

Battery Health Management System for Electric UAVs

Bhaskar Saha¹, Edwin Koshimoto², Cuong C. Quach³, Edward F. Hogge⁴, Thomas H. Strom⁵, Boyd L. Hill⁶, Sixto L. Vazquez³, Kai Goebel⁷

¹Mission Critical Technologies, Inc. (NASA ARC), 2041 Rosecrans Ave., Ste 220, El Segundo, CA 90245

²NASA Dryden Flight Research Center, Edwards, CA 93523

³NASA Langley Research Center, 1 S. Wright St., Hampton, VA 23681

⁴Lockheed Martin Corporation (NASA LaRC), 1 S. Wright St., Hampton, VA 23681

⁵ATK Space Systems (NASA LaRC), 1 S. Wright St., Hampton, VA 23681

⁶Vigyan Inc. (NASA LaRC), 30 Research Dr, Hampton VA 23681

⁷NASA Ames Research Center, Moffett Field, CA 94035

{bhaskar.saha,ed.t.koshimoto,cuong.c.quach,edward.f.hogge,thomas.h.strom,boyd.l.hill,sixto.l.vasquez,kai.goebel@nasa.gov}

Abstract—This paper presents a novel battery health management system for electric UAVs (unmanned aerial vehicles) based on a Bayesian inference driven prognostic framework. The aim is to be able to predict the end-of-discharge (EOD) event that indicates that the battery pack has run out of charge for any given flight of an electric UAV platform. The amount of usable charge of a battery for a given discharge profile is not only dependent on the starting state-of-charge (SOC), but also other factors like battery health and the discharge or load profile imposed. This problem is more pronounced in battery powered electric UAVs since different flight regimes like takeoff/landing and cruise have different power requirements and a dead stick condition (battery shut off in flight) can have catastrophic consequences.

Since UAVs deployments are relatively new, there is a lack of statistically significant flight data to motivate data-driven approaches. Consequently, we have developed a detailed discharge model for the batteries used and used it in a Bayesian inference based filtering (Particle Filtering) technique to generate remaining useful life (RUL) distributions for a given discharge. The results section presents the validation of this approach in hardware-in-the-loop tests.¹²

TABLE OF CONTENTS

| | |
|--|----------|
| 1. INTRODUCTION | 1 |
| 2. BATTERY BEHAVIOR | 2 |
| 3. PROGNOSTIC FRAMEWORK | 3 |
| 4. UAV PLATFORM | 4 |
| 5. BATTERY HEALTH MANAGEMENT HARDWARE | 4 |
| 6. RESULTS | 5 |
| 7. CONCLUSIONS | 8 |
| REFERENCES | 9 |
| BIOGRAPHY | 9 |

1. INTRODUCTION

Electric UAVs (unmanned aerial vehicles) have become the new face of green aviation. They are being increasingly deployed in military, civilian and scientific missions all over the globe. There are many benefits of electric propulsion. An electric UAV has no exhaust noise and can temporarily minimize prop noise to tactical advantage. The control bandwidth available to the motor can be used to improve roll axis stability for sensing and targeting. Electric propulsion systems usually have fewer moving parts than conventional internal combustion engine (ICE) based systems, and hence are more reliable.

However, like ground vehicles, battery powered electric UAVs suffer from uncertainties in estimating the remaining charge and hence most flight plans are highly conservative in nature. Usually ICE based powertrains run within narrow bands of RPMs with metered fuel delivery. This combined with a known volume fuel tank allows reasonably accurate predictions of remaining use time or travel distance. Batteries on the other hand, decrease in capacity with time and usage. Various factors like ambient storage temperatures and the state-of-charge (SOC) at which the battery was stored affects capacity fade.

Additionally, the amount of usable charge of a battery for a given discharge profile is not only dependent on the starting SOC, but also other factors like battery health and the discharge or load profile imposed. This is because in most battery powered propulsion systems, the battery shut off criteria are based on the terminal voltage. This voltage is related to the SOC of the battery, but it is a highly non-linear relation, which is further complicated by a sharp drop off of the terminal voltage as the battery SOC nears empty. This problem is more pronounced in battery powered electric UAVs since different flight regimes like takeoff/landing and cruise have different power requirements and a dead stick condition (battery shut off in flight) can have catastrophic consequences.

¹ U.S. Government work not protected by U.S. copyright

² IEEEAC paper#1484, Version 4, Updated January 11, 2011

In this paper our application platform is a subscale aerobatic UAV, the Edge 540, powered by four 18.5V 6000mAh Lithium-polymer (Li-Poly) battery packs and our goal is to develop prognostic algorithms to predict the end-of-discharge (EOD) event that indicates that the battery pack has run out of charge for any given flight. The Battery Health Management (BHM) system is fabricated out of Commercial-off-the-Shelf (COTS) systems. A TERN SensorCore board is used for digital signal acquisition, and a Gumstix OVERO with a SUMMIT expansion board is used for recording data and real-time implementation of the prognostic algorithm. An auxiliary signal conditioning board was custom designed and fabricated to handle inter-board signal voltage processing.

For prognostics, a detailed discharge model was developed for the Li-Poly cells and verified using hardware-in-the-loop tests of the Edge 540. This model was then used in a Particle Filter (PF) based prognostic framework that combines state estimation with model adaptation to accurately predict the remaining battery charge. This information is used in conjunction with stochastic estimates of future usage to give remaining run time for the UAV. This paper also discusses how these predictions may be used to increase operational safety, optimize mission plans and extend battery life.

2. BATTERY BEHAVIOR

Batteries are energy storage devices that facilitate the conversion, or *transduction*, of chemical energy into electrical energy, and vice versa [3]. They consist of a pair of *electrodes* (*anode* and *cathode*) immersed in an *electrolyte* and sometimes separated by a *separator*. The chemical driving force across the cell is due to the difference in the chemical potentials of its two electrodes, which is determined by the difference between the *standard Gibbs free energies* the products of the reaction and the reactants. However, the theoretical *open circuit voltage*, E^0 is not available during use. This is due to the various factors like Ohmic loss, work function at the solid electrolyte interface, resistance to ion transfer through the electrolyte, etc.

The output current plays a big role in determining the losses inside a battery and is an important parameter to consider when analyzing battery performance. The term most often used to indicate the rate at which a battery is discharged is the *C-Rate* [2]. The discharge rate of a battery is expressed as C/r , where r is the number of hours required to completely discharge its nominal capacity. So, a 2 Ah battery discharging at a rate of $C/10$ or 0.2 A would last for 10 hours. The terminal voltage of a battery, as also the charge delivered, can vary appreciably with changes in the C-Rate. Furthermore, the amount of energy supplied, related to the area under the discharge curve, is also strongly C-Rate dependent. Figure 1 shows the typical discharge of a battery and its variation with C-Rate. Each curve corresponds to a different C-Rate or C/r value (the lower the

r the higher the current) and assumes constant temperature conditions.

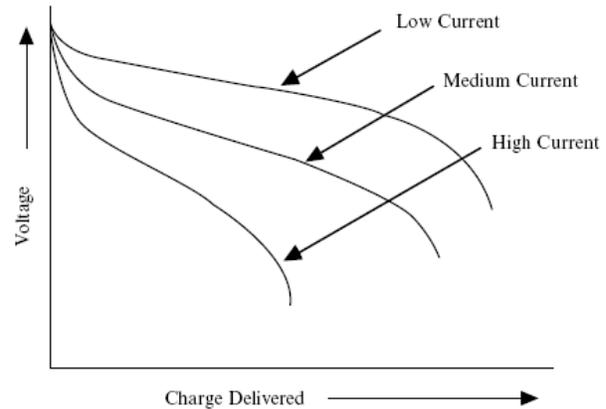


Figure 1 – Schematic drawing showing the influence of the current density upon the discharge curve (Reproduced from [2])

There are several rechargeable battery technologies available on the market right now, each having distinct characteristics. However, Li-ion batteries are becoming increasingly popular for a variety of applications, from consumer electronics to power tools, to electric vehicles and even to space applications. Li-ion batteries have a number of important advantages over competing technologies [2]. Some of these advantages are:

- The electrodes of a Li-ion battery are made of lightweight lithium and carbon. This translates into a very high energy density for Li-ion batteries as compared to other chemistries like lead-acid or NiCd (nickel-cadmium) or NiMH (nickel-metal hydride).
- They have a low *self-discharge* rate, meaning that they hold their charge for longer periods of time.
- Li-ion batteries have a long cycle life. They can handle hundreds of charge and discharge cycles without significant degradation of their capacity.

With these significant advantages in mind Li-Poly batteries were chosen as the power source for our application. Li-Poly battery packs have the same chemical processes as Li-ion cells, except that the electrolyte is in gel form. Modeling a Li-ion battery from the first principles of the internal electrochemical reactions can be very tedious and computationally intractable. Hence, we take the approach of analyzing the way the impedance parameters change with charge depletion during the discharge cycle.

Since the impedance parameters are essentially representations of electrochemical reactions and transport processes inside the battery, they are strongly affected by the internal temperature of the battery, the current load and the ionic concentrations of the reactants. We postulate that as discharge progresses the heat generated by the reactions

and the current flow causes the internal temperature to go up, effectively increasing the mobility of the ions in the electrolyte, thus decreasing the Warburg resistance, R_W . Decreasing R_W , however, increases the self-discharge rate, effectively increasing the electrolyte resistance, R_E , of the battery. Also, the increase in temperature results in faster consumption of the cell reactants resulting in an increase in the charge transfer resistance, R_{CT} , and a sharp drop in the cell voltage. EOD is reached when the output voltage hits the minimum safe voltage threshold, E_{EOD} . For a cell current I , the output voltage E is given by:

$$E = E^o - I(R_E + R_{CT} + R_W). \quad (1)$$

The variations in E^o with internal temperature [2] are not explicitly modeled, but accounted for by the adaptive powers of the PF framework described later. For the empirical charge depletion model considered here, we express the output voltage in terms of the effects of the changes in the internal parameters:

$$E(t) = E^o - \Delta E_{sd}(t) - \Delta E_{rd}(t) - \Delta E_{mt}(t), \quad (2)$$

where, t is the time variable during a discharge cycle, ΔE_{sd} is the drop due to self-discharge, ΔE_{rd} is the drop due to cell reactant depletion and ΔE_{mt} denotes the voltage drop due to internal resistance to mass transfer (diffusion of ions). These individual effects are modeled as:

$$\Delta E_{sd}(t) = \alpha_1 \exp(-\alpha_2/t), \quad (3)$$

$$\Delta E_{rd}(t) = \alpha_3 \exp(\alpha_4 t), \quad (4)$$

$$\Delta E_{mt}(t) = \Delta E_{init} - \alpha_5 t, \quad (5)$$

where, ΔE_{init} is the initial voltage drop when current I flows through the initial value of the internal resistance R_E at the start of the discharge cycle, and $\alpha = \{\alpha_1, \alpha_2, \alpha_3, \alpha_4, \alpha_5\}$ represents the set of model parameters to be estimated from the data. Figure 2 shows how the different voltage drop components defined in eqns. (3)-(5) combine to give the Li-ion discharge profile.

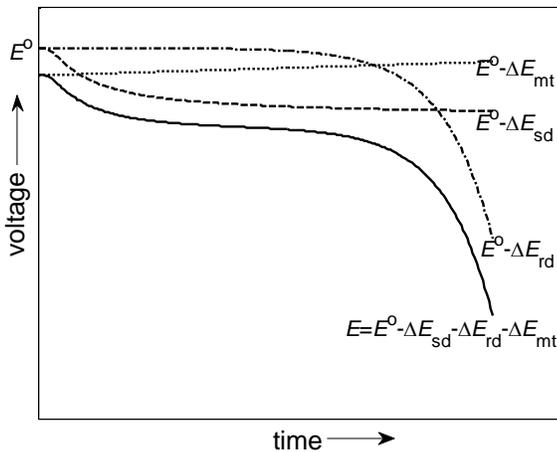


Figure 2 – Decomposition of the Li-ion discharge profile in to different components (Reproduced from [4])

3. PROGNOSTIC FRAMEWORK

The formulation of a model, though, is just a part of the solution. As mentioned above there are a number of unknown parameters that need to be identified. Even after identification, they may not be directly applicable to the test set since the values may differ from one battery to another, or for the same battery from one cycle to the next. Furthermore, for any given cycle the parameter values may be non-stationary. In general, given a model, the task of tracking a state variable and predicting future values is usually cast as a *filtering* problem. The variety of filtering techniques published in literature is enormous with each having performance advantages over others depending upon the application. For our task of battery prognostics, comprising the prediction of EOD and EOL (end-of-life), we need to reconcile our method with non-exact non-linear non-stationary models with non-Gaussian noise. Particle Filtering provides us a viable framework that allows us to explicitly represent and manage the uncertainties inherent to our problem.

Particle Filters [1] are a novel class of non-linear filters that combine Bayesian learning techniques with importance sampling to provide good state tracking performance while keeping the computational load tractable. The idea is to represent the system state (in this case the battery SOC or voltage or capacity) as a probability density function (pdf) that 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 from an *a priori* estimate of the state pdf, propagated through time using a nonlinear process model, and recursively updated from measurements through a measurement model. The main advantage of PFs here is that model parameters can be included as a part of the state vector to be tracked, thus performing model identification in conjunction with state estimation [4]. After the model has been tuned to reflect the dynamics of the specific system being tracked, it can then be used to propagate the particles till the failure (e.g. EOD or EOL) threshold to give the RUL pdf [4].

In the case of our application, the EOD estimation problem is cast in the PF framework as follows:

State transition model \equiv

$$\begin{aligned} \alpha_{j,i+1} &= \alpha_{j,i} + \omega_{j,i}, \forall j = 1, \dots, 5, \\ E_{i+1} &= E_i - \{ \alpha_{1,i} \alpha_{2,i} \exp(-\alpha_{2,i}/t_i) / t_i^2 \\ &\quad - \alpha_{3,i} \alpha_{4,i} \exp(\alpha_{4,i} t_i) - \alpha_{5,i} \} / f_s + \omega_i, \end{aligned} \quad (7)$$

$$\text{Measurement model} \equiv \tilde{E}_i = E_i + v_i, \quad (8)$$

where, i is the time index, f_s is the sampling frequency, \tilde{E}_i denotes the measured cell voltage at time index i , and $\omega_{j,i}$ ($\forall j = 1, \dots, 5$), ω_i and v_i are independent zero-mean Gaussian noise terms.

It must be noted that for a given discharge run, we are using terminal voltage as an indicator of battery life instead of the SOC. There are two main reasons for this. Firstly, most battery powered systems have a cut-off based on battery voltage, not SOC. For most prediction tasks, it is advantageous to track the variable that the system EOL (in this case the battery EOD) is expressed in. Secondly, the relation between terminal voltage and SOC, as given by the manufacturer, does not hold throughout the full life of the battery or under extreme load and temperature conditions. Other variables like the internal impedances can also indicate when a battery is depleted, however costly measurement techniques like electrochemical impedance spectroscopy (EIS) are required to estimate those impedances, which make them impractical for onboard deployment.

4. UAV PLATFORM

The test vehicle used for this research is a COTS 33% scale model of the Zivko Edge 540T. The Edge 540T, pictured in Figure 3, is 96 inches long with a 100 inch wing span. It has 1881 square inches of wing surface area. It is powered by dual tandem mounted electric out-runner motors capable of moving the aircraft up to 85 knots using a 26 inch propeller. It is instrumented primarily for structural fault diagnosis and mitigation research with accelerometers on the wings and fuselage as well as strain gauges on the wings and wing joiner tube. Other sensors include sensors to monitor power plant and navigation parameters. Its relatively low cost makes it suitable for use as a rapid evaluation prototype for airframe damage experiments.



Figure 3 – 1: 33% Edge 540 test aircraft

The gas engine in the original kit specification was replaced by two electric out runner motors which are mounted in tandem to power a single drive shaft. The motors are powered by a set of 4 Li-Poly rechargeable batteries. The batteries are each rated for 6000 mAh discharge current.

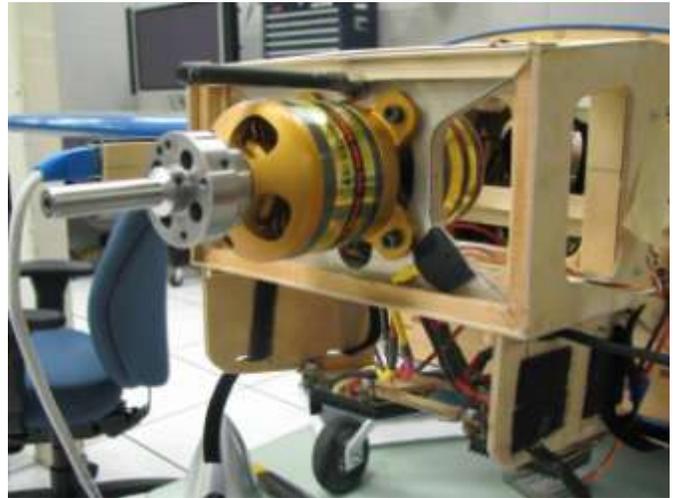


Figure 4 – Electric out runner motor

A 12 channel JR radio system is used to control the airplane. The system communicates in the 2.4 GHz band using a proprietary DSM2 protocol. Control surfaces are manipulated by seven actuators. The tandem motors are each controlled by separate motor controllers.

The airplane is equipped with a number of sensors to collect structure, propulsion, and navigation. The health of the structure is monitored using a series of strain gauges and accelerometers. Navigation data consist of GPS location, ground speed, altitude, true heading, and magnetic heading. Power plant data are intended to help assess adequacy of the thrust from the motors and 26 inch propeller, the relevant parameters being motor RPMs, currents, battery voltages and temperatures.

Data from sensors are logged by 2 separate data systems. The asynchronous systems are identified as RCATS and PC104. The RCATS is a turn-key system with proprietary software that creates an ASCII log of connected sensors. It also telemeters data to a laptop receiver which displays the data for callout to the pilot. It measures flight related parameters such as motor RPM, motor temperature, airspeed, z-axis acceleration at 10 Hz and interleaves GPS position and altitude data at 1 Hz. The PC104 stack consists of a CPU board, a DC/DC converter, an IO card, and a signal conditioning card for strain gauges. It runs MathWork's xPC Target operating system. The data in this release is acquired using a Simulink model that is compiled to an xPC target OS. This system records strain, accelerometer, battery temperature, and motor current at 200 Hz. It also outputs a 0.5 Hz sine wave to the RCATS system for synchronizing data.

5. BATTERY HEALTH MANAGEMENT HARDWARE

The BHM system is designed to be a relatively low cost analog-to-digital data acquisition system. The design

philosophy behind the first BHM system is to use COTS solutions which would have a light weight compact footprint. A second BHM system will then be designed around a custom printed circuit board (PCB) layout, which would yield an even smaller and lighter package.

There are three major elements to the BHM system - a signal conditioning board, an analog-to-digital acquisition board, and an embedded processor board. The signal conditioning board processes the analog sensor signals for the analog-to-digital acquisition board. The analog-to-digital acquisition board feeds the processed signal data into an embedded processor board, where the data are fed into the battery prognostics algorithm. Finally the embedded processor board outputs the battery prognostics algorithm results on an RS-232 data stream.

The signal conditioning board for the first BHM system was a board used in a past project, which was repurposed for the BHM system. As it was not specifically designed for the BHM purposes, these boards actually are larger than necessary, and cause the overall BHM system to be oversized and overweight. The second BHM system, with its custom PCB, would be about half the size.

The analog-to-digital acquisition board used in the BHM system is the SensorCore (SC) from Tern, Inc (www.tern.com). Several extra options were selected, such as 256KW SRAM, a Real-time clock, a CompactFlash interface, three 16 channel 24-bit ADC chips, a Low-drop regulator, and an 80-Mhz CPU upgrade. It was initially hoped that the SensorCore would be sufficient to run the battery prognostics algorithm, but it quickly became clear that a more powerful computing platform was necessary.

Gumstix, Inc provides an embedded processor system in its Overo™ computer-on-module (COM) which offers several attractive features for use in the BHM system – (i) it is very compact in size, (ii) it offers a reasonably powerful computing platform with its 600MHz Texas Instrument (TI) OMAP35x processor, (iii) it uses an open-source software development environment for low cost, and (iv) it can interface with the other BHM system components through industry standard RS-232 protocols. For the BHM system, the Overo Water COM and the Summit expansion board combination are used. Figure 5 shows the system installed onboard the Edge 540 airframe.

The Gumstix uses the openEmbedded development environment, which runs on an Ubuntu Linux laptop. All code for the Gumstix is developed in C. The Gumstix running environment is a custom boot image running on the onboard microSD card. The Gumstix is configured to be a turnkey environment, with data logging and processing occurring immediately after the Gumstix finishes booting up. The battery prognostic algorithm leverages POSIX threading for speed and efficiency. The Gumstix records raw data on the microSD card for post-run retrieval.



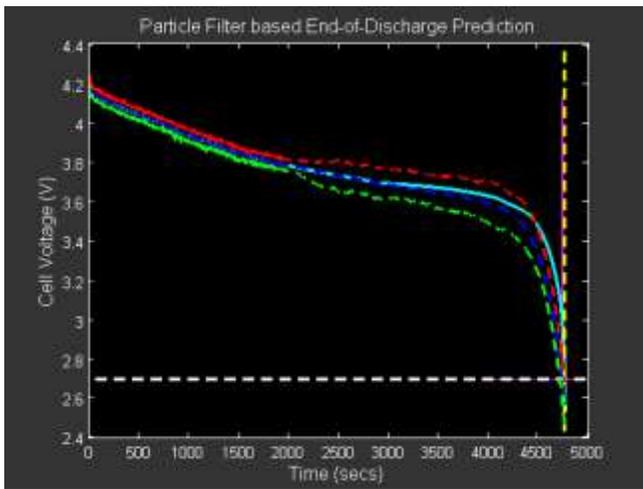
Figure 5 – The BHM system installed onboard

The first BHM system utilizes several small COTS boards to convert TTL signal voltages into PC RS-232 signal voltages. The second BHM system will incorporate these features onto its custom PCB, thereby reducing space and weight. The first BHM system has been used for the data used in this paper, recording 12 channels of data (4 battery voltages, 4 battery currents, and 4 battery temperatures) at rates up to 30 samples per second. The second BHM system is still being fabricated.

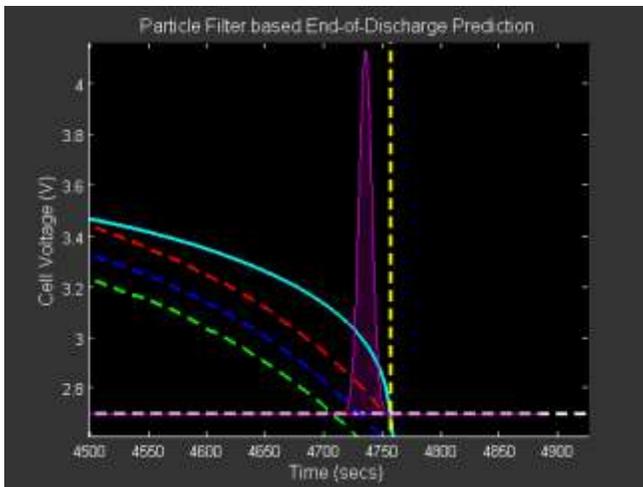
6. RESULTS

The tests to demonstrate the BHM methodology have been planned in three stages:

- (1) Battery discharge characterization tests on individual Li-Poly cells at the NASA Prognostics Center of Excellence battery prognostics testbed. Figure 6 shows the PF prediction of the EOD event of a Li-Poly cell. The solid cyan line shows the actual cell voltage. The aim is to predict when this curve crosses the EOD threshold 2.7V (dashed white line) shown by the vertical dashed yellow line. The PF mean is shown in blue and the upper and lower uncertainty limits are in red and green respectively. The tracking phase is shown in solid lines and the prediction curve is shown in dashed lines. The EOD RUL pdf is shown in magenta in Figure 6(b). The plots show the high accuracy and precision of the prediction when the prediction are made half way into the discharge.
- (2) The next stage of testing involved hardware-in-the-loop discharge runs with the Edge 540 restrained on the ground. The PF prediction is shown in Figure 7. The Edge 540 motors are run halfway through the discharge run before the prediction is made. The mean and lower and upper 95% confidence bounds for the EOD event are shown in magenta while the actual event is shown in yellow.



(a)



(b)

Figure 6 – (a) Particle Filter based EOD prediction; (b) zoomed in view of the EOD event

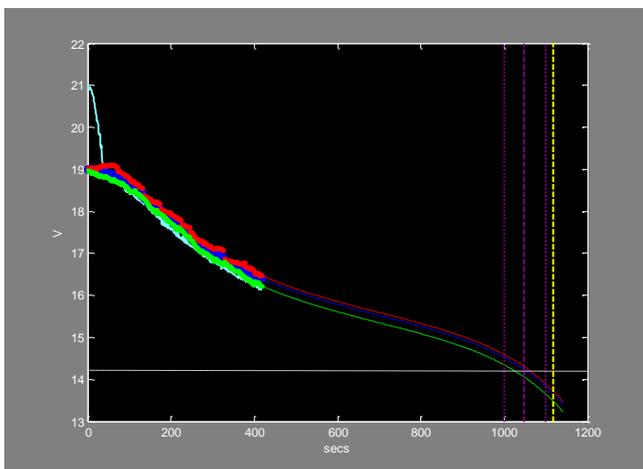


Figure 7 – Hardware-in-the-loop EOD prediction

- (3) The final test involved profiling the energy consumption in various flight regimes and making in-flight EOD predictions.

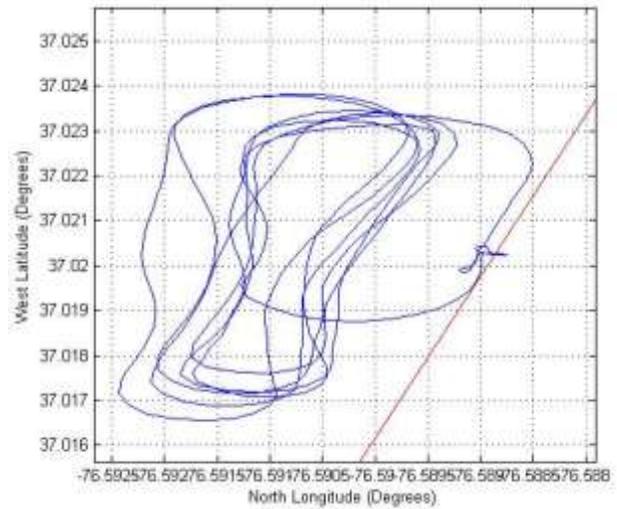


Figure 8 – Typical flight profile

Figure 8 shows a typical flight profile for the Edge 540, while Figure 9 shows the typical current loads in different flight regimes (x-axis in secs). This is a simulated flight test carried out on the ground using the UAV airframe. Actual flight tests are scheduled to be carried out in early 2011 dependent on weather. Figures 10 – 12 show further prediction results for this ground test.

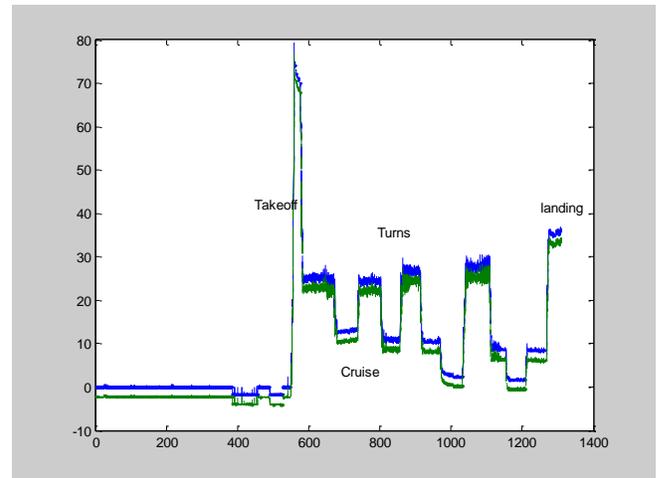


Figure 9 – Simulated flight current loads

Figure 10 shows the battery voltages measured during the ground test in solid lines, while the predicted voltage trajectories are shown in dotted lines. Although predictions can be made throughout the test, for the sake of clarity, we analyze one prediction made at the 13 minute mark. For this test, we found that the aft motor voltage channels had bias errors, so we used the forward (fwd) motor channels for making the EOD predictions.

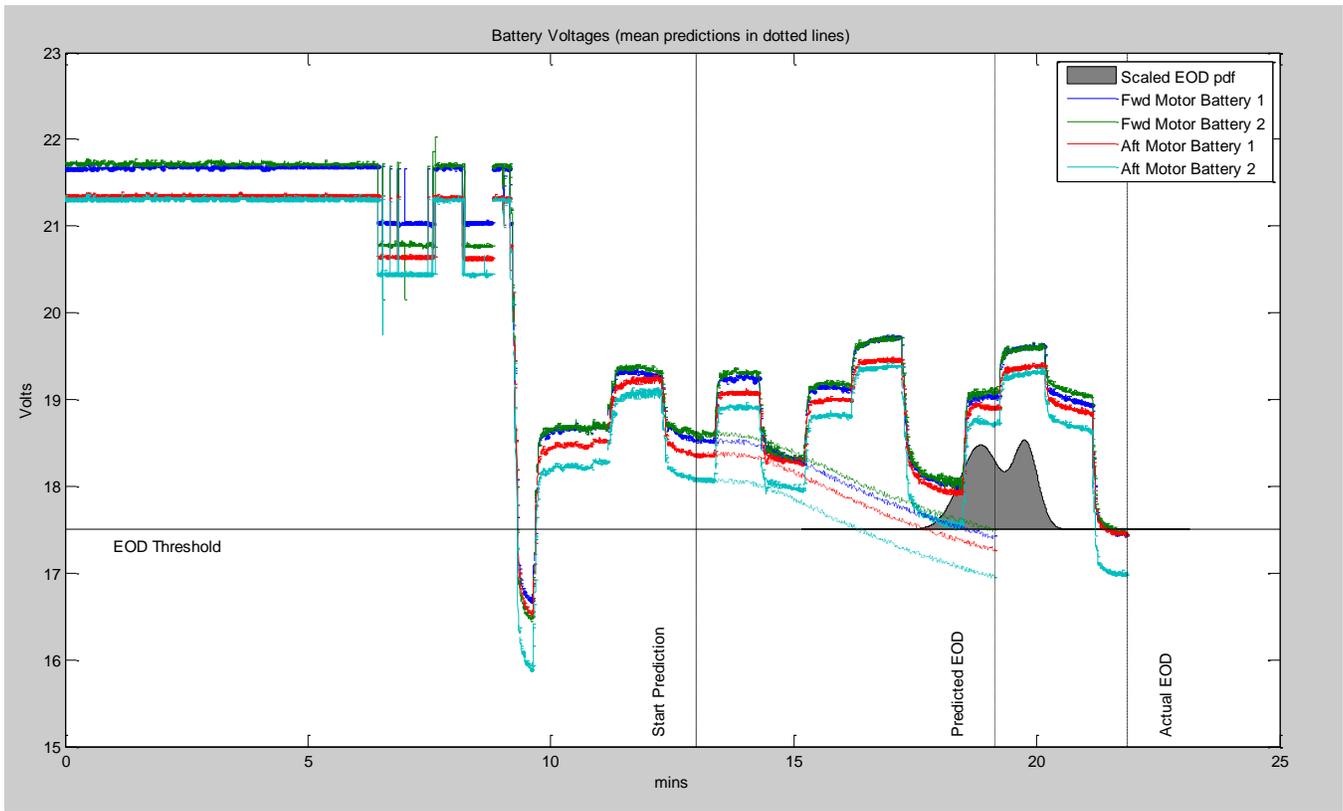


Figure 10 – Predicting EOD with average current loads

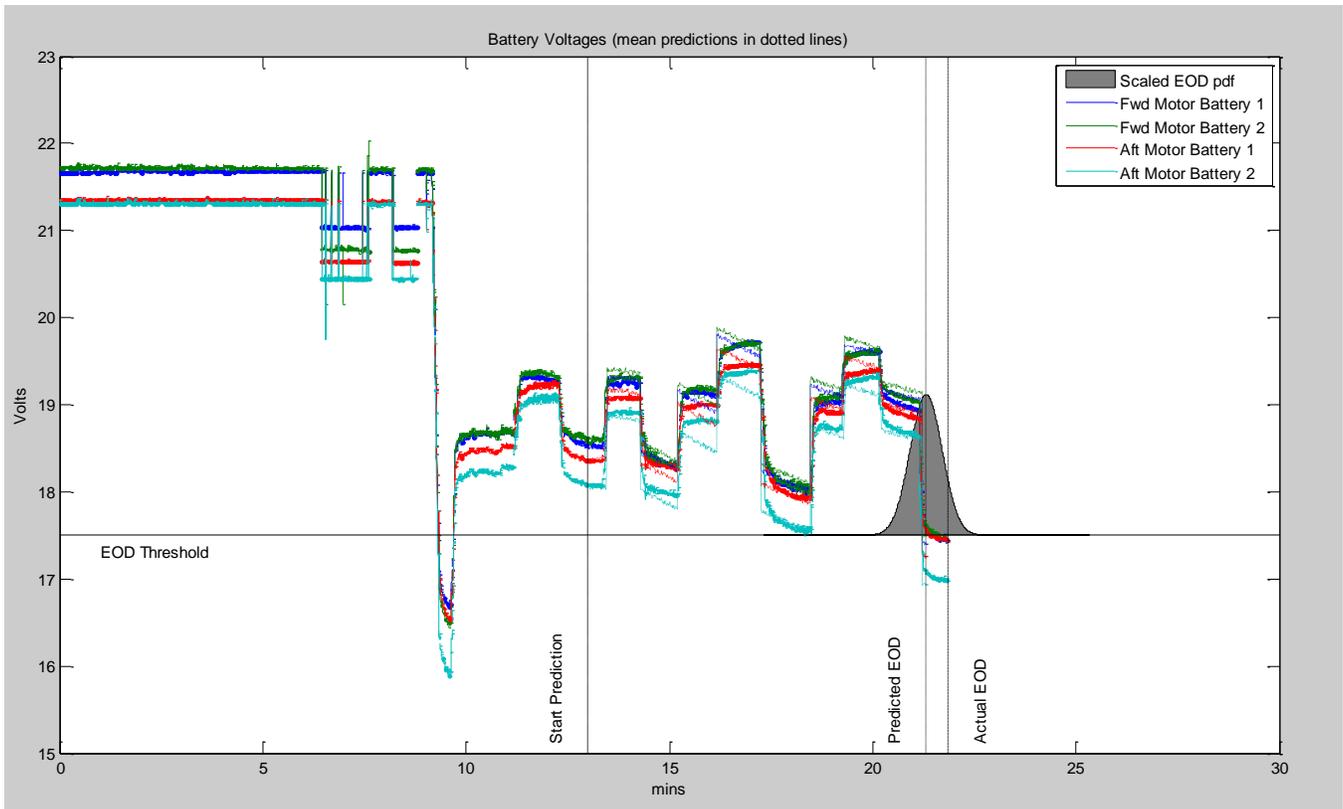


Figure 11 – Predicting EOD with known future current loads

For the predictions made in Figure 10, we use average flight currents loads and run the PF predictions until the forward motor batteries reach the cutoff voltage threshold of 17.5 V. This threshold is set by the electronic speed control (ESC) hardware in the UAV airframe. As seen from the plot in Figure 9, the test simulates some low power glide modes after the 13 minute mark. This means that the UAV has more flight time left than if it were consuming average power. Consequently, the EOD pdf shown in Figure 10 is conservative. On the other hand, if the future load profile is known, the PF model can account for the changes in load levels and make more accurate predictions, as shown in Figure 11.

It should be noted, however, that the future profile may not always be known, but even if a statistical distribution of future load levels is known, the PF framework can sample from that distribution to produce more accurate predictions than those produced by assuming average load levels. Also, if the intended future load profile is known, then the cumulative probability of the EOD occurring before the time that the intended mission profile ends (shown in red in Figure 12) gives the failure probability of the mission. This information may then be used for mission replanning purposes.

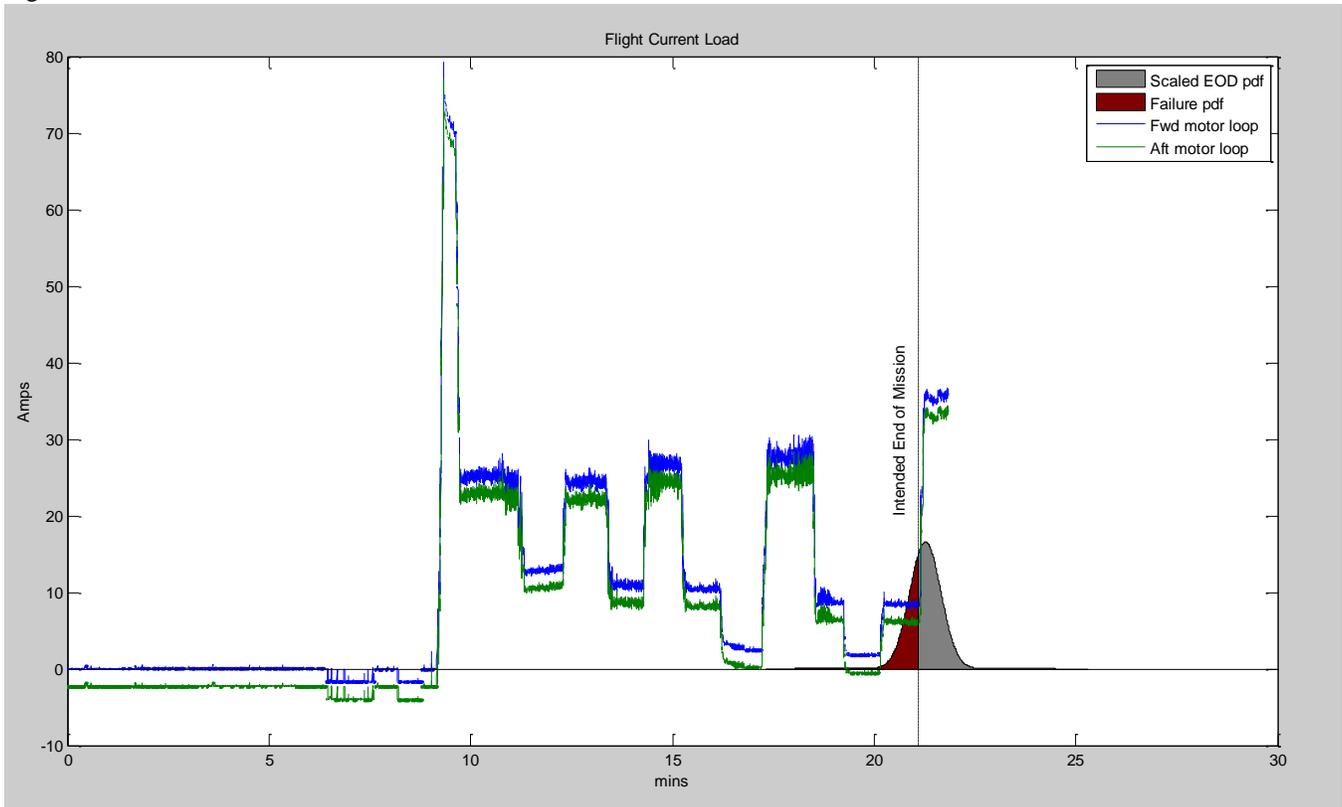


Figure 12 – Calculating mission failure probability from EOD pdf

7. CONCLUSIONS

In summary, this paper lays out a novel battery health management technique for application onboard an electric UAV. This technique is also applicable to other electric vehicles, like electric cars. Electric cars are not new; they were more popular and reliable at the turn of the twentieth century as compared to their gasoline powered rivals. However, the uncertainty in determining battery life plagued electric vehicles then as it does now. A recent report by the Consumer Electronics Association, “Electric Vehicles: The Future of Driving”, indicates that although these vehicles are increasing in popularity, running out of battery power on the road is the top concern for consumers (71% of adults

surveyed). Thus developing advanced battery health management techniques is very important.

The approach presented here is model-based where the form of the model has been linked to the internal processes of the battery and validated using experimental data. Subsequently, the model has been used in a PF framework to make predictions of EOD effectively. Since the prediction result is in the form of a pdf, it is easy to integrate the BHM routine into a higher level decisioning algorithm that can provide advance warning about when to land the aircraft.

By profiling the power required for different flight regimes like cruise segments, banked turns and landings, we can estimate the mission completion probability by calculating

the RUL pdf. This in turn can motivate and inform mission replanning activity. The PF prognostic framework can also be used to evaluate the effectiveness of possible mitigation actions, thus improving operational safety and optimizing battery power.

REFERENCES

- [1] 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.
- [2] R. L. Hartmann II, *An Aging Model for Lithium-Ion Cells*, PhD dissertation, University of Akron, 2008.
- [3] R. Huggins, *Advanced Batteries: Materials Science Aspects*, 1st ed., Springer, 2008.
- [4] B. Saha and K. Goebel, "Modeling Li-ion Battery Capacity Depletion in a Particle Filtering Framework", in *Proceedings of the Annual Conference of the Prognostics and Health Management Society 2009*, San Diego, CA.

BIOGRAPHY



Bhaskar Saha received his Ph.D. from the School of Electrical and Computer Engineering at Georgia Institute of Technology, Atlanta, GA, USA in 2008. He received his M.S. also from the same school and his B. Tech. (Bachelor of Technology) degree from the Department of Electrical Engineering, Indian Institute of Technology, Kharagpur, India. He is currently a Research Scientist with Mission Critical Technologies at the Prognostics Center of Excellence, NASA Ames Research Center. His research is focused on applying various classification, regression and state estimation techniques for predicting remaining useful life of systems and their components, as well as developing hardware-in-the-loop testbeds and prognostic metrics to evaluate their performance. He has been an IEEE member since 2008 and has published several papers on these topics.