accurate estimation on the future voltage trend and the domain knowledge of battery features, we can then obtain the battery lifetime and reserve time used for the battery allocation optimization in the BatAlloc framework, which will be evaluated next.

We then select one representative battery as a case study example to illustrate the voltage trend using different prediction methods as shown in Fig. 18. We can see that our approach captures the voltage varying trend more accurate-

ly and the predicted voltage is relatively close to the real data, while the ARIMA approach fails to capture the voltage varying trend. This is because ARIMA focuses on extracting the internal time series features of battery voltages, and does not consider the impacts of external events on battery working conditions.

6.3 Experiment on Battery Allocation

We next evaluate our BatAlloc framework on battery allocation results. For comparison, we extract the current battery deployment as a baseline from the real world dataset and use the *Original* allocation to represent it. Fig. 19 plots the annual average service interruption time with different budget limit \mathcal{B} . For ease of comparison, the budget limit is normalized by the baseline budget (i.e., 100% means the budget limit is equal to 100% of the original baseline budget). The minimum budget we need is 69% of baseline, which is actually the allocation result of the first stage in our optimization algorithm. Even with the 69% of baseline budget, our framework can still achieve a lower average service interruption compared to the original allocation, which is at least a 30% cost saving. The service interruption time drops observably as the budget limit increases and we can achieve nearly full service guarantee with only 85% of the baseline budget. These results demonstrate that our BatAlloc framework is capable of providing much more reliable service with a remarkably reduced cost.

To better understand the impact of different battery group numbers on base stations, we conduct a case study shown in Fig. 20, which plots our different metrics for a

typical base station when equipped with different numbers of battery groups. As the number of battery groups increases from 1 to 4, the emergent maintenance cost and the service interruption time decrease monotonously due to more sufficient backup power. On the other hand, the battery replacement cost achieves minimum when the number of battery groups is 2 since the additional battery group can drastically reduce the impact of overdischarging and prolong the battery lifetime. If we keep increasing battery groups, the extra battery power continues to reduce the service interruption time, while the battery replacement cost rises largely mostly due to the unavoidable battery degradation process.

Fig. 21 compares the different allocation results on battery group number between the original allocation scheme and our BatAlloc framework. The original battery allocation result is largely skewed that over 65% base stations are equipped with only one battery group. Our framework considers both the base station situations and battery features, allocating 2 battery groups to most base stations and 3 or 4 battery groups to those with long time power outages.

We also investigate the impact of different battery allocation strategies on battery lifetime. As shown in Fig. 22, in the original allocation the average battery lifetime is only around 1.5 years and far less than expected. After using BatAlloc to allocate suitable numbers of battery groups for base stations, the average battery lifetime has achieved to 4.3 years, roughly 1.8 times longer than that of the original allocation. The results indicate that our framework can also better protect base station batteries and significantly prolong their average lifetimes.

7 FURTHER DISCUSSION

Though each single power outage of one given base station is truly hard to predict precisely, the statistical long-term power outage trends (e.g., in every year) can have a very similar pattern (e.g., a base station built in cold area may suffer from several power outages due to the heavy snow every year). In this paper, restricted by the 1.5-year timespan of the data, the long-term characteristics of the power outages for each base station might still not be able to be captured with very high accuracy. In practical application, however, the service provider can have data for as long as tens of years and such data can be mined for better power outage prediction. From this perspective, our framework can still apply well in the practical battery allocation as long as the statistical long-term power outage trends can be predicted with more abundant data.

Different base stations may set different low voltage disconnect (LVD) value according to their practical situations. A high LVD value may potentially help extend the battery lifetime by avoiding deep discharging but will increase the service interruption time and result in high service interruption cost. Thus in reality, most mobile service providers would set LVD to a quite low value given the service interruption cost can usually be higher compared to the cost introduced by shortened battery lifetime due to deep discharge. In this paper we focus on studying the relationship between the power outage and the battery lifetime duration with the consideration of allocating multiple battery groups.

The trade-off between the setting of LVD, battery lifetime and service interruption cost can be an interesting future work for further exploring.

8 CONCLUSION

Current cellular communication base stations are facing serious problems due to the mismatch between the power outage situations and the backup battery supporting abilities. In this paper, we proposed BatAlloc, a battery allocation framework to address this issue. We first conducted a systematical analysis of a massive dataset of base stations and batteries. Then we built up a deep learning based model to precisely capture the battery conditions and further profile the battery features. With the profiling results, we formulated this battery allocation issue as a multi-objective optimization problem and designed an efficient algorithm to solve it. Our real trace-driven experiments showed that compared to the current practical deployment, our framework can remarkably reduce the service interruptions as well as the overall cost. It is worth noting that although our battery profiling model focuses on lead-acid batteries, the general allocation framework can still be well applied to other chemical battery scenarios (e.g., Li-ion batteries) once the battery degradation aspect is updated.

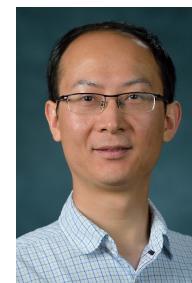
REFERENCES

- [1] F. Wang, F. Wang, X. Fan, and J. Liu, "Batalloc: Effective battery allocation against power outage for cellular base stations," in *Proceedings of the Eighth International Conference on Future Energy Systems*, pp. 234–241, ACM, 2017.
- [2] "Powering Your Cell Towers." http://www.westell.com/assets/MKTG-rA-Powering-Your-Cell-Tower_White-Paper.pdf, 2016/1/19. Accessed: 2017-01-20.
- [3] W. Balshe, "Power system considerations for cell tower applications," *Cummins Power Generation*, 2011.
- [4] "Power outage in southwestern connecticut." http://www.nj.com/news/index.ssf/2010/03/storm_brings_power_outages_roa.html, 2010. Accessed: 2017-01-20.
- [5] M. A. Marsan, G. Bucalo, A. Di Caro, M. Meo, and Y. Zhang, "Towards zero grid electricity networking: Powering bss with renewable energy sources," in *Communications Workshops (ICC), IEEE International Conference on*, pp. 596–601, IEEE, 2013.
- [6] "The power outage situations in okanaganlake area." http://www.okanaganlakebc.ca/community/personal/westside_road/bc_hydro_power.htm, 2017. Accessed: 2018-02-10.
- [7] F. Wang, C. Xu, L. Song, Q. Zhao, X. Wang, and Z. Han, "Energy-aware resource allocation for device-to-device underlay communication," in *2013 IEEE International Conference on Communications (ICC)*, pp. 6076–6080, IEEE, 2013.
- [8] H. Holtkamp, G. Auer, S. Bazzi, and H. Haas, "Minimizing base station power consumption," *IEEE Journal on Selected Areas in Communications*, vol. 32, no. 2, pp. 297–306, 2014.
- [9] R. Ramamonijs and V. K. Bhargava, "Energy allocation and cooperation for energy-efficient wireless two-tier networks," *IEEE Transactions on Wireless Communications*, vol. 15, no. 9, pp. 6434–6448, 2016.
- [10] V. Chamola, B. Sikdar, and B. Krishnamachari, "Delay aware resource management for grid energy savings in green cellular base stations with hybrid power supplies," *IEEE Transactions on Communications*, 2016.
- [11] J. Gong, J. S. Thompson, S. Zhou, and Z. Niu, "Base station sleeping and resource allocation in renewable energy powered cellular networks," *IEEE Transactions on Communications*, vol. 62, no. 11, pp. 3801–3813, 2014.
- [12] K. Kutluay, Y. Cadirci, Y. S. Ozkazanc, and I. Cadirci, "A new online state-of-charge estimation and monitoring system for sealed lead-acid batteries in telecommunication power supplies," *IEEE Transactions on Industrial Electronics*, vol. 52, no. 5, pp. 1315–1327, 2005.

- [13] A. H. Anbuky and P. E. Pascoe, "Vrla battery state-of-charge estimation in telecommunication power systems," *IEEE Transactions on Industrial Electronics*, vol. 47, no. 3, pp. 565–573, 2000.
- [14] P. E. Pascoe and A. H. Anbuky, "Vrla battery discharge reserve time estimation," *IEEE transactions on power electronics*, vol. 19, no. 6, pp. 1515–1522, 2004.
- [15] M. Coleman, C. K. Lee, C. Zhu, and W. G. Hurley, "State-of-charge determination from emf voltage estimation: Using impedance, terminal voltage, and current for lead-acid and lithium-ion batteries," *IEEE Transactions on industrial electronics*, vol. 54, no. 5, pp. 2550–2557, 2007.
- [16] B. S. Bhangu, P. Bentley, D. A. Stone, and C. M. Bingham, "Nonlinear observers for predicting state-of-charge and state-of-health of lead-acid batteries for hybrid-electric vehicles," *IEEE Transactions on Vehicular Technology*, vol. 54, no. 3, pp. 783–794, 2005.
- [17] P. E. Pascoe and A. H. Anbuky, "The behaviour of the coup de fouet of valve-regulated lead-acid batteries," *Journal of power sources*, vol. 111, no. 2, pp. 304–319, 2002.
- [18] I. Kurisawa and M. Iwate, "Capacity estimating method of lead-acid battery by short-time discharge-introduction of approximate expression of discharge curve obtained by compound expression of straight line and hyperbola," in *Telecommunications Energy Conference, 1997. INTELEC 97., 19th International*, pp. 483–490, IEEE, 1997.
- [19] A. Charnes, E. Frome, and P.-L. Yu, "The equivalence of generalized least squares and maximum likelihood estimates in the exponential family," *Journal of the American Statistical Association*, vol. 71, no. 353, pp. 169–171, 1976.
- [20] G. E. Box, G. M. Jenkins, and G. C. Reinsel, "Linear nonstationary models," *Time Series Analysis, Fourth Edition*, pp. 93–136, 1976.
- [21] X. Fan, F. Wang, and J. Liu, "On backup battery data in base stations of mobile networks: Measurement, analysis, and optimization," in *Proceedings of the 25th ACM International Conference on Information and Knowledge Management*, pp. 1513–1522, ACM, 2016.
- [22] W. D. Reeve, *DC power system design for telecommunications*, vol. 14. John Wiley & Sons, 2006.
- [23] A. H. Anbuky, P. E. Pascoe, and R. G. Lane, "Vrla battery capacity measurement and discharge reserve time prediction," in *Telecommunications Energy Conference, 1998. INTELEC. Twentieth International*, pp. 302–310, IEEE, 1998.
- [24] C.-L. Hwang and A. S. M. Masud, *Multiple objective decision making-methods and applications: a state-of-the-art survey*, vol. 164. Springer Science & Business Media, 2012.
- [25] F. Chollet *et al.*, "Keras." <https://github.com/fchollet/keras>, 2015.
- [26] M. Abadi, A. Agarwal, P. Barham, E. Brevdo, Z. Chen, C. Citro, G. S. Corrado, A. Davis, J. Dean, M. Devin, *et al.*, "Tensorflow: Large-scale machine learning on heterogeneous distributed systems," *arXiv preprint arXiv:1603.04467*, 2016.
- [27] Theano Development Team, "Theano: A Python framework for fast computation of mathematical expressions," *arXiv e-prints*, vol. abs/1605.02688, May 2016.
- [28] "Market price of 48v battery groups." <https://gb-battery.com/view-products?olsPage=t%2F48-volt-batteries>. Accessed: 2018-02-10.
- [29] "Labor cost on stationary engineers." http://www.payscale.com/research/CA/Job=Stationary_Engineer/Hourly_Rate, 2017. Accessed: 2017-01-20.



Xiaoyi Fan (S'14) received the B.E. degree from Beijing University of Posts and Telecommunications, Beijing, China, in 2013, and the M.Sc. degree from Simon Fraser University, Canada, in 2015. He is now a Ph.D. student in the School of Computing Science, Simon Fraser University, British Columbia, Canada. His areas of interest are Internet-of-Things, big data and deep learning.



Feng Wang (S'07-M'13) received both the Bachelor's degree and Master's degree in Computer Science and Technology from Tsinghua University, Beijing, China in 2002 and 2005, respectively. He received the PhD degree in Computing Science from Simon Fraser University, Burnaby, British Columbia, Canada in 2012. He is currently an Assistant Professor in the Department of Computer and Information Science at the University of Mississippi, University, MS, USA. He is a member of IEEE and a recipient of IEEE ICME Quality Reviewer Award (2011). He is a Technical Committee Member of Elsevier Computer Communications. He served as Program Vice Chair in International Conference on Internet of Vehicles (IOV) 2014, and as TPC co-chair in IEEE CloudCom 2017 for Internet of Things and Mobile on Cloud track. He also serves as TPC member in various international conferences such as IEEE INFOCOM, ICPP, IEEE/ACM IWQoS, ACM Multimedia, IEEE ICC, IEEE GLOBECOM and IEEE ICME.



Jiangchuan Liu (S'01-M'03-SM'08-F'17) received B.Eng. (Cum Laude) from Tsinghua University, Beijing, China, in 1999, and Ph.D. from The Hong Kong University of Science and Technology in 2003, both in Computer Science. He is currently a Full Professor (with University Professorship) in the School of Computing Science at Simon Fraser University, British Columbia, Canada. He is an IEEE Fellow and an NSERC E.W.R. Steacie Memorial Fellow.

He is a Steering Committee Member of IEEE Transactions on Mobile Computing, and Associate Editor of IEEE/ACM Transactions on Networking, IEEE Transactions on Big Data, and IEEE Transactions on Multimedia. He is a co-recipient of the Test of Time Paper Award of IEEE INFOCOM (2015), ACM TOMCCAP Nicolas D. Georganas Best Paper Award (2013), and ACM Multimedia Best Paper Award (2012).



Fangxin Wang received the B.S. degree from the Department of Computer Science of Technology, Beijing University of Post and Telecommunication, Beijing, China in 2013 and the M.S. degree in the Department of Computer Science and Technology, Beijing, China in 2016. He is currently pursuing the Ph.D. degree in the School of Computing Science, Simon Fraser University, Burnaby, B.C., Canada. His research interests include wireless networks, big data analysis and machine learning.