Improving Spacecraft Data Visualization Using Splunk

Final Report

Matthew Conte
Polytechnic Institute of New York University

Mentor: Victor Hwang

Definitions:

- DIF - Deep Impact Flyby Vehicle
- DMD pages – GDS interface
- EPOXI – Extrasolar Planet Observation and Deep Impact Extended Investigation
- GDS - Ground Data System
- JPL - Jet Propulsion Laboratory

Background

EPOXI, like all spacecraft missions, receives large volumes of telemetry data from its spacecraft, DIF. It is extremely important for this data to be updated quickly and presented in a readable manner so that the flight team can monitor the status of the spacecraft. Existing DMD pages for monitoring spacecraft telemetry, while functional, are limited and do not take advantage of modern search technology. For instance, they only display current data points from instruments on the spacecraft and have limited graphing capabilities, making it difficult to see historical data. The DMD pages have fixed refresh rates so the team must often wait several minutes to see the most recent data, even after it is received on the ground. The pages are also rigid and require an investment of time and money to update.

To more easily organize and visualize spacecraft telemetry, the EPOXI team has begun experimenting with Splunk, a commercially-available data mining system. Splunk can take data received from the spacecraft's different data channels, often in different formats, and index all the data into a common format. Splunk allows flight team members to search through the different data formats from a single interface and to filter results by time range and data field to make finding specific spacecraft events quick and easy. Furthermore, Splunk provides functions to create custom interfaces which help team members visualize the data in charts and graphs to show how the health of the spacecraft has changed over time.

One of the goals of my internship with my mentor, Victor Hwang, was to develop new Splunk interfaces to replace the DMD pages and give the spacecraft team access to historical data and visualizations that were previously unavailable. The specific requirements of these pages are discussed in the next section.

Objectives

Creating the Splunk pages required more than simply "translating" the DMD interfaces into Splunk views. Splunk's internal architecture and visualization style presented two main problems that needed to be addressed:

1. Although Splunk can present more information through graphs and charts than the DMD pages, it cannot fit as much data onto a single web page. Naturally, the chart visualizations take up much more display space than the single values of the old interfaces. The flight team

often needs a quick overview of the spacecraft's health so it is important to fit as much data onto a single screen as possible.

2. The DMD pages display data by piping data fields from stored telemetry into predefined fields on the pages. In order to display the same data, Splunk must search through its entire data index before sending the results to the display. In order for the Splunk interfaces to be useful, they need to be loaded quickly so the background searches needed to be optimized.

The front-end displays had to be designed so that they contain as much data as possible and waste as little screen space as possible. And for these displays to be rendered quickly, the back-end searches needed to be optimized.

Approach

The first step in this project was to identify the DMD interfaces that are used most often by the EPOXI flight team. There are dozens of DMD pages so in order to best help the team, it was important to concentrate on displaying the data referenced most often. By talking with team members I identified several interfaces to remake in Splunk.

The process for putting the views together was very similar for each page. First, I had to identify which data channel was the source of the necessary data. Once I had the channel name, I then constructed a Splunk search query to return the data. I would then check the results to make sure they were consistent with the DMD pages. Since the speed of the search is so important, writing these queries was an iterative process: I experimented with different search commands and structures to make sure the data was being returned as quickly as possible.

After the search query was written, tested, and optimized, I would use the provided Splunk tools for creating visualizations. For most of the data sources, the visualization consisted of a table showing the most recent values of a channel (for instance, the most recent voltage reading from each of the batteries) and a graph which would plot this reading over time. As with the search queries, I would experiment with different display methods to find the fastest way to render the view.

Results:

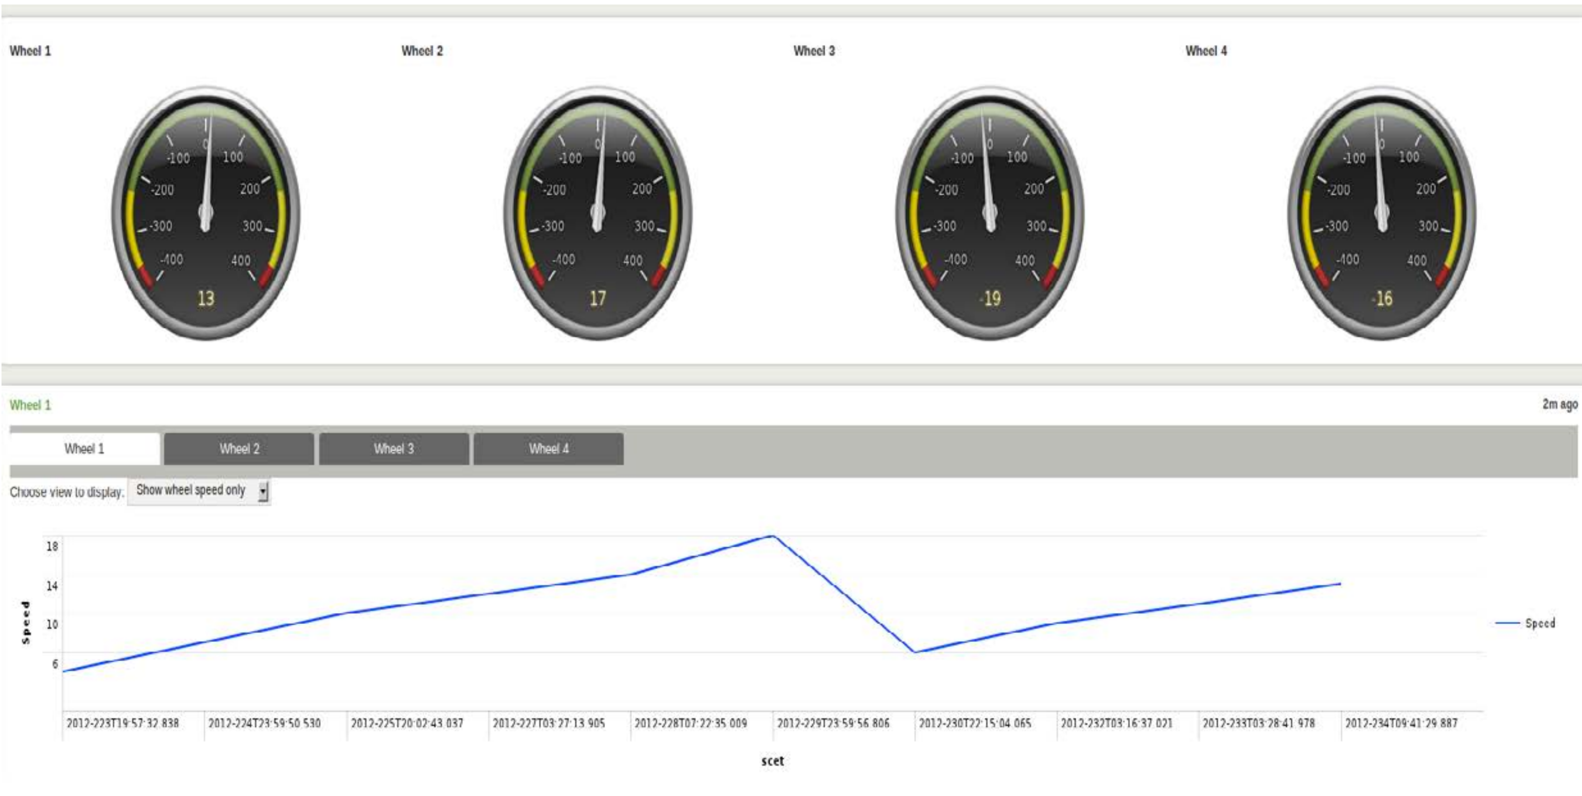

Figure 1: Reaction wheel status page in Splunk

The status page shown above displays both the current velocities of the reaction wheels (devices which control the spacecraft's momentum) as well as historical data. Although this page cannot display as much data as the DMD page for the reaction wheels, it presents the most important and most commonly-used data in a more helpful way. The Splunk page also reflects the most recent telemetry as opposed to the DMD page which is only updated once every 10 minutes.

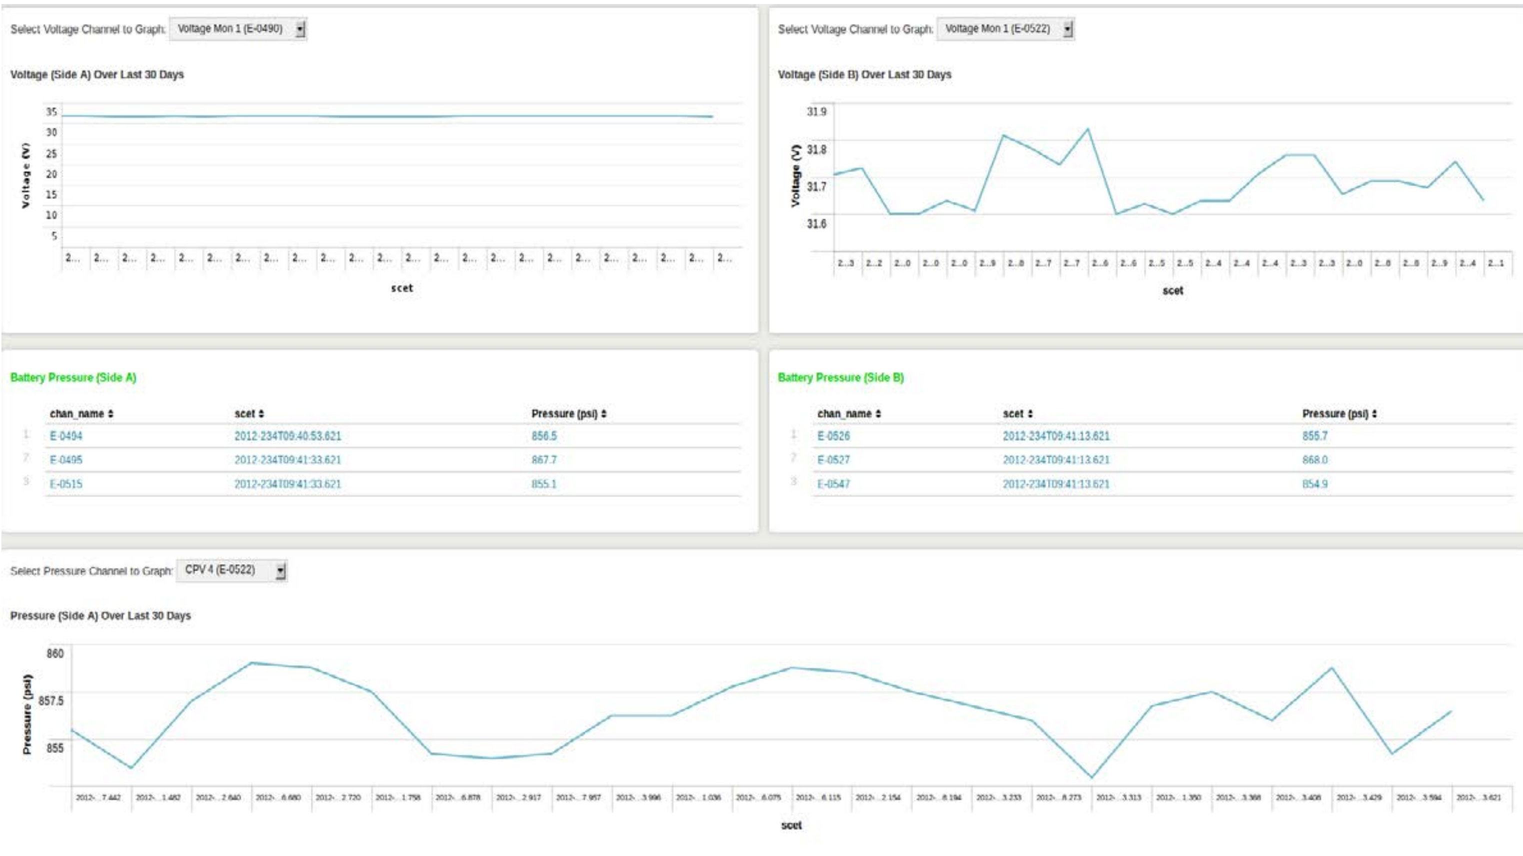

Figure 2: Battery status page in Splunk

Similarly to the reaction wheel page, the battery status page shows the current voltage and pressure readings from the spacecraft's batteries as well as a graph showing these readings over time. In case of an anomalous reading, these graphs will help the team see recent trends in the reading to help diagnose the problem.

Conclusion:

I believe that the Splunk interfaces will be helpful for the EPOXI flight team in monitoring the state of the spacecraft. The pages that I created fulfill the requirements of displaying all the necessary data on a single screen and rendering the visualizations quickly.

However, I do not think that Splunk will be able to completely replace the DMD pages. Because of the architecture of Splunk's search and display systems, it would be impossible to include all of the data from a DMD page in a single Splunk page without the page rendering too slowly. Displaying all of the data using Splunk would require reorganizing each DMD page into several sub-pages in Splunk. This would make all of the data available in Splunk but the spacecraft team members find it most useful to have all data from a particular subsystem on a single display. Thus, Splunk will likely be used best by providing fast access and helpful visualizations for the data that is needed most often. The team would then be able to primarily use the Splunk interfaces and only need to refer to the DMD pages on occasion when they are interested in a more obscure data channel.

Although Splunk will not be able to completely replace the old interfaces, I believe that it will still be a helpful tool for the spacecraft team. It is more interactive than the DMD pages, allowing the team to explore the telemetry data instead of simply reading data points. It provides visualizations that can monitor the spacecraft's health historically and in real time. And it is fully customizable, so the team members can quickly adapt it to their needs. Splunk is a powerful solution for EPOXI's telemetry data and would be a useful addition to other spacecraft teams' toolkits.

Acknowledgements:

This work was carried out at the Jet Propulsion Laboratory, California Institute of Technology, and was sponsored by the JPL Summer Internship Program and the National Aeronautics and Space Administration.