

A Li-ion battery charge protocol with optimal aging-quality of service trade-off

Original

A Li-ion battery charge protocol with optimal aging-quality of service trade-off / Chen, Yukai; Bocca, Alberto; Macii, Alberto; Macii, Enrico; Poncino, Massimo. - ELETTRONICO. - (2016), pp. 40-45. (Intervento presentato al convegno 2016 International Symposium on Low Power Electronics and Design tenutosi a San Francisco Airport, CA, USA nel August 8-10,2016) [10.1145/2934583.2934591].

Availability:

This version is available at: 11583/2647385 since: 2020-02-24T12:21:14Z

Publisher:

ACM

Published

DOI:10.1145/2934583.2934591

Terms of use:

This article is made available under terms and conditions as specified in the corresponding bibliographic description in the repository

Publisher copyright

(Article begins on next page)

A Li-Ion Battery Charge Protocol with Optimal Aging-Quality of Service Trade-off

Yukai Chen Alberto Bocca Alberto Macii Enrico Macii Massimo Poncino

Politecnico di Torino, Corso Duca degli Abruzzi 24, 10129 Torino, Italy

{yukai.chen, alberto.bocca, alberto.macii, enrico.macii, massimo.poncino}@polito.it

ABSTRACT

The reduction of usable capacity of rechargeable batteries can be mitigated during the charge process by acting on some stress factors, namely, the average state-of-charge (SOC) and the charge current. Larger values of these quantities cause an increased degradation of battery capacity, so it would be desirable to keep both as low as possible, which is obviously in contrast with the objective of a fast charge.

However, by exploiting the fact that in most battery-powered systems the time during which it is plugged for charging largely exceeds the time required to charge, it is possible to devise appropriate *charge protocols* that achieve a good balance between fast charge and aging.

In this paper we propose a charge protocol that, using an accurate estimate of the charging time of a battery and the statistical properties of the charge/discharge patterns, yields an optimal trade-off between aging and quality of service. The latter is measured in terms of the distance of the actual SOC from 100% at the end of the charge phase. Results show that the present method improves significantly over other similar protocols proposed in the literature.

1. INTRODUCTION

The aging of a Li-Ion battery, intended as the loss of usable capacity over time, depends on quantities such as temperature, depth of discharge, average state of charge (SOC), and charge/discharge current [1,2]. While the values of these quantities during the discharge phase cannot be controlled because they are determined by the user habits, the charge phase offers some margin for their control.

Firstly, battery charge is usually constrained to some standardized schemes with pre-defined current and voltage charge profiles. In the case of Li-ion batteries, the charge is based on the Constant Current-Constant Voltage (CC-CV) protocol, characterized by a well-defined charge process that cannot be altered. The CC-CV constraint is motivated by cost (it requires simple hardware to be implemented) and by safety reasons. While the CC-CV constraint may appear as a limitation, it has some degrees of freedom (namely, the charge current) that can be used to mitigate the aging.

Secondly, users do control the *charge patterns*; devices are often connected to a plug (*plug-in time*) for a time much larger than the time needed to charge the battery (*charge time*). This condition impacts on the average SOC of the battery, which affects the aging [1]. There exist therefore an evident trade-off between (fast) charge time and battery aging: the former requires larger charge currents, which, however, degrade battery capacity in two ways: directly, because aging depends on charge current; indirectly, because long plug-in times imply a larger average SOC in the battery. Previous works have explored this trade-off without however taking all the required variables and constraints into account. In [3,4] only charge current or average SOC are considered, respectively. In [5], both quantities are appropriately taken into account, but the

analysis is limited to a single cycle for which the actual plug-in time is assumed to be known. Finally, [6,7] propose solutions that do not stick to the CC-CV scheme.

In this work we propose a CC-CV compliant charge protocol, which, by taking into account **all** the relevant parameters, calculates an optimal charge current based on a simple (and thus easily implementable in hardware) prediction of the plug-in time, in order to achieve a “just-in-time” charge (i.e., with the smallest average SOC and with the smallest possible current). While the key parameters of the protocols are borrowed from [5], the introduction of a prediction requires the introduction of a novel Quality of Service (QoS) metric, calculated as the deviation from the 100% charge level. In this way we can define an aging/QoS space in which the trade-off between these two quantities can be explored.

Results in this Aging/QoS space show that the minimum current just-in-time protocol that we propose (called ASAP), provides the best trade-off under various user charge/discharge pattern statistics. It is worth mentioning that the proposed protocol is general and not tied to a particular type of battery-powered device. Although we present some data relative to smartphones, the same protocol is obviously suitable also for hybrid electric vehicles (HEVs), where aging is a primary metric due to the relative cost of the energy storage devices.

2. BACKGROUND AND RELATED WORK

2.1 Battery charging

Charging a battery is an operation that has a great impact on battery lifetime, even more than discharging [2]. Selecting the appropriate charge protocol is thus a critical step, for both limiting the drift in battery performance and for avoiding dangerous side effects such as overtemperature and overcharging. The above mentioned CC-CV protocol is considered the de-facto charging protocol for Li-Ion batteries. It consists of two phases. In the first phase (CC), the battery is charged at a constant current until its voltage reaches a pre-determined limit; in the second phase (CV), the battery is charged at a constant voltage until the current drops to a pre-defined value.

The need to adhere to a standard protocol clearly rules out the many complex schemes suggested for both aging-aware slow and fast charge (e.g., [2,6–8]). However, it is still possible to stick to CC-CV and alter some of its characteristic parameters (namely, the charging current and its distribution over time) in order to reduce the degradation of batteries [8]. Such minor variants of the basic CC-CV protocol would not violate its main electrical and safety properties, nor would significantly increase its cost.

2.2 Battery aging

Life degradation of rechargeable batteries depends on four main factors: (i) temperature, (ii) depth of discharge (DOD) at each cycle, (iii) average state of charge (SOC_{avg}) [1,9], and (iv) charge/dis-

charge current [1, 5]. Aging worsens with an increase of any of these quantities.

While temperature can be assumed as more or less constant during charge and DOD is meaningful only for discharge, the other two parameters can be managed during the charge process. As a general rule-of-thumb, the charge phase should ideally reach 100% by the end of the plug-in time - this will yield the smallest average SOC [4]. Some previous efforts for aging-aware charging schemes have tried to achieve this objective in different ways. However, batteries are often left in plug-in mode although fully charged, even for a long time. Besides energy waste, this has a huge impact on battery degradation, which is strongly affected by the SOC_{avg} [1]. These may endanger the battery *SOH* (defined as ratio of the capacity of an aged battery with respect to the nominal capacity) and consequently accelerate its permanent degradation. whose value, for a generic time interval from t_0 to t_1 , is given by the following expression:

$$SOC_{avg} = \int_{t_0}^{t_1} SOC(t)dt / (t_1 - t_0) \quad (1)$$

For instance, [9] reports that, for $LiFePO_4$ batteries, SOC should be less than 60% in average for maintaining battery life acceptable. In [3], it mitigates battery aging by considering only the charge current, by calculating a minimum current that ensures a fully charged battery at the end of a predicted plug-in time whenever it is greater than the charge time needed in standard CC-CV mode. However the non-linear relation *charge current vs. charging time* [7] required an analysis. On the other hand, [4] proposed a delayed start time for charging batteries as late as possible in order to minimize the SOC_{avg} only. [5] both charge current and SOC_{avg} were analyzed, reporting that the aging-optimal charge current is more related to battery usage rather than plug-in time, and that *capacity loss vs. charge current* characteristic is not a monotonic function. In [10], the authors proposed a method for automatically quantifying battery aging in mobile devices without using external equipment, by considering charge time and focusing on the middle region of the battery level where the plug-in time and SOC have a nearly linear characteristic in the CC phase.

Aging is usually evaluated through the aggregate metric of the State-of-Health (SOH), defined as the ratio of the capacity of an aged battery and the nominal capacity. In this work, since we focus on multiple charge cycles, we need a model that expresses the aging for a single cycle. To this purpose we use the classical model of [9] augmented as in [5] to account for charge and discharge current. Battery degradation (L) in the i -th cycle is determined by the following expression:

$$L_i = L_{0,i} \cdot e^{(K_{ic,i} \cdot I_{ch,i} + K_{id,i} \cdot I_{dis,i})}$$

where $L_{0,i}$ is the degradation factor provided by the Millner's model [9], which accounts for temperature, DOD, and SOC_{avg} in the i -th cycle, while the second term accounts for the charge/discharge current ($I_{ch,i}$ and $I_{dis,i}$) in that cycle; $K_{ic,i}$ and $K_{id,i}$ are empirical coefficients extracted from datasheet information [5] or experimental data. By summing L_i over M cycles we get the total loss of capacity L_M . L_M and SOH are both normalized, so they are simply related as $SOH = 1 - L_M$.

CC/CV requires relatively simple hardware to be implemented: the charge begins with a constant current charge to some voltage level, at which point the voltage caps and the current begins to decrease until some current level is reached. Besides simplicity, CC/CV is also driven by safety reasons, since it effectively manages the risk of overcharging, which is quite dangerous in Li-Ion batteries.

that accounts for all relevant parameters, namely, average SOC, deviation of the SOC (here considered as depth-of-discharge), discharge/charge current, and temperature (e.g., [9]), (b) an analytical macro-model of CC-CV charge time based on a subset of these quantities [11]. Thanks to the use of models, it is also possible to explore the overall parameter space and identify the charging solution that best fits, from both battery aging and QoS point of view, a specific user behavior that determines the discharge patterns.

3. METHODOLOGY

3.1 Characterization of the Charge Period

Figure 1 shows a generic charge-discharge cycle of a typical device, and, specifically, the evolution of the battery SOC over time. Since our focus is on the charge phase, in the figure the discharge time has been compressed to better emphasize the charge phase.




Figure 1: Generic Charge-Discharge Cycle.

Starting from some value SOC_{eoc} (end-of-charge) at the end of the previous cycle ¹ the discharge will evolve according to the user activity on the device and will reach a value SOC_{eod} (end-of-discharge) at some point in time T_{charge}^i (i -th charge cycle) when the user will start charging the device. The difference $\Delta SOC = SOC_{eoc} - SOC_{eod}$ represents the depth of discharge DOD_i of the i cycle. The charge curve identifies the conventional CC-CV charge implemented by typical chargers: as soon as the device is plugged, the standard CC-CV protocol described in Section 2 is applied; in the diagram, we assume that the CC and CV phase lasts t_{cc} seconds and t_{cv} , respectively. Let $T_{charge} = t_{cc} + t_{cv}$. It is clear that if the time the device is plugged in (the *plug-in time*) exactly matches T_{charge} (e.g., the user polls the charge termination and disconnects as soon as it is charged) there is no better way to charge: the device will be 100% charged in the smallest possible time. Needless to say, this situation seldom occurs, and in most case the battery is 100% charged well before the device is unplugged, i.e. $T_{plug-in}$ is in general much larger than T_{charge} . In these cases, the battery stays plenty of charge for an amount of time equal to $T_{plug-in} - T_{charge}$ (the time t_{stdby} in the figure). Given the dependency of aging on the average SOC, this is obviously detrimental for the aging of the battery. This is particularly critical for mobile devices or HEVs, where charging frequently occurs overnight.

3.2 Alternative Charge Protocols

Whenever plug-in time exceeds charge time, better solutions are possible, as shown in Figure 2.

The options lie between two extremes:

- Delay the charge as much as possible, (dotted curve), using the same nominal current as the standard charge so that the charge ends just in time when the device is plugged out. We called this schedule **ALAP**.

¹ SOC_{eoc} will not necessarily be 100%




Figure 2: Alternative Charge Protocols.

- The other extreme (solid curve) starts the charge as soon as the device is plugged in and modulates the current so as to reach 100% charge at plug-out time. This implies clearly using a smaller current during the CC phase. We called this schedule **ASAP**.

Any scheme between these two extremes will use some delay in the charge and a current $I_{cc,min} < I_{cc} < I_{cc,nom}$. The ALAP scheme obviously yields the smallest SOC (a smaller area below the curve), whereas ASAP uses the smallest possible charge current. Which one is best for aging depends on the relative weight of the two quantities (SOC vs. charge current) on battery life degradation. Therefore, it is possible to derive, for a given $T_{plug-in}$, an aging-optimal schedule (i.e., by [5]) with a profile that lies in between ASAP and ALAP. These four CC-CV based protocols (Standard, ASAP, ALAP and the Aging-optimal) will be compared in this work.

3.3 Effect of Estimated Plug-In Times

The comparative analysis of the charge protocols would be a relatively straightforward exercise if we focus on a single cycle and rely on the knowledge of the plug-in time $T_{plug-in}$; this was the analysis done in [5], where all the $T_{plug-in}$ were assumed to be known. However, the exact plug-in time is not generally known in advance; for this reason we improve the algorithm in [5] by including a predictor. In fact, if $T_{plug-in}$ in Figure 1 is actually an *estimate* of the actual plug-in time, the analysis becomes non-trivial due to the intrinsic inaccuracy on the estimation. Two are the main consequences of this inaccuracy, which we discuss in the following sections.

3.3.1 Quality of Service

One first impact of imprecise prediction is that the protocol might not succeed in fully charging the battery. In other terms, we need to include in the analysis the user perception of the **quality** of the charge process. This quality of service (QoS) is intuitively measured as the percentage of charge at plug-out time. QoS is obviously in contrast with aging: any “fast” charge that aims at increasing the chance of full charge will stress the battery both in terms of high charge current and average SOC. In terms of QoS alone, it is immediate to see that, as a rule-of-thumb, early-starting protocols will reach faster SOC close to the 100% even in the case of under-predicted plug-in times. Therefore, for a given value of $T_{plug-in}$, the standard protocol will provide the best quality, followed by the ASAP charge; the aging-optimal one will lie between ASAP and ALAP, which yields the worst charge quality. However, depending on the accuracy of the estimates and the actual length of plug-in time, the difference between the standard protocols and the other ones might result in being negligible, achieving similar QoS while keeping the above mentioned benefits in terms of aging. Because it yields a higher SOC in the cases where the actual plug-in time is

shorter than the average: the ASAP SOC curve always lies above the ALAP one. However, the difference is mainly evident at the beginning (corresponding to large underestimations of $T_{plug-in}$, whereas the two curves become indistinguishable towards the end of the charge. Again, the optimal QoS/aging tradeoff might lie between these two extremes. (and the CV time is normally larger than the CC one). An accurate evaluation of these effects requires an appropriate simulation framework and the relative models in order to determine the best charge protocol, for a set of representative “workload” models, i.e., user plug-in patterns.

3.3.2 Impact of Mispredictions

From the previous section, it is evident that the actual ranking among the various protocols in the QoS/aging space depend on the accuracy of the plug-in time estimates. Figure 3 shows the impact of mispredictions on both QoS and aging.

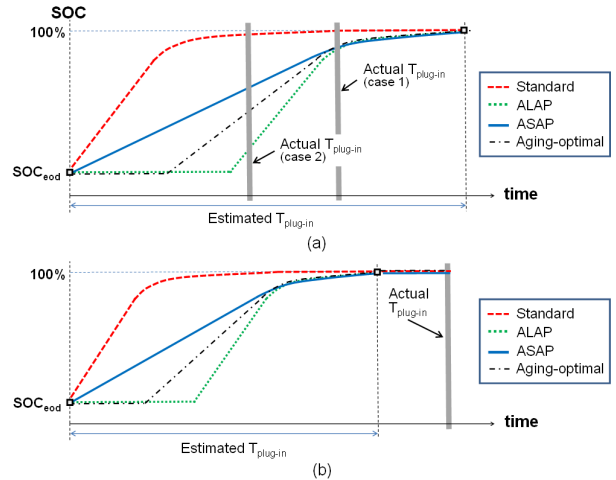


Figure 3: Impact of Mispredictions: Overprediction (a) and Underprediction (b).

The case of an overprediction (Figure 3-(a)), i.e., the predicted $T_{plug-in}$ is longer than the actual one, is the most critical case. If the overprediction is modest (Case 1), the impact is limited; all the considered protocols follow the CC-CV template, such misprediction will likely occur during the CV phase, which takes a significant portion of the overall charge time but contributes only to approximately 10% of the total SOC for a standard charge current [7]. In these cases, mispredictions will therefore result in an moderately incomplete charge. However, when the overprediction is sizable (Case 2), as a limit case, the ALAP protocol might even reach a charge close to 0%, which is of course not acceptable. As already discussed, in general early-starting protocols such as the standard and ASAP tend to be less sensitive to the accuracy prediction. The case of underprediction (Figure 3-(b)), conversely, is less critical. It has no impact on the QoS but only on the aging, in terms of (i) an extra period of time in which the battery stays fully charged, and (ii) an unnecessarily large current used to charge it. From the above discussion, it is evident that the most important requirement for the predictor is *to avoid large overpredictions*. Section 3.5 will discuss the possible strategies to implement the predictions.

3.4 Metrics

In order to consistently explore the QoS/aging space, we define normalized metrics for the two quantities. Concerning aging, the metric is naturally tracked by the battery $SOH = 1 - L_M$, as already discussed in Section 2,

which measures the overall capacity loss. The aging model of Section 2 accounts for multiple cycles and is also normalized, so we can use $SOH = 1 - L(m)$ as a normalized aging factor, where m is a reference number of cycles.

QoS, as already mentioned, is intuitively related to the percentage of full charge achieved in the charge process. That is, the user ideally wants to achieve 100% in every cycle. Although we want a normalized metric, we can envision two slightly different variants of quality metric:

- *Average Quality* Q_1 , defined as the average (over M charge cycles) deviation from the full charge condition. Defining $\Delta SOC_c^i = 1 - SOC_{eoc}^i$ for cycle i , Q_1 is obtained by averaging ΔSOC_c^i over the M cycles:

$$Q_1 = 1 - \frac{\sum_{i=1}^M \Delta SOC_c^i}{M}$$

Q_1 is a normalized term; it equals 1 only if 100% of the cycles reach 100% charge. Therefore, it might be smaller than 1 even for standard charging. Notice that it suffices to consider average SOC in the metric since $\Delta SOC_c^i \geq 0$ (i.e., it is not possible to overcharge), that is, there is no substantial difference between two sequences of ΔSOC_c^i with same average and different standard deviations.

- *Average Maximum Error* Q_2 ; this metric is parameterized with respect to a value ϵ , representing the maximum desired deviation from the full charge. It is simply defined as the percentage of charge cycles in which SOC_{eoc} is within the bound, i.e.:

$$Q_2(\epsilon) = \frac{\# \text{ of cycles } SOC_{eoc}^i > (1 - \epsilon)}{M}$$

Q_2 is also a normalized quantity; it equals 1 only if 100% of the cycles reach a $100 - \epsilon$ charge. Although Q_2 is still an average metric, it emphasizes the cycle-by-cycle quality rather than the overall average quality represented by Q_1 .

For the exploration of the aging/QoS tradeoff, we will assess the various protocols in a bi-dimensional space SOH/Q , where Q is either Q_1 or $Q_2(\epsilon)$.

3.5 Plug-In Time Prediction

As discussed in Section 3.3.2, accurate predictions are key for the definition of a protocol with an acceptable QoS. Predicting the future values of a quantity (in our case, plug-in time) based on its history falls under the widely studied statistical problem of *time series prediction*, which has applications basically in any discipline. Solutions to this problem range from simple regressions to state-space models and learning-based solutions based on neural networks [12]. In our specific context, the prediction of plug-in times has both constraints and specificities that push towards simple prediction strategies. The *constraints* are related to the overhead of the implementation of the prediction; computationally expensive algorithms imply a relevant energy overhead and force the use of a software implementation on the device, and prevent thus the use of a pure hardware, device-independent solution (e.g., in the charger). The *specificities* concern the fact that the charge process for popular battery-powered “devices” such as mobile devices and HEVs exhibit some structure that allow a statistical a priori characterization of the plug-in times. The latter tend in fact to follow two well defined patterns [13, 14]:

- *Unimodal charge*: A typical value (usually between 6 and 8 hours) with relatively narrow tails (± 2 hrs), that can be

statistically described by a bell-shaped distribution such as a Gaussian or a log-normal one; this case is representative of a nightly charge pattern (Figure 4-(a)).

- *Bimodal charge*: two typical values (one around 2 to 3 hours, the other as in the unimodal case), which can be described by a bimodal distribution, representative of a mix of a shorter daily charge and a longer nightly charge (Figure 4-(b)).




Figure 4: Unimodal (a) and Bi-Modal (b) Charge Patterns.

In the unimodal case, a very simple predictor that track the average value of like the moving average (simple or exponential) performs well enough. In this case, in fact, the average value of the distribution is representative of the universe as it denotes the most typical value. For the bimodal case, where the average is not representative, a different mechanism is required. In the example of Figure 4-(b), for instance, the average will lie around 350-400 minutes, a value with a very low probability of occurrence. One possibility is that of fitting the data to the bimodal distribution using a maximum likelihood estimator (e.g., [4]), a relatively expensive computation. In our case, we leverage another property of the charge patterns, that is, the dependence of charge duration on the time of day. It has been observed that the majority of short charge cycles occur during daytime, whereas the long charges occur during night time. This is conceptually shown in Figure 5-(a): most charge events lie in the dark shaded areas, while only a few outliers fall outside them (light shaded areas).

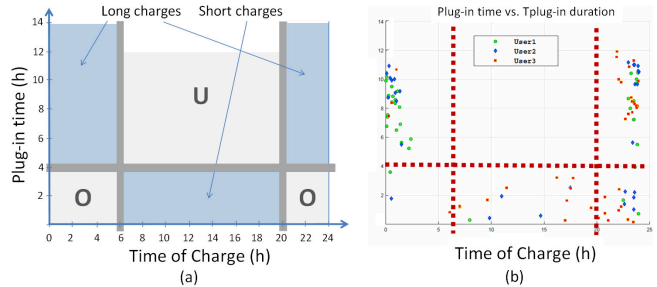


Figure 5: Time vs. Duration of Plug-In: Conceptual regions (a) and an Example Pattern of three Users (b).

The outliers lying in the region denoted with “U” correspond to an underprediction (long charges during daytime), which we saw being less critical for QoS. Those in the region denoted with “O” are instead overpredictions (short charges during nightly hours), which affect QoS more seriously. An example falling in this region for a portable device could be a short charge after working hours and before going out at night time. Figure 5-(b) shows the example of 30 days of charge events relative to three users for a smartphone. We notice that the events map quite well to the expected two-region pattern; in fact, similar data and distribution absolutely comply with those reported in [4, 15]. Generally, the plug-in and charging time are limited to a few hour during daytime, while they are much longer during the night.

Based on these observations our plug-in time predictor uses the information about the time of day and the history as follows. We keep two separate moving averages MA_{day} and MA_{night} for daily and nightly history, respectively; each MA averages the last five plug-in times in the respective hourly intervals defining “day” and “night”. These intervals have been defined by calibration based on the behavior of different users: “day” is between 6am and 7pm, whereas “night” covers the remaining hours. The prediction of the next plug-in time is simply obtained as $T_{plug-in} = MA_{day}$ if the start of charge occurs during the “day”, and $T_{plug-in} = MA_{night}$ otherwise. This simple scheme provides quite good accuracy resulting in a QoS comparable to that of the regular CC-CV charge with the nominal current.

4. SIMULATION RESULTS

4.1 Simulation Setup

We implemented a simulator of battery cycles in Matlab that evaluates the battery life by considering the relevant aging parameters. It can generate synthetic charge/discharge cycles according to predefined distributions of their relevant parameters (DOD, discharge time, and plug-in time).

For the discharge phase, we consider three different battery usage patterns corresponding to three DOD values: C (conservative, with an average DOD = 25%), M (moderate, average DOD = 50%), and A (aggressive, average DOD = 75%). Since the focus of this paper is on the charge phase, the discharge current profile is always the same for a given DOD value (A,M,C) in order to avoid that the results are affected by the discharge phase thus altering the evaluation of the charge protocols. The simulator can generate $T_{plug-in}$ traces according to Gaussian and bimodal distributions.

For each charge phase, we calculate the aging and QoS for the charge protocols discussed in Section 3 (**Standard**, **Aging-Optimal**, **ALAP**, and **ASAP**) according to the estimated $T_{plug-in}$. The SOC at the end of the charging phase SOC_{eoc} is used as initial SOC for the next discharge cycle.

The battery cell used in our experiments is the A123 Systems Li-Ion ANR26650m1; it has a 2.5Ah nominal capacity and a 3.3V nominal voltage. The operating temperature is set to 35°C. We extract the parameters for building the aging model from its datasheet [16].

4.2 Prediction Accuracy

A preliminary experiment is concerned with the accuracy of the Gaussian and bimodal predictor. Figure 6 shows the distribution of the prediction error for the $T_{plug-in}$ distributions. The figure shows that the predictor achieves good performance for both distributions: 64% (84%) and 68% (87%) of the predictions are within one hour (two hours) error for the Gaussian and bimodal distribution, respectively.

4.3 Charge Protocols Comparison

4.3.1 Impact of DOD

Our first experiment shows how different discharging profiles affect the final aging and QoS results. We generated four traces: three for each of the A, M, C profiles, and an extra one (“mixed”) containing an equal amount of A, M, C cycles. All the traces have the same $T_{plug-in}$ with bimodal distribution in the charging phase. The simulation cycle life in this experiment is 1000 cycles, equivalent to about a 2-year lifetime. Due to lack of space, in this experiment we show only results for the metric Q_1 . Q_2 has an almost identical layout in the SOH/QoS space, in which (1, 1) represents the ideal point where the aging has the smallest value and there is a 100% charge quality.




Figure 6: Error Distribution of the Predictor.




Figure 7: SOH vs. Q_1 with Different DOD Types.

Figure 7 illustrates the effect of the different DOD levels. The **Standard** protocol always reaches the best quality value (near 1), but it has the worst SOH (due to large average SOC). Expectedly, larger discharge depth exacerbates the aging. **ALAP** and **Aging-optimal** exhibit better SOH thanks to a decreased SOC_{avg} ; however, since the sequence of $T_{plug-in}$ values is the same for the four DOD profiles, as DOD increases the SOC at the start of a charge also decreases, and therefore a larger number of charge events does not achieve a 100% SOC. **ASAP** exhibits the best trade-off, as it is the closest to (1,1); moreover, it is also almost insensitive to DOD: all the points are very clustered.

4.3.2 Impact of $T_{plug-in}$ Distribution

A second experiment aims at showing how different $T_{plug-in}$ distributions affect the aging and QoS metrics when considering different charge protocols. We set as discharging profile mixed (A,M,C) DOD trace in order to average out different user behaviors. Five traces of 1000 cycles were simulated for both Gaussian and bimodal distributions.

Figures 8 and 9 show the results for Q_1 and $Q_2(10\%)$, respectively. Again, **ASAP** yields the optimal SOH/QoS trade-off under both $T_{plug-in}$ distributions, consistently with the results of Figure 7. Q_2 appears to be more demanding than Q_1 as a quality metric for **ALAP** and **Aging-Optimal**, which exhibit a quite low QoS. The **Standard** charge exhibits a 50% loss of capacity, while **ASAP** results in only a 20% loss of capacity with a marginal loss in quality (1% to 3% depending on the chosen metric). Expectedly, the




Figure 8: SOH vs. Q_1 for Different $T_{plug-in}$ Distributions.




Figure 9: SOH vs. Q_2 with Different $T_{plug-in}$ Distributions.

Aging-Optimal is the best in terms of aging (about 15% loss), but it achieves this at the price of a $Q_1 < 90\%$ and $Q_2 < 75\%$. Concerning sensitivity to the distribution, for **ASAP** it is only slightly larger than the **Standard** one, yet much smaller to the sensitivity of the other two protocols.

4.3.3 Results On Real Traces

The third experiment refers to the real traces from three users (U1, U2, and U3) logged from a smartphone. The simulation time here is only about 40 days, so the aging is less sizable; however, the results are consistent with those obtained with longer synthetic traces. Figure 10 shows the results for both Q_1 and Q_2 (10%), for the three users and the four charge protocols, visually grouped in the figure for the sake of clarity. The results are consistent with the previous section: **ASAP** provides the best trade-off, with a very similar quality to **Standard** and a marginally lower SOH than **Aging-Optimal**, while **ALAP** yields to worst tradeoff. Notice again that Q_2 is always a more stringent metric than Q_1 .

5. CONCLUSIONS

Battery aging is a phenomenon which must be kept under control because it affects the usable capacity.

In this paper we propose a battery charge protocol that, using an accurate estimate of the charging time of a device and the statistical characteristics of the discharge patterns, can reach an optimal trade-off between aging and QoS, by improving significantly over other similar protocols proposed in the literature.

The results in this Aging/QoS space show that the minimum current




Figure 10: SOH vs. Q_1 & Q_2 with Real Traces.

just-in-time protocol we proposed, called **ASAP**, provides the best trade-off under various user charge/discharge pattern statistics.

6. ACKNOWLEDGMENTS

This work was supported in part by the EC JTI ARTEMIS AR-ROWHEAD project under Grant Agreement No. 332987.

7. REFERENCES

- [1] J. Vetter, P. Novák, M. R. Wagner, C. Veit, K.-C. Möller, J. O. Besenhard, M. Winter, M. Wohlfahrt-Mehrens, C. Vogler, and A. Hammouche. Ageing mechanisms in lithium-ion batteries. *Journal of Power Sources*, 147(1-2):269–281, September 2005.
- [2] S. Bashash, S.J. Moura, J.C. Forman, and H.K. Fathy. Plug-in hybrid electric vehicle charge pattern optimization for energy cost and battery longevity. *Journal of Power Sources*, 196(1):541–549, January 2011.
- [3] N. Matsumura, N. Otani, and K. Hamaji. Intelligent battery charging rate management. *US Patent App. 12/059,967*. US 2009/0243549 A1, 1 October 2009.
- [4] A. Pröbstl, P. Kindt, E. Regnath, and S. Chakraborty. Smart2: Smart Charging for Smart Phones. In *RTCSA'15 Conference Proceedings*, pages 41–50. IEEE, August 2015.
- [5] A. Bocca, A. Sassone, A. Macii, E. Macii, and M. Poncino. An Aging-Aware Battery Charge Scheme for Mobile Devices Exploiting Plug-in Time Patterns. In *IEEE ICCD'15 Conference Proceedings*, pages 407–410. IEEE, October 2015.
- [6] R. Klein, N.A. Chaturvedi, J. Christensen, J. Ahmed, R. Findeisen, and A. Kojic. Optimal Charging Strategies in Lithium-Ion Battery. In *ACC'11 Conference Proceedings*, pages 382–387. IEEE, 2011.
- [7] Z. Guo, B.Y. Liaw, X. Qiu, L. Gao, and C. Zhang. Optimal charging method for lithium ion batteries using a universal voltage protocol accommodating aging. *Journal of Power Sources*, 274:957–964, January 2015.
- [8] W. Shen, T.T. Vo, and A. Kapoor. Charging Algorithms of Lithium-Ion Batteries: an Overview. In *ICIEA'12 Conference Proceedings*, pages 1567–1572. IEEE, July 2012.
- [9] A. Millner. Modeling Lithium Ion Battery Degradation in Electric Vehicles. In *CITRES'10 Conference Proceedings*, pages 349–356. IEEE, September 2010.
- [10] J. Lee, Y. Chon, and H. Cha. Evaluating Battery Aging on Mobile Devices. In *DAC'15 Conference Proceedings*, pages 1–6. ACM, June 2015.
- [11] D. Shin et al. , Alessandro Sassone, Alberto Bocca, Alberto Macii, Enrico Macii, Massimo Poncino A compact macromodel for the charge phase of a battery with typical charging protocol. *2014 International Symposium on Low Power Electronics and Design (ISLPED)*. In *ISLPED'14 Conference Proceedings*, pages 267–270. ACM, August 2014.
- [12] P. J. Brockwell. *Introduction to Time Series and Forecasting*, 2nd Edition. Springer, 2010.
- [13] D. Ferreira, A. K. Dey, and V. Kostakos. Understanding human-smartphone concerns: a study of battery life. In *Pervasive 2011 Conference Proceedings*, pages 19–33. Springer-Verlag Berlin Heidelberg, 2011.
- [14] A. Duran, A. Ragatz, K. Kelly, and K. Walkowicz. Characterization of In-Use Medium Duty Electric Vehicle Driving and Charging Behavior. In *IEVC'14 Conference Proceedings*, pages 1–8. IEEE, December 2014.
- [15] E.A. Oliver and S. Keshav. An empirical approach to smartphone energy level prediction. In *UbiComp'11 Conference Proceedings*, pages 345–354. ACM, September 2011.
- [16] A123 Systems, Inc. High Power Lithium Ion ANR26650m1-B. *Technical datasheet*, 2013.