Prognostics Methods for Battery Health Monitoring Using a Bayesian Framework

Bhaskar Saha, *Member, IEEE*, Kai Goebel, Scott Poll, and Jon Christophersen

***Abstract*—This paper explores how the remaining useful life (RUL) can be assessed for complex systems whose internal state variables are either inaccessible to sensors or hard to measure under operational conditions. Consequently, inference and estimation techniques need to be applied on indirect measurements, anticipated operational conditions, and historical data for which a Bayesian statistical approach is suitable. Models of electrochemical processes in the form of equivalent electric circuit parameters were combined with statistical models of state transitions, aging processes, and measurement fidelity in a formal framework. Relevance vector machines (RVMs) and several different particle filters (PFs) are examined for remaining life prediction and for providing uncertainty bounds. Results are shown on battery data.**[1]

***Index Terms*—Battery health, Bayesian learning, particle filter, prognostics, relevance vector machine, remaining useful life.**

I. Introduction

PROGNOSTICS is the emerging science of predicting the health condition of a system and/or its components based upon knowledge of past usage, current state, and future conditions. Past usage may be in the form of archived or historical data of similar system operations, an operational log of the system under study, or a combination of the two. The current state is usually a function of the sensor feedback or some feature derived from it. Future conditions encompass the operational and environmental parameters under which one wants to evaluate the health of the system. Such usage scenarios may be prespecified, derived from statistical analyses of archived instances, or based on expert opinion. In the case of battery health, the important questions boil down to whether the battery will provide the required power during the current discharge cycle, what kind of degradation processes are at work, and how many more missions this battery can support. These concepts are encapsulated in the terms state-of-charge (SOC), state-of-health (SOH), and state-of-life (SOL), respectively. While SOH is predominantly a diagnostics issue, SOC and SOL are primary prognostic concerns. A good battery prognostic algorithm should be able to accurately predict the remaining battery capacity for current and future cycles while also maintaining a reasonable tradeoff with respect to computational burden.

Traditional approaches to battery health management have mostly concentrated on addressing the SOC issue with limited attention to SOH and SOL. Commercial techniques include voltage monitoring, Coulomb counters, and internal impedance measurements. These methods often suffer from model and measurement inaccuracies, leading to erroneous estimates of battery health, the results of which may be catastrophic (e.g., the Mars Global Surveyor satellite). Such uncertainties have also prevented the adoption of cost-effective and environment-friendly rechargeable batteries over primary batteries in critical applications [4]. In recent years, significant advances have been made in the modeling of batteries [5] as well as the application of advanced mathematical tools in predicting battery health [8]. However, the problem of uncertainty management in battery health prediction under conditions that are different from the training data set is still largely untouched.

The Bayesian learning framework introduced in [10] is a first attempt at explicitly incorporating and propagating uncertainty in battery aging models. The relevance vector machine (RVM)–particle filter (PF) approach presented provides a probability density function (pdf) for the end-of-life (EOL) instead of just a mean time to failure. The results show good accuracy and precision in the prognostic predictions as well as their improvement with more measurements. An extension of this paper [6] discusses the comparative benefits of using a model-based (PF) approach, as opposed to data-driven techniques like neural networks (NNs) and Gaussian process regression (GPR) in dealing with model uncertainties and nonnominal operating conditions. In this paper, we extend our work in [10] to reduce the uncertainty in the prediction result using a Rao–Blackwellized PF (RBPF) framework. The idea behind the RBPF is to divide the state space into deterministic and probabilistic parts and analytically solve for the former while using PF for the latter, thus reducing the variance in the state estimate [2]. It is to be noted that some overlaps in content with [6] and [10], pertaining to the literature survey, data description, RVM–PF methodology, and its results, have been maintained to ensure completeness and readability.

Manuscript received November 1, 2007; revised August 14, 2008. First published October 21, 2008; current version published January 5, 2009. The Associate Editor coordinating the review process for this paper was Dr. John Sheppard.

B. Saha is with Mission Critical Technologies, Inc. (NASA ARC), El Segundo, CA 90245 USA (e-mail: bhaskar.saha-1@nasa.gov).

K. Goebel and S. Poll are with NASA Ames Research Center, Moffett Field, CA 94035 USA (e-mail: kai.goebel@nasa.gov; scott.poll@nasa.gov).

J. Christophersen is with the Idaho National Laboratory, Idaho Falls, ID 83415 USA (e-mail: Jon.Christophersen@inl.gov).

Color versions of one or more of the figures in this paper are available online at http://ieeexplore.ieee.org.

Digital Object Identifier 10.1109/TIM.2008.2005965

[1]Based on "An integrated approach to battery health monitoring using Bayesian regression and state estimation" by B. Saha, K. Goebel, S. Poll, and J. Christophersen, in the *Proceedings of IEEE AUTOTESTCON 2007*, pp. 646–653, September 2007. Several figures in this paper have been revised from [6] and [10].

II. Methodology

A. RVM

Support vector machines (SVMs) [12] are a set of related supervised learning methods used for classification and regression that belong to a family of generalized linear classifiers. The goal is to separate an n-dimensional data space (transformed using nonlinear kernels) by a $n-1$-dimensional *hyperplane* that creates the maximum separation (margin) between two classes. This technique can be extended to regression problems in the form of support vector regression [3]. Regression is essentially an inverse classification problem where, instead of searching for a maximum margin classifier, a minimum margin fit needs to be found. However, SVMs are not well suited to diagnostic applications due to the lack of probabilistic outputs. The RVM is a Bayesian form representing a generalized linear model of identical functional form to the SVM [11]. In addition to the probabilistic interpretation of its output, it uses far fewer kernel functions for comparable performance [11].

This type of supervised machine learning starts with a set of input vectors $\{\mathbf{t}_n\}_{n=1,\ldots,N}$ and their corresponding targets $\{\theta_n\}_{n=1,\ldots,N}$. The aim is to learn a model of the dependency of the targets on the inputs to make accurate predictions of θ for unseen values of $\mathbf{t}$. Typically, the predictions are based on some function $F(\mathbf{t})$ defined over the input space, and *learning* is the process of inferring the parameters of this function. In the context of SVM, this function takes the form

$$F(\mathbf{t};\mathbf{w}) = \sum_{i=1}^{N} w_i K(\mathbf{t},\mathbf{t}_i) + w_0 \tag{1}$$

where $\mathbf{w} = (w_1, w_2, \ldots, w_M)^T$ is a weight vector, and $K(\mathbf{t},\mathbf{t}_i)$ is a *kernel* function.

In the case of RVM, the targets are assumed to be samples from the model with additive noise

$$\theta_n = F(\mathbf{t}_n;\mathbf{w}) + \varepsilon_n \tag{2}$$

where ε_n's are independent samples from some noise process (Gaussian with mean 0 and variance σ^2). Assuming the independence of θ_n, the likelihood of the complete data set can be written as

$$p(\theta|\mathbf{w},\sigma^2) = (2\pi\sigma^2)^{-N/2}\exp\left\{-\frac{1}{2\sigma^2}\|\theta-\Phi\mathbf{w}\|^2\right\} \tag{3}$$

where Φ is the $N\times(N+1)$ design matrix with $\Phi = [\varphi(\mathbf{t}_1), \varphi(\mathbf{t}_2),\ldots,\varphi(\mathbf{t}_N)]^T$, wherein $\varphi(\mathbf{t}_N) = [1, K(\mathbf{t}_n,\mathbf{t}_1), K(\mathbf{t}_n, \mathbf{t}_2),\ldots,K(\mathbf{t}_n,\mathbf{t}_N),]^T$.

To prevent overfitting, a preference for smoother functions is encoded by choosing a zero-mean Gaussian prior distribution P_G over $\mathbf{w}$ as

$$p(\mathbf{w}|\eta) = \prod_{i=1}^{N} P_G(w_i|0,\eta_i^{-1}) \tag{4}$$

where η is a vector of $N+1$ *hyperparameters*. To complete the specification of this hierarchical prior, we must define *hyperpriors* over η as well as over the noise variance σ^2.

Having defined the prior, Bayesian inference proceeds by computing the *posterior* over all unknowns given the data from Bayes' rule, i.e.,

$$p(\mathbf{w},\eta,\sigma^2|\theta) = \frac{p(\theta|\mathbf{w},\eta,\sigma^2)p(\mathbf{w},\eta,\sigma^2)}{p(\theta)}. \tag{5}$$

Since this form is difficult to analytically handle, the *hyperpriors* over η and σ^2 are approximated as delta functions at their most probable values η_{MP} and σ^2_{MP}. Predictions for new data are then made according to

$$p(\theta_*|\theta) = \int p(\theta_*|\mathbf{w},\sigma^2_{\mathrm{MP}})p(\mathbf{w}|\theta,\eta_{\mathrm{MP}},\sigma^2_{\mathrm{MP}})d\mathbf{w}. \tag{6}$$

B. PFs

Bayesian techniques also provide a general rigorous framework for dynamic state-estimation problems. The core idea is to construct a pdf of the state based on all available information. For a linear system with Gaussian noise, the method reduces to the Kalman filter. For nonlinear systems or non-Gaussian noise, there is no general analytic (closed-form) solution for the state-space pdf. The extended Kalman filter (EKF) is the most popular solution to the recursive nonlinear state estimation problem [7]. In this case, the desired pdf is approximated by a Gaussian, which may have significant deviation from the true distribution causing the filter to diverge. In contrast, for the PF approach [1], the pdf is approximated by a set of particles (points) representing sampled values from the unknown state space and a set of associated weights denoting discrete probability masses. The particles are generated and recursively updated based on a probabilistic model as well as measurements. In other words, PF is a technique for implementing a recursive Bayesian filter using Monte Carlo (MC) simulations and, as such, is known as a sequential MC (SMC) method.

Particle methods assume that the state equations can be modeled as a first-order Markov process with the outputs being conditionally independent. This can be written as

$$\begin{aligned}\mathbf{x}_k &= f(\mathbf{x}_{k-1}) + \omega_k\\ \mathbf{y}_k &= h(\mathbf{x}_k) + v_k\end{aligned} \tag{7}$$

where k is the time index, $\mathbf{x}$ denotes the state, $\mathbf{y}$ is the output or measurements, and ω and ν are samples from independent noise distributions.

Sampling importance resampling (SIR) is a very commonly used PF algorithm that approximates the filtering distribution denoted as $p(\mathbf{x}_k|\mathbf{y}_{0:k})$ by a set of P weighted particles $\{(w_k^{(i)},\mathbf{x}_k^{(i)}) : i = 1,\ldots,P\}$. The *importance weights* $w_k^{(i)}$ are approximations to the relative posterior probabilities of the particles such that

$$\int f(\mathbf{x}_k)p(\mathbf{x}_k|\mathbf{y}_{0:k})dx_k \approx \sum_{i=1}^{P} w_k^{(i)} f\left(\mathbf{x}_k^{(i)}\right)$$

$$\sum_{i=1}^{P} w_k^{(i)} = 1. \tag{8}$$

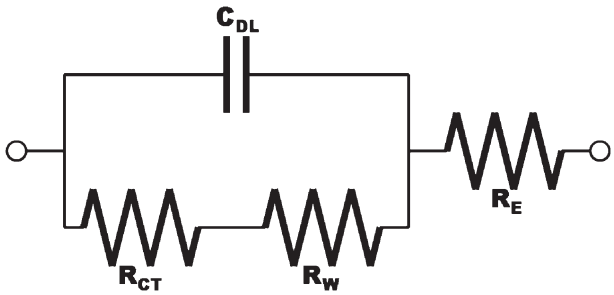

Fig. 1. Lumped parameter model of a cell (reproduced from [6, Fig. 3] and [10, Fig. 1]).

The weight update is given by

$$w_k^{(i)} = w_{k-1}^{(i)} \frac{p(\mathbf{y}_k|\mathbf{x}_k)p(\mathbf{x}_k|\mathbf{x}_{k-1})}{\pi(\mathbf{x}_k|\mathbf{x}_{0:k-1}, \mathbf{y}_{1:k})} \tag{9}$$

where the importance distribution $\pi(\mathbf{x}_k|\mathbf{x}_{0:k-1}, \mathbf{y}_{0:k})$ is approximated as $p(\mathbf{x}_k|\mathbf{x}_{k-1})$.

Resampling is used to avoid the problem of degeneracy of the PF algorithm, that is, avoiding the situation in which all but one of the importance weights are close to zero. This step needs to be performed when the effective number of particles $P_{\text{eff}} < P_{\text{threshold}}$, where P_{eff} is computed as the inverse of the sum of squared normalized particle weights given in (9). Resampling is performed by drawing P particles from the current set with probabilities proportional to their weights and then simply replacing the current set with the new one and assigning the same weight $1/P$ to all.

C. RBPF

In the case of models with large state vectors, resampling may not be sufficient in reducing the variance of particle errors. In such cases, if a part $\mathbf{u}_k$ of the state space $\mathbf{x}_k$ can be used to analytically compute the remaining part, then $\mathbf{x}'_k = E(\mathbf{x}_k|\mathbf{u}_k)$, which is known as the Rao–Blackwellized version of $\mathbf{x}_k$, can be used as the state estimator with the same mean as $\mathbf{x}_k$ but a variance that is smaller by $E\{\text{var}(\mathbf{x}_k|\mathbf{u}_k)\}$ [9], where E and var denote the statistical operators expectation and variance, respectively. Hence, it statistically makes sense to use the RBPF whenever possible.

III. Implementation

A. Model Development

To tie in the aforementioned techniques, namely, RVM and PF, with the battery health monitoring problem, the process is broken down into offline and online parts. During offline analysis, the battery/cell operation is expressed in the form of structural and functional models, which aid in the construction of the "physics of failure mechanism" model. The features extracted from electrochemical impedance spectroscopy (EIS) are used to estimate the internal parameters (like the double layer capacitance C_{DL}, the charge transfer resistance R_{CT}, the Warburg impedance R_W, and the electrolyte resistance R_E) of the battery model shown in Fig. 1. The values of these internal parameters change with various aging and fault processes like plate sulfation, passivation, and corrosion. The parameters of interest for our application are R_E and R_{CT} since their values show significant change due to the aging processes observed.

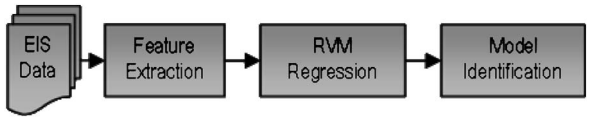

Fig. 2. Schematic of model development (revised from [6, Fig. 8] and [10, Fig. 2]).

RVM regression is performed on parametric data collected from a group of cells over a long period of time to find representative aging curves. The use of probabilistic kernels in RVM helps to reject the effects of outliers and the varying number of data points at different time steps, which can bias conventional least-square-based, model-fitting methods. Since we want to learn the dependency of the parameters with time, the RVM input vector $\mathbf{t}$ is time, whereas the target vector θ is given by the inferred parametric values. Exponential growth models, as shown in (10), are then fitted on these curves to identify relevant aging parameters like C_χ and λ_χ, e.g.,

$$\tilde{\chi} = C_\chi \exp(\lambda_\chi t) \tag{10}$$

where $\tilde{\chi}$ is the model predicted value of an internal battery parameter χ like R_{CT} or R_E. The overall model development scheme is depicted in the flowchart of Fig. 2.

B. Diagnosis and Prognosis

The system description model developed offline is fed into the online PF process. Features extracted from sensor data are used to estimate the SOC and SOH. Once the diagnostic module detects a fault, it triggers the prognosis routine. The PF incorporates the aging parameters λ_{R_E} and $\lambda_{R_{\text{CT}}}$ [from (10)], the internal battery parameters R_E and R_{CT}, and C/1 (capacity at rated current) as components of the state vector $\mathbf{x}$. Thus, it performs parameter identification in parallel with state estimation. The measurement vector $\mathbf{y}$ comprises the battery parameters R_E^* and R_{CT}^* extracted from measured data. The system model is

$$\mathbf{x}_k = \begin{cases} \lambda_{R_E} : x_{1,k} = x_{1,k-1} + \omega_{1,k} \\ \lambda_{R_{\text{CT}}} : x_{2,k} = x_{2,k-1} + \omega_{2,k} \\ R_E : x_{3,k} = x_{3,k-1} \exp(x_{1,k} \cdot \Delta t) + \omega_{3,k} \\ R_{\text{CT}} : x_{4,k} = x_{4,k-1} \exp(x_{2,k} \cdot \Delta t) + \omega_{4,k} \\ \mathbf{C}/1 : x_{5,k} = \alpha(x_{3,k} + x_{4,k}) + \beta + \omega_{5,k} \end{cases}$$

$$\mathbf{y}_k = \begin{cases} R_E^* : y_{1,k} = x_{3,k} + v_{1,k} \\ R_{\text{CT}}^* : y_{2,k} = x_{4,k} + v_{2,k}. \end{cases} \tag{11}$$

The values of λ_{R_E}, $\lambda_{R_{\text{CT}}}$, C_{R_E}, and $C_{R_{\text{CT}}}$ learned from RVM regression are used as estimates for $x_{1,0}$, $x_{2,0}$, $x_{3,0}$, and $x_{4,0}$, respectively, whereas $x_{5,0}$ is not required since it is defined in terms of $x_{3,0}$ and $x_{4,0}$. The current capacity estimate $x_{5,k}$ is used to compute the SOC, whereas the future predictions are compared against EOL thresholds to estimate SOL (as shown in Fig. 3). The SOH analysis is performed based on cause-and-effect studies published in the literature.

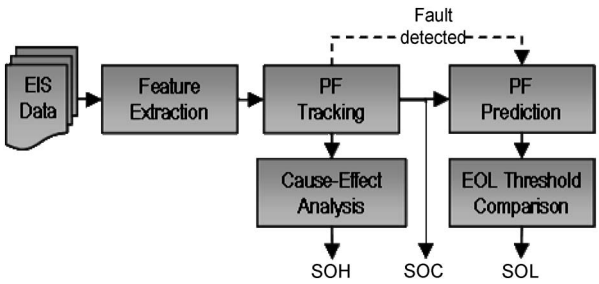

Fig. 3. PF framework (revised from [6, Fig. 9] and [10, Fig. 3]).

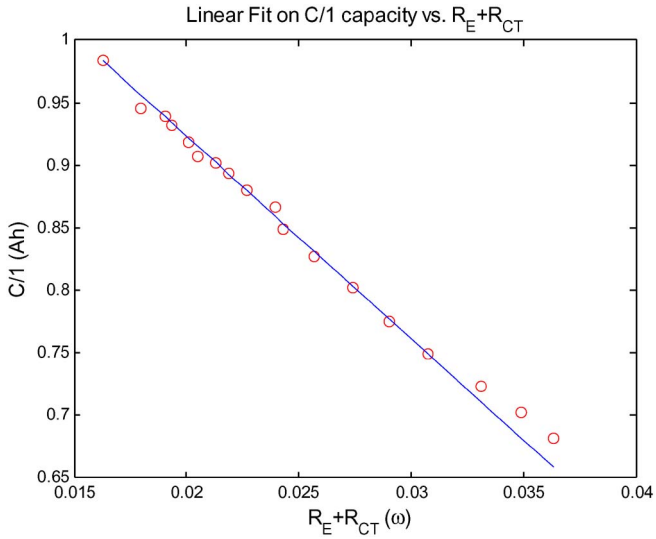

Fig. 4. Dependency between state variables (reproduced from [6, Fig. 5] and [10, Fig. 8]).

In the context of our application, $R_{\mathrm{CT}} + R_E$ was observed to be highly linearly correlated to C/1 capacity (as shown in Fig. 4). To take advantage of this property, an RBPF framework is used, where $x_{1:4,k}$ constitutes the part $\mathbf{u}_k$ that can be used to analytically determine the remaining part $x_{5,k}$ of the state vector $\mathbf{x}_k$. This reduces the variance in the particle values without changing their mean, resulting in more precise (narrower) pdfs.

IV. Results

The data used had been collected from second-generation, 18650-size, Li-ion cells (i.e., Gen 2 cells) that were cycle-life tested at the Idaho National Laboratory. The cells were aged at 60% SOC and various temperatures (25 °C and 45 °C). It is to be noted that, in the application scope of this paper, all data were collected beforehand, and hence, all analyses are effectively offline. We use the 25 °C data for training purposes and the 45 °C data for testing to determine the robustness of our approach to unmodeled operational conditions. Fig. 5 shows a zoomed view of the shift in EIS data of one of the test cells aged at 25 °C (full plot shown in [10, Fig. 4]).

Since the expected frequency plot of a resistance and a capacitance in parallel (Fig. 1) is a semicircle, we fit semicircular curves to the central sections of the data in a least-square sense, which are shown by black dashed lines in Fig. 5. The left intercept of the semicircles give the R_E values, whereas the

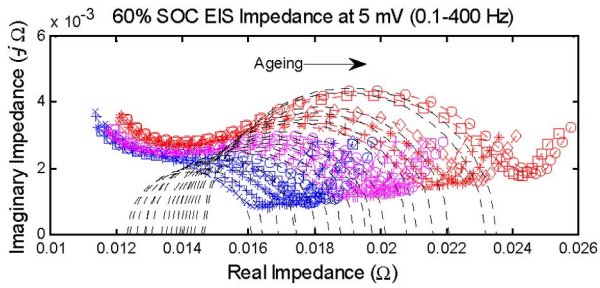

Fig. 5. Zoomed EIS plot with battery model parameter identification (reproduced from [6, Fig. 4] and [10, Fig. 5]).

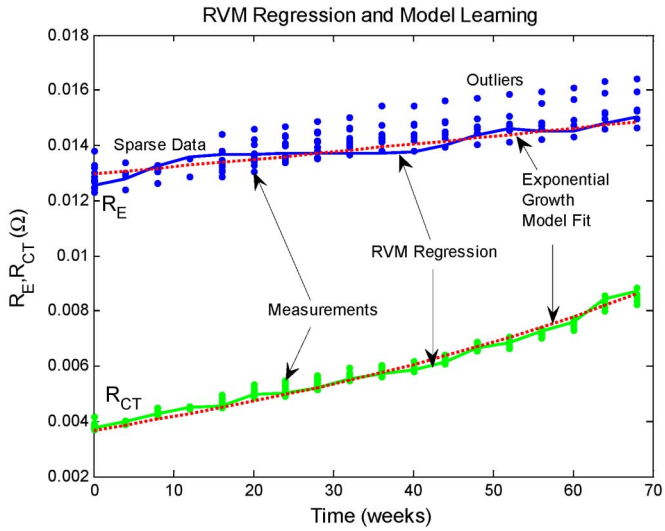

Fig. 6. RVM regression and growth model fit (revised from [10, Fig. 6]).

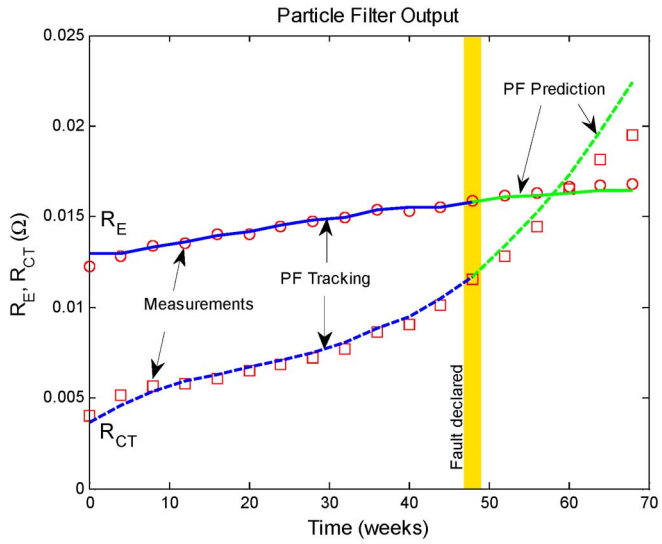

Fig. 7. PF output (revised from [10, Fig. 7]).

diameters of the semicircles give the R_{CT} values. Fig. 6 shows the output of the RVM regression along with the exponential growth model fits for R_E and R_{CT}. The advantage of RVMs in rejecting outliers and nonuniform sparseness of data can be seen in the R_E regression fit of Fig. 6.

Fig. 7 shows both the state tracking and future state prediction plots for test data. The threshold for fault declaration has

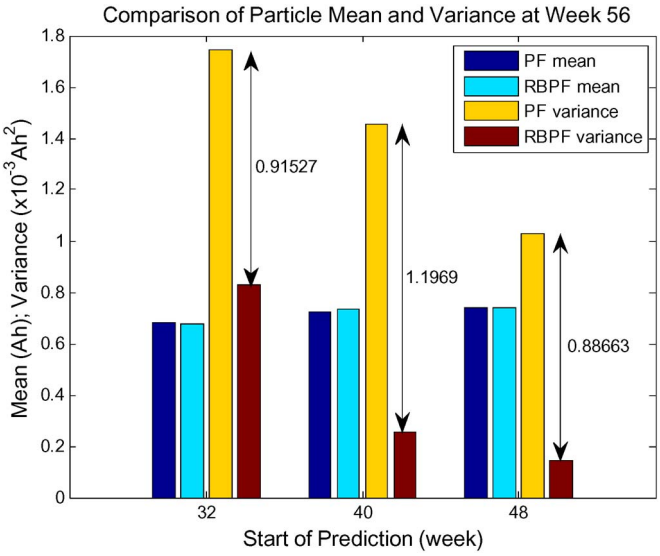

Fig. 8. PF and RBPF particle means and variances.

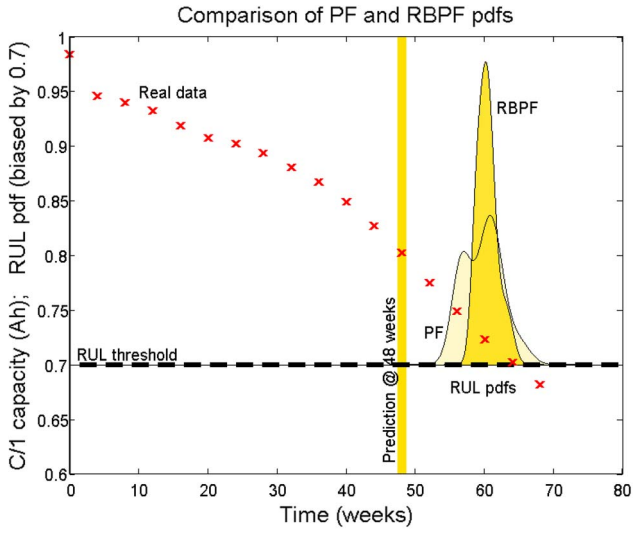

Fig. 9. Comparison of PF and RBPF RUL pdfs.

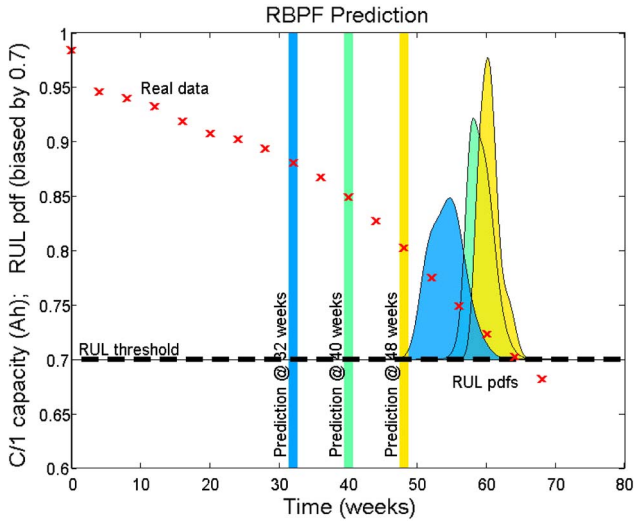

Fig. 10. RBPF RUL predictions.

arbitrarily been chosen. The growth rate of R_{CT} is considerably larger than the training data, suggesting that the cell has undergone rapid passivation due to the elevated temperatures (SOH). The SOC, in any aging cycle, is derived by subtracting the amount of charge drawn expressed as the integral of current over time $\int I \cdot dt$ from the corresponding estimated C/1 capacity $x_{5,k}$.

Remaining useful life (RUL) is used as the relevant metric for determining SOL. This is derived by projecting out the $x_{5,k}$ estimates of each particle from the time of prediction t_P into the future until they hit the predetermined EOL threshold of 0.7 Ah at time t_{EOL}. Both PF and RBPF prognostic frameworks are implemented, with the latter having $x_{5,k}$ as the deterministic state variable and the former assigning an additive zero-mean Gaussian noise of variance 0.001 to it. Although RUL is defined as $t_{\mathrm{EOL}} - t_P$, we plot the estimated EOL values as RUL for ease of visualization. The RUL pdf is computed by fitting a mixture of Gaussians to the RUL values generated by the particle population. As expected from theoretical considerations, the RBPF variance for $x_{5,k}$ is less than the PF value by approximately 0.001 in three separate cases with different t_P, whereas the means are similar (shown in Fig. 8).

Fig. 9 shows the corresponding improvement in precision (spread of the pdf over time) of the RUL pdf with similar accuracy (centering of the pdf over the actual EOL). Fig. 10 shows that the RBPF still maintains the advantages of Bayesian learning, as depicted by the RUL prediction improving in both accuracy and precision with the inclusion of more measurements before prediction.

V. Conclusion

The most significant challenges for making prognostic predictions are the inherent uncertainties in the system model, external and internal noise, and sensor errors. In [10], we explored the possibility of posing the battery prognostic problem in a Bayesian learning (RVM–PF) framework to encapsulate the randomness of RUL and improve its prediction with increasing information (measurements). We presented the advantages of this model-based approach over data-driven techniques capable of handling uncertainties like NN and GPR in [6]. In this paper, we have implemented a refinement of the PF prognostic framework to further reduce the prediction uncertainty.

We exploit the correlations between battery performance (C/1 capacity) and model parameters (R_E and R_{CT}) by posing the RUL prediction task in an RBPF formulation. The presented results show that it is possible to quantitatively analyze the particle distribution representing the system state pdf in terms of its contributing factors. This property can then be used, in the presence of deterministic relations in the system model, to significantly reduce the spread of the RUL distribution while still maintaining the convergence properties of the basic PF (improvement in accuracy and precision with more data). We conclude that the RBPF is thus a suitable candidate for handling prognostics for many engineered systems, where there is high-fidelity expert knowledge about a part of the state space and uncertainty in the rest.

REFERENCES

[1] S. Arulampalam, S. Maskell, N. J. Gordon, and T. Clapp, "A tutorial on particle filters for online nonlinear/non-Gaussian Bayesian tracking," *IEEE Trans. Signal Process.*, vol. 50, no. 2, pp. 174–188, Feb. 2002.

[2] G. Castella and C. P. Robert, "Rao–Blackwellization of sampling schemes," *Biometrika*, vol. 83, pp. 81–94, 1996.

[3] H. Drucker, C. J. C. Burges, L. Kaufman, A. J. Smola, and V. Vapnik, "Support vector regression machines," in *Advances in Neural Information Processing Systems*, vol. 9, M. Mozer, M. Jordan, and T. Petsche, Eds. Cambridge, MA: MIT Press, 1997, pp. 155–161.

[4] G. S. Fein, "Battery supplies ran dangerously low in Iraq," *Nat. Defense Mag.*, pp. 16–19, Sep. 2003. [Online]. Available: http://www.nationaldefensemagazine.org/archive/2003/September/Pages/Batter_Supplies3776.aspx

[5] L. Gao, S. Liu, and R. A. Dougal, "Dynamic lithium-ion battery model for system simulation," *IEEE Trans. Compon. Packag. Technol.*, vol. 25, no. 3, pp. 495–505, Sep. 2002.

[6] K. Goebel, B. Saha, A. Saxena, J. R. Celaya, and J. Christophersen, "Prognostics in battery health management," *IEEE Instrum. Meas. Mag.*, vol. 11, no. 4, pp. 33–40, Aug. 2008.

[7] A. H. Jazwinski, *Stochastic Processes and Filtering Theory*. New York: Academic, 1970.

[8] J. D. Kozlowski, "Electrochemical cell prognostics using online impedance measurements and model-based data fusion techniques," in *Proc. IEEE Aerosp. Conf.*, Mar. 2003, vol. 7, pp. 3257–3270.

[9] B. Ristic, S. Arulampalam, and N. Gordon, *Beyond the Kalman Filter*. Boston, MA: Artech House, 2004.

[10] B. Saha, K. Goebel, S. Poll, and J. Christophersen, "An integrated approach to battery health monitoring using Bayesian regression and state estimation," in *Proc. IEEE AUTOTESTCON*, Sep. 2007, pp. 646–653.

[11] M. E. Tipping, "The relevance vector machine," in *Advances in Neural Information Processing Systems*, vol. 12. Cambridge, MA: MIT Press, 2000, pp. 652–658.

[12] V. N. Vapnik, *The Nature of Statistical Learning*. Berlin, Germany: Springer-Verlag, 1995.

Bhaskar Saha (M'08) received the B.Tech. and M.S. degrees from the Indian Institute of Technology, Kharagpur, India, and the Ph.D. degree from the Georgia Institute of Technology, Atlanta, in 2008.

He is currently a Research Programmer with Mission Critical Technologies (National Aeronautics and Space Administration (NASA) Ames Research Center), Moffett Field, CA. He has developed an integrated Bayesian framework for estimating present damage and damage growth rates as a function of operational parameters to determine time-to-failure probability densities. He has authored a fair number of papers on these topics. His research interests include the application of various classification, regression, and state estimation techniques for predicting the remaining useful life of systems and their components.

Kai Goebel received the Dipl.-Ing. degree from Technische Universität München, München, Germany, in 1990 and the M.S. and Ph.D. degrees from the University of California, Berkeley, in 1993 and 1996, respectively.

He is currently a Senior Scientist with the National Aeronautics and Space Administration (NASA) Ames Research Center, Moffett Field, CA, where he leads the Diagnostics and Prognostics Groups in the Intelligent Systems Division. He is also directing the Prognostics Center of Excellence. Prior to that, he was with General Electric's Global Research Center, Niskayuna, NY, from 1997 to 2006, as a Senior Research Scientist. He is the author of more than 75 papers. He is the holder of nine patents. He has carried out applied research in the areas of artificial intelligence, soft computing, and information fusion. His research interests include advancing these techniques for real-time monitoring, diagnostics, and prognostics.

Scott Poll received the B.S.E. degree in aerospace engineering from the University of Michigan, Ann Arbor, in 1994 and the M.S. degree in aeronautical engineering from the California Institute of Technology, Pasadena, in 1995.

He is currently an Associate Principal Investigator for prognostics with the Integrated Vehicle Health Management Project under the National Aeronautics and Space Administration (NASA)'s Aviation Safety Program. He is also co-leading the evolution of a laboratory at the NASA Ames Research Center, Moffett Field, CA, which is designed to enable the development, maturation, and benchmarking of diagnostic, prognostic, and decision technologies for system health-management applications.

Jon Christophersen received the B.S. and M.S. degrees in electrical engineering from the University of Idaho, Moscow, in 1999 and 2005, respectively.

He is currently a Research Engineer with the Energy Storage and Transportation Systems Department, Idaho National Laboratory, Idaho Falls. He has lead responsibility for all high-power cell testing, analyses, and reporting under the Department of Energy's Advanced Technology Development Program, as well as various USABC deliverables. This includes the investigation and successful implementation of novel testing profiles and procedures, the development of new analysis techniques, the development and validation of various life-predictive modeling and prognostic tools, and the coordination of complex testing and analyses between various national laboratories. He has authored more than 20 publications and has numerous conference presentations related to energy storage.