

Verification of Prognostic Algorithms to Predict Remaining Flying Time for Electric Unmanned Vehicles

Edward F. Hogge¹, Brian M. Bole², Sixto L. Vazquez³, Chetan S. Kulkarni⁴, Thomas H. Strom⁵, Boyd L. Hill⁶, Kyle M. Smalling⁷, and Cuong C. Quach⁸

^{1,5,7}*National Institute of Aerospace, Hampton, Virginia 23666*

edward.f.hogge@nasa.gov

thomas.h.strom@nasa.gov

kyle.m.smalling@nasa.gov

²*ARMUS Corporation, San Mateo, California 94404*

⁴*Stinger Ghaffarian Technologies, Inc., NASA Ames Research Center, Moffett Field, California 94035*

chetan.s.kulkarni@nasa.gov

^{3,8}*NASA Langley Research Center, Hampton, Virginia 23681*

sixto.l.vazquez@nasa.gov

cuong.c.quach@nasa.gov

⁶*Analytical Mechanics Associates, Inc., NASA Langley Research Center, Hampton, Virginia 23681*

boyd.l.hill@nasa.gov

ABSTRACT

This paper addresses the problem of building trust in the online prediction of a eUAV's remaining available flying time powered by lithium-ion polymer batteries. A series of ground tests are described that make use of an electric unmanned aerial vehicle (eUAV) to verify the performance of remaining flying time predictions. The algorithm verification procedure described is implemented on a fully functional vehicle that is restrained to a platform for repeated run-to-functional-failure (charge depletion) experiments. The vehicle under test is commanded to follow a predefined propeller RPM profile in order to create battery demand profiles similar to those expected during flight. The eUAV is repeatedly operated until the charge stored in powertrain batteries falls below a specified limit threshold. The time at which the limit threshold on battery charge is crossed is then used to measure the accuracy of the remaining flying time prediction. In our earlier work battery aging was not included. In this work we take into account aging of the batteries where the parameters were updated to make predictions. Accuracy requirements are considered for an alarm that warns operators when remaining flying time is

estimated to fall below the specified limit threshold.

1. INTRODUCTION

Improvements in battery storage capacity have made it possible for general aviation vehicle manufacturers to consider electrically-powered solutions. The development of trust in battery remaining operating time estimates, however, is currently a significant obstacle to be overcome when considering adoption of electrical propulsion systems in aircraft (Patterson, German & Moore, 2012). There are several ways in which predicting remaining operating time is more complicated for battery-powered vehicles than it is for vehicles with a conventionally-powered liquid-fueled combustion system. Unlike a liquid-fueled system, where the fuel tank's volume remains unchanged over successive refueling procedures, a battery's charge storage capacity will diminish over time. Another complicating feature of a battery system is the time-varying relationship between battery output power and battery current draw. Whereas a conventional liquid combustion system uses an approximately constant amount of liquid fuel to produce a given motive power, the power from a battery system is equal to the product of battery voltage and current. Thus, as batteries are discharged, their voltages drop lower, and they will lose charge at a faster rate.

Edward Hogge et al. This is an open-access article distributed under the terms of the Creative Commons Attribution 3.0 United States License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.

There are a number of methods used to model the voltage drop of batteries as they are discharged. A survey of the theoretical issues using the extended Kalman filter method for nonlinear state of charge estimation is found in Wang, Fang, Zhou and Wada (2017). For the small electric aircraft use case, our previous papers introduced several tools for battery discharge prediction onboard. One paper described a battery equivalent circuit model to simulate battery state (Bole, Teubert, Quach, Hogge, Vazquez & Goebel, 2013). The model's battery capacity, internal resistance and other parameters were identified through two laboratory experiments that used a programmable load. In the first experiment, batteries were slowly discharged at 1/50 C constant current to an open circuit voltage of 16V. (1C is the current necessary to completely discharge the battery in one hour). In the other experiment a repeated pulsed loading was done. Current and voltage profiles logged during flights of a small electric airplane further tuned the battery model (Quach, Bole, Hogge, Vazquez, Daigle, Celaya, Weber & Goebel, 2013). The use of a flight plan to define the energy required to complete the mission as well as upper and lower uncertainty bounds was presented along with an approach to identify additional parasitic battery loads (Bole, Daigle & Gorospe, 2014). Another paper introduced a verification testing procedure that is intended to build trust in predictions of remaining flying time prior to actual flight testing (Hogge, Bole, Vazquez, Celaya, Strom, Hill, Smalling & Quach, 2015). This article is a further discussion of the verification testing of remaining flying time prior to flight testing. The philosophy behind the testing procedure described here is to translate system performance and safety goals into requirements for an alarm that warns system operators when the estimated remaining flying time falls below a certain threshold. Ground testing of the actual vehicle provides the closest possible testing conditions short of actual flight and captures some of the variation that the powertrain hardware and that the pilot introduce while avoiding the risks inherent in flight. For instance, the batteries may be drained to a lower capacity during testing of the remaining flying time prediction without danger of vehicle loss.

A large electric unmanned aerial vehicle (eUAV) was used in this study. The eUAV is a 33 percent sub-scale version of the Zivko Aeronautics Inc. Edge 540 T tandem seat aerobatic aircraft (Fig. 1). This vehicle has been actively used by researchers at NASA LaRC to facilitate the rapid deployment and evaluation of remaining flying time prediction algorithms for electric aircraft since 2010. Examples of prior works using this platform are found in the following papers: (Saha, Koshimoto, Quach, Hogge, Strom, Hill, Vazquez & Goebel, 2011), (Hogge, Quach, Vazquez & Hill, 2011), (Daigle, Saxena & Goebel, 2012), and (Bole et al., 2013).

Remaining flying time prediction algorithms focus on the prediction of battery charge depletion over a eUAV flight. A



Figure 1. The Edge 540 T Rapid Evaluation eUAV

lower-bound on the battery state of charge (SOC) that is considered safe for flight is set at 30 percent in this work. Flying the vehicle with batteries below 30 percent SOC is considered to be a high-risk mode of operation. Policy and guidelines are set according to the rulings and the engineering judgment of the NASA Langley UAS Operations Office and the NASA Langley Airworthiness and Safety Review Board. Such violations of operating guidelines are referred to here as a functional failure of the vehicle's assigned mission. The primary use case for remaining flying time prediction is to warn system operators when landing procedures must be initiated to avoid aircraft batteries depleting below a set threshold limit. It was determined that initiating landing procedures when the eUAV batteries reach 30 percent SOC would provide a sufficient energy buffer for at least two "missed approach" maneuvers without risk of exceeding the battery current limits and any associated excessive heating. This was based upon operator's experience and upon ground tests. The predictive element to be tested in this work is an alarm that warns system operators when the powertrain batteries are two minutes from reaching the 30 percent SOC threshold under normal operating conditions. This should allow the pilot sufficient time to prepare the eUAV for landing without exceeding a moderate work load.

An equivalent circuit model is used in our work. Equivalent circuit models are computationally efficient and are popular for engineering applications. Examples include Ceralo (2000) and Chen & Rincon-Mora (2008). They have the disadvantage of limited accuracy due to approximations to battery internal chemical mechanisms, and as a result do not account for changes due to aging. An analysis of accuracy issues in the battery capacity and internal resistance estimation process can be found in Lin & Stefanopoulou (2015). The accuracy of the onboard remaining flying time estimation algorithms is tested in this work. A series of controlled run-to-functional-failure (charge depletion) experiments were conducted. The vehicle under test was strapped down to a platform and commanded to follow an RPM profile that created battery demand profiles similar to

those expected for flight while a ground station operator monitored the battery health parameters.

The time it takes for powertrain batteries to reach 30 percent SOC establishes a truth value for the functional failure time. Unlike actual flight tests, powertrain batteries can be repeatedly run down to their lower-limits in the ground-based testing described here.

The defined performance requirements are then verified by repeating ground tests a specified number of times. The performance requirement testing procedure used here was originally introduced in Saxena, Roychoudhury, Lin and Goebel (2013).

Section 2 of this paper provides an overview of the Edge 540T powertrain. Algorithms used for onboard battery state estimation and remaining flying time predictions are summarized in Section 3. The process used to verify onboard remaining flying time predictions through ground testing and experimental results are described in Section 4. Finally, concluding remarks are given in Section 5.

2. OVERVIEW OF EDGE 540T POWERTRAIN

A wiring diagram for the vehicle powertrain is shown in Fig. 2. The aircraft has two 3-phase tandem motors that are mechanically coupled to the aircraft propeller. Powertrain batteries are arranged in two pairs of series connected battery packs. A switchable parasitic load R_p injects a fault to test the robustness of the remaining flying time estimation algorithms to changes in battery loading demand.

Remaining flying time predictions are generated by propagating a number of estimates of the battery charge forward. Forward propagation of the present battery state estimate is performed using an estimate of the future powertrain demand that will occur over the known flight plan. These future loads include propeller loads and parasitic loads. The prognostic tools make use of the known

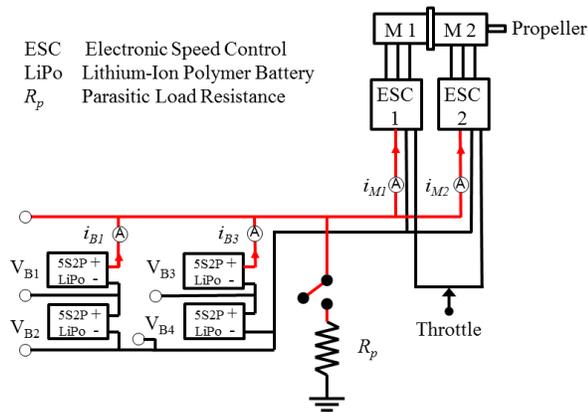


Figure 2. Schematic of electric powertrain

flight plan to inform future load predictions, but no prior information is assumed to be available regarding when a parasitic load may be injected.

3. REMAINING FLYING TIME PREDICTION

Battery discharge prediction is described here in terms of the following components; (i) online battery state estimation; (ii) prediction of future battery power demand as a function of an aircraft flight plan; (iii) online estimation of additional parasitic battery loads; and (iv) prediction of battery discharge over the future flight plan. The assumptions and algorithms used for each of these steps are summarized in this section.

3.1. Online Battery State Estimation

Our previous papers (Quach et al., 2013) and (Bole et al., 2014), described the use of an equivalent circuit model and an unscented Kalman filter (UKF) (Julier & Uhlmann, 1997, 2004) to update battery state estimates based on observations of current and voltage at the battery output terminals. This approach is also summarized here for convenience. An equivalent circuit battery model in Fig. 3 is used to represent battery terminal voltage dynamics as a function of battery current. It is similar to models presented in (Chen & Rincon-Mora, 2006, and Ceralo, 2000). The model is based on Thevenin's theorem to model the current and voltage profile of the battery as a black box input-output device. A first-approximation assumption is made such that the battery state can match a linear electrical network with voltage and current sources and only resistances. Thevenin states that the black box can be replaced at the input output terminals by an equivalent voltage source in series connection with an equivalent resistance. To better match standard battery phenomenon, such as internal resistance voltage drops and hysteresis effects, additional pairs of series connected RC parallel circuits are added to the model. The R_s, C_s pair are added for the internal resistance drop and the R_{cp}, C_{cp} pair are added for the concentration polarization effect. The correspondence of these RC circuits to actual

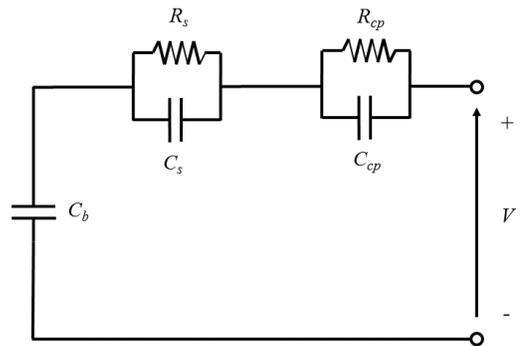


Figure 3. Lithium-Ion battery equivalent circuit model

battery chemical phenomena is only notional. Models that better account for the electrochemical behavior and aging effects are being considered for future work. (Daigle and Kulkarni, 2013). See section 4.4 for further discussion. In the equivalent circuit model, some of the components were made to vary according to the bulk charge stored in C_b as described in (Zhang and Chow, 2010). The State of Charge (SOC) is a battery charge estimate of the bulk charge. The battery input-output voltage dynamics will change as a function of this bulk charge estimate. Battery SOC is defined here as:

$$SOC = 1 - \frac{q_{max} - q_b}{C_{max}} \quad (1)$$

Where q_b represents the charge stored in capacitor C_b , q_{max} is the maximum charge that the battery can hold, and C_{max} is the maximum charge that can be drawn from the battery in practice. This battery model contains six electrical components that are tuned to recreate the observed current-voltage dynamics of the Edge-540T battery packs. The following SOC parameterizations (Bole et al., 2014) were used to model the bulk charge influence on the C_b , C_{cp} , and R_{cp} circuit elements of Fig. 3:

$$C_b = C_{Cb3} + C_{Cb2}SOC + C_{Cb1}SOC^2 + C_{Cb0}SOC^3 \quad (2)$$

$$C_{cp} = C_{Cp0} + C_{Cp1} \exp(C_{Cp2}(SOC)) \quad (3)$$

$$R_{cp} = R_{Cp0} + R_{Cp1} \exp(R_{Cp2}(SOC)) \quad (4)$$

These parameters (C_{max} , R_s , R_{Cp0} , R_{Cp1} , R_{Cp2} , C_{Cp0} , C_{Cp1} , and C_{Cp2}) were identified by fitting a pulsed discharge laboratory experiment voltage profile shown in Fig. 4 with a Nelder-Mead downhill simplex method solution search that minimizes the error between the modeled and actual voltage profile (Nelder & Mead, 1965).

These identified parameters are associated with a selected battery from a batch of batteries of a given chemical formulation. These parameters are assumed to be unvaried across all similar battery packs of a given batch. Any differences in individual batteries due to manufacturing variation is accounted for by adaptation of the battery charge capacity term C_{max} of the C_b capacitor in the equivalent circuit model. C_{max} is identified by running a 1/50 C discharge cycle for each battery pack as shown in Fig. 5. During this low current discharge cycle, the voltage across the C_b capacitor plays a dominant role. Thus, this experiment allows the C_{max} parameter in the equivalent circuit model to be fitted in isolation, also through use of the Nelder-Mead simplex method (Bole et al., 2014). According to the SOC definition (equation 1) C_{max} will always be less than q_{max} , due to electrochemical side-reactions that make some portion of a battery's charge carriers unavailable.

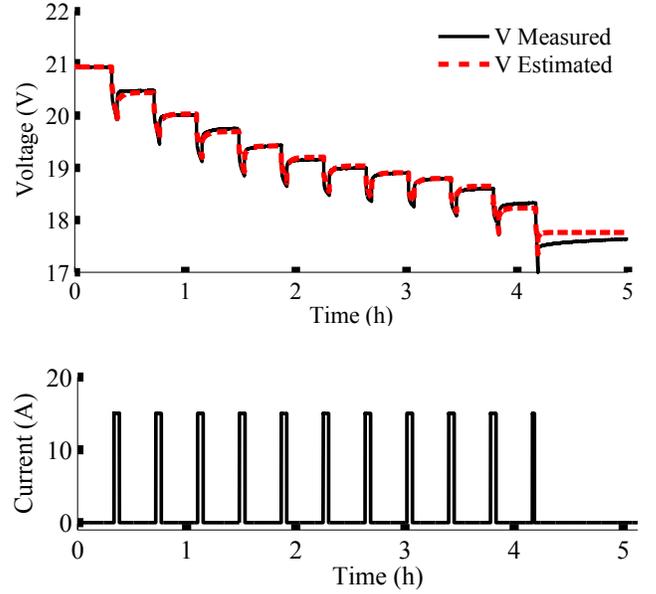


Figure 4. Comparison between measured and predicted battery voltage over a pulsed current discharge (Bole et al., 2014) used by permission.

As the battery ages more of its internal charge will become unavailable because of these side reactions. The C_{max} parameter must be refitted periodically to capture this effect (we use 10 recharge cycles between refits). The C_{max} and R_s parameters obtained were used in a simulation run using the current profile from a chamber run. The SOC battery plots were examined to see if the SOC estimates remained constant while the battery voltage recovered in the recorded rest period after the run. If they remained nearly constant, they were used. If estimated SOC was seen to rise during the rest period, it indicated that R_s was too low and needed

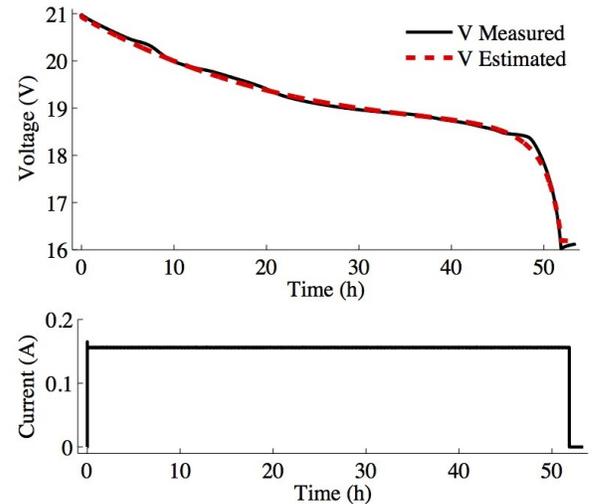


Figure 5. Comparison between measured and predicted battery voltage over a low current discharge. (Bole et al., 2014) used by permission.

to be increased. Likewise, if SOC was seen to fall, then R_s was too high and needed to be decreased. This R_s selection was repeated until the SOC estimates remained constant after a run. This empirical process was necessary because the constant current discharge poorly identified R_s . An observability issue that may be the root-cause of this is suggested in Lin & Stefanopoulou, 2015. Their analysis indicates that the accuracy of the R_s measurement depends having a high current for the test condition. The initial characterization of a set of batteries at purchase used a relative high current spike discharge as shown in Fig. 4 for a representative sample. Use of Maccor 4000 battery cycling equipment made this high-current measurement easy. At the eUAV operations location all the batteries need routine refitting after 10 recharges, a logging battery charger which has an effective limit of only 4A for the high current logged discharge is available. A special bench setup would be a way to perform this R_s measurement.

3.2. Unscented Kalman Filter

Once all battery parameters are fitted, the UKF is used to update model projections of the battery output voltage with past measurement data in a tracking mode. The UKF takes in the measured battery current and voltage, and gives probability distributions for the charge states of each of the three capacitors in the equivalent circuit model as components of a state vector.

The UKF is a tool for computing probabilistic belief in system state estimates based on stochastic models of the system's dynamics. The UKF assumes a general nonlinear form of the state and output equations, and efficiently propagates model and state uncertainties. The UKF employs an unscented transform (UT) using a minimal set of weighted samples, called sigma points, whose mean and covariance are preserved once transformed by the unscented transform and the nonlinear battery model function (Julier & Uhlmann, 1997, 2004). The UKF takes battery power demand (current) as a controlling input to the system, and the measured battery voltage from the previous time step. The UKF gives a probability distribution for charge state variables and the future voltage state output distribution. The SOC distribution can be directly derived from this. The number of sigma points required is minimal as compared to particle filters that require an order of magnitude more random variable sampling. The future system state must be simulated until a given cut-off threshold is reached for each sampled particle. (Daigle et al., 2012). Readers interested in the application of UKF to the estimation of battery SOC are referred to our previous papers (Bole et al., 2013; Daigle et al., 2012) and the references therein. The Bayesian use of actual past system behavior makes the model-based filtering approaches such as UKF much less susceptible to initialization and measurement errors than the coulomb counting method currently used in many battery monitoring systems (Dai et al., 2006).

3.3. Prediction of Motor Power Demand as a Function of Aircraft Flight Plan

After estimating battery state, the next step towards predicting remaining flying time is the estimation of motor power demand over the remainder of a given flight plan. The aircraft's flight plan is assumed here to be specified in advance in terms of fixed airspeed segments. Each segment includes a desired vehicle airspeed along with an expected duration or other ending condition. An example flight plan is defined here as:

1. Takeoff and climb to 200 m:
Set airspeed to 25 m/s, hold for 1.0 min
2. Maintain altitude, airspeed:
Set airspeed to 23 m/s, hold for 3.0 min
3. Maintain altitude, increase airspeed:
Set airspeed to 25 m/s, hold for 2.0 min
4. Maintain altitude, decrease airspeed:
Set airspeed to 20 m/s, hold for 2.0 min
5. Maintain altitude, increase airspeed:
Set airspeed to 23 m/s, hold until landing is called by monitors on the ground.
6. Remote control landing: airspeed and duration may vary widely depending on pilot behavior and environmental conditions.

The energy required for an aircraft to fly the remainder of a given flight plan will necessarily be uncertain due to random variation in pilot behavior and environmental conditions. Data from previous flights of a similar vehicle was used to estimate the mean motor power (current) required for the flight plan. A plus or minus 30 percent variation about this mean value seemed to contain most of the flight plan maneuver power variation (Bole et al., 2014).

The motor power demand is estimated for each flight plan segment using a previously developed reduced order powertrain and aerodynamic drag model, discussed in Bole et al. (2013) and in Bole et al. (2014). The motor power demand was assumed to be distributed uniformly between 30 percent variation limits. Future motor power (current demand) can be described by a parameterized family of functions based upon the future flight plan, drawn from the uniform distribution about the mean power value. The equivalent circuit model, the powertrain and aerodynamic drag model, and UKF are used to simulate the flight plan demand for each input function into the future until an assumed SOC cut-off threshold is reached (Daigle et al., 2012). As an alternative to exhaustive sampling or random sampling, the UT can also be used to sample the input power distribution random variable $u \in \mathbb{R}^{n_u}$ that has mean of \bar{u} and covariance \mathbf{P}_{uu} . If the UT is used to sample the input power demand trajectories, the sigma points computed happen to coincide with the input power distribution's

maximum and minimum bounds and its mean value. Since the most likely value of remaining flying time and its bounds are of primary interest for this application, this is a significant saving in computation. The UT uses a minimal set of weighted samples, sigma points U^i whose mean and covariance are preserved once transformed. Daigle et al., 2012, states that there is no guidance for choosing a value for the sigma point free parameter κ when the distribution is uniform. The κ for a Gaussian distribution was used for this one-dimensional input space of parameterized flight plan power train input functions \mathbb{R}^{n_u} which is a κ of 2. The symmetric unscented transform method is used in our case (Daigle & Goebel, 2010), (Julier & Uhlmann, 2004). In the symmetric unscented transform, the $2n_u + 1$ sigma points are selected about the mean in the following way:

$$w^i = \begin{cases} \frac{\kappa}{n_u + \kappa}, & i = 0 \\ \frac{\kappa}{2(n_u + \kappa)}, & i = 1, \dots, 2n_u \end{cases} \quad (5)$$

$$U^i = \begin{cases} \bar{u}, & i = 0 \\ \bar{u} + \left(\sqrt{(n_u + \kappa)\mathbf{P}_{uu}}\right)^i, & i = 1, \dots, n_u \\ \bar{u} - \left(\sqrt{(n_u + \kappa)\mathbf{P}_{uu}}\right)^i, & i = n + 1, \dots, 2n_u \end{cases} \quad (6)$$

Where $\left(\sqrt{(n_u + \kappa)\mathbf{P}_{uu}}\right)^i$ refers to the i th column of the matrix square root of $(n_u + \kappa)\mathbf{P}_{uu}$. In addition, Daigle et al. 2012 states in the equivalent circuit battery model case study considered, the model was more sensitive to changes in the input than to changes in the process noise, concluding that the process noise could be assumed to be zero without significant effect. This method was used to generate the minimum, maximum, and mean power demand predictions shown in Fig. 6. These three power estimates can then be integrated to form predictions of the minimum, maximum, and mean motor energy consumption over the remaining flight plan. Figure 6 shows sample predictions of future motor power and energy demand over segments 1-5 of the given flight plan. Here, segment 5 of the flight plan is shown to extend out indefinitely (20 min.), representing the intent to continue flying until the ground team calls for a landing.

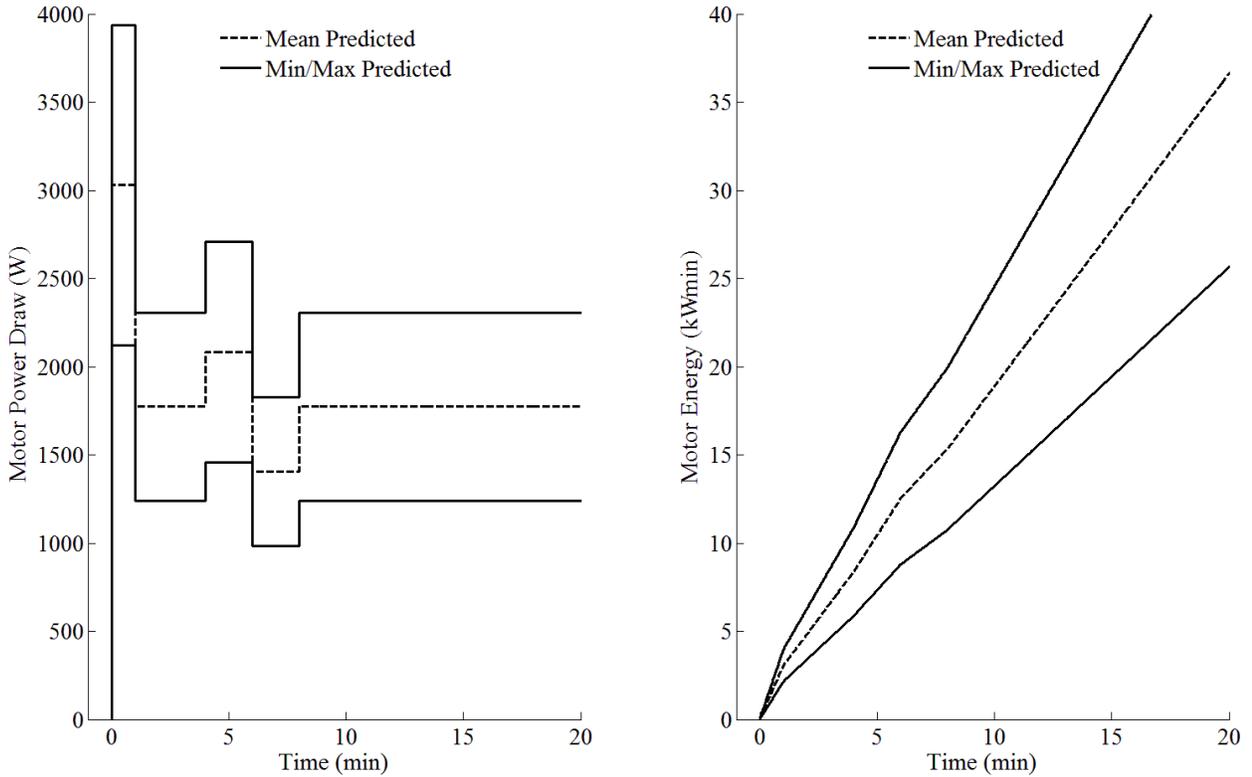


Figure 6. Uncertain predictions of motor power and energy draw over the sample flight plan

The minimum battery SOC required to safely land the aircraft with adequate reserve energy remaining is considered to limit the aircraft's maximum safe flying time. Prediction of available flying time remaining can thus be considered as the time until the battery SOC reaches 30 percent, assuming that a landing will not be called until the last possible moment. A triplet of minimum, maximum, and median remaining flying time estimates will ultimately be produced by estimating when the battery SOC threshold would be reached for each of the minimum, maximum, and mean motor power profiles.

3.4. Online Estimation of Additional Parasitic Battery Loads

Parasitic demands on the battery system that cannot be known in advance are simulated with a resistive load that may be injected in parallel with the aircraft batteries at any time during flight. Let R_p be the unknown parasitic load. The parasitic current, i_p , is the difference in the current i measured at the battery and the current i_m measured at the motor controller. The locations of the battery current sensors i_{B1} and i_{B2} for battery current i and the motor current sensors i_{M1} and i_{M2} for motor current i_m are found in Fig. 2. A residual, defined as the difference between an observed signal and its model-predicted value, can be defined for the parasitic fault detection based on the measured values of i and i_m . In the nominal case, our model for i is $i = i_m$. We can then define a residual, r_i , as $r_i = i^* - i_m^*$, where the *superscript indicates a measured value. Nominally, $r_i = 0$, and we can define a simple threshold-based fault detector that triggers when $r_i = 0$ for some threshold T . Once a fault is detected, we can estimate the parasitic current at time k using

$$\hat{i}_p(k) = i^*(k) - i_m^*(k). \quad (7)$$

The parasitic resistance can then be estimated with Ohm's Law

$$\hat{R}_p(k) = \frac{v_b^*(k)}{\hat{i}_p}. \quad (8)$$

The estimate $\hat{R}_p(k)$ will be noisy, since it is computed based on measured values. Assuming that R_p is constant, we take the median of all computed values to provide a robust estimate of R_p , i.e.

$$R_p(k) = \text{median}(\{\hat{R}_p(k_j) : k_d \geq k_j \geq k\}), \quad (9)$$

where k_d is the time of fault detection (and the time that fault identification begins). This online filtering routine is described further in Bole et al. (2014).

A battery current profile and parasitic load estimates from a sample aircraft data set is shown in Fig. 7. Here, a 5.5 Ω parasitic load is injected in parallel with the aircraft batteries at 5 minutes into the run. The time at which the parasitic load is injected is shown with a dashed line on the third

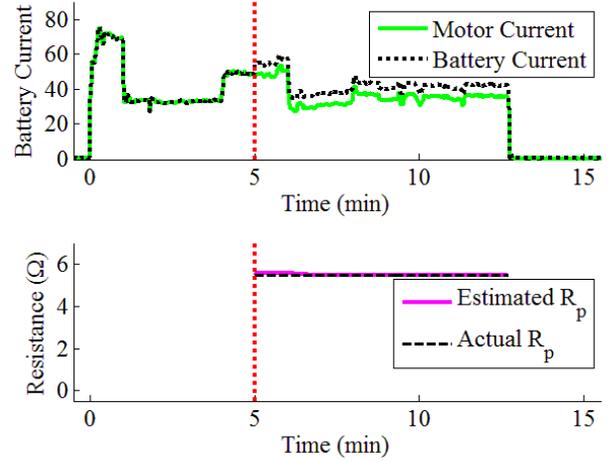


Figure 7. Sample motor and battery current profiles (top), along with parasitic load estimates (bottom)

column of plots in Fig. 8. At the time the load is injected, the battery current is seen to become notably higher than the motor current. The estimated parasitic load is then seen to rapidly converge to approximately 5.5 Ω . Online parasitic load estimates are directly incorporated into battery discharge predictions. This results in an immediate shift in battery discharge predictions each time the parasitic load estimate is updated. This immediate shift in discharge predictions is demonstrated in the following subsection.

3.5. Prediction of Battery Discharge Over a Flight Plan

Figure 8 shows plots of measured and predicted battery current, voltage, and SOC at three sample times over the battery discharge run. The minimum, median, and maximum predictions are plotted from each sample time until the predicted SOC reaches 30 percent. The predictions made at the first two sample times occur prior to injection of the parasitic load. These predictions are seen to overestimate the future battery current loads, resulting in underestimation of future battery voltage and SOC. The parasitic load has been detected by the third sample time, and the predictions at that time are seen to increase the degree of over-estimation of current demand which the model uses to underestimate the battery voltage, and SOC.

Figure 9 shows predictions of remaining flying time for the example run shown in Fig. 8. The solid line in Fig. 9 indicates the true flying time remaining from ground truth done after the run. The dashed line in Fig. 9 represents the median remaining time prediction. A median rather than mean is used to estimate the remaining flying time in order to be consistent with assessment when there is an ensemble of many simulated estimates. When distributions containing outliers are used, the median is a more robust predictor than the mean (Hoaglin, Mosteller and Tukey, 1983). The

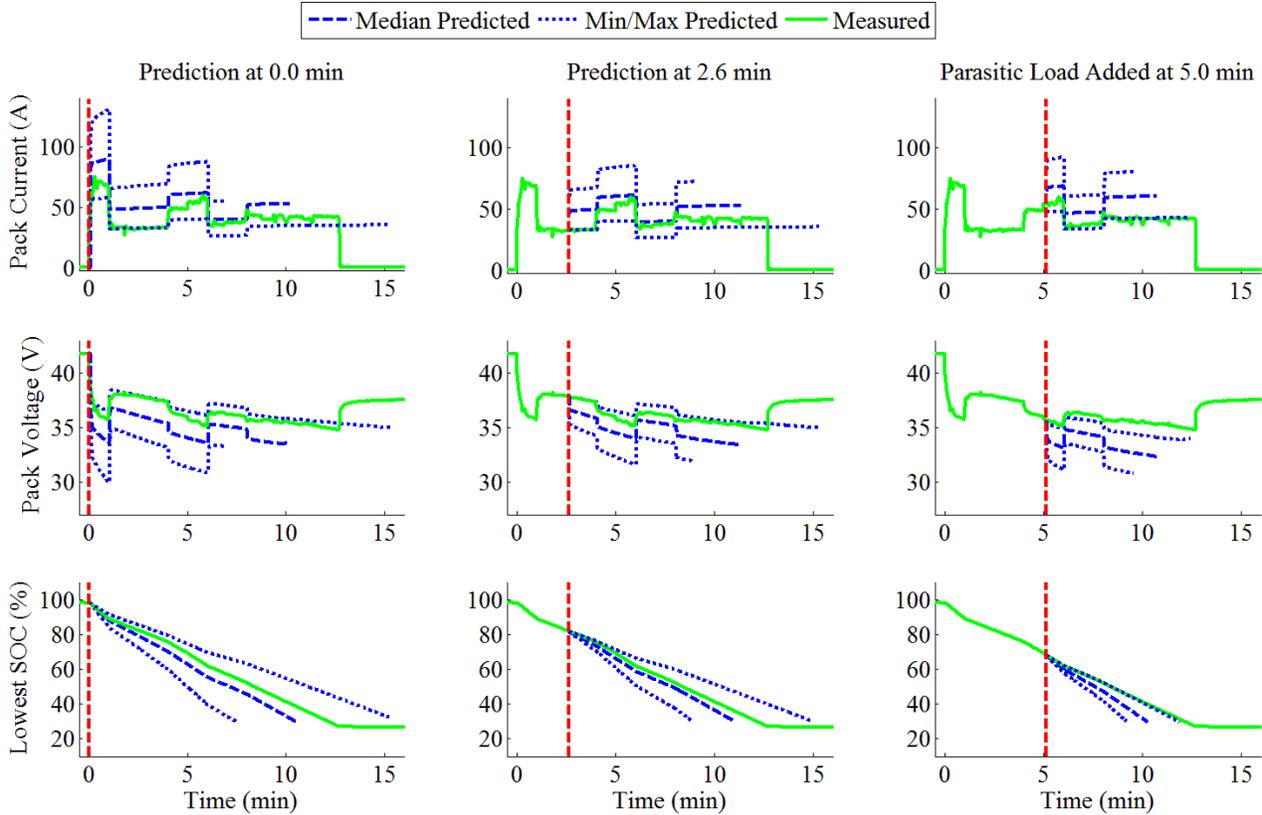


Figure 8. Example plot of measured and predicted battery current (top) and voltage (bottom) shown at three sample times over a trial battery discharge run

vertical solid bars occur at each calculation of the remaining flying time. They represent the spread of the predictions given the maximum, minimum and mean power future flight plan demand inputs. These map to the maximum, minimum and median remaining time prediction. The true

flying time remaining is found by subtracting the current time from the time at which the lowest battery SOC crossed 30 percent. The predictions are seen to underestimate remaining flying time until the parasitic load is detected at about 5 minutes into the run. After the parasitic load is detected the remaining flying time predictions are immediately shifted down increasing the degree of underestimation. The α^+ and α^- accuracy cone bounds are for many applications specified symmetrically about the ground truth value. However, in our case, the bounds are biased to be well on the low side because of the hazard present in overestimating the remaining flying time. The α - λ metric is from (Saxena, Celaya, Saha, Saha, and Goebel, 2010).

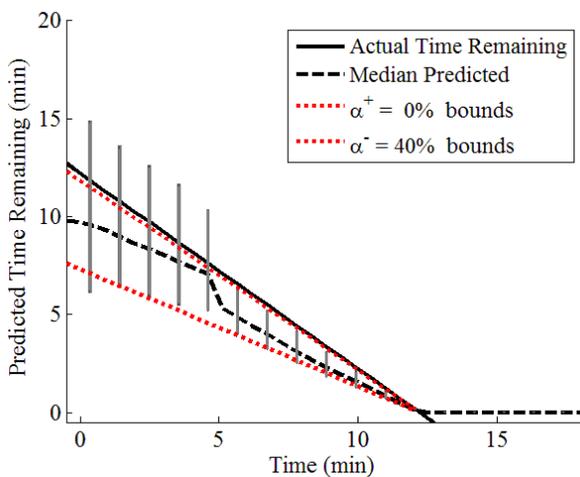


Figure 9. Predicted remaining flying time

4. GROUND TEST VERIFICATION OF REMAINING FLYING TIME PREDICTION

The ground-based verification testing of the Edge 540 T hardware and software was performed by strapping the vehicle down in the LaRC Electromagnetics and Sensors Branch High Intensity Radiated Fields (HIRF) test chamber. More information about the HIRF Chamber can be found in a report of an earlier UAS radio frequency emissions test in (Ely, Koppen, Nguyen, Dudley, Szatkowski, Quach, Vazquez, Mielnik, Hogge, Hill & Strom, 2011). The



Figure 10. Ground test chamber setup for active motor simulated flight

airplane was placed upon expanded-polystyrene blocks centered within the chamber, as seen in Fig. 10. The aircraft powertrain with propeller was operated with the vehicle anchored using a steel cable to the chamber wall. Its motor and actuators were operated from another room using the same remote control radio that would be used in subsequent flight tests.

Measured aircraft states, battery SOC estimates, and remaining flying time estimates were transmitted to a ground station over a wireless downlink. The ground station also had an uplink interface that enables the aircraft's autopilot to autonomously follow a given flight plan in chamber testing. This autopilot hardware-in-the-loop interfacing capability is discussed in (Bole et. al., 2013).

Only manual control of the throttle was used for the test results described in this paper. Aircraft propeller RPM, estimated battery SOC, and predicted remaining flying time were displayed on the ground station for the system operators in near real-time. The motor throttle was commanded using the control radio by a manual operator, who read the RPM display from the ground station. The operator adjusted the remote control throttle to maintain the target values for the time duration as determined by the flight plan described in Section 3.3. The test proceeded until a 28 percent SOC condition was indicated on the ground station display for the lowest battery. The throttle was set to zero, stopping the battery current draw. The motor batteries were allowed to rest for approximately one

hour. The battery terminal voltages at rest were used to compute an empirical approximation of battery SOC at the end of the experimental run. Onboard data logging during the experiment runs was performed by the data system described in (Hogge, 2011).

4.1. Test experience and lessons learned.

During analysis of the test results it was noticed that the chamber runs had been run at too low an equivalent energy to match flight. When the RPM was increased to match the electrical power drawn during past flight maneuvers, motor heating became an issue. Our model aircraft subject matter expert (SME) indicated that the motor was overheating because the static test air flow is not representative of that of flight. The propeller-induced air flow was not enough to keep the area near the motor cool. Temperature monitoring was done to verify that motor winding temperature limits were not exceeded. This is a limitation of this technique. The motor, the electronic speed controller (ESC), and the batteries all can be damaged by this overheating during static testing. Use of a temperature monitor was found to be important to prevent this type of damage.

Discrepancies were noticed in the SOC and remaining flying time estimation between the offline, object-oriented, battery model simulation code and the online, data-flow, real-time, operating system code. Suspected implementation errors were sought when there was a disagreement. A search for coding discrepancies revealed the following:

- The simulation input current value differed by one iteration step that was driving the UKF filter sigma point logic.
- The simulation input current value differed by one iteration step that was driving the battery state space model output logic.
- The temperature model code was planned for future work and not intended for use with the actual battery temperatures encountered because the equivalent circuit model parameters had been fitted only at a constant 30 degrees C in the laboratory.

Once the online and offline code matched these conditions, the online SOC matched the offline SOC check case value instead of being one-third more conservative.

4.2. Performance Requirements

The specification of performance requirements for ground verification of remaining flying time predictions is described next. The predictive element to be tested in this work is an alarm that warns system operators when the powertrain batteries are two minutes from reaching 30 percent SOC under normal operations. These requirements were drawn from a case study of a eUAV with battery operated propulsion (Saxena, Roychoudhury, Celaya, Saha, Saha, & Goebel, 2012). The highest level goal was stated “...complete a specified research mission by a given milestone deadline in a cost effective and safe manner.” The performance goals derived from that were translated into functional requirements to ensure that adequate battery power remains until the aircraft lands safely. The requirements were further deconstructed to suggest specific requirements for a prognostic algorithm. The decision lead time λ was specified based upon consultation with the operators to give what would be an adequate lead time for a decision and action to be taken to not put the eUAV to risk. The target λ was set at two minutes. The λ metric is related to the prediction horizon (PH) which is the time when the algorithm’s performance converges to within desired specifications. Accuracy requirements for the two-minute warning were specified as in Saxena et al., 2012, with a number of extensions added derived from the flight testing rules, operational experience, and the NASA Langley UAS Operations Office:

1. *The prognostic algorithm shall raise an alarm no later than two minutes before the lowest battery SOC estimate falls below 30 percent for at least 90 percent of verification trial runs.*
2. *The prognostic algorithm shall raise an alarm no earlier than three minutes before the lowest battery SOC estimate falls below 30 percent for at least 90 percent of verification trial runs.*
3. *Verification trial statistics shall be computed using at least 20 experimental runs.*

4. *At the end of the two-minute warning period the pilot shall have an option of doing two go arounds before the aircraft must land.*
5. *After the two-minute warning the pilot is advised not to increase altitude significantly.*
6. *The ending SOC estimation error as identified from the resting battery voltage shall be less than five percent for at least 90 percent of verification trial runs.*

To meet requirement one, “raise an alarm no later than two minutes”, the algorithm is tuned to raise the two-minute alarm early rather than late since landing becomes unsafe if not enough charge reserve is present. This is shown in Fig. 9 where the $\pm\alpha$ accuracy cone is below the actual time remaining and biases the estimated remaining flying time to be shorter than the actual remaining flying time. Requirement two, “raise an alarm no earlier than three minutes”, limits the “opportunity cost” of unnecessarily denied flying time (Saxena et al., 2012). This is to utilize the flying time resource to the utmost possible. The additional number of shorter flights come with an increased vehicle loss risk from the elevated hazard from the additional takeoffs and landings performed. Requirement three, “statistics shall be computed using at least 20 experimental runs”, is an attempt to define a target number of experimental trials needed to give a desired confidence limit. There is always a cost trade-off between the increased number of trials required for high confidence and the acceptance of risk from low confidence predictions resulting from fewer trials. Requirement four, “pilot shall have an option of doing two go arounds”, is an energy reserve safety requirement to allow two landing attempts before battery exhaustion. It was initially based upon operators’ experience and engineering judgment. A landing overrun incident required ground-test verification of the energy reserve required to accomplish two repeated landing attempts. This confirmed the already established 30 percent SOC time-to-land threshold. Requirement five, “pilot is advised not to increase altitude significantly”, grew out of variability in maneuvers chosen by the pilot just before the time of landing, and is a constraint on unplanned maneuvers close to the 30 percent SOC minimum energy threshold. When the battery is close to the 30 percent SOC, it is operating in the non-linear region, thus the SOC falls at a faster rate. The linear assumptions behind the aircraft’s powertrain and aerodynamics model (Bole et al., 2013) are also at odds with this non-linear behavior. Requirement six, “SOC estimation error... shall be less than five percent for...90 percent trials” is an accuracy goal for the method of estimating the SOC. There is a stable, empirical, relationship between resting battery voltage and SOC that can then be used to compute the ending SOC error between the resting SOC and that estimated by the prognostic algorithm at the end of the flight.

4.3. Initial experimental results and corrected results

Figure 11 is a cumulative plot of 26 verification test chamber runs that shows the difference between the times at which the two minutes remaining alarm was raised and the time at which the lowest battery SOC estimate crossed the 30 percent threshold. The remaining flying time can be thought of as the Remaining Useful Life (RUL) for the specific flying cycle. This becomes zero by definition when the lowest battery SOC estimate is below 30 percent. The RUL is plotted on the horizontal axis. Note that the positive scale increases right-to-left, the reverse of normal convention. Vertical dashed lines mark this axis and denote the two-minute and three-minute alarm time limit criteria. The criteria stated in section 4.2 are the design requirements indicating the acceptable range that the two-minute alarm is early or late. The vertical axis measures the run sequence number. This number happens to increase with calendar time when the run was performed. There was a one-year hiatus between the first 26 runs and the runs after 26 appearing on later plots and is denoted by the horizontal solid line with the year 2014 on the lower side for the initial runs and 2015 on the upper side for the later runs. Runs that were performed with and without parasitic load injection are identified by triangle and circle symbols respectively. In Fig. 11, one verification run out of the 26 is seen to violate the two-minute warning requirement late prediction limit. This was originally reported in Hogge et al. (2015). It was later discovered that there was a problem with the method used to characterize the battery capacity parameters that caused the algorithm to bias the estimate of the remaining flying time of Fig. 11. Figure 12 is the same plot made using corrected battery capacity values. The incorrect capacity came from a fit to the voltage profile using the high current chamber run data instead of from a fit using a laboratory low current discharge (2A) battery characterization cycle.

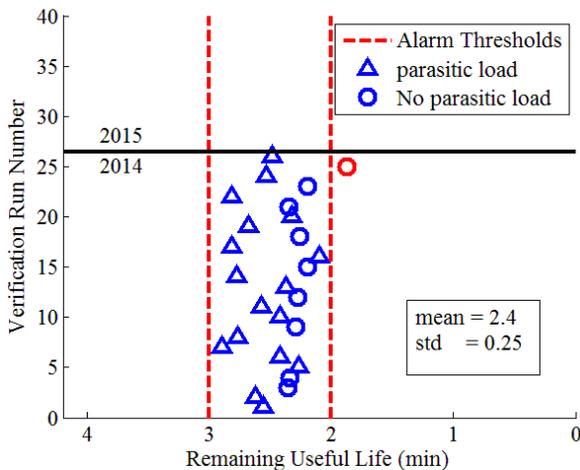


Figure 11. Original plot of two-minute alarms for 26 runs done in 2014.

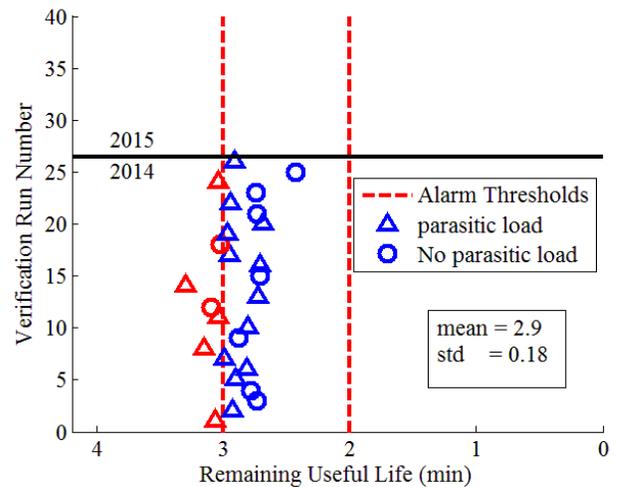


Figure 12. Two-minute alarms for 26 runs using corrected battery capacity parameters.

Violation of the assumed condition of chemical equilibrium made it impossible to identify the capacity parameter accurately with this method. This was corrected through use of the low current discharge capacity values. In practice there were additional aging effects that were not captured accurately because the laboratory 2A slow discharge experiment was not the best method to identify them. A trial and error process based upon logging the battery voltage after the run was used to work around this limitation as discussed in section 3.1.

The two-minute alarm algorithm tuning was made more conservative by using a coefficient of 1.09 to increase the velocity input to the drag model. This was based on preliminary flight experience to compensate for wind gusts and pilot variation in the trim of the aircraft that led to late biases in the estimated remaining flying time. In addition, not all flights were performed under the same weather conditions. When a coefficient of 1.0 was used to adjust the velocity input to the drag model, the pilot's increased use of speed during gusty winds and the pilot's difficulty in trimming the eUAV for optimum flight led draws higher current leading to faster charge depletion resulting in unsafe operational limits. The resulting late alarm predictions violated the safety criteria set by the operators. Uncertainty in those is not incorporated.

4.4. Battery parameter deterioration with age

Evidence of deterioration in the results is seen in Fig. 13 after a one-year gap in test runs. Ten additional ground test runs were added to Fig. 12 plot. The increase seen in the dispersion of the remaining useful life estimate for the 2015 series of runs has some runs failing early and some runs failing late. Another issue discovered was that the battery

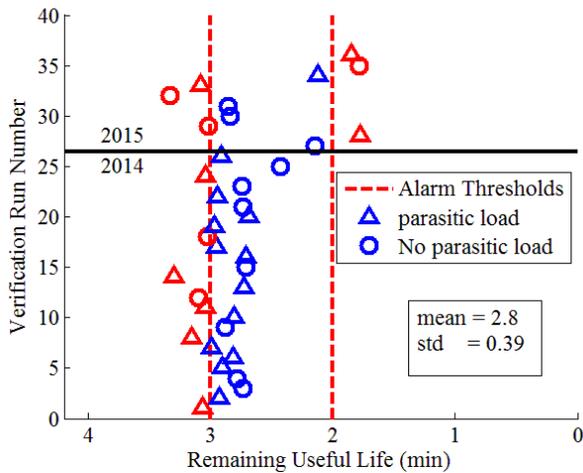


Figure 13. Two-minute alarms for additional runs done a year later using out-of-date battery capacity parameters.

aging had changed the capacity parameters significantly, but the year-old characterization runs had overestimated the battery capacity C_{max} used in the later runs (2015). This deficiency was mitigated through use of revised battery capacities (June 2015) to recalculate the data (Fig. 14). Even though the improvements were made to the capacity parameter estimates, the number of runs that exceeded the two-minute alarm bound has an increasing trend. The increased failures in the later runs may be due to the batteries approaching the end of their service life as stated by the manufacturer as being a 20 percent reduction in their initial capacity (Thunder Power RC, 2013). The batteries were approaching this value during the 2015 ground tests.

This issue resulted in an additional operational requirement

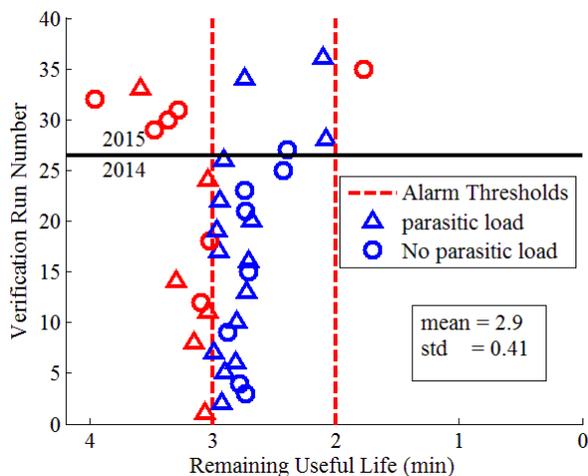


Figure 14. Two-minute alarms for additional runs done a year later using revised battery capacity parameters.

to repeat the bench capacity characterization after 10 recharge cycles to update the C_{max} parameter for each battery. The equivalent circuit models are not able to track any battery degradation unless the battery parameters are updated. This is one of the main drawbacks of the model where it cannot track aging efficiently. A slow discharge (2A) characterization experiment (Fig. 5) was attempted after 10 recharge cycles using our logging battery chargers. However, this process poorly identified R_s . Analysis by Lin & Stefanopoulou, 2015 indicates that the accuracy of the R_s measurement depends having a high current for the test condition as was done at initial purchase for one of the family of batteries (Fig. 4). This additional experiment may be necessary for all batteries to keep the model accurate. Additionally, other methods that use more detailed models of the electrochemical processes or can use dynamic in-use demand profiles could be developed and implemented (Daigle & Kulkarni, 2013).

With updated battery parameters subject to the limitations mentioned above, there are fewer runs that fail by predicting the two-minute alarm late as seen in Fig. 14. Only one (run 35) fails by predicting the alarm late. 13 verification runs out of the 36 performed are seen to violate requirement two's "...raise an alarm no earlier than three minutes..." accuracy bound. The bulk of the predictions fail by predicting the two-minute alarm slightly over one minute early. The last eight runs have the poorest performance. This could be due to the battery parameters not being identified by the 2A slow discharge experiment, or it could be due to underlying divergence of the battery cell's properties from aging. Consequently, the requirement that 90 percent of trials pass this benchmark is seen not to be satisfied because of the excessively early predictions, but this is better than failing late. Also there is increasing deviation from the mean trend indicating that there may be a systemic problem with the tracking of the battery internal state. This problem can be narrowed down to be age-related since the runs failing the two-minute to three-minute acceptance limits increase markedly after the 2014 runs, and the algorithm is very conservative.

4.5. SOC estimation error results

Now, let's examine the SOC estimation error that underlies the two-minute alarm plots. The SOC error estimation of the first 26 runs performed during the 2014 test series and previously reported uncorrected in (Hogge et al., 2015) is shown in Fig. 15 with corrections made to the battery capacity and SOC. The 2014 portion of this data set correspond to the two-minute alarm plot of Fig. 12. When the battery SOC errors from the ten additional runs made in 2015 are added, not much change is seen. Figure 16 shows box plots from the 2014 and 2015 combined data set. A total of 36 verification runs were performed if both years are included. Both plots use the empirical process "... R_s selection was repeated until the SOC estimates remained

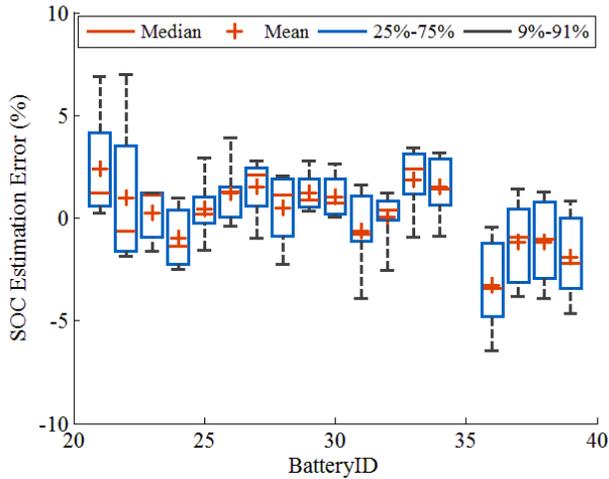


Figure 15. SOC estimation error from verification runs in 2014 (26 runs that each use 4 batteries)

constant after a run...“, which was described at the end of section 3.1. The R_s parameter was corrected through this process which then corrected the SOC error.

Because each verification run requires four powertrain batteries, 144 measurements of SOC estimation error are produced. The overall aspect of the errors are consistent. Only four of these measurements fall outside of the five percent error tolerance specified. Requirement six that 90 percent of trials pass this benchmark is thus seen to be satisfied. More measurements were outside the five percent

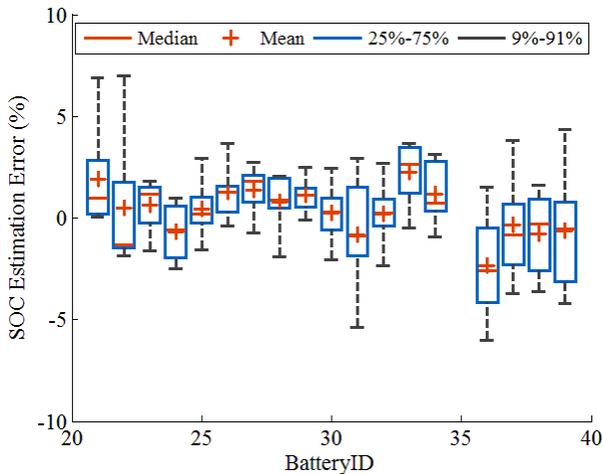


Figure 16. SOC estimation error from 10 additional verification runs in 2015 (36 runs that each use 4 batteries)

error tolerance before this correction was implemented.

5. CONCLUSION

A procedure for verifying the performance of remaining flying time predictions for a small electric aircraft was demonstrated. Aircraft battery packs reaching 30 percent SOC in flight was defined as a high risk operation for our experimental flying vehicle, to be avoided if possible. Ground-based simulated flight testing was shown to enable a safe means of running the aircraft power train to 30 percent SOC in order to obtain an empirical measurement of the aircraft's available safe operating time if motor temperature is monitored. Battery parameter identification from use and aging remains a problem to be solved for verifiably safe operations.

ACKNOWLEDGEMENT

This work was funded by the NASA System-wide Safety Assurance Technologies (SSAT) project under the Aviation Safety (AvSafe) Program of the Aeronautics Research Mission Directorate (ARMD), Prognostics and Decision Making. Patrick Tae is thanked for his contribution to the testing of the parasitic load apparatus as part of his intern assignment. Samuel Bibelhauser is thanked for his contribution of a test chamber data logger used for ground tests as part of his intern assignment. Dr. Brian Bole provided the analysis framework and previous papers that made this possible. His untimely death is a great loss and he will continue to be missed.

REFERENCES

- Bole, B., Daigle, M., Gorospe, G., (2014). Online prediction of battery discharge and estimation of parasitic loads for an electric aircraft. *European Conference of the Prognostics and Health Management Society*.
- Bole, B., Teubert, C., Quach, C., Hogge, E., Vazquez, S., Goebel, K., & Vachtsevanos, G. (2013). SIL/HIL replication of electric aircraft powertrain dynamics and inner-loop control for V&V of system health management routines. *Annual Conference of the Prognostics and Health Management Society*.
- Ceraolo, M. (2000, November). New dynamical models of lead-acid batteries. *IEEE Transactions on Power Systems*, 15(4), 1184-1190.
- Chen, M., & Rincon-Mora, G.A. (2006, June). Accurate electrical battery model capable of predicting runtime and I-V performance. *IEEE Transactions on Energy Conversion*. 21(2), 504-511.
- Dai, H., Wei, X., & Sun, Z. (2006). Online SOC estimation of high power lithium-ion batteries used on HEVs. In *IEEE international conference on vehicular electronics and safety*.
- Daigle, M., Goebel, K. (2010). Improving computational efficiency of prediction in model-based prognostics

- using the unscented transform. *Annual Conference of the Prognostics and Health Management Society, 2010*.
- Daigle, M., Kulkarni, C. (2013). Electrochemistry-based battery modeling for prognostics. *Annual Conference of the Prognostics and Health Management Society, 2013*.
- Daigle, M., Saxena A. & Goebel, K. (2012). An efficient deterministic approach to model-based prediction uncertainty estimation. *Annual Conference of the Prognostics and Health Management Society, 2012*.
- Ely, J., Koppen, S., Nguyen, T., Dudley, K., Szatkowski, G., Quach, C., Vazquez, S., Mielnik, J., Hogge, E., Hill, B. & Strom, T. (2011). *Radiated Emissions From a Remote-Controlled Airplane - Measured in a Reverberation Chamber*. NASA/TM-2011-217146.
- Hoaglin, D. C., Mosteller, F., Tukey, J. W. (1983). Analysis of two-way tables by medians. In Emerson, J. & Hoaglin, D. (Eds.), *Understanding Robust and Exploratory Data Analysis* (176-182).
- Hogge, E., Bole, B., Vazquez, S., Celaya, J., Strom, T., Hill, B., Smalling, K. & Quach, C. (2015). Verification of a remaining flying time prediction system for small electric aircraft. *Annual Conference of the Prognostics and Health Management Society 2015*.
- Hogge, E., Quach, C., Vazquez, S. & Hill, B. (2011). *A Data System for a Rapid Evaluation Class of Subscale Aerial Vehicle*. NASA/TM-2011-217145.
- Julier, S. J., & Uhlmann, J. K. (1997). A new extension of the Kalman filter to nonlinear systems. In *Proceedings of the 11th international symposium on aerospace/defense sensing, simulation, and controls* (pp. 182-193)
- Julier, S. & Uhlmann, J. (2004, March). Unscented filtering and nonlinear estimation. *Proceedings of the IEEE, 92(3), 401-422*.
- Lin, X. & Stefanopoulou, A. (2015). Analytic bound on accuracy of battery state and parameter estimation. *Journal of the Electrochemical Society, 162 (9) A1879-A1891*.
- Nelder, J. & Mead, R. (1965). A simplex method for function minimization. *Computer Journal 1965; 7 (4), 308-313*.
- Patterson, N., German, B. J. & Moore, M. D. (2012). Performance analysis and design of on-demand electric aircraft concepts. *12th AIAA Aviation Technology, Integration, and Operations (ATIO) Conference and 14th AIAA/ISSM* (p. 27). Reston, VA: AIAA.
- Quach, C., Bole, B., Hogge, E., Vazquez, S., Daigle, M., Celaya, J., Weber, A. & Goebel, K. (2013). Battery charge depletion prediction on an electric aircraft. *Annual Conference of the Prognostics and Health Management Society 2013*.
- Saha, B., Koshimoto, E., Quach, C., Hogge, E., Strom, T., Hill, B., Vasquez, S. & Goebel, K. (2011). Battery health management system for electric UAV's. *IEEE Aerospace Conference*. Big Sky, MT: IEEE.
- Saxena, A., Celaya, J., Saha, B., Saha, S., & Goebel, K. (2010). Metrics for offline evaluation of prognostic performance. *International Journal of Prognostics and Health (IJPHM)*, Vol. 1, 2010.
- Saxena, A., Roychoudhury, I., Celaya, J., Saha B., Saha S., & Goebel, K. (2012). Requirements flowdown for prognostics and health management. *Infotech@Aerospace, 2013*. AIAA, Garden Grove, CA.
- Saxena, A., Roychoudhury, I., Lin, W. & Goebel, K. (2013). Towards requirements in systems engineering for aerospace IVHM design. *AIAA Conference, 2013*. AIAA, Reston, VA.
- Thunder Power RC, 2013 Safety instructions and warnings, revision 3 (December 10, 2013). www.ThunderPowerRC.com (2).
- Wang, Y., Fang, H., Zhou, L. & Wada, T. (2017). Revisiting the state-of-charge estimation for lithium-ion batteries. *IEEE Control Systems*, Vol. 37 (4), 73-96. doi: 10.1109/MCS.2017.2696761
- Zhang, H. & Chow, M.-Y. (2010). Comprehensive dynamic battery modeling for PHEV applications. *Power and Energy Society General Meeting, 2010, IEEE*. Minneapolis, MN: IEEE.

BIOGRAPHIES



Edward F. Hogge received a B.S. in Physics from the College of William and Mary in 1977. He has provided engineering services to the government and currently is employed by National Institute of Aerospace. He has recently been supporting aviation safety research through the implementation of electronic systems for subscale remotely piloted aircraft and through commercial aircraft simulation. He is a member of the American Institute of Aeronautics and Astronautics and the Prognostics and Health Management Society.



Brian M. Bole graduated from the FSU-FAMU School of Engineering in 2008 with a B.S. in Electrical and Computer Engineering and a B.S. in Applied Math. He received M.S. and Ph.D. degrees in Electrical Engineering from the Georgia Institute of Technology. Dr. Bole was recently deceased. Before his death his research interests included: analysis of stochastic processes, risk analysis, and optimization of stochastic systems. He had been investigating the use of risk management and stochastic optimization techniques for prognostics and prognostics-informed decision making in robotic and aviation applications. From 2011 to 2013 he performed joint research with the Prognostic Center of Excellence at NASA Ames under the NASA graduate student research fellowship.



Sixto L. Vazquez Mr. Vazquez obtained MSEE from Old Dominion University in 1990 and BSEE from the University of Puerto Rico in 1983. He has developed real-time 3D graphical simulations to aid in the visualization and analysis of complex sensory data. He has developed techniques to interactively process, analyze, and integrate sensory data from multiple complex, state-of-the-art sensing technologies, i.e. FMCW Coherent Laser Radar range measuring system, Bragg grating Fiber Optic Strain Sensing system, etc., into simulation. In recent years, he has developed software for the Ardupilot and associated ground station.



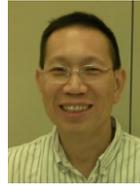
Chetan S. Kulkarni received the B.E. (Bachelor of Engineering) degree in Electronics and Electrical Engineering from University of Pune, India in 2002 and the M.S. and Ph.D. degrees in Electrical Engineering from Vanderbilt University, Nashville, TN, in 2009 and 2013, respectively. He was a Senior Project Engineer with Honeywell Automation India Limited (HAIL) from 2003 till April 2006. From May 2006 to August 2007 he was a Research Fellow at the Indian Institute of Technology (IIT) Bombay with the Department of Electrical Engineering. From Aug 2007 to Dec 2012, he was a Graduate Research Assistant with the Institute for Software Integrated Systems and Department of Electrical Engineering and Computer Science, Vanderbilt University, Nashville, TN. Since Jan 2013 he has been a Staff Researcher with SGT Inc. at the Prognostics Center of Excellence, NASA Ames Research Center. His current research interests include physics-based modeling, model-based diagnosis and prognosis. Dr. Kulkarni is a member of SMAIAA and SMIEEE.



Thomas H. Strom was born in Aberdeen, WA, 1924. He graduated from Hoquiam High School in 1942, and attended Seattle University and the University of Washington 1953-1960. He served in the U.S. Navy during World War II as a Radar Technician. He was an employee of the Boeing Corp. Wind Tunnel, Seattle, WA 1947-1972. He was Senior Engineer with principal expertise in flutter and aeroelastic wind tunnel modeling. He was the founder and president of Dynamic Engineering, Inc., Newport News, VA, until 1997. He has served as a consultant to the Aeroelasticity Group of NASA Langley Research Center (LaRC) and continues to provide engineering and technical services to LaRC through the National Institute of Aerospace.



Kyle M. Smalling Kyle Smalling obtained his B.S. in Aerospace Engineering from Cal Poly Pomona in 2013. He is an avid remote control vehicle enthusiast both personally and professionally. His areas of research include Health Prognostics and developing Safety Critical hardware and software. He is employed by the National Institute of Aerospace.



Cuong C. Quach received his M.S. from the School of Physics and Computer Sciences at Christopher Newport University in 1997. He is a staff researcher in the Safety Critical Avionics Systems Branch at NASA Langley Research Center. His research areas include development and testing of software for airframe diagnosis and strategic flight path conflict detection