DUK
OSI

WHITE
PAPER

Sensitivity of Lithium-ion Battery SOC and SOH Estimates to Sensor Measurement Error and Latency

Sensitivity of Lithium-Ion Battery SOC and SOH Estimates to Sensor Measurement Error and Latency[1]

Prepared by Dr. Gregory L. Plett
Battery Control LLC, Colorado Springs, CO 80920
In association with Dukosi

Abstract

Highly accurate and highly confident estimates of state of charge (SOC) and state of health (SOH) are crucial prerequisites to maximizing the performance and safety achieved from a battery pack. Neither SOC nor SOH are directly measurable, so algorithms of varying complexity and computational cost are employed to provide estimates of the values. Required inputs to the algorithms are the measurable quantities of the cell, whose performances are defined by precision, accuracy and synchronicity. This white paper provides an evaluation of the impact on the performance of SOC and SOH estimation based on the integrity of these inputs through execution of model-based simulation. It considers typical usage scenarios in electric-vehicle and ESS applications, cell chemistry, estimation method and measurement performance. The cell measurements under examination are: cell temperature, cell voltage and cell current. The cell-model derivation and methods used in SOC and SOH estimation are described.

1 Introduction

Battery-management systems must be able to compute estimates of state of charge (SOC) and state of health (SOH). There exist many methods that create these estimates, each having its own characteristics. However, the methods commonly rely on measurements of battery-cell voltage, current, and temperature. We seek to address two questions in this white paper:

1. *How do SOC estimates depend on sensor total measurement error and synchronization error when using coulomb counting or sigma-point Kalman filters to make the estimates?*
2. *How do SOH estimates depend on sensor total measurement error and synchronization error when using total-least-squares methods to make the estimates*?

We use a simulation-based approach. First, we create a synthetic "truth" dataset by simulating an equivalent-circuit model (ECM) for different simulation scenarios. Then, we add measurement and synchronization errors to the truth outputs. This modified dataset is used as input to the SOC-estimation algorithms. The output of these algorithms is used as input to the SOH methods, which seek to estimate cell total capacity and equivalent-series resistance (ESR). We evaluate results by comparing the estimates of SOC and SOH to the truth values from the original simulation.

In this white paper, Sect. 2 describes the simulation approach in detail. Sect. 3 shares the different SOC-estimation algorithms and their characteristics and Sect. 4 presents SOC-estimation results. Sects. 5 and 6 share the different SOH-estimation algorithms and Sect. 7 presents SOH-estimation results. Finally we offer some summary observations.

[1]This white paper expands on the results reported in: [1, 2].

2 Simulation approach

This section presents our approach in more detail. Since the study depends on a synthetic data set, we first describe the mathematical lithium-ion battery-cell model used to create this data set. Then, we present the simulation scenarios that are used to evaluate the effect of different levels of sensor measurement error and latencies on the algorithms.

2.1 The enhanced-self-correcting (ESC) cell model

The synthetic datasets were generated using simulations of cells described using an "enhanced-self-correcting" (ESC) ECM [3]. Fig. 1 illustrates this model.

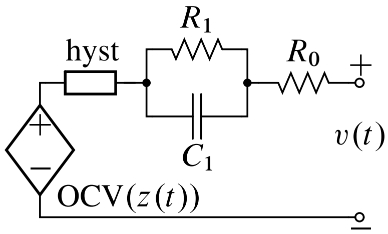

Fig. 1: An ESC-type ECM having a single resistor–capacitor "Voigt" network.

In the ESC model, cell SOC $z(t)$ is modeled as:

$$\frac{\mathrm{d}z(t)}{\mathrm{d}t} = -\frac{\eta(t)i(t)}{3600Q},$$

where $\eta(t)$ is coulombic efficiency ($\eta(t) = 1$ on discharge and $\eta(t) \leq 1$ on charge) and Q is cell total capacity measured in ampere-hours. In this sign convention, discharge current is positive and charge current is negative. Capacitor voltages are modeled as:

$$\frac{\mathrm{d}v_c(t)}{\mathrm{d}t} = -\frac{1}{\tau}v_c(t) + \frac{R}{\tau}i(t),$$

where τ is the R–C time constant in seconds and R is the value of the resistor in the R–C pair in ohms. Dynamic hysteresis fraction is modeled as:

$$\frac{\mathrm{d}h(t)}{\mathrm{d}t} = -\left|\frac{\eta(t)i(t)\gamma}{3600Q}\right|(h(t) + \mathrm{sgn}(i(t))),$$

where γ is a hysteresis decay-rate factor in seconds. Instantaneous hysteresis is modeled as:

$$s(t) = \begin{cases} -\mathrm{sgn}(s(t)), & |i(t)| > 0; \\ s(t-\varepsilon), & \text{otherwise.} \end{cases}$$

Cell voltage is then predicted as:

$$v(t) = \mathrm{OCV}(z(t)) + Mh(t) + M_0 s(t) - v_c(t) - R_0 i(t),$$

where OCV$(\cdot)$ is the open-circuit voltage of the cell at a given SOC and where R_0 is the cell's ESR. Note that the model can have multiple R–C branches ("Voigt" networks) if we implement individual ODEs for each capacitor voltage and if we subtract the sum of capacitor voltages when computing cell voltage.

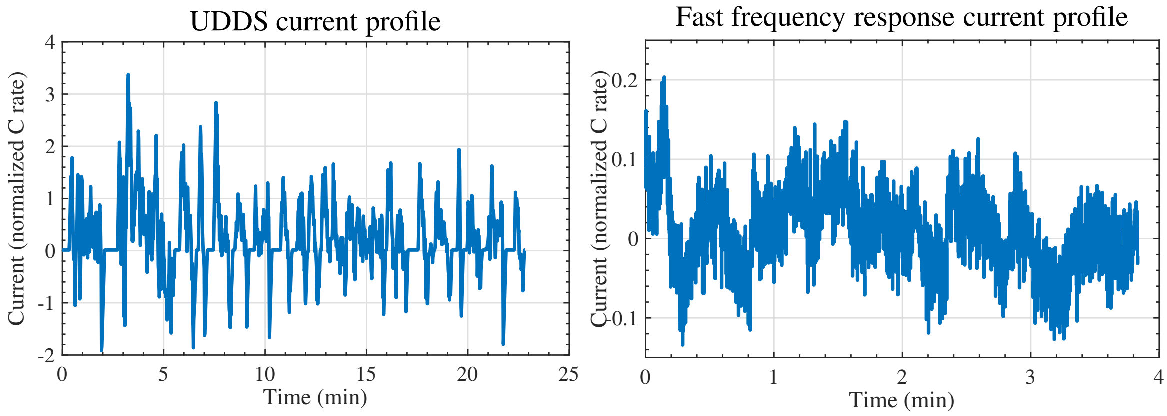

Fig. 2: UDDS drive profile (left); FFR profile (right).

2.2 Simulation-study framework

Cell type SOC-estimate quality depends on the chemistry and design of the cells. Lithium-iron-phosphate (LFP) cells are commonly used for stationary storage and are increasingly being used in automotive applications. From an algorithm perspective, they have a flat OCV relationship and significant hysteresis, which can make estimating SOC well very challenging. Nickel-manganese-cobalt-oxide (NMC) cells are commonly used in automotive and other applications requiring high energy density. The models we used in this study were generated using data collected in the laboratory from A123 LFP cells and Panasonic NMC cells using the methods presented in [3].

Scenarios Automotive applications are represented using an urban dynamometer driving schedule (UDDS) profile for both LFP and NMC cells. Energy-storage applications are represented using a fast frequency response (FFR) profile for LFP cells. For brevity, all of the results presented here were for truth cell temperatures of 25 °C.

The left frame of Fig. 2 presents a single normalized UDDS current profile. In the simulations, the true initial cell SOC was 95 %. This profile was repeated 10 times to discharge the cell to a final SOC of 5 %.[2] The right frame of the figure presents a single normalized FFR current profile. The true initial cell SOC was 50 % and the profile was repeated 94 times for a total duration somewhat longer than 6 h. The FFR profile is nearly charge-neutral, so the final SOC was approximately 44 %.

Cases Eight test cases investigated the sensitivity of results to algorithm initialization and parameterization. We considered cases where cell SOC was incorrectly initialized by ± 1 % and where cell capacity was incorrectly initialized by ± 1 %. The synthetic dataset was generated using ESC models having three Voigt networks and the SOC-estimation algorithms were implemented using ESC models of the same cells using models having two Voigt networks to account for realistic truth/model mismatch. The simulation cases also considered different levels of sensor dc bias, noise, and timing latency between the voltage and current measurements.

[2]This is quite a wide operating window. If we were to use a narrower window—for example from 90 % to 10 %, we would expect that the root-mean-squared errors on SOC estimates would be lower (it is easier to estimate SOC over a narrower operating range) but that RMSE on total-capacity estimates would be higher (it is important to move SOC as much as possible between updates to a total-capacity estimate). This study did not investigate the tradeoffs in detail.

Table 1: Simulation-study parameter settings.

Test parameter	Parameter settings
Scenarios	{UDDS for NMC and LFP; FFR for LFP}
Temperature	{25 °C (ambient)}
Cases	{z_0 initialization error of ±1 %}, {Q initialization error of ±1 %}, {truth data simulated using model having 3 R–C pairs and SPKF SOC estimates made based using model having 2 R–C pairs} The eight cases of Table 2 span these conditions.
Sensor models	dc bias and random error specified for seven sensors in Table 3.
Sensor latencies	{0 ms, ±100 ms} between voltage and current measurements.
SOC algorithm	{Coulomb counting, SPKF}
SOH algorithm	{WLS and AWTLS-based methods}

Table 2: The eight pseudo-random simulation cases.

Category \ Case	1	2	3	4	5	6	7	8
Init. SOC-est. error	−1 %	−1 %	−1 %	−1 %	+1 %	+1 %	+1 %	+1 %
Capacity estimate	$Q/0.99$	$Q/0.99$	$Q/0.99$	$Q/0.99$	$0.99Q$	$0.99Q$	$0.99Q$	$0.99Q$
Current-sensor bias	− max	− max	− max	− max	+ max	+ max	+ max	+ max
Voltage-sensor bias	+ max	+ max	− max	− max	+ max	+ max	− max	− max
Temp.-sensor bias	+ max	− max	+ max	− max	+ max	− max	+ max	− max

SOC-estimation algorithm The quality of SOC estimates depends on the estimation algorithm used. Of the plethora of available algorithms [4], we select coulomb counting (a crude but commonly used approach) and sigma-point Kalman filters (SPKF, a very good and generally robust approach).

2.3 Implementing noise-free simulation scenarios

Table 1 summarizes the simulation-study parameter settings. Table 2 itemizes the eight simulation cases considered for each scenario. Table 3 lists the assumed sensor specifications.

For each combination of scenario and case from Tables 1 and 2, we simulated an ESC model using noise-free inputs: we discretized the ESC model equations using a sampling frequency of 100 Hz [3] and implemented them in Simulink.

3 SOC estimation methods

3.1 Coulomb counting

In discrete time, cell SOC can be expressed as:

$$z[k] = z[0] - \frac{\Delta t}{3600Q} \sum_{j=0}^{k-1} i_{\text{app}}[k],$$

Table 3: Sensor TME settings for the simulation study.

Sensor	Label	Parameter	Vehicle setting	ESS setting
Voltage	SV1	dc bias	0.35 mV	0.35 mV
		noise 1σ	142 μV	142 μV
	SV2	dc bias	3.5 mV	3.5 mV
		noise 1σ	1.5 mV	1.5 mV
Current (C-rate)	SI1	dc bias	0.284×10^{-3}	0.071×10^{-3}
		noise 1σ	0.130×10^{-3}	0.0325×10^{-3}
	SI2	dc bias	4.2×10^{-3}	1.05×10^{-3}
		noise 1σ	0.07×10^{-3}	0.0175×10^{-3}
	SI3	dc bias	18×10^{-3}	4.5×10^{-3}
		noise 1σ	24×10^{-3}	6×10^{-3}
Temp.	ST1	dc bias	0.4 °C	0.4 °C
		noise 1σ	0.013 °C	0.013 °C
	ST2	dc bias	5 °C	5 °C
		noise 1σ	0.013 °C	0.013 °C

if we assume that coulombic efficiency is perfect, where $i_{\text{app}}[k]$ is the true applied current and Δt is the sampling period. Coulomb counting estimates cell SOC as:

$$\hat{z}[k] = \hat{z}[0] - \frac{\Delta t}{3600\widehat{Q}} \sum_{j=0}^{k-1} i_{\text{meas}}[k],$$

where $\hat{z}[0]$ is the initial SOC estimate, $\widehat{Q}$ is the total-capacity estimate, and $i_{\text{meas}}[k] = i_{\text{app}}[k] + i_{\text{bias}}[k] + i_{\text{noise}}[k]$. SOC-estimation error can then be expressed as:

$$\begin{aligned}\tilde{z}[k] &= z[k] - \hat{z}[k] \\ &= (z[0] - \hat{z}[0]) - \left(\frac{1}{Q} - \frac{1}{\widehat{Q}}\right)\left(\frac{\Delta t}{3600} \sum_{j=0}^{k-1} i_{\text{app}}[k]\right) + k\frac{i_{\text{bias}}\Delta t}{3600\widehat{Q}} + \frac{\Delta t}{3600\widehat{Q}} \sum_{j=0}^{k-1} i_{\text{noise}}[k], \quad (1)\end{aligned}$$

assuming that the dc-bias error is constant. The first term is the SOC-estimation error due to an offset in the initial SOC estimate—this offset is never corrected since coulomb counting does not have a feedback mechanism that might be able to make a correction. The second term is a slope error due to an incorrect estimate of cell total capacity, which is also never corrected but tends to cancel out over a charge-neutral cycle (error grows as a cell is discharged or charged, but then decreases as the cell is subsequently charged or discharged, respectively). The third term is a linear function of time that grows due to a current-sensor dc bias and is never corrected. The final term is a zero-mean error caused by measurement noise; its actual value will grow and shrink randomly, but its variance (uncertainty) increases linearly over time (assuming that noises are uncorrelated). Because of the random sensor noise, every execution of the coulomb-counting method will give a different result; however, we can compute the standard deviation of the noise-induced SOC-estimate error as:

$$\sigma_{\tilde{z}_{\text{noise}}}[k] = \frac{\sigma_i \cdot \Delta t \cdot \sqrt{k}}{3600\widehat{Q}}, \quad (2)$$

Table 4: SOC-estimation results—using coulomb counting and SPKF—for the 13 scenarios and three latencies considered in this work.

						Latency = −100 ms			Latency = 0 ms			Latency = 100 ms		
SI	SV	ST	Profile	Cell(s)	RMSE CC (%)	RMSE KF (%)	RMSB KF (%)	BNDX KF (%)	RMSE KF (%)	RMSB KF (%)	BNDX KF (%)	RMSE KF (%)	RMSB KF (%)	BNDX KF (%)
1	1	1	UDDS	NMC	1.54	0.27	0.56	0.00	0.27	0.56	0.00	0.27	0.56	0.00
			UDDS	LFP	1.54	0.38	1.29	0.00	0.38	1.29	0.00	0.38	1.29	0.00
			FFR	LFP	1.05	0.65	1.06	0.00	0.65	1.06	0.00	0.65	1.06	0.00
1	2	2	UDDS	NMC	1.54	0.63	1.57	0.00	0.63	1.57	0.00	0.63	1.57	0.00
			UDDS	LFP	1.54	1.32	2.38	0.00	1.32	2.38	0.00	1.32	2.38	0.00
			FFR	LFP	1.05	1.05	1.13	0.00	1.05	1.13	0.00	1.05	1.13	0.00
2	1	1	FFR	LFP	1.36	1.31	1.76	0.00	1.31	1.76	0.00	1.31	1.76	0.00
2	2	2	FFR	LFP	1.36	1.33	1.85	0.00	1.33	1.85	0.00	1.33	1.85	0.00
3	1	1	UDDS	NMC	5.43	0.28	0.57	0.00	0.28	0.57	0.00	0.28	0.57	0.00
			UDDS	LFP	5.43	4.82	8.94	0.00	4.82	8.94	0.00	4.82	8.94	0.00
3	2	2	UDDS	NMC	5.43	0.77	1.78	0.00	0.77	1.78	0.00	0.77	1.78	0.00
			UDDS	LFP	5.43	2.21	7.11	0.00	2.21	7.11	0.00	2.22	7.11	0.00
			FFR	LFP	2.51	1.99	4.14	0.00	1.99	4.14	0.00	1.99	4.14	0.00

where σ_i is the standard deviation of the current-sensor noise.

Note that only the current-sensor noise specifications are factors in the coulomb-counting error. Coulomb counting does not use sensed voltage and no variables in the coulomb-counting expression are functions of temperature.

3.2 Sigma point Kalman filter

Kalman filters implement a model-based approach to estimating the state of a dynamic system. When the system whose state is being estimated is nonlinear, the steps of the Kalman filter must be adjusted to approximate the intent of the original Kalman filter for the nonlinear model. One approach to doing so uses an extended Kalman filter (EKF) [5, 6]. An alternate approach that can give better results for highly nonlinear models is the sigma-point Kalman filter (SPKF) [7, 6]. Note that unscented Kalman filters (UKF) and cubature Kalman filters (CKF) are specific examples of the more general SPKF. In this work, we implemented a central-difference Kalman filter (CDKF), which is also a specific example of SPKF.

The SPKF repeatedly executes two substeps every measurement interval. First, it uses the ESC model equations to predict the model state and cell voltage. Then, it compares predicted to measured voltage and adjusts its state estimate based on this feedback. The details of the SPKF are presented in the above references, but the critical observation is that voltage feedback is used to improve the estimates versus those produced by coulomb counting, which does not incorporate feedback. Whenever the measured voltage contains quality information regarding cell SOC, we would expect SPKF to outperform coulomb counting. But, when the cell model or sensed voltage are poor, or when the cell voltage is a weak function of SOC, we might not expect dramatic improvements.

One feature of SPKF is that the algorithm automatically computes a 3σ confidence bound on its SOC estimates. We consider the estimator to be robust if the true SOC lies within this confidence bound with high certainty.

4 SOC estimation results and discussion

4.1 Generating the noisy datasets

Noise-free truth datasets were simulated for each of the three basic scenarios listed in Table 1. Then, the SOC estimators were executed for each of the eight simulation cases of Table 2 and three latency levels for each of these scenarios and a subset of possible sensor combinations. To do so, pseudo-random Gaussian noises were added to the current, voltage, and temperature per the sensor specifications listed in Table 3. The "max" (and "min") entries in Table 2 refer to adding (or subtracting) the dc-bias value defined by the particular sensor's specification to (from) the true value.

Cases 1–4 bias the initial SOC estimate 1 % above the true value, bias the capacity at $\approx$ 1 % above its true value, and bias the current-sensor in the charging direction. These have the effect of trying to force the SOC estimate to be greater than the truth at all times (since all profiles start in the discharge direction). Cases 5–8 bias the estimate in the other direction. These three variables are always changed together to provide $\pm$ worst-case results.

Sensor timing latency is assumed to impact only the voltage signal. If the latency is negative, then changes to cell voltage precede changes to current and SOC (the voltage signal is effectively shifted left compared with the other signals). If the latency is positive, then the voltage changes lag behind changes to current and SOC (the voltage signal is effectively shifted right compared with the other signals). Since all of the latencies we considered are integer multiples of the sampling period (which is 10 ms), these shifts are easily accomplished without resimulating the synthetic datasets.

This modified dataset is used as input to the SOC-estimation algorithms. We evaluate results by comparing the estimated SOC to the truth SOC from the original simulation.

4.2 Tuning the SPKF

Note that the SPKF method requires specifying the initial-state covariance matrix $\Sigma_{\tilde{x},0}$, process-noise covariance matrix $\Sigma_{\widetilde{w}}$, and sensor-noise covariance matrix $\Sigma_{\widetilde{v}}$. The process of optimizing these variables is often called 'tuning the filter.' If we constrain $\Sigma_{\tilde{x},0}$ to be diagonal, then it has four elements that must be chosen (the initial uncertainties of the two R–C states, the hysteresis state, and the SOC state). Each of $\Sigma_{\widetilde{w}}$ and $\Sigma_{\widetilde{v}}$ have one element to be chosen. In total, six 'tuning values' must be selected. Different sets of tuning values will work best for different applications, so these six values were optimized independently for each of the simulation scenarios and sensor-package combinations. (Thirteen sets of tuning values were optimized for the results presented in this white paper, corresponding to the thirteen rows of Tables 4–6.)

When tuning the SPKF via an optimization routine, the input data must be consistent in order for the objective function being optimized also to be consistent. Since the sensor configurations involve random noise, we must somehow convert that randomness into a deterministic set of battery-cell input/output data to be applied to the SPKF while optimizing its tuning. We chose to optimize the tuning across the eight pseudo-random 'simulation cases,' applied to each operating-cycle and temperature combination, as listed in Table 2. In all cases, pseudo-random Gaussian noises are added to the current, voltage, and temperature per the specifications of the sensors (the random-number-generator seed was reset before input data for the first case were computed, so that the same pseudo-random values would be generated for each simulation run).

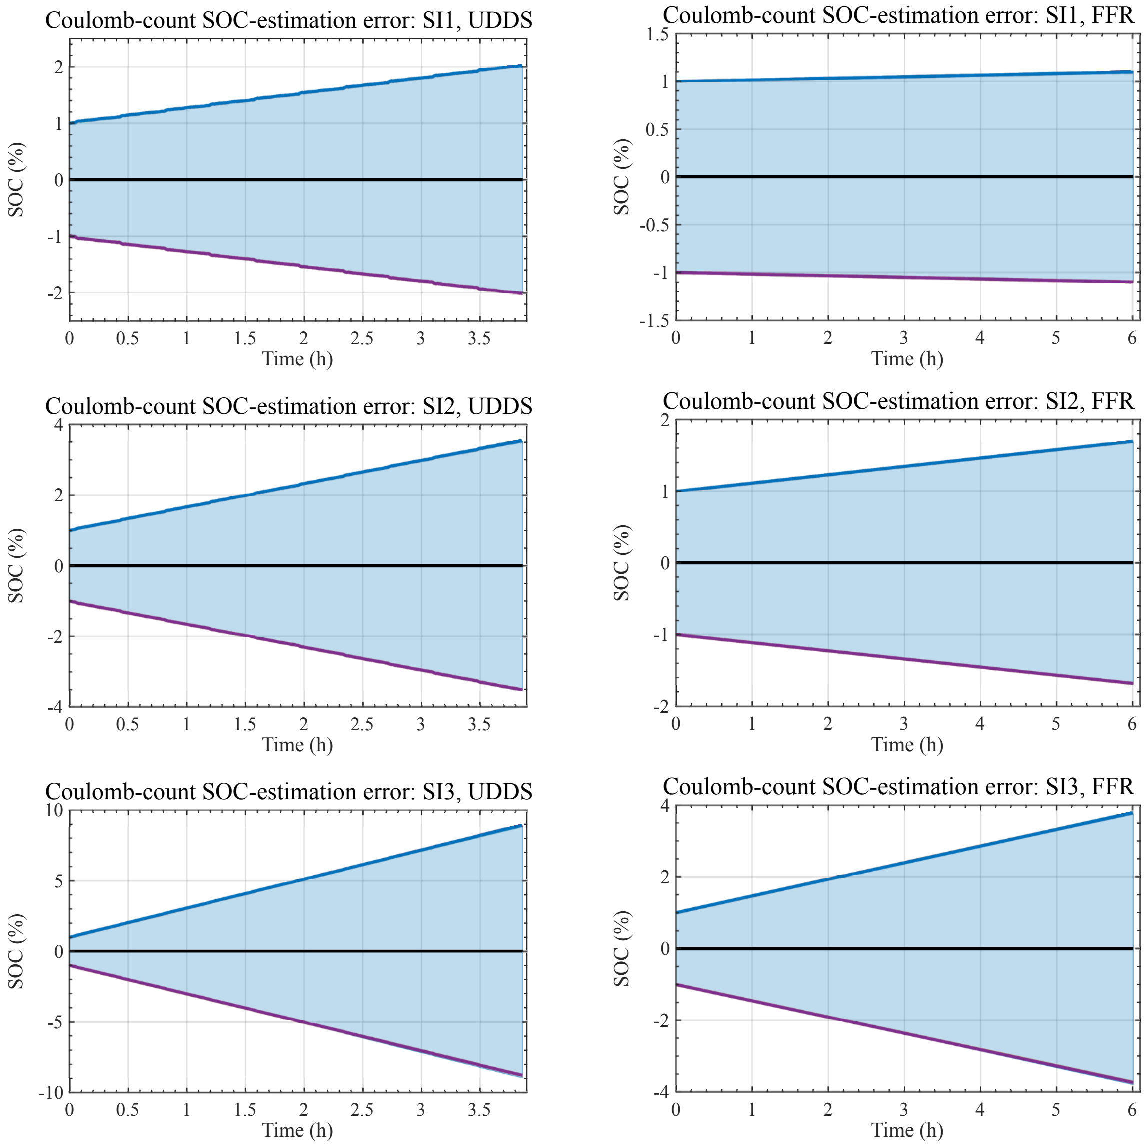

Fig. 3: Coulomb-counting SOC estimation-error results. The notation 'SIx' in the titles indicates that current sensor 'SIx' from Table 3 was used when producing that plot's results.

The 'max' (and 'min') entries in the table refer to adding (or subtracting) the worst-case dc-bias value defined by the particular sensor's specification to (from) the true value.

This tuning process is very important for the SPKF to produce reliable results. Notice however, that the tuning is accomplished using the worst-case sensor specifications and so the SPKF does not need to be retuned by every BMS for its own sensor package.

4.3 Results and discussion

Table 4 lists numeric SOC-estimation results. The "RMSE CC" column lists the root-mean-squared SOC-estimation error using coulomb counting, in percent. This is statistically similar to a 1σ error value. The "RMSE KF" columns list the corresponding result when estimating SOC using the SPKF.

The "RMSB KF" columns list the root-mean-squared value of the confidence bounds for that scenario. We would prefer this number to be small, because it indicates that the filter has high

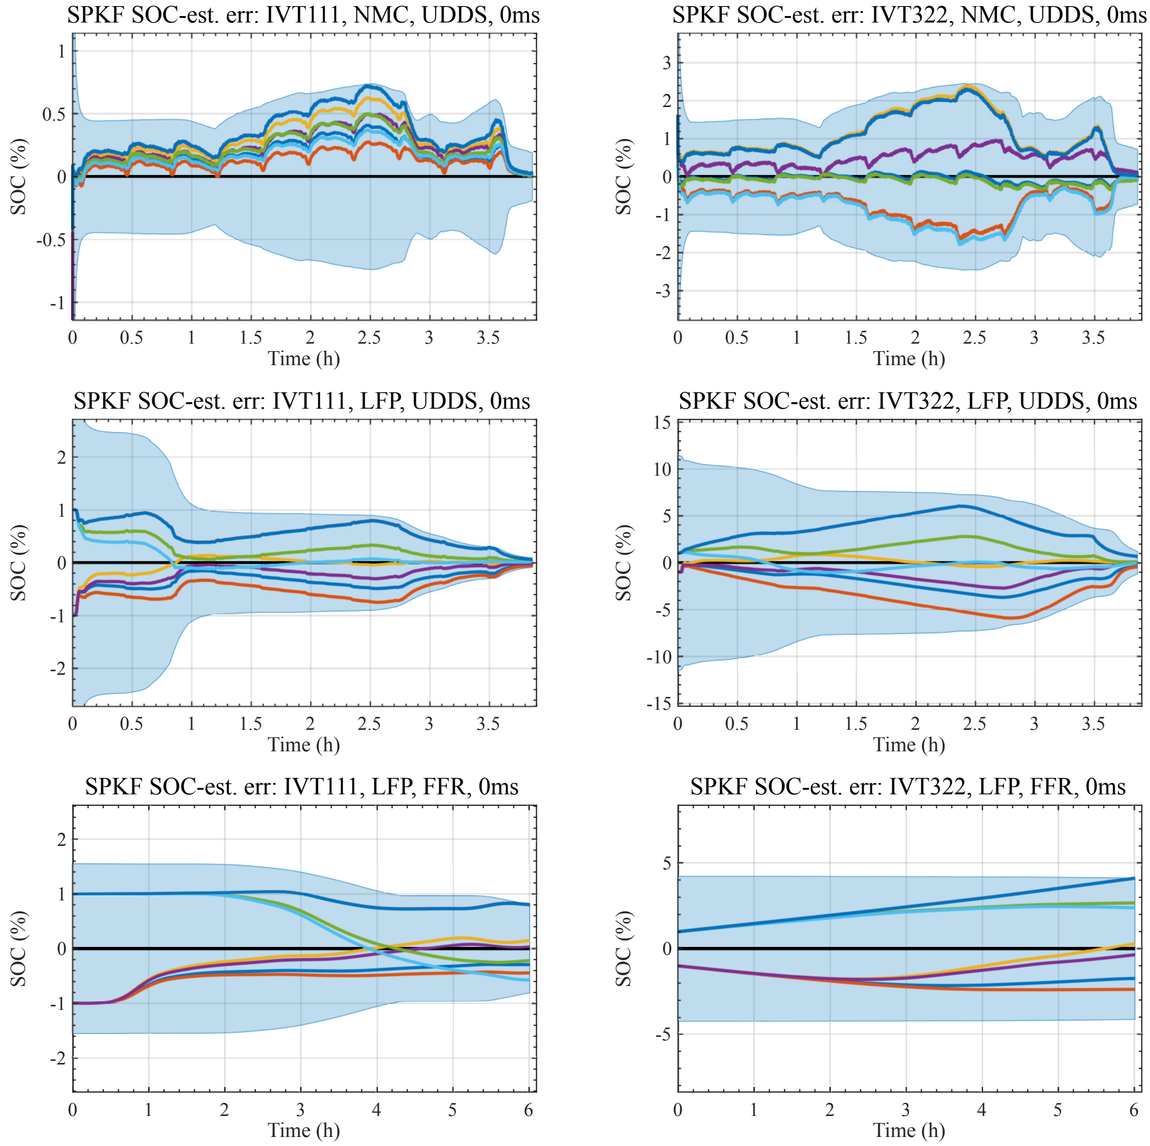

Fig. 4: Subset of SPKF SOC estimation-error results. The notation 'IVTxyz' in the titles indicates that current sensor 'SIx', voltage sensor 'SVy', and temperature sensor 'STz' from Table 3 were used when producing the results of that plot.

confidence in its estimates. However, we also wish for it to be reliable—if the true error is ever outside the confidence bounds, then we lose faith in those bounds. The "BNDX KF" column presents the percent of time that the true error is outside the confidence bounds. Notice that this is always zero in this table, indicating that we have high confidence that the SPKF method is performing robustly.

Fig. 3 presents coulomb-counting results in a graphical format. The blue shaded regions are the confidence regions of the estimate. The black solid line indicates the desired SOC-estimation error of zero. The colored lines (eight in total) correspond to the eight test cases specified in Table 2. We notice that the eight test cases divide into two categories since only two colored lines are apparent. In fact, the top line (which appears as blue) is a combination of four test cases whose SOC-estimation errors are indistinguishable at this scale, and the bottom line (which appears as purple) is also a combination of four test cases whose results are indistinguishable. This verifies that the design objective used to formulate the cases of Table 2 was met.

It also shows that the confidence region of the coulomb-counting estimate from Eqs. (1)-(2) is a tight/achievable region.

Fig. 4 presents a subset of the SPKF results. The blue shaded regions again represent the (union of the) confidence regions of the estimates for the eight simulation cases. The individual-case responses are usually sufficiently distinct to be distinguished from one another. The voltage feedback used by the SPKF allows the estimator to recover from initial SOC-estimation errors, so that the final SOC-estimation error was often lower than the initial error.

Studying the numeric results of Table 4, we make several observations. All else being equal, the quality of the current sensor directly impacts the RMSE of the coulomb-count estimate. If an application specifies a maximum permitted SOC-estimation error, this will impose requirements on the measurement error of the current sensor. The quality of the voltage and temperature sensors have no impact on coulomb counting, since there is no feedback in the coulomb-counting method and the cell-model quantities involved when estimating SOC using coulomb counting are not temperature-dependent.

We also notice that the quality of the current sensor has a direct impact on the RMSE of the SPKF estimate. However, since the SPKF uses voltage feedback, it is relatively insensitive to errors in the value of total-capacity used in its processing. So, nearly always, the SPKF estimation errors are lower in magnitude than the corresponding coulomb-counting errors. Since the LFP chemistry has a very flat OCV curve, the value of voltage feedback for LFP applications is lower than when using nickel-based chemistries such as NMC. However, there still is some benefit in using the feedback as compared with open-loop coulomb counting. Even so, the LFP absolute errors tend to be larger than the NMC absolute errors.

Since the SPKF relies on the temperature sensor to select the model parameter values used in its predictions and since it relies on the voltage sensor for feedback, it is sensitive to the quality of both of these sensors. All else being equal, as the voltage and current sensors degrade, the quality of the SPKF estimates also degrade. So, again, if an application specifies a maximum permitted SOC-estimation error, this will impose requirements on the measurement error of the current, voltage, and temperature sensors. Perhaps surprisingly, sensor latency as large as 100 ms has very little impact on the SPKF results.

5 SOH-Q: Total-capacity estimation methods

SOH is usually quantified by evaluating changes in cell total capacity and ESR as the cell ages. We consider these two effects separately.

Total capacity is estimated in a number of ways in the literature. Some methods first estimate series resistance (or impedance) and then correlate changes in resistance to presumed changes in capacity (most machine-learning methods follow this approach). They do this because estimating resistance/impedance well is relatively simple to do, since resistance/impedance is a direct contributor to the voltage measurement. However, the assumption that is made in doing so is that resistance/impedance changes in a deterministic way when capacity changes, which is not physically true in general.

A physically sound method bases the capacity estimate directly on the cell-model equations. For example, consider the change in SOC $z(t)$ between two points in time t_1 and t_2:

$$z(t_2) = z(t_1) - \frac{1}{Q} \int_{t_1}^{t_2} i(\tau)\, \mathrm{d}\tau,$$

where $i(t)$ is applied current and Q is cell total capacity (and we are assuming that $\eta \approx 1$). We can rearrange this expression to find:

$$\underbrace{-\int_{t_1}^{t_2} i(\tau)\,\mathrm{d}\tau}_{y} = Q\,\underbrace{(z(t_2) - z(t_1))}_{x},$$

where the linear form $y = Qx$ becomes apparent. This observation allows us to apply methods of linear regression to datasets comprising $\{x, y\}$ pairs to find an estimate of the total capacity.

We must be careful when doing so: because current sensors are imperfect, the coulomb count y will have errors; because SOC-estimators produce imperfect results, x will also have errors. Therefore, it is mathematically incorrect to use standard ordinary least squares regression to estimate Q; instead, we must use total least squares. Total-least-squares methods are not recursive in the general case and so are impractical to implement on an embedded system like a BMS. So, approximations or simplifying assumptions must be made. In this white paper, we compare the standard weighted ordinary least squares method (WLS, which is biased by noises on the x variable and so should not be used) to the approximate weighted total least squares method (AWTLS) [8, 6].

In order to implement the methods, we need datasets of $\{x, y\}$ pairs along with the uncertainties/variances of x and y (σ_x^2 and σ_y^2) for each of these scenarios. The variable x is the difference in SOC estimates between two points in a profile. The quality of the total-capacity estimate is generally highest when the magnitude of x is maximized. So, for each profile, we chose the times $t_1 = 0$ and t_2 equal to the final time in the profile to maximize the true SOC difference.

The uncertainty of x is described by σ_x^2. If we assume that the errors on the SOC estimates at the beginning and end of the profiles are uncorrelated, then $\sigma_x^2 = \sigma_{z(t_1)}^2 + \sigma_{z(t_2)}^2$. We compute the values of σ_z^2 from the "RMSE KF" column in Table 4 as $\sigma_z^2 = (\text{rmse in percent}/100)^2$.

The total-capacity estimates will be best if we can minimize σ_x^2. If there is ever a point in a profile where there is no uncertainty regarding the cell's SOC, we can use that time as t_1 and that SOC as $z(t_1)$, which sets $\sigma_{z(t_1)}^2 = 0$. The UDDS profile has an initial point where the cell begins the profile with a fixed SOC that we might assume was achieved after a prescribed charging profile. Therefore, for the UDDS profile we assume $\sigma_{z(t_1)}^2 = 0$. However, the FFR profile is intended to execute continuously and so we do not assume that $\sigma_{z(t_1)}^2 = 0$. In sum, $\sigma_x^2 = \sigma_z^2$ for UDDS and $\sigma_x^2 = 2\sigma_z^2$ for the FFR profile.

Then, to compute x we compute an SOC estimate at time t_1 as the true SOC at that time plus a Gaussian random variable having variance $\sigma_{z(t_1)}^2$ and another SOC estimate at time t_2 as the true SOC at that time plus a Gaussian random variable having variance $\sigma_{z(t_2)}^2$. Then, x is the difference between these two SOC estimates.

The variable y is (-1 times) the measurement of accumulated ampere hours that pass through the cell between times t_1 and t_2. This value is found by coulomb counting.[3] The variable σ_y^2 is the variance of y.

To find σ_y^2, we use the uncertainty bounds from the final point in the simulation profile when using coulomb-counting as the SOC-estimation method. We assume that this bound is equal to $3\sigma_y$, which allows computing σ_y^2.

[3]Note, for this reason we cannot use coulomb counting to find $z(t_1)$ and $z(t_2)$ in this method. We will end up with the result that $1 = 1$ and capacity will cancel from the equations. The SPKF method appears to work quite well for finding estimates of $z(t_1)$ and $z(t_2)$, however.

Then, to compute y, we use the true change in SOC, multiplied by the true capacity, multiplied by -1, and then we add to it a Gaussian random variable having variance σ_y^2.

Simulation setup and results For each of the scenarios, we set up a simulation where that scenario is executed 1000 times. For every execution, we compute random values of x, y, σ_x^2, and σ_y^2 using the approach just described. We then execute the WLS and AWTLS methods to compute an updated capacity estimate for this iteration (plus confidence bounds). For the purpose of these simulations, we initialized the capacity estimates at $1.01 \times$ the true capacity, and did not degrade the true capacity over time.

6 SOH-R0: Resistance estimation methods

6.1 Method 1: The "simple method"

Estimating a cell's ESR turns out to be relatively simple to do well and quickly with a high-quality sensor package because it is highly observable from voltage measurements. There are different approaches in the literature; here, we consider two approaches based on total least squares.

Consider the cell's discrete-time voltage equation:

$$v_k = \mathrm{OCV}(z_k) + Mh_k - \sum_i v_{c_i,k} - i_k R_0.$$

One approach to estimating R_0 is to subtract voltages at two adjacent time samples:

$$\begin{array}{r}
v_k = \mathrm{OCV}(z_k) + Mh_k - \sum_i v_{c_i,k} - i_k R_0 \\
v_{k-1} = \mathrm{OCV}(z_{k-1}) + Mh_{k-1} - \sum_i v_{c_i,k-1} - i_{k-1} R_0 \\
\hline
v_k - v_{k-1} \approx R_0 \left(i_{k-1} - i_k \right),
\end{array}$$

where we use our knowledge that that cell state of charge z_k, diffusion voltages $v_{c_i,k}$, and hysteresis h_k change relatively slowly compared to how quickly i_k changes. So, we can formulate the ESR estimation problem as:

$$\underbrace{v_k - v_{k-1}}_{y} = R_0 \underbrace{i_{k-1} - i_k}_{x},$$

which has the linear form $y = R_0 x$. By noting that both y and x have measurement errors, we recognize that this is again a total-least-squares problem and not a standard ordinary-least-squares problem. So, we use AWTLS method to estimate ESR based on this model. We call this the "simple" method.

To implement the method, the variable x at any timestep is simply the difference in current between the prior and present timestep and the variable y is the difference between the present and prior voltages. Assuming that measurement noise is white, σ_x^2 is set to twice the variance of the current-sensor random error and σ_y^2 is set to twice the variance of the voltage-sensor random error. Measurement dc bias cancels from x and y due to the subtractions used when computing x and y so does not need to be considered.

Table 5: RMS total-capacity estimation error of the final 200 timesteps of each simulation, expressed as a percentage of true capacity. For each method, one column tabulates the normalized RMS error in percent and a second column tabulates the percentage of time that the true total capacity is outside of the estimate's confidence bounds.

					Latency = −100 ms (values in (%))				Latency = 0 ms (values in (%))				Latency = 100 ms (values in (%))			
SI	SV	ST	Profile	Cell	WLS RMSE	WLS Bounds	AWTLS RMSE	AWTLS Bounds	WLS RMSE	WLS Bounds	AWTLS RMSE	AWTLS Bounds	WLS RMSE	WLS Bounds	AWTLS RMSE	AWTLS Bounds
1	1	1	UDDS	NMC	0.01	0.00	0.01	0.00	0.01	0.00	0.01	0.00	0.01	0.00	0.01	0.00
			UDDS	LFP	0.01	0.00	0.01	0.00	0.01	0.00	0.01	0.00	0.01	0.00	0.01	0.00
			FFR	LFP	2.02	85.80	0.11	0.00	2.02	85.90	0.11	0.00	2.02	85.70	0.11	0.00
1	2	2	UDDS	NMC	0.01	0.00	0.01	0.00	0.01	0.00	0.01	0.00	0.01	0.00	0.01	0.00
			UDDS	LFP	0.06	33.30	0.04	0.00	0.06	33.30	0.04	0.00	0.06	33.30	0.04	0.00
			FFR	LFP	4.97	94.20	0.11	0.00	4.97	94.20	0.11	0.00	4.97	94.20	0.11	0.00
2	1	1	FFR	LFP	7.69	94.50	0.46	0.00	7.69	94.50	0.46	0.00	7.69	94.50	0.46	0.00
2	2	2	FFR	LFP	7.97	94.50	0.48	0.00	7.97	94.50	0.48	0.00	7.97	94.50	0.48	0.00
3	1	1	UDDS	NMC	0.08	0.00	0.08	0.00	0.08	0.00	0.08	0.00	0.08	0.00	0.08	0.00
			UDDS	LFP	0.42	89.60	0.13	0.00	0.42	89.60	0.13	0.00	0.42	89.60	0.13	0.00
3	2	2	UDDS	NMC	0.05	0.00	0.06	0.00	0.05	0.00	0.06	0.00	0.05	0.00	0.06	0.00
			UDDS	LFP	0.08	0.00	0.03	0.00	0.08	0.00	0.03	0.00	0.08	0.00	0.03	0.00
			FFR	LFP	16.57	96.40	1.10	0.00	16.57	96.40	1.10	0.00	16.57	96.40	1.10	0.00

6.2 Method 2: The "SPKF method"

The performance of the "simple method" just presented is expected to degrade when the input current is constant since the magnitude of $i_{k-1} - i_k$ will then be very small since it will depend only on measurement noise and not on actual changes to the load current. AWTLS attempts to provide an intelligent smoothed version of what we might naïvely compute as:

$$R_0 \approx \frac{v_k - v_{k-1}}{i_{k-1} - i_k},$$

which loses quality as $i_{k-1} - i_k \rightarrow 0$. Further, we might expect estimates of R_0 to become meaningless if synchronization between the voltage and current sensors is not exact. Anticipating this problem, we implemented a second TLS-based method to estimate R_0. This method relies on signals produced by the SPKF state estimator.

To develop the SPKF method, we start with the cell voltage equation, replacing true quantities with estimates of those quantities from the SPKF, denoted with the "hat" symbol $(\hat{\cdot})$:

$$\hat{v}_k = \mathrm{OCV}(\hat{z}_k) + M\hat{h}_k - \sum_i \hat{v}_{c_i,k} - i_k^{\mathrm{meas}} R_0,$$

where $\hat{v}_{c_i,k} = R_i \hat{i}_{R_i,k}$ from the estimator. Rearranging,

$$\underbrace{-\left(\hat{v}_k - \mathrm{OCV}(\hat{z}_k) - M\hat{h}_k + \sum_i \hat{v}_{c_i,k}\right)}_{y} = R_0 \underbrace{i_k^{\mathrm{meas}}}_{x},$$

where we once again see the linear structure $y = R_0 x$, and where both y and x have noises associated.

To implement a TLS estimator using this relationship, we exported the $\hat{z}_k$, $\hat{h}_k$, and $\hat{v}_{c_i,k}$ signals from the SPKF state estimator, along with their covariances $\sigma^2_{z,k}$, $\sigma^2_{h,k}$, and $\sigma^2_{v_{c_i},k}$. Computing y is straightforward and x is simply the measured current and no computation is needed to find its value.

The variance σ_x^2 needed by the TLS method is set to the current-sensor variance $\Sigma_{\widetilde{w}}$ tuned to optimize the SPKF. The time-varying variance $\sigma^2_{y,k}$ is computed as:

$$\sigma^2_{y,k} = \sigma_v^2 + \left(\frac{\mathrm{d}\,\mathrm{OCV}(\hat{z}_k)}{\mathrm{d}\hat{z}_k}\right)^2 \sigma^2_{z,k} + M^2\sigma^2_{h,k} + \sum_i \sigma^2_{v_{c_i},k},$$

where $\sigma_v^2 = \Sigma_{\widetilde{v}}$, the voltage-sensor variance tune to optimize the SPKF. This computation assumes that all the uncertainties are uncorrelated (which is unlikely to be true in general).

In summary, we have a new way to present x, y, σ_x^2, and σ_y^2 to a TLS algorithm to estimate R_0. Again, we used AWTLS in the results presented next.

7 SOH-Q and SOH-R0 estimation results and discussion

This modified dataset from Sect. 4.1 was used as input to the SOC-estimation algorithms. Output SOC estimates and statistics were used as inputs to the SOH estimates. We evaluate results by comparing the estimated SOH (total capacity or ESR) to the truth SOH from the original simulation.

Table 6: Numeric results for R_0 estimation using the simple and SPKF methods. For each method, one column tabulates the normalized RMS error in percent and a second column tabulates the percentage of time that the true resistance is outside of the estimate confidence bounds. Entries denoted '—' represent values that are too large to print in the table.

					Latency = −100 ms (values in (%))				Latency = 0 ms (values in (%))				Latency = 100 ms (values in (%))			
SI	SV	ST	Profile	Cell	Simple RMSE	Simple Bounds	SPKF RMSE	SPKF Bounds	Simple RMSE	Simple Bounds	SPKF RMSE	SPKF Bounds	Simple RMSE	Simple Bounds	SPKF RMSE	SPKF Bounds
1	1	1	UDDS	NMC	99.94	100.00	0.96	0.00	1.00	100.00	0.82	0.00	100.02	100.00	1.37	0.00
			UDDS	LFP	99.87	100.00	1.85	0.85	3.02	100.00	1.69	0.77	99.97	100.00	2.35	1.21
			FFR	LFP	145.60	100.00	34.16	93.84	3.06	100.00	2.40	0.00	145.71	100.00	35.13	94.40
1	2	2	UDDS	NMC	99.95	100.00	7.51	0.00	1.02	60.75	7.67	0.00	100.04	100.00	7.38	0.00
			UDDS	LFP	99.90	100.00	7.78	0.00	3.05	94.29	7.77	0.00	99.99	100.00	7.90	0.00
			FFR	LFP	145.28	100.00	38.69	0.00	3.28	81.94	23.54	0.00	145.84	100.00	39.54	0.00
2	1	1	FFR	LFP	145.60	100.00	27.91	0.00	3.06	100.00	12.57	0.00	145.70	100.00	28.82	0.00
2	2	2	FFR	LFP	145.28	100.00	37.98	0.00	3.28	81.94	25.10	0.00	145.84	100.00	38.78	0.00
3	1	1	UDDS	NMC	—	3.86	1.77	0.00	0.99	94.81	1.72	0.00	—	0.06	1.99	0.00
			UDDS	LFP	—	2.81	26.09	0.00	2.97	100.00	26.46	0.00	—	5.71	25.34	0.00
3	2	2	UDDS	NMC	99.79	100.00	8.53	0.66	0.98	42.03	8.44	0.00	99.95	100.00	8.59	0.00
			UDDS	LFP	99.94	100.00	12.42	0.00	3.03	92.38	12.19	0.00	99.95	100.00	12.91	0.00
			FFR	LFP	145.89	100.00	50.55	0.00	3.52	83.94	30.67	0.00	146.45	100.00	51.36	0.00

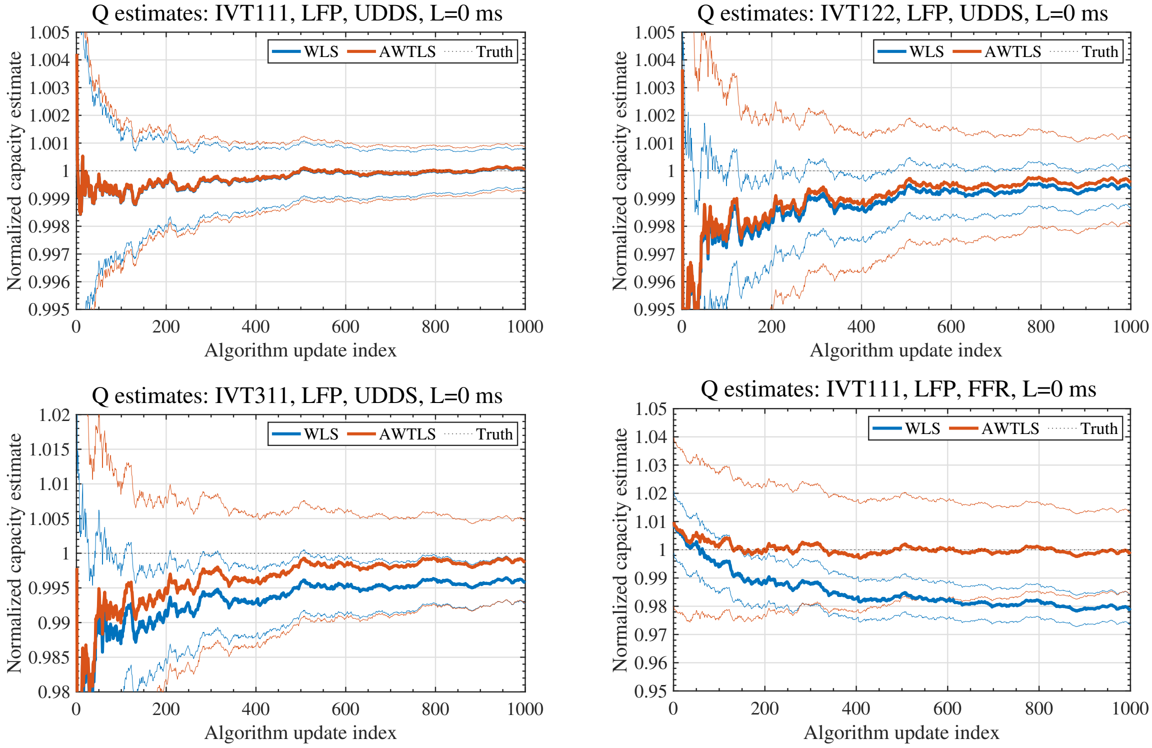

Fig. 5: Subset of capacity-estimation results. The notation 'IVTxyz' in the titles indicates that current sensor 'SIx', voltage sensor 'SVy', and temperature sensor 'STz' from Table 3 were used when producing the results of that plot.

7.1 SOH-Q: SOH represented by total-capacity estimate

Table 5 lists the numeric total-capacity-estimation results. It presents RMSE between true and estimated total capacity for the WLS and AWTLS methods, computed over the final 200 timesteps of each simulation. The table also lists the percentage of time that the total-capacity confidence bounds do not encompass the true total capacity. We would like both of these quantities to be low. From the table, we see that sensor synchronization latency has negligible effect on the quality of the total-capacity estimates.

Fig. 5 shows the evolution of the total-capacity estimates plus their confidence bounds over time for some interesting cases. The plots present normalized estimates, $\widehat{Q}/Q$, so the ideal result is always 1.0. From the top row, we observe that the estimates degrade when the voltage and temperature sensors are less accurate; from the left column we observe a similar trend when the current sensor is less accurate. The illustrated cases also show examples where the WLS method fails (either it converges to the wrong value or its confidence interval is too narrow). AWTLS should be used instead.

7.2 SOH-R0: SOH represented by equivalent-series-resistance estimate

Table 6 lists the numeric ESR-estimation results. It presents RMSE between true and estimated resistance for the simple and SPKF methods and also lists the percentage of time that the estimator confidence bounds do not encompass the true ESR.

Fig. 6 shows the evolution of the resistance estimates over time for two example cases that use the simple method. Each frame shows eight thick lines illustrating the R_0 estimates evolving

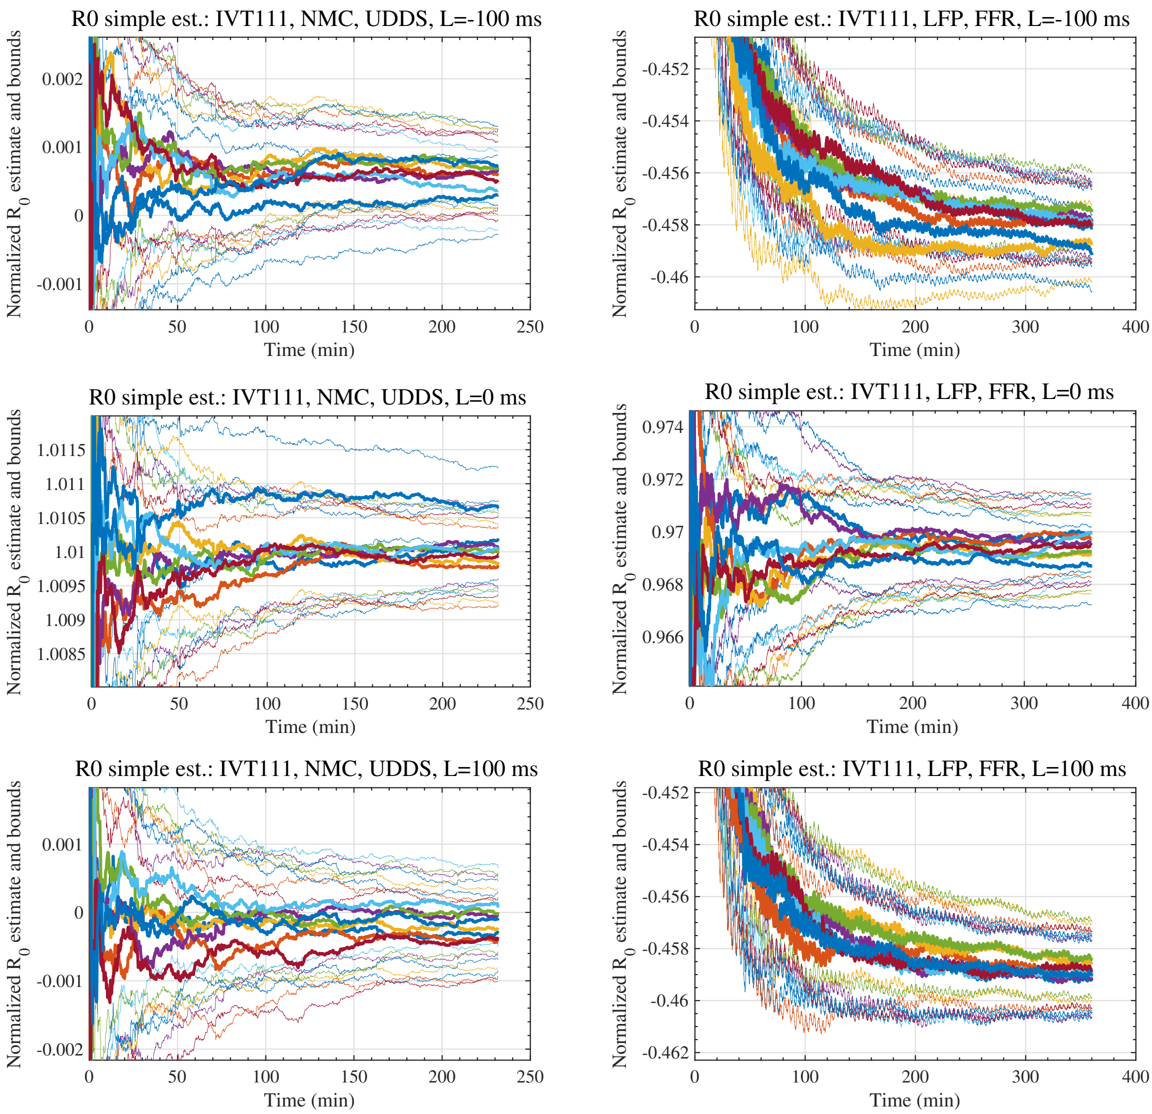

Fig. 6: Subset of R_0 estimation results using the simple method. The notation 'IVTxyz' in the titles indicates that current sensor 'SIx', voltage sensor 'SVy', and temperature sensor 'STz' from Table 3 were used when producing the results of that plot.

over time for the eight cases of Table 2. Thin lines of the same color draw the confidence bounds predicted by the estimator. Note that the plots are for normalized estimates, $\widehat{R}_0/R_0$, so the ideal result is always 1.0. We observe that the simple method consistently fails when there is latency between the voltage and current sensors—its estimates are close to zero since the asynchronous voltage and current signals have essentially no correlation. The simple method even struggles when the measurements are synchronized; its estimates have low RMSE but its confidence intervals are too narrow.

Fig. 7 illustrates several results using the SPKF method. As with the simple method, the SPKF method performs best in the no-latency case but performance degrades gracefully and more robustly when the magnitude of latency increases. Because the SPKF is a filter, it provides estimates of $\hat{z}_k$, $\hat{h}_k$, and $\hat{v}_{c_i,k}$ that are reasonable, despite the timing latencies. Therefore, the overall net voltage depression caused by average current multiplied by R_0 becomes observable, whereas the simple method relies on unfiltered instantaneous changes in voltage and current. The biggest shortcoming of the SPKF method is that its confidence intervals are very wide

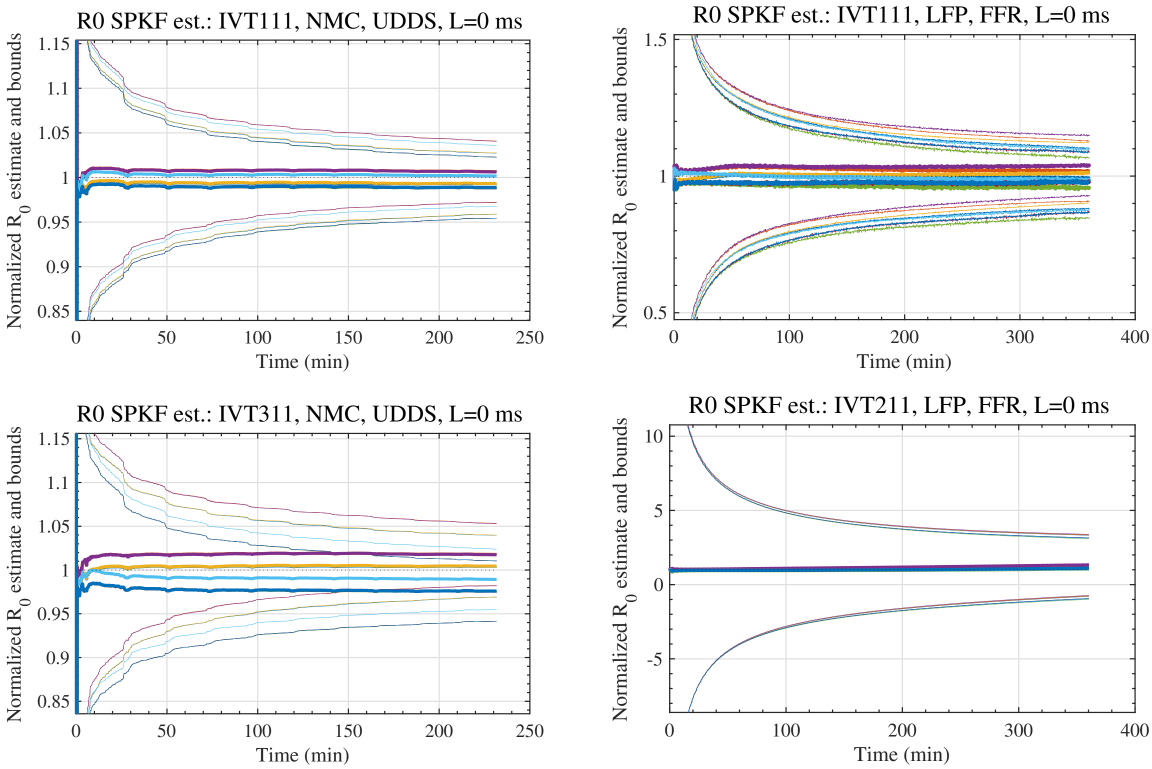

Fig. 7: Subset of R_0 estimation results using the SPKF method. The notation 'IVTxyz' in the titles indicates that current sensor 'SIx', voltage sensor 'SVy', and temperature sensor 'STz' from Table 3 were used when producing the results of that plot.

and do not converge quickly over the course of the simulations. Given longer datasets, we believe that the SPKF method would improve its estimates and confidence windows, although this remains a topic of future work.

7.3 Discussion

Studying the numeric results of Tables 5 and 6, we make several observations. All else being equal, the quality of the sensors directly impacts the RMSE of the SOH estimates. High-accuracy synchronized sensors enable high-accuracy SOH estimates with reliable confidence intervals. Degrading any of the sensors causes a resulting degradation in the quality of the SOH estimates. Losing synchronization between voltage and current measurements is likely to degrade resistance estimates unless advanced methods (such as the SPKF approach presented here) are implemented. As a general rule, if an application specifies a maximum permitted SOH-estimation error, this will impose requirements on the measurement error and synchronization of the sensors.

8 Summary

This white paper has presented a framework for evaluating the impact of sensor measurement error and synchronization latency on SOC- and SOH-estimation results. It has applied that framework to a combination of scenarios encompassing automotive and energy-storage use cases for different sensor error levels and latencies. The study has confirmed in a quantitive way the intuitive expectation that better sensors enable better results. This implies that applications which specify maximum permitted SOC- and SOH-estimation error will have referred requirements imposed on the quality of the sensing subsystem of the BMS. Some high-level direct observations and conclusions are:

- SOC estimates made using coulomb counting do not incorporate voltage or temperature measurements, so are independent of the quality of the voltage and temperature sensors and are not affected by timing latency between the voltage and current measurements.
- SOC estimates made using SPKF use measured voltage via a (temperature-dependent) model-based feedback mechanism and so are affected by the quality of voltage and temperature sensors. Comparing the results from the '111' and '322' sensor packages listed in Table 4, we see that the RMSEs for the high-quality sensor package are on the order of 1/3 of those for the low-quality sensor package. Error bounds are also much tighter for the 111 case, indicating that the SPFK had much more confidence in its estimates when using the high-quality sensor package. The SOC estimates are not affected by timing latencies as high as ± 100 ms between voltage and current measurements.
- SOC estimates made using SPKF were generally better than those made using coulomb counting, but the degree of improvement depended on the flatness of the cell's open-circuit-voltage curve. A disadvantage of SPKF is that it is important that the filters be tuned for the expected application and sensor systems, as described in Sect. 4.2. However, this tuning can be done off-line using worst-case sensor specifications and does not need to be done online by every BMS for its specific sensor package.
- Total-capacity estimates made using ordinary weighted least squares (WLS) ignore the presence of uncertainty on the SOC estimates used as inputs to the method; therefore, WLS fails to provide robust and reliable SOH-Q.
- Total-capacity estimates made using AWTLS incorporate the uncertainty of the SOC estimates in their computations and so succeed in providing robust and reliable SOH-Q. Since the quality of the total-capacity estimate depends directly on the quality of the SOC estimates used as input, a high-quality sensor package enables better estimates of cell total capacity than a low-quality sensor package. Total-capacity estimates are not affected by timing latencies as high as ± 100 ms between voltage and current measurements.
- Unlike SOC and SOH-Q, estimates of SOH-R0 (i.e., ESR) are strongly affected by timing latencies between the voltage and current measurements when using the 'simple method.' Synchronization is very important if we are to achieve meaningful ESR estimates.
- When estimating SOH-R0 using the 'SPKF method,' the level of synchronization between voltage and current measurements is not as critical, but the confidence bounds are much wider than those from the 'simple method' and so the value of the estimate—even though robust—is questionable. We notice again from Table 6 that the quality of the ESR estimate is highly dependent on the quality of the sensor package that is being used.

Some inferred conclusions are:

- As a general rule, if an application specifies a maximum permitted SOC- or SOH-estimation error, this will impose requirements on the measurement error and synchronization of the sensors. If it specifies a maximum permitted 3σ confidence interval on the estimates, this additionally places requirements on the quality of the sensors.

- Since battery available energy is a function of SOC and total capacity, the quality of estimates of available energy will depend on the quality of the SOC and SOH-Q estimates. If a BMS is computing available energy along with a confidence on that computation, then we would like for the SOC and SOH-Q estimates to be very good and to have tight confidence intervals. This requires a high-quality sensor subsystem.
- Also, since battery available power is a function of SOC and ESR, the quality of estimates of available power will depend on the quality of the SOC and SOH-R0 estimates. If a BMS is computing available power along with a confidence on that computation, then we would like for the SOC and SOH-R0 estimates to be very good and to have tight confidence intervals. This requires a high-quality sensor subsystem, preferably with little or no timing latency between the voltage and current measurements.
- Battery phyical health is a complicated function of the specific electrochemical and mechanical degradation factors that it has encountered over its lifetime. Often, SOH-Q and SOH-R0 are used as surrogates for summarizing important aspects of battery health. Therefore, it is important that estimates of both SOH-Q and SOH-R0 be as accurate and robust as possible. As already mentioned, this requires a high-quality sensor system.

References

[1] G. L. Plett and G. McVeigh, "Sensitivity of lithium-ion battery SOC estimates to sensor measurement error and latency," in *Proc. 4th International Conference on Electrical, Computer, and Energy Technologies (ICECET 2024)*, (Sydney, Australia), July 2024.

[2] G. L. Plett and G. McVeigh, "Sensitivity of lithium-ion battery SOH estimates to sensor measurement error and latency," in *Proc. 4th International Conference on Electrical, Computer, and Energy Technologies (ICECET 2024)*, (Sydney, Australia), July 2024.

[3] G. L. Plett, *Battery Management Systems, Vol. 1: Battery Modeling*. Artech House, 2015.

[4] G. L. Plett, "Review and some perspectives on different methods to estimate state of charge of lithium-ion batteries," *J. Automotive Safety and Energy*, vol. 10, no. 3, pp. 249–72, 2019.

[5] G. L. Plett, "Extended Kalman filtering for battery management systems of LiPB-based HEV battery packs—Part 2: Modeling and identification," *Journal of Power Sources*, vol. 134, pp. 262–76, August 2004.

[6] G. L. Plett, *Battery Management Systems, Vol. 2: Equivalent-Circuit Methods*. Artech House, 2015.

[7] G. L. Plett, "Sigma-point Kalman filtering for battery management systems of LiPB-based HEV battery packs: Part 1. Introduction and state estimation," *J. of Power Sources*, vol. 152, no. 2, pp. 1356–1368, 2006.

[8] G. L. Plett, "Recursive approximate weighted total least squares estimation of battery cell total capacity," *Journal of Power Sources*, vol. 196, pp. 2319–2331, 2011.

Distribution

Contact Dukosi

www.dukosi.com

info@dukosi.com