EVS26
Los Angeles, California, May 6 – 9, 2012

Mission Skyline: Design, implementation, and benefits of a cloud-based electric powertrain data acquisition and visualization system

Jit Bhattacharya[1], Seth LaForge[2], Jon Wagner[3]
[1]*CEO of Mission Motors, jit@ridemission.com*
[2]*Sr. Software Architect, Mission Motors, seth@ridemission.com*
[3]*Chief Technology Officer, Mission Motors, jon.wagner@ridemission.com*

Abstract

The optimization of electric drive systems for electric and hybrid vehicles depends on detailed data on vehicle use and powertrain performance. With the evolution of modern data acquisition and wireless communication technology, electric vehicles are able to capture and transmit large amounts of raw data in the field. However, these large, complex datasets are difficult to manage, process, and filter. Vehicle data is often sampled at different time scales, resolutions, and sample rates, each with its own degree of relevance. Engineers must spend excess time processing vehicle datasets in order to identify time segments and channels of importance. Mission Motors created a cloud-based data acquisition and visualization platform for streaming vehicle data, named Skyline. The Skyline platform catalogs data into a pre-structured database and then accesses that data through a novel visualization platform on the web. Through the Skyline platform, vehicle developers are able to visualize any combination of data acquisition channels at any resolution. Engineers are also able to flag segments of data in order to efficiently recall them at another time. Initial results from use of the Skyline system show that engineers are able to more efficiently process and identify critical data streaming from vehicles. The software technology within Skyline addresses the gap within current data acquisition and processing systems by providing engineers with a tool to handle vehicle specific, time resolved data.

Keywords: wireless, data acquisition, data visualization, system modeling, powertrain optimization

1 Introduction

The development of a reliable electric/hybrid powertrain system for a consumer vehicle demands a detailed understanding of component interactions across a broad range of driver use cases and environmental conditions. The collection of low-level system data, including data on individual battery cell performance in a large array, can enable more optimized system designs during vehicle development while also enabling better risk management for production electric vehicles in the field.

Vehicle developers currently obtain the data needed to inform system designs through a combination of laboratory testing and onboard data acquisition systems. In addition to these methods, modern acquisition technology and wireless communications has empowered engineers to stream extensive low-level data from vehicles operating in the field. For electric vehicle designers, these detailed data sets provide key insights into drive cycles, battery degradation, and long-term use, while also providing high-resolution data on powertrain performance and reliability under a variety of use cases.

While vehicle field data has become more common, these large, complex data sets are difficult to manage, process, and filter. Vehicle system data is often sampled at different time scales, resolutions, and sample rates, each with its own degree of relevance. Engineers are forced to spend excess time processing vehicle data sets, often through tools such as Matlab, Excel, or custom software, in order to identify time segments and channels of importance.

To unlock the power of wireless, real-time data acquisition in electric vehicle development, tools are required that 1) Process data as it arrives, 2) Handle data at a variety of resolutions and sample rates, and 3) Enable users to efficiently access and visualize data in order to derive value from large datasets.

2 About Mission Motors

Mission Motors is an advanced electric powertrain technology company. Headquartered in San Francisco, California, the company develops advanced lithium-ion battery systems, power electronics, electric motors, and vehicle software for electric and hybrid vehicles. Mission Motors has developed powertrain systems for a range of vehicle segments including hybrid cars, electric cars, electric powersports, and hybrid heavy-duty vehicles.

3 Mission Skyline

To meet the need for better data processing and visualization in electric and hybrid vehicle development, Mission Motors developed the Skyline platform. Skyline is a cloud-based platform for processing and visualization of streaming vehicle system data. The Skyline software platform catalogs data into a pre-structured database as it is streamed from a vehicle. Skyline enables access and visualization of vehicle data through a novel web-based visualization platform. Through the Skyline platform, vehicle developers are able to visualize any combination of data acquisition channels at any time resolution.

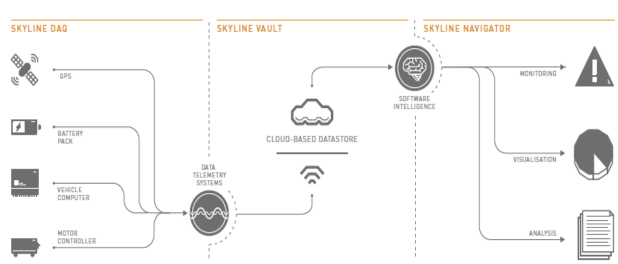

Figure 1: Skyline Flowchart

The Skyline platform aimed to solve three principal challenges with remote vehicle data acquisition

1) Processing large datasets streaming in real-time at multiple sample rates.
2) Enabling secure access and visualization of raw data via the web
3) Flagging of specific time segments and channels for efficient identification of important data

3.1 Database Structure – Managing data acquired at multiple sample rates

Skyline represents samples in a form that allows for logging of any data – even channels that have not been pre-structured into the database. In addition, the sample representation does not require fixed sample rates. Conceptually, each *real sample* consists of:

- Numeric vehicle ID.
- Textual channel name (*e.g.* "motor/positionSensor.speed_rpm"_
- Sample start time, in microseconds since the epoch (January 1st, 1970 UTC)
- Sample end time, in microseconds since the epoch
- Sample value

Each sample's time range represents the time period over which that sample is relevant – typically the time since the previous sample in the channel. Using absolute time stamps allows arbitrary or variable sampling rates, and allows easy combining of data from different sampling sessions. Channel names describe the origin of the data in the vehicle. Sample values can be floating point or integral numbers, strings, or structured data.

To support efficient web-based display of high-resolution data over long timeframes, Skyline implements a multi-resolution data store. When samples are added to the database, lower-resolution *synthetic samples* are also created. Synthetic sample rates have been specified based on the expected data resolutions. The synthetic sample stores the mean, minimum, and maximum values of the raw-data across the synthetic sample range.

Samples are stored using the MongoDB nosql database. For storage efficiency, both raw-data and synthetic samples are grouped into buckets, such that on average 50 to 100 samples are stored in a single database row, avoiding excessive overhead. MongoDB collections are used to segregate samples by duration – for example, there are separate collections for each synthetic sample rate as well as a collection for raw-data in its native rate. The use of collections maintains the integrity of the raw dataset.

When data is inserted into the database, the affected synthetic samples are automatically generated. Special techniques are used to allow parallel processing across servers to enable scaling.

To make the Skyline database an effective tool, both raw and synthetic data had to effectively integrate into the web-based client where a user will view the data. The web-based client requests samples for a time range at a given resolution. The server queries to find synthetic samples at the requested resolution, and samples of raw data at durations greater than the requested resolution. Raw-data and synthetic samples are combined to return a simple list of sample values and times to the client.

3.2 Accessing and Displaying Vehicle Data through the Web

Engineers use acquired vehicle data to evaluate vehicle performance, understand system faults, and optimize system designs and parameters. To achieve these goals, engineers must process numerous channels of data taken at high sample rates over long time periods. Vehicle engineers often use Excel, Matlab scripts, or custom developed software programs for data processing and visualization. These systems often fail to handle large volumes of data, require extensive technical training, and are expensive and available only to a few team members.

To address the need for better tools, Mission Motors developed and tested new methods for visualizing and exploring stored vehicle data using the latest in web-based application technology. The Skyline interface is designed as a web application, rather than as a native application residing on the PC. As a web application, the Skyline interface can give teams in multiple locations the flexibility to access from anywhere in the world and from any machine with a web browser. This eliminates the requirement for installing specialized software for all users who need access to data.

The web client initially presents a list of tracked vehicles to which the user has access. Once the user selects a vehicle, the app displays a hierarchical tree of collected data channels, one or more graphs of selected channels, events of interest, and a map with the GPS path of the vehicle. Channels can be dragged from the channel tree to a graph, which then requests data from the server and displays it at an appropriate level of resolution. The user can then fluidly drag the graph or use the mouse scroll wheel to zoom in or out of the data, with the web app loading additional data in the background as necessary.

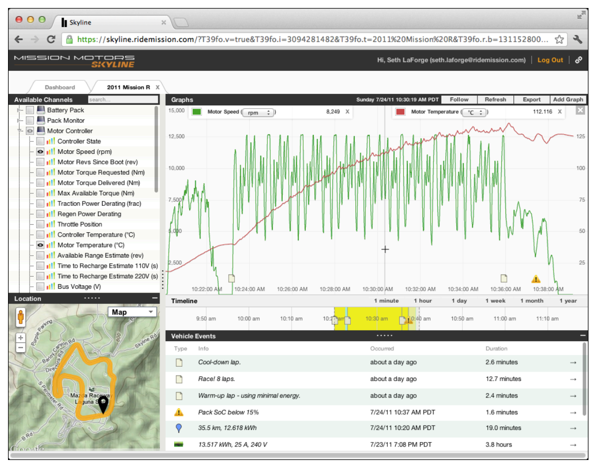

Figure 2: Skyline Navigator UI

Since data is being transmitted over the Internet, both access and transmission of data must be kept secure, without affecting the performance of the interface. The software encrypts the connection to the server through standard SSL/TLS protocols in accordance with IETF RFC 5246. Skyline employs the OAuth 2.0 protocol to authenticate using Google Apps accounts.

To effectively enable users to efficiently browse large amounts of high-resolution data, the Skyline platform enables a user to smoothly zoom from millisecond level resolution to multiple days, without any lag in the interface or excessive data transfer. By displaying minimum, maximum, and mean values from synthetic samples when zoomed out, Skyline enables users to spot trends and outliers even when the data resolution significantly exceeds the pixel density of the user's display, as in Figure 3a. As the user zooms in, Skyline smoothly fetches higher-resolution data until the user can see every sample, as in Figure 3b.

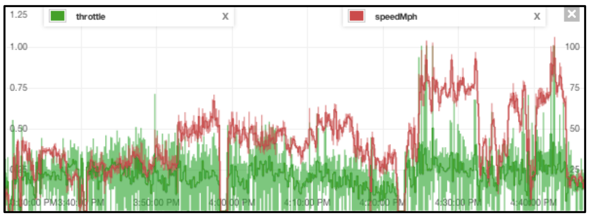

Figure 3a: Skyline graph, zoomed out

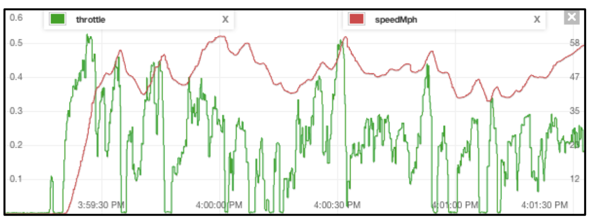

Figure 3b: Skyline graph, zoomed in

3.3 Events

When navigating large datasets, it is crucial to have landmarks to put the data into context and make it easier to locate significant events. Additionally, allowing multiple user comments and collaborative communication greatly increases the value of the data to a development team. Thus Skyline provides methods for creating events, where each event is related to a specific vehicle, time range, and set of data channels. Skyline Navigator displays a list of events relevant to the current vehicle. Users can immediately view the related data for an event through a single click.

Skyline implements two methods for creating events: automatic events and user-generated notes. The server generates automatic events based on detection of system-specific fault codes, or user-defined data conditions. For example, figure 4 shows an automatically generated warning event based on low State of Charge for the vehicle battery pack. By appearing in the event list, a user is able to navigate to the time segment of data surrounding that warning with a single click. Automatic event types include a vehicle drive cycle, a charge cycle (plug-in vehicle only), warnings, and errors.

Figure 4: Automatic warning event

"User-generated notes" are created by users as they browse vehicle data. The user selects a time range and set of channels, and types a textual comment (Figure 5). All users who access the vehicle will see the created note, and can reply to it. Options for email notifications ensure that all participants in the conversation can remain involved. User-generated notes improve collaboration for diverse teams interacting with and evaluating vehicle data.

Figure 5: Note with replies

3.3 Exporting Data

Skyline does not attempt to implement processor intensive data analysis tools within the web-application. Skyline seeks to improve data processing efficiency, including interoperability with tools such as MatLab and Excel, which are better suited to processor-intensive analysis. In order to make interoperability with other tools easy, Skyline implements a data export feature.

Only the time range and channels visible in the graph are exported, with optional resampling available to reduce the quantity of exported data and/or to time-align the data samples.

Data collected from vehicle systems is often sampled at different rates. Comparing and analyzing such diverse data streams can be difficult with traditional tools. As discussed in Section 3.1, Skyline's database structure allows multiple sample rates and timing by storing samples with absolute time stamps, rather than enforcing a fixed sampling interval. In order to simplify working with such data, export with resampling effectively coerces all channels to the same sample rate, as seen in Figure 7. Data can also be exported without resampling to retain the original timing of the raw data.

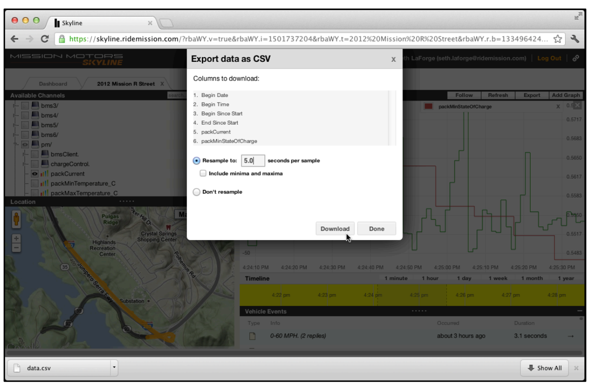

Figure 6: Exporting data

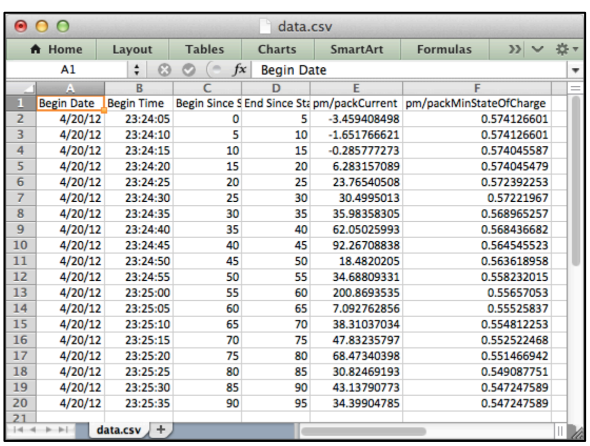

	A	B	C	D	E	F
1	Begin Date	Begin Time	Begin Since S	End Since Sta	pm/packCurrent	pm/packMinStateOfCharge
2	4/20/12	23:24:05	0	5	-3.459408498	0.574126601
3	4/20/12	23:24:10	5	10	-1.651766621	0.574126601
4	4/20/12	23:24:15	10	15	-0.285777273	0.574045587
5	4/20/12	23:24:20	15	20	6.283157089	0.574045479
6	4/20/12	23:24:25	20	25	23.76540508	0.572392253
7	4/20/12	23:24:30	25	30	30.4995013	0.57221967
8	4/20/12	23:24:35	30	35	35.98358305	0.568965257
9	4/20/12	23:24:40	35	40	62.05025993	0.568436682
10	4/20/12	23:24:45	40	45	92.26708838	0.564545523
11	4/20/12	23:24:50	45	50	18.4820205	0.563618958
12	4/20/12	23:24:55	50	55	34.68809331	0.558232015
13	4/20/12	23:25:00	55	60	200.8693535	0.55657053
14	4/20/12	23:25:05	60	65	7.092762856	0.55525837
15	4/20/12	23:25:10	65	70	38.31037034	0.554812253
16	4/20/12	23:25:15	70	75	47.83235797	0.552522468
17	4/20/12	23:25:20	75	80	68.47340398	0.551466942
18	4/20/12	23:25:25	80	85	30.82469193	0.549087751
19	4/20/12	23:25:30	85	90	43.13790773	0.547247589
20	4/20/12	23:25:35	90	95	34.39904785	0.547247589

Figure 7: Exported data in Microsoft Excel

4 Results: Examples of using the Skyline System in EV/Hybrid Development

Mission Motors has tested the Skyline database and web visualization system as part of the Mission's development of advanced powertrains for electric and hybrid vehicles. Skyline has been used to visualize data from internal vehicle development as well as monitoring of prototype vehicles developed with OEM's.

Initial use of the Skyline software platform has demonstrated four primary benefits:

1) Vehicle data is effectively catalogued as it is streamed from vehicles in the field
2) The process of visualizing and filtering to important data is streamlined
3) Multiple users are able to access data on multiple vehicles, increasing specialist interaction with raw data
4) Users use Skyline to identify raw data for further analysis and then export the data to other software tools.

The following case studies provide two examples of how Mission Motors uses Skyline to advance electric/hybrid vehicle development.

4.1 Optimizing Battery System Design by Monitoring Real-World Degradation

Battery systems remain the principal challenge in any advanced electric or hybrid vehicle. As the most expensive part of the electric drive system, the battery pack must be optimally sized to meet vehicle needs. When selecting battery chemistry, sizing a battery system, and appropriately selecting a cooling system, vehicle manufacturers must balance between competing goals for vehicle performance, vehicle range, and battery life.

Real-world battery life is inherently difficult to predict, adding risk and cost to vehicle manufacturers deploying a new electric or hybrid vehicle. Battery performance varies with drive cycle, driving style, environmental conditions, and calendar age of the battery. In order to ensure battery life meets warranty targets, vehicle manufacturers must be conservative in battery pack design, typically by oversizing a pack or the cooling system. Although vehicle and cell manufacturers conduct significant simulation and laboratory cycling, data from vehicles in the field would improve manufacturers' understanding of battery performance.

Mission Motors initial deployment of Skyline explored how the Skyline system could be used for monitoring of battery performance from vehicles in operation. For all Mission electric powertrain systems deployed in the field, Mission integrated a tablet computer running on the Android operating system and with a 3G cellular data connection. Mission designed software on the Android computer to collect powertrain data via the vehicle's CAN network. This included individual channels for voltage data from each battery within the battery pack, collected through Mission's Battery Management System. The computer also collected battery temperature data, and battery current. The tablet computer then streamed this data to the Skyline server where it was catalogued and archived according to the Skyline structure.

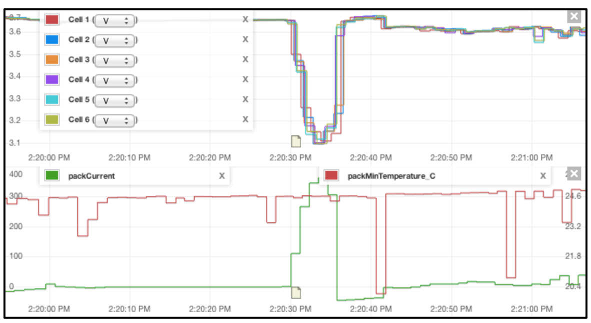

Figure 8: Cell states over 60 seconds

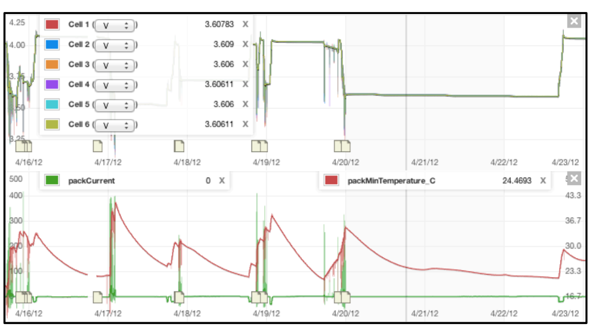

Figure 9: Cell states over 7 days

The resolution of battery data archived in Skyline resembles the level of data collected in a battery laboratory. The data collected through Skyline enables a number of key analyses for battery pack developers. Temperature and performance gradients for cells can be evaluated across the pack. Pack usage can be compared accurately against individual driver duty cycles. Finally, patterns of battery capacity fade can be evaluated from across multiple vehicles.

Although Mission's current battery data in Skyline represents vehicles that have been in the field less than 6 months, the long-term potential of a detailed battery pack archive has been determined. The ability to handle and process detailed battery data can provide a vehicle manufacturer with key insights for the design of next generation battery packs for future vehicle models, optimizing system sizing and cooling requirements. Trend-based analysis of cell level data can potentially identify battery warranty issues before they arise, and can further inform uncertain warranty models.

4.2 Validating and Informing Simulation Models

Typical design process for a powertrain begins with a simulation of every component based on a use profile and a range of environmental conditions. The system designer uses available information from bench testing to inform the simulation model in initial stages of development.

Data collected through Skyline can exercise the simulation model against real-world results in close to real-time. The system can be used to identify weaknesses or areas for further characterization, creating an immediate feedback loop for vehicle models and powertrain developers. Currently such data is collected on prototype vehicles through complex onboard data collection systems and computers that travel with prototypes. Skyline helps engineers overcome this challenge. Skyline can also enable more detailed vehicle model development by encompassing more vehicle and powertrain variables. The more data acquired from vehicles in the field, the more informed engineers can make electric and hybrid vehicle simulations, for both short-term performance and long-term life.

4.3 Drive Cycle Development

Understanding the use case of a vehicle or a battery system is crucial to optimizing the cell selection, pack sizing, and thermal systems; which ultimately drives the cost and performance of an electric vehicle. Standardized drive cycles exist for many vehicle types, and many vehicle manufacturers have gone to extensive efforts to create internally validated drive cycles. However, these drive cycles depend on past data – usually collected with traditional ICE vehicles. With the advent of HEV, PHEV, and BEV architectures, drivers and users may change their behavior due to the vehicle architecture. It is therefore important for a developer to collect drive cycle information specific to the platform architecture.

Skyline is an excellent tool for drive cycle analysis and creation. By exporting data for speed, distance, and altitude across the entire vehicle or fleet's history, an engineer can run statistical analysis on a range of users including the lightest drivers, the most aggressive drivers, and the 90th percentile users. Generating representative drive cycles that correspond directly to vehicles in operation becomes straightforward.

5 Conclusions

The acquisition and analysis of low-level data on batteries and powertrain system performance within an electric or hybrid vehicle can help to create more optimized system designs that better meet the needs of drivers. While vehicle data acquisition and telematics have improved to provide engineers with more vehicle data than ever before, engineers must still overcome the challenge of how to handle and process the data efficiently in order to inform vehicle specification and system design.

Mission Motors developed Skyline as a cloud-based platform for processing and visualization of streaming vehicle data. Skyline is unique in the way it catalogs and stores data streamed to the server from a vehicle. The Skyline database architecture enables data to be processed at multiple sample rates and from multiple vehicles. The Skyline web client enables multiple users to easily view historical or real-time vehicle data through the web.

Skyline has the potential to significantly improve the process of electric/hybrid powertrain development, optimization, and monitoring. Electric/hybrid vehicle design and optimization is a complex process balancing multiple variables affecting vehicle performance, reliability, and life. To properly understand electric powertrain performance, system data must be analyzed at both short time periods (seconds) and long time periods (days and months). The Skyline interface enables engineers to compare multiple channels of data at any time resolution and efficiently identify the most important data for use. This capability can be used to better understand complex problems such as vehicle battery life, vehicle-specific drive cycles, or real-time powertrain performance.

The Skyline software platform can unlock the power of data from electric and hybrid powertrains, enabling us to create more optimized, higher performance, and more efficient drive systems. In the long-term, systems such as Skyline can form the basis for trend-based preventative maintenance, warranty risk mitigation, and battery financing.

6 About the Authors

Seth LaForge is a Sr. Software Architect at Mission Motors and served as the chief software architect for the Skyline system. Prior to Mission Motors, he worked as a software engineer at Google, focusing on the Google Maps application. He has also served as Director of Software Engineering at Newton Research Labs. Mr. LaForge earned a degree in Applied Science and Engineering from the California Institute of Technology.

Jon Wagner Jon Wagner has served as Chief Technology Officer at Mission Motors since 2009. Mr. Wagner has led development of high performance motor and control systems for over 10 years, including development of Mission Motors' suite of electric powertrain products. Prior to Mission, Mr. Wagner was the Sr. Electro-Mechanical Engineer at Velocity11 (now part of Agilent). Mr. Wagner earned his Master's degree in electro-mechanical engineering from Stanford University, and is one of the foremost experts in the Silicon Valley in motor control and the dynamics of complex electro-mechanical systems.

Jit Bhattacharya has served as Mission Motors' CEO since 2009. Mr. Bhattacharya has been studying lithium-ion battery systems for electric and hybrid vehicles since 2007, working as a venture capital analyst for Spring Ventures. Mr. Bhattacharya holds a Master's in mechanical engineering from Stanford and an MBA from the Haas School of Business at UC Berkeley.