Uncertainty Management for Diagnostics and Prognostics of Batteries using Bayesian Techniques

Bhaskar Saha
Georgia Institute of Technology
777 Atlantic Drive NW
Atlanta, GA 30332
404-894-4130
bsaha@ece.gatech.edu

Kai Goebel
Research Institute for Advanced Computer Science
NASA Ames Research Center
Moffett Field, CA 94035
650-604-4204
goebel@email.arc.nasa.gov

Abstract—Uncertainty management has always been the key hurdle faced by diagnostics and prognostics algorithms. A Bayesian treatment of this problem provides an elegant and theoretically sound approach to the modern Condition-Based Maintenance (CBM)/Prognostic Health Management (PHM) paradigm. The application of the Bayesian techniques to regression and classification in the form of *Relevance Vector Machine* (RVM), and to state estimation as in *Particle Filters* (PF), provides a powerful tool to integrate the diagnosis and prognosis of battery health. The RVM, which is a Bayesian treatment of the Support Vector Machine (SVM), is used for model identification, while the PF framework uses the learnt model, statistical estimates of noise and anticipated operational conditions to provide estimates of remaining useful life (RUL) in the form of a probability density function (PDF). This type of prognostics generates a significant value addition to the management of any operation involving electrical systems.

TABLE OF CONTENTS

1. INTRODUCTION

Uncertainty management is the most significant challenge faced by state-of-the-art health monitoring systems in these times of ever increasing system autonomy. As the onus of decision making in complex engineered systems shifts towards highly evolved algorithms, simple threshold based diagnostic decisions or single valued model or data driven prognostics are insufficient. The various sources of uncertainty inherent to the field of diagnostics and prognostics must be accounted for in a probabilistic fashion for the approach to make sense.

Batteries form a core component of many machines and are often times critical to the well being and functional capabilities of the overall system. Failure of a battery could lead to reduced performance, operational impairment and even catastrophic failure, especially in aerospace systems. A case in point is NASA's Mars Global Surveyor which stopped operating in November 2006. Preliminary investigations revealed that the spacecraft was commanded to go into a safe mode, after which the radiator for the batteries was oriented towards the sun. This increased the temperature of the batteries and they lost their charge capacity in short order. This scenario, although drastic, is not the only one of its kind in aerospace applications. The AFRL ARGOS satellite and the Vikings 2 Mars Lander are both examples of systems shut down by battery failure. An efficient method for battery monitoring would greatly improve the reliability of such systems.

The phrase "battery health monitoring" has a wide variety of connotations, ranging from intermittent manual measurements of voltage and electrolyte specific gravity to fully automated online supervision of various measured and estimated battery parameters. In the aerospace application domain, researchers have looked at the various failure modes of the battery subsystems. Different diagnostic methods have been evaluated, like discharge to a fixed cut-off voltage, open circuit voltage, voltage under load and electrochemical impedance spectrometry (EIS) [16]. In the field of telecommunications, people have looked to combine conductance technology with other measured parameters like battery temperature/differential information and the amount of float charge [5].

Other works have concentrated more on the prognostic perspective rather than the diagnostic one. Statistical parametric models have been built to predict time to failure [9]. Electric and hybrid vehicles have been another fertile area for battery health monitoring [12]. Impedance spectrometry has been used to build battery models for cranking capability prognosis [3]. State estimation techniques, like the Extended Kalman Filter (EKF), have been applied for real-time diagnosis of automotive batteries [2].

As the popular cell chemistries changed from lead acid to nickel metal hydride to lithium ion, cell characterization efforts have kept pace. Dynamic models for the lithium ion

---

1

batteries that take into consideration nonlinear equilibrium potentials, rate and temperature dependencies, thermal effects and transient power response have been built [7]. Automated reasoning schemes based on neuro-fuzzy and decision theoretic methods have been applied to fused feature vectors derived from battery health sensor data to arrive at estimates of battery life [11]. However, not withstanding the body of work done before, it still remains notoriously difficult to accurately predict the end-of-life of a under environmental and load conditions different from training data sets. This is where advanced regression, classification and state estimation algorithms have an important role to play. The work presented here is an extension of the Bayesian diagnostic-prognostic framework presented in [13]. In this paper, the various sources of uncertainty are analyzed and their mitigation using Bayesian techniques is presented. We aim to do a more rigorous quantitative analysis of uncertainty management in battery health prognostics in the future.

2. SOURCES OF UNCERTAINTY

Prognostic predictions needs to contend with multiple sources of error like modeling inconsistencies, system noise and degraded sensor fidelity. Irrespective of whether the diagnostic/prognostic algorithms are model-driven or data-driven it is not feasible to eliminate all of the above error factors. The following subsections discuss these uncertainty sources in more details.

Modeling Error

Modeling error arises due to the inability to create an analytical model that represents the actual system exactly. Often this is because the system dynamics are too complex to be accurately modeled in a computationally feasible framework. Also there might be insufficient knowledge about the system itself or its response to various environmental stimuli.

Noise

Noise in a system can be generated from a multitude of sources. Some may be internal to the system while others can infiltrate the system dynamics through noisy external inputs. In complex systems, the noise may electrical, mechanical or even thermal in nature. Electrical disturbances can result from things like faulty connectors, relay chatter, wiring crosstalk, electromagnetic interference and electrostatic discharge. Most mechanical disturbances result from vibrations or component degradation due to ageing, overloading or corrosion. Thermal noise can result from insufficient cooling or uneven heat distribution and can be highly disruptive to electrochemical systems like batteries.

Sensors

With ever increasing sensor suites being used to monitor today's systems, sensor fidelity has assumed critical importance. Sensor noise can also result from a variety of sources like electrical interference, digitization error, sensor bias, deadband, backlash and nonlinearity in response.

3. UNCERTAINTY MANAGEMENT

Early development of diagnostic/prognostic algorithms concentrated on logical systems, which interacted with the world through "if and then" statements. The importance of probabilities rose with the realization that logical systems could not anticipate all possible contingencies. Consequently Bayesian techniques (amongst others) started to be assimilated in newer approaches. Simply put, Bayes' theory defines the concept of probability as the degree of belief that a proposition is true. Furthermore, it also suggests that Bayes' theorem can be used as a rule to infer or update the degree of belief in light of new information or data – the more the data, the better the predictions. An additional advantage is that Bayesian models are self-correcting, meaning that the predictions change with change in data trends.

Support Vector Machines (SVMs) [15] are a set of related supervised learning methods used for classification and regression that belong to a family of generalized linear classifiers. The *Relevance Vector Machine* (RVM) [14] is a Bayesian form representing a generalized linear model of identical functional form of the SVM. Bayesian techniques also provide a general rigorous framework for dynamic state estimation problems. The core idea is to construct a probability density function (PDF) of the state based on all available information. For a linear system with Gaussian noise, the method reduces to the Kalman filter. The state space PDF remains Gaussian at every iteration and the filter equations propagate and update the mean and covariance of the distribution.

Relevance Vector Machine

In a given classification problem, the data points may be multidimensional (say n_{dim}). The task is to separate them by an n_{dim}-1 dimensional hyperplane. This is a typical form of linear classifier. There are many linear classifiers that might satisfy this property. However, an optimal classifier would additionally create the maximum separation (margin) between the two classes. Such a hyperplane is known as the maximum-margin hyperplane and such a linear classifier is known as a maximum-margin classifier. Nonlinear kernel functions can be used to create nonlinear classifiers [4]. This allows the algorithm to fit the maximum-margin hyperplane in the transformed feature space, though the classifier may be nonlinear in the original input space.

This technique was also extended to regression problems in the form of support vector regression (SVR) [6]. Regression can essentially be posed as an inverse classification problem where, instead of searching for a maximum margin classifier, a minimum margin fit needs to be found. Although, SVM is a state-of-the-art technique for classification and regression, it suffers from a number of disadvantages, one of which is the lack of probabilistic outputs that make more sense in health monitoring applications. The RVM attempts to address these very issues in a Bayesian framework. Besides the probabilistic interpretation of its output, it uses a lot fewer kernel functions for comparable generalization performance.

This type of supervised machine learning starts with a set of input vectors $\{\mathbf{t}\}_{n=1\ldots N}$ and their corresponding targets $\{\theta\}_{n=1\ldots N}$. The aim is to learn a model of the dependency of the targets on the inputs in order 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:

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

(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:

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

(2)

where, ε_n 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:

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

(3)

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

To prevent over-fitting a preference for smoother functions is encoded by choosing a zero-mean Gaussian prior distribution $\wp$ over $\mathbf{w}$:

$$p(\mathbf{w} \mid \eta) = \prod_{i=1}^{N} \wp(w_i \mid 0, \eta_i^{-1}),$$

(4)

with η 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. This accounts for various sources of uncertainty like noise and unmodeled effects not encountered during training. Having defined the prior, Bayesian inference proceeds by computing the posterior over all unknowns given the data from Bayes' rule:

$$p(\mathbf{w},\eta,\sigma^2 \mid \theta) = \frac{p(\theta \mid \mathbf{w},\eta,\sigma^2)p(\mathbf{w},\eta,\sigma^2)}{p(\theta)}$$

(5)

Since this form is difficult to handle analytically, 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:

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

(6)

Particle Filter

In the Particle Filter (PF) approach [1], [8] the state 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 from a nonlinear process model that describes the evolution in time of the system under analysis, a measurement model, a set of available measurements and an a priori estimate of the state PDF. 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 [10].

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:

$$\mathbf{x}_k = f(\mathbf{x}_{k-1}) + \boldsymbol{\omega}_k$$
$$\mathbf{y}_k = h(\mathbf{x}_k) + \boldsymbol{v}_k$$

(7)

where, $\mathbf{x}$ denotes the state, $\mathbf{y}$ is the output or measurements, and $\boldsymbol{\omega}_k$ and $\boldsymbol{v}_k$ are samples from a noise distribution. Sampling importance resampling (SIR) is a very commonly used particle filtering algorithm, which approximates the filtering distribution denoted as $p(\mathbf{x}_k|\mathbf{y}_0,\ldots,\mathbf{y}_k)$ 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

$$\int f(\mathbf{x}_k)p(\mathbf{x}_k \mid \mathbf{y}_0,...,\mathbf{y}_k)dx_k \approx \sum_{i=1}^{P} w_k^{(i)} f(\mathbf{x}_k^{(i)})$$

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

At each time step k, the algorithm proceeds in two parts – first, by calculating the prior distribution $p(\mathbf{x}_k|\mathbf{x}_{k-1})$ according to the system model; and second, by updating the particle weights using the posterior distribution after the measurement at time k. The weight update is given by:

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

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

Although the inclusion of several particles to track the state variables is a significant step towards managing uncertainty in the model-driven PF framework, additional steps like resampling of the particles must be carried out to ensure that the state tracking is robust enough to handle sudden deviations caused by unmodeled effects. Resampling is used to avoid the problem of degeneracy of the algorithm, that is, avoiding the situation that all but one of the importance weights are close to zero. This step needs to be performed when the effective number of particles $P_{eff} < P$, 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 of them.

4. APPLICATION DOMAIN

The application domain chosen to validate the above described approach is a batch of second generation 18650-size lithium-ion cells (i.e., Gen 2 cells) that were cycle-life tested at the Idaho National Laboratory under the Advanced Technology Development (ATD) Program. This program was initiated in 1998 by the U.S. Department of Energy in order to find solutions to the barriers that limit the commercialization of high-power lithium-ion batteries. The cells were aged at 60% state-of-charge (SOC) and various
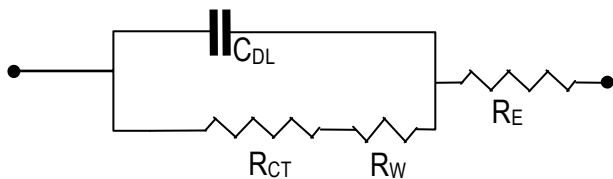
temperatures (25°C and 45°C). The 25°C data was used as the baseline, while the 45°C data was designated as the faulty sequence.

Model Development

The first step in model development is to extract features from sensor data comprising of voltage, current, power, impedance electro-chemical impedance spectrometry (EIS), frequency and temperature readings. These features are used to estimate the internal parameters of the battery model shown in Figure 1. The parameters of interest are the double layer capacitance C_{DL}, the charge transfer resistance R_{CT}, the Warburg impedance R_W and the electrolyte resistance R_E. Other variants of the lumped parameter battery model are not investigated since they all essentially consist of a resistance and capacitance in parallel with another resistance in series, although the nomenclature of the parameters varies.

The values of these internal parameters change with various ageing and fault processes like plate sulfation, passivation and corrosion. Figure 2 shows the shift in electro-chemical impedance spectrometry (EIS) data of one of the test cells with ageing at 25°C. Figure 3 shows a zoomed in section of the data presented above in Figure 2 with feature extraction shown by the dashed curves. Since the expected frequency plot of a resistance and a capacitance in parallel is a semicircle, we fit semicircular curves to the central sections of the data in a least-square sense. The left intercept of the semicircles give the R_E values while the diameters of the semicircles give the R_{CT} values. Other internal parameters like R_W and C_{DL} are not plotted since they showed negligible change over the ageing process and are excluded from further analysis.

Baseline data consists of parametric time series extracted from a group of cells aged at 25°C over a long period. RVM
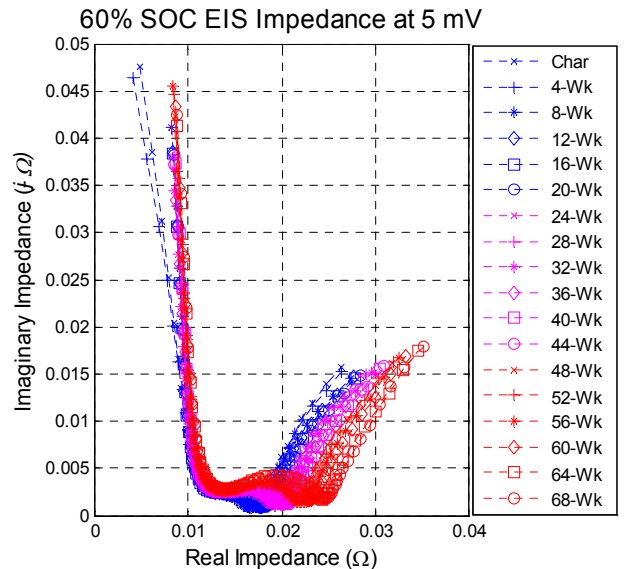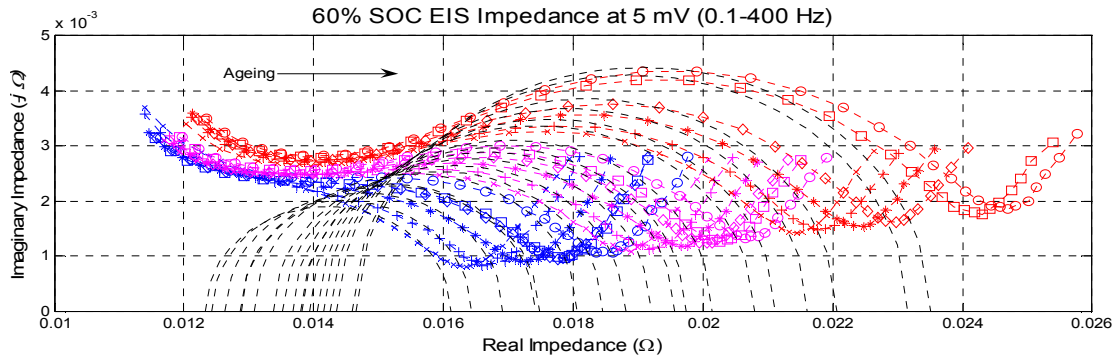

Figure 1 – Lumped Parameter Model of a Cell

Figure 2 – Shift in EIS Data with Ageing

Figure 3 – Zoomed EIS Plot with Internal Battery Model Parameter Identification
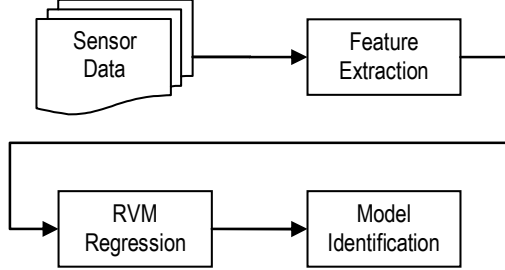

Figure 4 – Schematic of Model Development

Figure 5 – Schematic of RUL Prediction

regression is performed on this data so as to find the representative ageing curves for the different parameters. The RVM input vector $\mathbf{t}$ is time, while 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 the relevant decay parameters like C and λ:

$$\widetilde{\theta} = C \exp(\lambda t), \tag{10}$$

where, $\tilde{\theta}$ is the model predicted value of an internal battery parameter like R_{CT} or R_E. The overall model development scheme is depicted in the flowchart of Figure 4.

Diagnostics and Prognostics

The model developed in the previous section is fed into the Particle Filter (PF) framework. Data from the system sensors are mapped into system parameters which are subsequently used for further analysis. The diagnostic thresholds on the parameters are arbitrarily chosen based on the baseline data sets. Once the diagnostics module detects a fault, it triggers the particle filtering prognosis framework. The PF uses the parameterized exponential growth model, described in (10), for the propagation of the particles in time. The algorithm incorporates the model parameters C and λ as well as the internal battery parameters R_E and R_{CT} as components of the state vector $\mathbf{x}$, and thus, performs parameter identification in parallel with state estimation, thus accounting for more sources of uncertainty. Taking advantage of the highly linear correlation (as derived from data) between $R_{CT}+R_E$ and $C/1$ capacity, i.e. battery capacity at rated current, predicted values of the internal battery model parameters are used to calculate expected charge capacities of the battery. Future
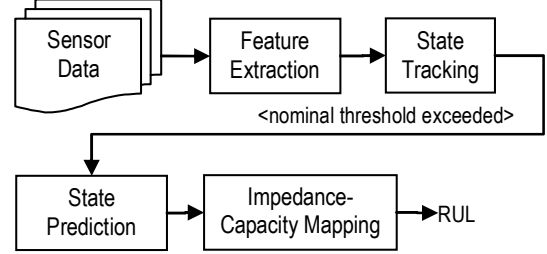
predictions are compared against end-of-life thresholds to derive remaining-useful-life (RUL) estimates. Figure 5 shows a simplified schematic of the process described above.

The state and measurement equations that describe the battery model are given below:

$$\mathbf{z}_0 = \mathbf{C} \; ; \; \boldsymbol{\Lambda}_0 = \boldsymbol{\Lambda}$$

$$\mathbf{z}_k = \mathbf{z}_{k-1}.\exp\boldsymbol{\Lambda}_k + \boldsymbol{\omega}_k$$

$$\boldsymbol{\Lambda}_k = \boldsymbol{\Lambda}_{k-1} + \boldsymbol{v}_k$$

$$\mathbf{x}_k = [\mathbf{z}_k \; ; \; \boldsymbol{\Lambda}_k]$$

$$\mathbf{y}_k = \mathbf{z}_k + \boldsymbol{v}_k \tag{11}$$

where, the vector $\mathbf{z}$ comprises of R_E and R_{CT}, and $\mathbf{C}$ and $\boldsymbol{\Lambda}$ contain their C and λ values respectively. The $\mathbf{z}$ and $\boldsymbol{\Lambda}$ vectors are combined to form the state vector $\mathbf{x}$. The measurement vector $\mathbf{y}$ comprises of the battery parameters inferred from measured data. The values of the $\mathbf{C}$ and $\boldsymbol{\Lambda}$ vectors (for both R_E and R_{CT}) learnt from RVM regression are used as initial estimates for the particle filter. The noise samples $\boldsymbol{\omega}$, $\boldsymbol{v}$ and $\boldsymbol{v}$ are picked from zero mean Gaussian distributions whose standard deviations are derived from the given training data, thus accommodating for the sources of uncertainty in feature extraction, regression modeling and measurement. Resampling of the particles is carried out in each iteration so as to reduce the degeneracy of particle weights. This helps in maintaining track of the state vector even under the presence of disruptive effects like unmodeled operational conditions (in our case, high temperature).
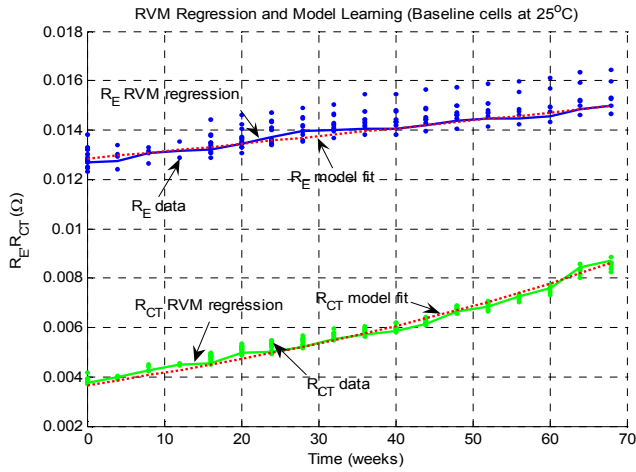
5

Figure 7 – RVM Regression and Growth Model Fit

Figure 8 – Correlation between Capacity and Impedance Parameters

5. RESULTS

The output of the RVM regression along with the exponential growth model fits for R_E and R_{CT} are shown in Figure 6. 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.

Figure 7 shows both the state tracking and future state prediction plots for data collected at 45°C. The threshold for fault declaration has been arbitrarily chosen. The estimated λ value for the R_{CT} growth model (10) is considerably larger than of the training data (collected at 25°C) primarily due to rapid passivation at elevated temperatures.

Figure 8 shows the high degree of linear correlation between the C/1 capacity and the internal impedance parameter R_E+R_{CT}. We exploit this relationship to estimate the C/1
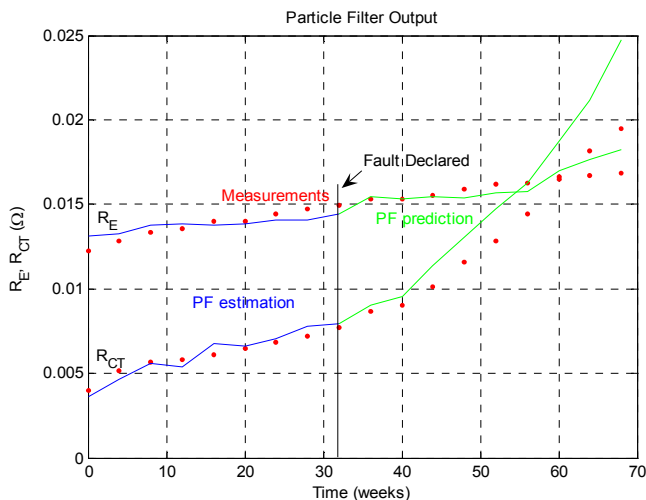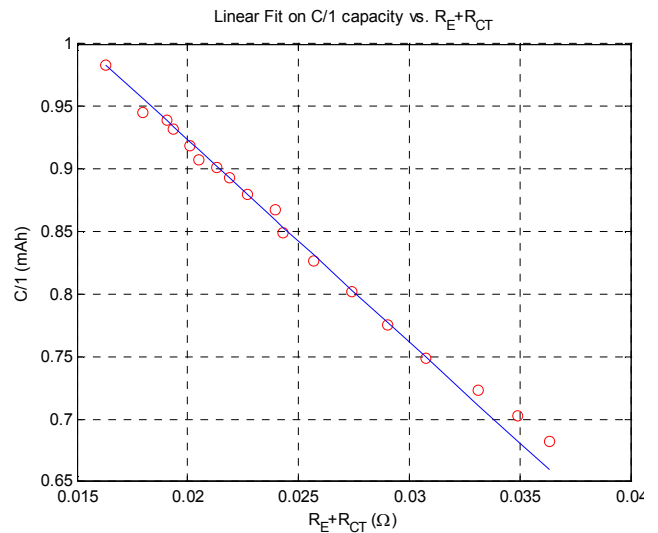
capacities of the cells.

Remaining-useful-life (RUL) or time-to-failure (TTF) is used as the relevant metric for prognostics. This is derived by projecting out the capacity estimates into the future (Figure 9) until expected capacity hits a certain predetermined end-of-life threshold. The particle distribution is used to calculate the RUL PDF by fitting a mixture of Gaussians in a least-squares sense. As shown in Figure 9, the RUL PDF improves in both accuracy (centering of the PDF over the actual failure point) and precision (spread of the PDF over time) with the inclusion of more measurements before prediction.

6. CONCLUSIONS

The combined Bayesian regression-estimation approach implemented as a RVM-PF framework has significant
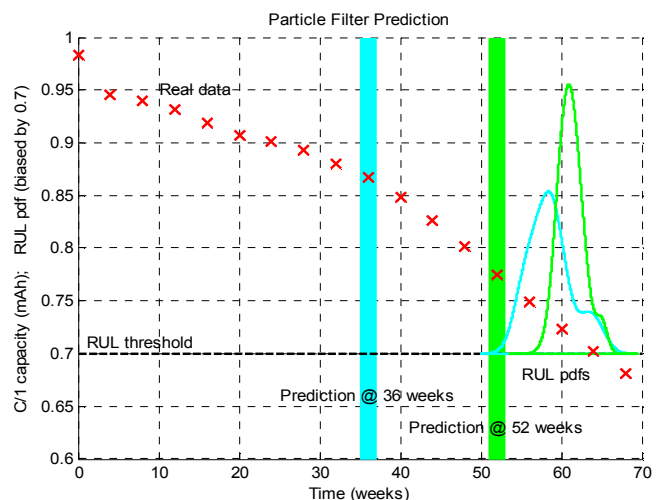

Figure 9 – Particle Filter Output

Figure 6 – Particle Filter Prediction

advantages over conventional methods of RUL estimation like Autoregressive Integrated Moving Average (ARIMA) and Extended Kalman Filter (EKF) [11]. ARIMA, being a purely data-driven method that assumes system stationarity, does not incorporate any physics of the process into the computation, making it unsuitable for long-term predictions. Additionally, it may not be possible to eliminate all non-stationarity from a dataset even after repeated differencing, thus adding to prediction inaccuracy. EKF, though robust against non-stationarity, suffers from the inability to accommodate un-modeled effects and can diverge quickly in the presence of unanticipated operating conditions. A Bayesian statistical approach, on the other hand, is well suited to handle various sources of uncertainties since it defines probability distributions over both parameters and variables and integrates out the nuisance terms. Also, it does not simply provide a mean estimate of the time-to-failure; rather it generates a probability distribution over time that best encapsulates the uncertainties inherent in the system model and measurements and in the core concept of failure prediction.

REFERENCES

[1] S. Arulampalam, S. Maskell, N. J. Gordon and T. Clapp, "A Tutorial on Particle Filters for On-Line Non-Linear/Non-Gaussian Bayesian Tracking", IEEE Trans. on Signal Processing 50(2), 174–188, 2002.

[2] B. S. Bhangu, P. Bentley, D. A. Stone and C. M. Bingham, "Nonlinear Observers for Predicting State-of-Charge and State-of-Health of Lead-Acid Batteries for Hybrid-Electric Vehicles", IEEE Trans. on Vehicular Technology 54(3), 783–794, May 2005.

[3] H. Blanke, O. Bohlen, S. Buller, R. W. De Doncker, B. Fricke, A. Hammouche, D. Linzen, M. Thele and D. U. Sauer, "Impedance Measurements on Lead-Acid Batteries for State-of-charge, State-of-health and Cranking Capability Prognosis in Electric and Hybrid Electric Vehicles", Journal of Power Sources 144(2), 418–425, 2005.

[4] B. E. Boser, I. M. Guyon and V. N. Vapnik, "A Training Algorithm for Optimal Margin Classifiers", in D. Haussler, ed., 5th Annual ACM Workshop on COLT, 144–152, ACM Press, Pittsburgh, PA, 1992.

[5] D. C. Cox and R. Perez-Kite, "Battery State of Health Monitoring, Combining Conductance Technology with other Measurement Parameters for Real-time Battery Performance Analysis", Twenty-second International Telecommunications Energy Conference, 2000, INTELEC, 342–347, Sep. 2000.

[6] H. Drucker, C. J. C. Burges, L. Kaufman, A. J. Smola and V. Vapnik, "Support Vector Regression Machines" in M. Mozer, M. Jordan and T. Petsche, eds., Advances in Neural Information Processing Systems 9, 155–161, Cambridge, MA, MIT Press, 1997.

[7] L. Gao, S. A. Liu and R. A. Dougal, "Dynamic Lithium-Ion Battery Model for System Simulation", IEEE Trans. on Components and Packaging Technologies 25(3), 495–505, Sep. 2002.

[8] N. J. Gordon, D. J. Salmond and A. F. M. Smith, "Novel Approach to Nonlinear/Non-Gaussian Bayesian State Estimation", Radar and Signal Processing, IEE Proceedings F 140(2), 107–113, April 1993.

[9] R. K. Jaworski, "Statistical Parameters Model for Predicting Time to Failure of Telecommunications Batteries", The 21st International Telecommunications Energy Conference, INTELEC '99, June 1999.

[10] A. H. Jazwinski, Stochastic Processes and Filtering Theory, Academic Press, New York, 1970.

[11] J. D. Kozlowski, "Electrochemical Cell Prognostics Using Online Impedance Measurements and Model-Based Data Fusion Techniques", in Proc. IEEE Aerospace Conference 7, 3257–3270, March 2003.

[12] E. Meissner and G. Richter, "Battery Monitoring and Electrical Energy Management – Precondition for Future Vehicle Electric Power Systems", Journal of Power Sources 116(1), 79–98(20), July 2003.

[13] B. Saha, K. Goebel, S. Poll and J. Christophersen, "An Integrated Approach to Battery Health Monitoring using Bayesian Regression and State Estimation", in Proc. AUTOTESTCON, 2007 IEEE, 646-653, 2007.

[14] M. E. Tipping, "The Relevance Vector Machine", in Advances in Neural Information Processing Systems 12, 652–658, Cambridge MIT Press, 2000.

[15] V. N. Vapnik, The Nature of Statistical Learning, Springer, Berlin, 1995.

[16] D. G. Vutetakis and V. V. Viswanathan, "Determining the State-of-Health of Maintenance-Free Aircraft Batteries", in Proc. of the Tenth Annual Battery Conference on Applications and Advances, 13–18, Jan. 1995.

BIOGRAPHY

Bhaskar Saha is a Ph.D. candidate in the School of Electrical and Computer Engineering at Georgia Institute of Technology. He received his B. Tech. (Bachelor of Technology) degree from the Department of Electrical Engineering, Indian Institute of Technology, Kharagpur in 2001. His research interests include applying Bayesian classification, regression and state estimation techniques for estimating remaining useful life, using model-based reasoning for fault diagnosis of complex engineered systems, and formulating verification methodologies for such models.

Kai Goebel received the degree of Diplom-Ingenieur from the Technische Universität München, Germany in 1990. He received the M.S. and Ph.D. from the University of California at Berkeley in 1993 and 1996, respectively.

Dr. Goebel is a senior scientist at NASA Ames Research Center where he is coordinator of the Prognostics Center of Excellence. Prior to that, he worked at General Electric's Global Research Center in Niskayuna, NY from 1997 to 2006 as a senior research scientist. He has carried out applied research in the areas of artificial intelligence, soft computing, and information fusion. His research interest lies in advancing these techniques for real time monitoring, diagnostics, and prognostics. He has fielded numerous applications for aircraft engines, transportation systems, medical systems, and manufacturing systems. He holds half a dozen patents and has published more than 75 papers.