A New Model for Power Plant and Health Estimation

A Technical Paper prepared for SCTE by

Kang Lin, PhD
P&E RF Lab Manager
Comcast
1002 Cornerstone Blvd, Downingtown, PA 19335
480 430-8254
Kang_lin@comcast.com

Michael Nispel
Senior Principal Engineer
Comcast
1002 Cornerstone Blvd, Downingtown, PA 19335
610 952-3783
Michael_nispel@comcast.com

Table of Contents

List of Figures

1. Introduction

Backup batteries are an essential part of the reliability equation for the cable outside plant architecture. When utility power fails, batteries are relied upon to power the network and helps ensure customers remain on-line. Being able to predict how long a battery can support the network is an essential planning and maintenance tool.

As one of the largest industrial Internet-of-Things (IOT) applications, Comcast has developed the unique ability to automatically test and monitor over a million installed batteries at well over a quarter of a million installations. This paper describes possible ways to use this continuous sensing of metrics from the power supplies of our outside plant network to begin developing a sophisticated diagnostic and planning tool. When finalized, this tool will use machine learning and artificial intelligence to be able to predict the expected runtime of a battery during a utility outage based on its actual load, and then to adjust this runtime prediction based on multiple key factors. The final desired result is the ability to know whether a battery is performing as expected, and to be able to track its degradation for preemptive maintenance purposes. This includes unexpected loss of battery performance from unknown factors as well as the expected losses from known factors. The status of this work is presented, showing the current predictive model, along with a brief discussion of the future work planned in this area.

2. Background

Battery systems have been used for many years to power the Cable network during a utility outage. Utilities themselves are undergoing a significant transformation, with increases in renewable generation, flexible load programs, decarbonization of commercial buildings and major capital deferments. Throughout and sometimes due to these major strategic shifts, energy reliability and disruptions remain a significant issue that network providers must deal with. The disturbing and growing trend of electrical disturbances for the past 20 years can be clearly seen in Figures 1 and 2.

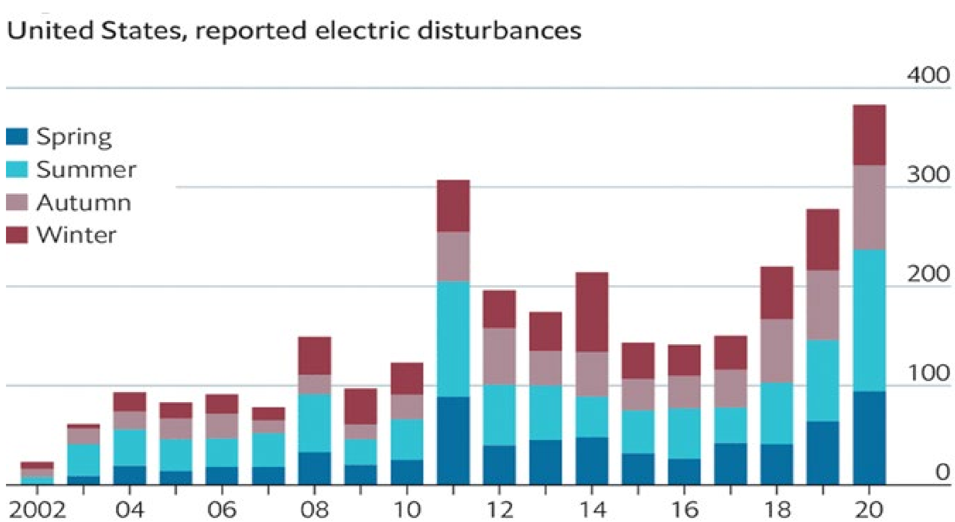

Figure 1 – Department of Energy OE417 – Annual Summary

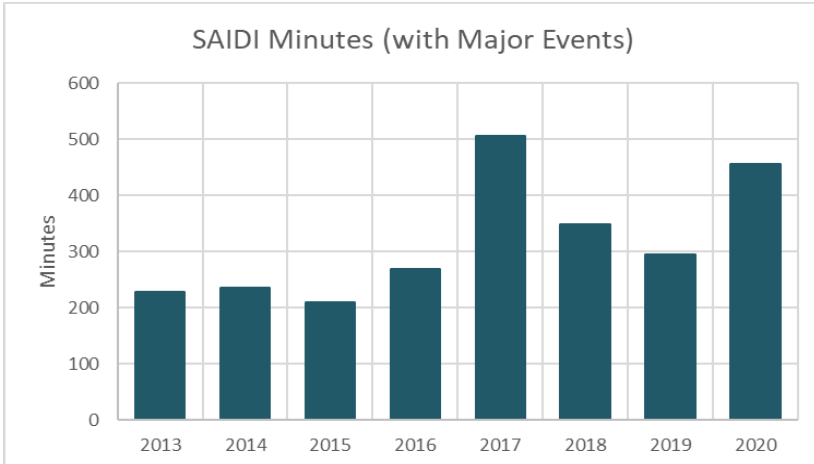

Figure 2 – System Average Interruption Duration Index (SAIDI) – EIA-0861, Annual Report

It is relatively easy to detect an outage and to report on the duration of a battery's discharge event. However, this in-and-of itself is not sufficient to determine the health of the network. Battery manufacturers provide tables and curves for the expected discharge runtime vs. load of a new battery. Over time as the battery ages, the runtime performance of a battery will be expectedly reduced by numerous factors and conditions, while the battery remains healthy. The goal of this work was to use the existing metrics available from the power supplies to calculate the runtime of a new Outside Plant (OSP) battery, and to then adjusted the runtime based on the effects of the major, known factors. Future work will focus on identifying unhealthy batteries that have degraded faster than would be expected. This tool is expected to ultimately result in fewer truck rolls as well as enable a more effective OSP battery (see Figure 3) replacement strategy.

Figure 3 – Pad mounted Outside Plant with Batteries Shown

3. Approach

The existing set of metrics that are reported from the power supply was first examined to determine what gaps exist. This set includes the following key items.

- Power output of the power supply
- Current output of the power supply
- Individual battery voltage
- Ambient temperature
- Timestamps of all items
- Inverter status (charging, discharging, +)

Additional items known include the age of the battery, the manufacturer and the model, the number of battery strings, and the total number of batteries in the power supply. Using all of these metrics, we were able to craft a working predictive model.

The first step was to determine how the battery was expected to perform as new. Manufacturers generally provide a limited performance data set. This runtime data can be based on a constant current discharge or a constant power discharge. Using this data set, the discharge performance for each battery model in the network was plotted and curve fit to allow a calculated runtime for any load imposed on the battery. This generally follows the well-known Peukert's relationship, which is a generalized relationship between discharge load and runtime. Since the performance curve of all batteries vary based on their design, it was decided to individually curve fit the performance data for each battery model in the network. Using a standard exponential equation, a very accurate curve equation was then created for each battery model. An example of the manufacturer's data provided and the resultant curve equation for one of the battery models are shown in Figure 3.

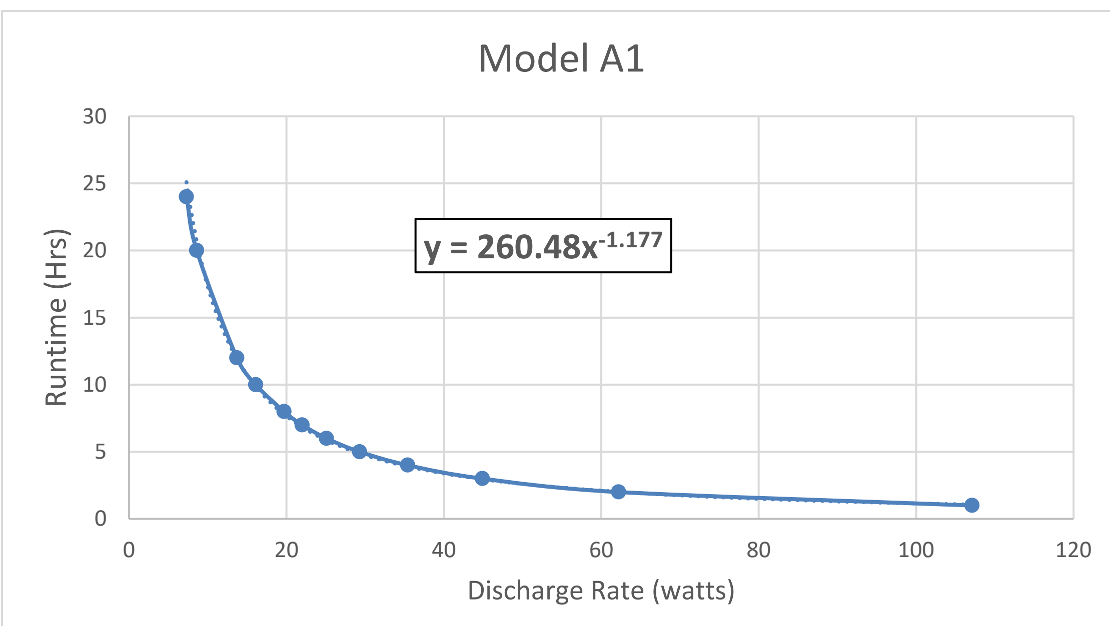

Figure 4 - Manufacturer's Published Battery Performance Data – with curve fit equation

To use this curve equation, the actual direct current (DC) load placed on the battery is required. Unfortunately, this is not one of the measurements of the power supply. To resolve this, the efficiency of the power supply inverter was calculated using the measured alternating current (AC) voltage and current output of the power supply. By using this calculated inverter efficiency, the DC load on the battery could then be calculated. With this, we were then able to calculate a predicted battery runtime for any utility outage where the power supply transferred to battery backup.

The limitations of this preliminary equation were immediately evident. Using the performance curves from the manufacturers predicts their runtime as a brand new battery. All batteries will expectedly decline in performance over time as they naturally age. This degradation is not accounted for in the curve fit of the performance data. Additionally, temperature has a major effect on battery performance. Without accommodating for these factors, one is unable to ascertain if a battery that runs for a shorter than expected period of time is prematurely failing. A short-running battery may be defective, or it may be perfectly healthy but several years old. It could also be healthy but with a reduced runtime because it is being used in the winter in the Minneapolis area.

4. Temperature Effect

The power supply cabinets are environmentally uncontrolled, and the temperatures vary widely depending on the latitude and the season. As all students of chemistry are aware, virtually all chemical reaction rates are affected by temperature as described by Arrhenius many years ago. To complicate this temperature effect on the reaction rate, a battery's performance is the result of a two-phase reaction. The discharge reaction occurs in the liquid electrolyte phase, and then deposits onto the solid surface of the plates. Temperature will affect both rates differently depending on the battery design. The liquid electrolyte phase is further complicated due to temperature induced convective movements during extended runtimes. Finally, the surface reaction is affected by double-layer capacitive effects. For this first level model, a test program was initiated within the Comcast labs as shown in Figure 4 to experimentally quantify the total effect of temperature on the battery performance for each major battery model in the network. Using controlled environmental chambers, measured performance discharges were conducted on battery models under varying ambient temperatures.

Figure 5 – Tested Batteries for Temperature Effect

Based on these results, the Temperature Correction Factor (TCF) was determined to follow the general polynomial equation as shown in equation 1. Applying the TCF to the predicted runtime allows the overall effect of temperature for each battery model to be included in the runtime prediction.

$$TCF = A_1 * \exp(-05) * (\text{Temp, degrees Celsius})^2 + A_2 * (\text{Temp, degrees Celsius}) + A_3 \qquad \text{Eqn. 1}$$

[A_n are experimentally derived constants]

5. Aging Effect

The second major effect on batteries is the aging effect. Industrial batteries are considered at end-of-life when their performance has degraded to 80% of their initial capacity. Beyond 80%, the change in degradation accelerates significantly, and the battery is in danger of unpredictably and sharply falling below the minimum requirements. There are two separate types of aging – calendar and cycle effects. Either of these will independently degrade the performance of a battery.

A battery has a predetermined calendar life. Even if never discharged, the internal side reactions that occur within the battery will corrode the positive grids and cause electrolyte loss. Numerous factors will affect the corrosion rate, including grid alloy, temperature, float voltage, electrolyte strength and AC ripple. These factors are generally balanced by the battery manufacturer, so that in an ideal setting, the battery will degrade to 80% capacity at the end of its published design life. Life claims vary by manufacturer, but a seven-year design life was used as an initial estimate. This is known to be a gross estimate, as the corrosion rate is known to be strongly affected by latitude/temperature. For our initial model, a straight-line estimation was used, assuming the battery loses 2.8% of its capacity each year, thus hitting its 80% end-of-life capacity at the end of 7 years.

$$\text{Calendar Aging Factor (AF)} = -(1 - 0.80) * (\text{year})/7 + 1 \qquad \text{Eqn. 2}$$

The second major aging effect is due to cycling. In addition to the grid corrosion occurring during calendar aging, any discharge/recharge cycle will degrade the positive and negative plates within the battery and reduce its capacity. This is understood, but not yet implemented. The issue is that a generic model cannot be used, as the ability to cycle is dependent upon the specific design of each battery model. The types of batteries used in the network commonly are designed for approximately 200 ‘deep’ cycles, but can easily vary from 75 to 500 depending upon the internal components and method of construction. Additionally, the ability of a battery to withstand shallow vs. deep discharges varies tremendously in a non-linear manner, making it necessary to know the cycle-life curve for every battery model. This information is typically not provided by manufacturers in sufficient detail and lab testing is expected to be required. Once characterized in our labs, the number and depth of the discharge information will be collected and compiled from the field, with the intent of adding this element to a future revision of this model.

$$\text{Cycle Aging Factor} = \text{function(depth of discharge)} + \text{function(cycle quantity)} \qquad \text{Eqn. 3}$$

6. Discussion/Results

The predictive model as described takes the known informational points of the power supply listed below. It calculates the predicted runtime of the power supply, based on the output load, the battery model and battery quantity. This calculation is then modified based on the temperature and the calendar age of the battery to provide a final predicted runtime.

- Battery manufacturer - recorded
- Battery model - recorded
- Number of battery strings - recorded
- Total number of batteries - recorded
- Date code - recorded
- AC voltage output - measured
- AC current output - measured
- AC power output - measured
- Temperature - measured

Actual field data was collected and compiled over a 6 week period and plotted in Figure 5. The predicted, adjusted runtimes were calculated and plotted vs. the actual runtimes. Obvious sources of error were excluded, such as missing dates, errant currents and runtime that were too short or too long to be meaningful. This is an area of ongoing work as the errors found continue to be cleaned and corrected. In an ideal world, all points would lie on a straight, 45 degree line (shown as the thin black line in Figure 5), which would indicate the model was perfectly accurate in predicting battery performance over a wide range of runtimes.

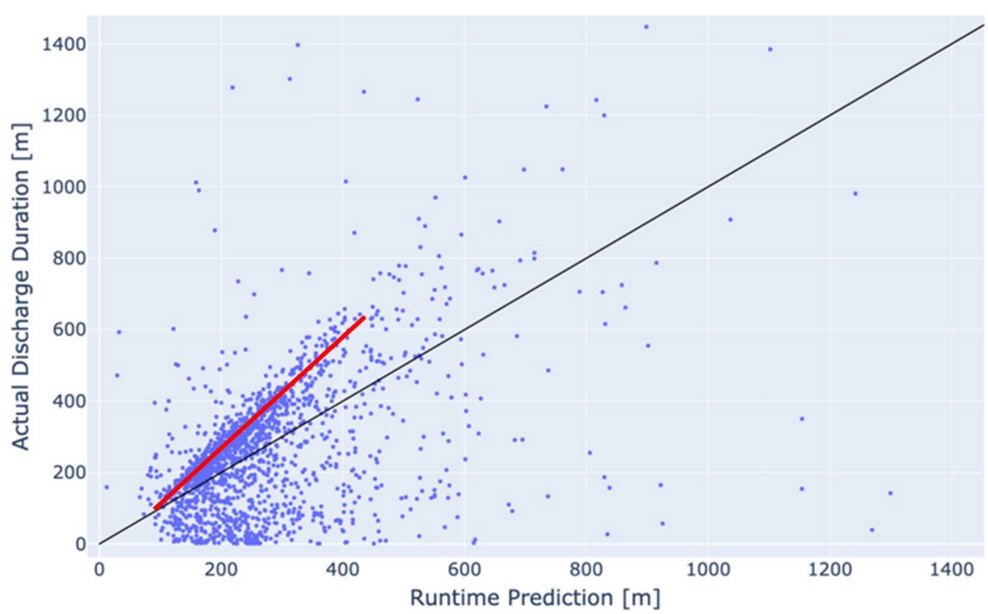

Figure 6 - Predicted Runtimes vs. Actual Runtimes

This is a very preliminary snapshot of the data, and refinements and adjustments are continuing. However, despite the substantial scatter in the data, even this preliminary snapshot is extremely encouraging. A distinct trendline shown in red in the plot shows that a correlation is evident in the runtime calculation. The deviation from the ideal 45 degree slope indicates that as the runtimes extend longer, the deviation grows from the expected. Since the model is based on the manufacturer's published data, the prediction is only as accurate as this data. One possible cause of this error shift is based on the author's experience that the published manufacturers data is often progressively more conservative for shorter runtimes. This will be verified in future planned lab testing to see if the actual performance data does indeed deviate from the published data. Another possible cause of error may be the temperature rise in the battery from ohmic heating during a discharge. During longer discharges, it is possible this internal heating could progressively improve the performance of a battery. The current model calculator neglects this effect as it corrects for only the initial temperature measurement.

In Figure 5, all points below the 45 degree line indicate batteries that ran shorter than predicted. The points very close to the horizontal axis indicate batteries that failed immediately. This could be from batteries with defects, improper connections, broken cables, etc. There currently are no predictive elements in the present calculator to identify such defects. However, from knowledge of battery failure modes, there are suspected metrics that are currently being reviewed that may identify defective batteries prior to a discharge. For instance, a large standard deviation for the batteries within a string could conceivably indicate one prematurely failing battery and could thus be a predictor of a shorter runtime. Another example is that a battery not fully charged will perform poorly. To identify batteries that are not at a full state-of-charge, it is planned to track the time and depth of previous discharges to determine the recharge efficiency. Additionally, it is expected that a battery in this condition could be identifiable by its high internal resistance, which could conceivably be seen in a low initial voltage at the onset of the outage. For this example and other suspected causes of poor performance, there are now efforts underway to use machine learning as a tool. This will allow correlation of possible groups and trends with identifiable causes.

7. Conclusion

Comcast has developed the significant ability to automatically test and monitor over a million installed batteries through its Outside Plant network. Using the data collected from this immense installed base, a runtime calculator was created to predict the runtime of outside plant batteries. This calculator is based on Peukert's law, which approximates the non-linear change in battery capacity due to changes in discharge rate. An exponential equation was fitted to each battery model's performance curve to allow an accurate prediction of the runtime of a new battery at a constant, 77 degrees F temperature. This model was then adjusted for the two major, known causes of variations in battery performance - temperature and calendar aging. With this new model, preliminary comparisons to actual field outage data shown a promising correlation between predicted and actual runtimes. Refinement of this preliminary model is continuing to identify and reduce the sources of error present. Future work will focus on additional laboratory testing of batteries to more fully characterize their performance over the range of field conditions. This includes the addition of a cycling degradation factor due to accumulated discharge events of different depths. In addition, future efforts are also planned on machine learning to associate patterns in the data with identifiable causes. (see Acknowledgements) It is optimistically expected that as this model evolves and become increasingly sophisticated and accurate, the ability to quantify and predict battery degradation will prove to be an immensely valuable tool for predictive maintenance and asset planning.

Acknowledgements

The authors would like to acknowledge our colleagues that were major contributors to the success of this effort. Matt Stehman and Chris D'Andrea were instrumental in setting up and developing the tools used in the data collection and analysis of this project. Their work has been greatly appreciated and they are expected to continue to play a major role in the machine learning efforts as this exercise continues into the future. Interested readers are encouraged to read the details of their machine learning efforts "Machine Learning and Telemetry Improves Outside Plant Power Resiliency for More Reliable Networks", by Stephanie Ohnmacht and Matt Stehman, presented at this Expo.

Abbreviations

IOT	Internet of things
OSP	Outside Plant
DC	Direct current
AC	Alternating current
TCF	Temperature correction factor
AF	Aging factor