JOURNAL OF AEROSPACE INFORMATION SYSTEMS
Vol. 15, No. 5, May 2018

Check for updates

Telemetry Fault-Detection Algorithms: Applications for Spacecraft Monitoring and Space Environment Sensing

Ashley Carlton,* Rachel Morgan,† Whitney Lohmeyer,* and Kerri Cahoy‡
Massachusetts Institute of Technology, Cambridge, Massachusetts 02139

DOI: 10.2514/1.I010587

Algorithms have been developed that identify unusual behavior in satellite health telemetry. Telemetry from solid-state power amplifiers and amplifier thermistors from 32 geostationary Earth orbit communications satellites from 1991 to 2015 are examined. Transient event detection and change-point event detection techniques that use a sliding window-based median are used, statistically evaluating the telemetry stream compared to the local norm. This approach allows application of the algorithms to any spacecraft platform because there is no reliance in the algorithms on satellite- or component-specific parameters, and it does not require a priori knowledge about the data distribution. Individual telemetry data streams are analyzed with the event detection algorithms, resulting in a compiled list of unusual events for each satellite. This approach identifies up to six events of up to six events that affect 51 of 53 telemetry streams at once, indicative of a spacecraft system-level event. In two satellites, the same top event date (4 December 2008) occurs over more than 10 years of telemetry from both satellites. Of the five spacecraft with known maneuvers, the algorithms identify the maneuvers in all cases. Event dates are compared to known operational activities, space weather events, and available anomaly lists to assess the use of event detection algorithms for spacecraft monitoring and sensing of the space environment.

Downloaded by 73.149.22.231 on May 18, 2018 | http://arc.aiaa.org | DOI: 10.2514/1.I010587

I. Introduction

A. Motivation

SATELLITES provide communications and media distribution, meteorological information, global positioning and timing, reconnaissance, and intelligence services. The reliability of geostationary Earth orbit (GEO) communications satellites (ComSats) is critical to many industries worldwide, such as organizations remotely operating offshore oil and gas drilling facilities, where reliable connectivity is key for safety and efficiency. GEO ComSats make up over 50% of the satellites on orbit (with more than 600 ComSats reported on orbit in 2015), totaling over $208 billion in the services they provide in 2015 [1].

Given societal reliance on satellite services, satellite failures and other events that degrade performance can be highly disruptive [2,3]. Almost all satellites will suffer some unexpected or unusual problems, often referred to as “anomalies”, during their lifetime due to a variety of sources such as the space environment, human error in design or manufacturing, age of the component, or unanticipated software modifications or operating scenarios. The key impacts on satellites due to space weather are 1) surface charging, 2) internal charging, 3) single-event effects or single-event upsets (SEUs), and 4) total ionizing dose (TID) effects. Surface charging occurs due to a difference between ambient electron and ion fluxes, resulting in possible electrostatic discharges. GEO and near-GEO satellites are most at risk due to their movement in and out of the plasmasphere [4]. This can lead to anomalies such as component failures, degradation of sensors and solar panels, and serious physical damage to materials. For a more detailed introduction to the ionizing radiation environment, space environment hazards, and how these hazards lead to anomalies, the reader is referred to [4–6]. Anomalies can cause harmful interference or interruptions in service, reduce spacecraft functionality, or cause complete spacecraft failure. Anomalies can be difficult to diagnose and even more difficult to resolve, draining operator and manufacturer resources (i.e., time, money) in both the short and long term [7]. Quick anomaly identification and resolution can mitigate the effects of anomalies. If there is a “soft” anomaly, like a single-event upset in a field-programmable gate array, the operator can power cycle the device, and nominal operations can resume. For a “hard” anomaly, such as an electrostatic discharge in a power amplifier, the nominal device operation is not recoverable, leading to outages and interruptions, the need to switch to a redundant unit, or a change in operations.

Spacecraft telemetry provides the primary source of health information available from a spacecraft and contains key information that can be used in trying to recover from problems. The word “telemetry” has Greek etymology that means “measurement at a distance”. In the case of a spacecraft, this means downlinking electrical signals proportional to the quantity being measured [8,9]. Spacecraft telemetry originates from sensors from each subsystem, the payload (if applicable), and the attitude control system. Housekeeping data are monitored to check the health and operating status of onboard components in subsystems. Housekeeping telemetry can be in the form of operational or redundancy on/off statuses; sampled temperature, current, or pressure measurements; or deployment of mechanisms [10]. For a modern, large communications satellite, there are typically many thousands of telemetry streams. Monitoring each telemetry stream individually provides the “pulse” of a component or of a subsystem. In this study, we analyze GEO ComSat current and temperature telemetry from solid-state power amplifiers (SSPAs) and thermistors.

B. Fault-Detection Techniques

Longer mission durations, with evolving component performance and reduced dependence on ground control, drive the need for updated fault-detection and diagnosis techniques compared to those currently used on spacecraft. Additionally, the evolving component technology and decrease in lithographic feature size for chips renders components potentially more susceptible to faults [11]. On the component-level, simple

Received 23 June 2017; revision received 31 October 2017; accepted for publication 2 March 2018; published online Open Access 6 April 2018. Copyright © 2018 by Ashley K. Carlton. Published by the American Institute of Aeronautics and Astronautics, Inc., with permission. All requests for copying and permission to reprint should be submitted to CCC at www.copyright.com; employ the ISSN 2327-3097 (online) to initiate your request. See also AIAA Rights and Permissions www.aiaa.org/randp.

*Graduate Researcher, Department of Aeronautics and Astronautics, 77 Massachusetts Avenue. Student Member AIAA.
†Undergraduate Researcher, Department of Aeronautics and Astronautics, 77 Massachusetts Avenue. Student Member AIAA.
‡Associate Professor of Aeronautics and Astronautics, Department of Aeronautics and Astronautics, Earth Atmospheric Planetary and Science Department, 77 Massachusetts Avenue. Member AIAA.

limit-checking methods called "thresholding" are used. Upper and lower numerical limits are set for a particular component. Current thresholding techniques have evolved to include hard and soft limits (e.g., failures and warnings, respectively) and thresholds that are applicable for different situations or operational modes (often called "rule-based" methods) [12,13]. However, component performance may change over time, making certain thresholds no longer applicable. In addition, expert knowledge is necessary to decide the appropriate thresholds for each spacecraft component.

On the system level, polling (or voting) schemes are used on spacecraft. These techniques include "consecutive occurrence counters" and "persistence filters" [12,14], which are designed to quickly assess systemwide faults. Polling schemes also include rule-based methods, using "if–then" rules encoded by system experts. Fault trees are commonly used to determine the actions that should be taken, given the occurrence of a particular fault.

New fault-detection techniques are being developed for aerospace applications but are not yet widely employed on active spacecraft due to the computational resource requirements and limitations of current models that are required for the techniques. Analytical model-based methods compute residuals between measured and estimated values. Current work focuses on developing better dynamical models of aerospace systems and components [15,16]. The European Space Agency's SMART-FDIR study in 2003 uses the Gravity Field and Steady-State Ocean Circulation Explorer satellite simulation environment for validation purposes [17]. For detection of faults, no training data are used. The SMART-FDIR approach relies on fuzzy inductive reasoning using a model-based framework. The system behavioral model makes decisions using possibilistic logic theory, meaning that the behavioral model represents the system knowledge about causal dependencies between inputs and outputs, which are represented by logical formulation. Meitinger and Shulte have developed a technique that uses a cognitive automation approach, which mimics human cognition. Goal-directed planning is implemented while considering the current situation. The approach was tested on a successful unmanned aerial vehicle flight experiment [18]. There are still challenges with the Meitinger and Shulte cognitive automation approach because it depends on a priori knowledge (environment models, models of every subsystem, etc.), increasing the system complexity.

Downloaded by 73.149.22.231 on May 18, 2018 | http://arc.aiaa.org | DOI: 10.2514/1.I010587

C. Spacecraft as a Sensor

There are ongoing efforts to use the "spacecraft as a sensor", most notably by the Aerospace Corporation in El Segundo, California, starting in 1997. Bowman and DeSieno patented their work on the subject [19]. They have developed neural network-based methods that are able to adapt to the dataset but rely on historically labeled training data. Training data instances must be classified as anomalous or nominal. Called "supervised learning", using labeled training data has advantages, such that basic trends and ranges of allowable values are captured and incorporated in the algorithm.

With the algorithms we develop in this work, we aim to detect anomalous events without the use of training data due to lack of availability or applicability. Often, there is not an identical (or similar) mission for reference training data, making labeled training data nearly impossible to obtain or generate. Functional test data could be used from ground testing, but differences in the operational environments and procedures make the performance degradation unpredictable. Although GEO ComSat components come from a small number of manufacturers, on-orbit performance is not always well characterized. "Identical" satellites have varied workmanship, test, environmental exposure, performance parameters, and settings, such as the linear channel amplification (preamp). In addition, "layouts" and possible influences, such as thermal effects, are unique for each spacecraft (no two harnesses are the same, for example). Therefore, training data from a previous "identical" mission may not be directly applicable. Dependence on labeled training data can potentially exclude the system detecting events that do not also exist in the training data.

II. Approach

A. Overview

Using only housekeeping telemetry, our approach is to develop algorithms to extract unusual events. We develop transient detection and change-point detection algorithms that statistically evaluate the telemetry stream compared to a local norm. The algorithms can be applied to telemetry from any spacecraft platform because there is no reliance on satellite- or component-specific parameters, and the algorithms do not require a priori knowledge of the data distribution. A high-level diagram of our approach can be found in Fig. 1.

Events in a spacecraft system are identified, and the number of telemetry streams is increased at each level. First, the detection algorithms are applied to each individual telemetry stream, identifying events in single telemetry streams at the component level. Then, the events detected from all telemetry streams of a certain type or subsystem are examined, such as all of the temperature measurements from thermistors in the amplifier system. Similar performance might be expected from components of the same type, or in the same subsystem, and so directly comparing and compiling the events from that component type may reveal an event that affects the subsystem level.

At the system level, the detection of events from the telemetry streams for many subsystems (or the entire spacecraft system) is considered. The intersection of the events may be indicative of an external (environmental) or internal (spacecraft-level) event. When combining and comparing events in datasets from multiple satellites, the intersection of system-level events could be an environmental event (natural or manmade space environment events that could affect a local environment). Detected events are compared with known operational activities, space weather events, and publicly available anomaly lists, such as SpaceTrak from Serdata.

B. Data Used in this Analysis

GEO ComSat telemetry is evaluated in this analysis for two primary reasons. First, ComSats provide important services, requiring high reliability that can be improved by fault-detection algorithms and an understanding of hazardous space environments. Second, ComSat operators have decades of telemetry, providing a large quantity of data for analysis. For a modern, large communications satellite, there are typically several thousand telemetry streams.

The design lifetime for a modern ComSat is 10–15 years, and many satellites operate beyond this nominal lifetime for return-on-investment reasons, with increasing interest and development in technologies to extend satellite lifetime [20]. The long mission durations of the satellites, which span the temporal variations in the space environment (mostly importantly, the 11-year solar cycle) and thousands of housekeeping telemetry streams provide a long baseline of telemetry. The telemetry is typically archived by operators and only analyzed in the event of an unexplained issue. Collaborating with GEO ComSat operators who maintain these telemetry archives enables scientific and statistical assessment of events, trends, and relationships to known space and operational environments.

Downloaded by 73.149.22.231 on May 18, 2018 | http://arc.aiaa.org | DOI: 10.2514/1.I010587

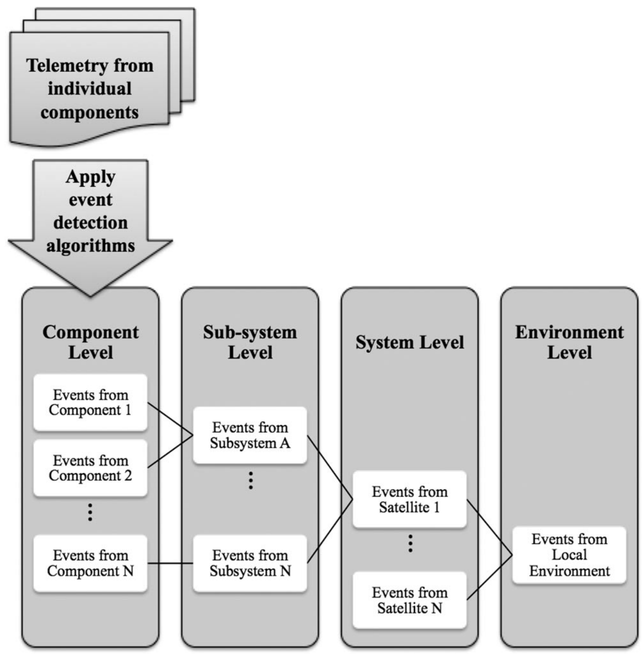

Fig. 1 Approach for telemetry event detection algorithms. Telemetry from individual streams is input to the algorithms, and events are evaluated at different levels: component, subsystem, system, and environment.

C. Telemetry Characteristics and Algorithm Goals

To determine how the event detection algorithms should be designed and operate, we assess the characteristics of our telemetry dataset. We survey anomaly detection techniques for time-series data and identify the algorithm characteristics relevant for our application to GEO ComSats. A key challenge is that it is difficult to apply techniques in one domain to another; the exact notion of an anomaly is often different when applied in another domain. As a result, most existing anomaly techniques solve a specific formulation of the problem. Applications of anomaly detection that are similar or relevant to the spacecraft telemetry event detection problem include intrusion detection, fraud detection, medical and public health anomaly detection, and sensor networks. Common techniques for anomaly detection in time-series data include classification-based, nearest-neighbor, and cluster anomaly detection. Detailed surveys of anomaly detection for time-series data can be found in [21–23].

In telemetry, as a univariate, time-series dataset, point anomalies and contextual anomalies can occur. Point anomalies are when individual data instances can be considered anomalous with respect to the rest of the data. Thresholding techniques can often detect point anomalies, given that the point is a sharp departure from the rest of the dataset. Point anomalies for sequence data (linearly ordered), where data instances are related, can often appear as a set or as collective anomalies [24,25], which are a more difficult detection problem.

Contextual anomalies occur if a data instance is only anomalous in a certain context [26,27]. A satellite temperature measurement of 50°C may be allowable during certain times of the mission but may be contextually anomalous during eclipse periods, when the spacecraft is not directly irradiated by the sun, for example. Context is important because component performance may change over time due to damage or degradation. Defining a "normal" region or time period in the series that encompasses all normal behavior is difficult because boundaries between normal and anomalous are often not well defined. Also, normal behavior can evolve; a current notion of normal behavior might not be sufficiently representative in the future, leading for the need to detect contextual anomalies in time series.

Labeled training data, where dataset instances are labeled normal or anomalous, are difficult or impossible to obtain. In the case of satellite telemetry, it is not guaranteed that even identical satellites will yield the same telemetry or performance. Even if a satellite is designed and built as part of a fleet of identical satellites from one manufacturer, the operational demands and spacecraft environment are often different, making labeled training data inaccurate or nonexistent. As a result, we have chosen to develop event detection algorithms that are unsupervised, meaning that techniques do not require training data. Techniques in this category make the implicit assumption that normal data instances are far more frequent than anomalous ones; if this assumption is not true, then the technique may suffer from a high false alarm rate [21].

Noise in the data is often difficult to identify, distinguish, and remove from the signal. Onboard processor interruptions for task-management reasons, corruption of data, error during communication access, etc., can all affect the telemetry signal. For the telemetry in this analysis, there are dropouts in the data where the values are zeros or a discontinuous jump to later time step. The data dropouts typically lasted between 12 and 24 h in duration with a frequency of once or twice a year.

Because of the size of the dataset, the computational complexity of the detection algorithms is a key consideration, such as the time to run the algorithm, memory available to run the algorithm, and onboard storage for the results. Complexity of the algorithms will affect whether real-time (or near-real-time) algorithms can be used in the future or if the events will be detected retroactively in batch algorithms [23].

The desired output of an anomaly detection technique for telemetry is an "event score", indicative of the degree to which a data instance is anomalous. The scoring enables a ranked list of anomalies from which a user can select the scoring threshold for detection. Scoring can include preferential weighting for telemetry streams that monitor critical payload or mission functions.

D. Spacecraft Telemetry

The GEO ComSat telemetry used in this analysis is from Massachussets Institute of Technology (MIT) collaborations with two commercial operators: Intelsat and Inmarsat, two of the world's leading providers of global telecommunications. Headquartered in Washington, D.C., Intelsat has operated satellites since 1965 and currently operates over 50 satellites.[§] Since 1979, Inmarsat (headquartered in London) has operated communications satellites, with a current fleet of 11 satellites.[¶] A summary of the acquired telemetry can be found in Table 1.

[§]Data available online at http://www.intelsat.com/about-us/overview [retrieved 19 June 2017].

[¶]Data available online at http://www.inmarsat.com/about-us [retrieved 19 June 2017].

Table 1 Summary of telemetry acquired and analyzed from Intelsat and Inmarsat

Company	Intelsat	Inmarsat
Headquarters	Washington, D.C.	London
Number of satellites	21	10
Number of bus types	4	3
Time range	1996–2015	1991–2012
Years of data	20	22
Telemetry sampling rate	Every hour, minute, minor frame (less than 1 min)	Every hour
Data quantity	>0.5 TB	>500 MB
Telemetry obtained	TWTA and SSPA current and temperature telemetry; solar panel current, total bus power, shunt loads; magnetometer measurements	SSPA current and temperature telemetry; solar panel current, total bus power; anomaly and SEU lists

Collaboration with Inmarsat started in 2011. We have analyzed over 22 years of archived telemetry from 10 satellites in three fleets (from different manufacturers, not identified here for proprietary reasons). The data acquired include solid-state power amplifier (SSPA) current and temperature telemetry, solar array power and total bus loads, and anomaly and single-event upset (SEU) lists, totaling approximately 1 GB of data.

In 2014, we began a collaboration with Intelsat. We have analyzed over 20 years of archived telemetry from 22 satellites in four fleets (from different manufacturers, not identified here for proprietary reasons). We acquired current and temperature telemetry from SSPAs and traveling-wave tube amplifiers (TWTAs), solar panel current, total bus power, shunt loads, and magnetometer measurements, totaling over 0.5 TB of data. For this analysis, we focus on the SSPA current and temperature measurements. We also have SEU lists from four satellites.

Because of the proprietary nature of the data, the satellites and components have the following naming convention: COMPONENT-OPERATOR-UNIQUEID. A plot of a nominally performing SSPA from Inmarsat is shown in Fig. 2, called SSPA-INMARSAT-1 in this analysis. A plot of thermistor telemetry from the same amplifier can be found in Fig. 3 (THERM-INMARSAT-1). The sampling time for both figures is hourly and had no reported anomalies (SEUs, etc.). For housekeeping data used in this study, the sampling rates are typically 30 s to 1 min

Downloaded by 73.149.22.231 on May 18, 2018 | http://arc.aiaa.org | DOI: 10.2514/1.010587

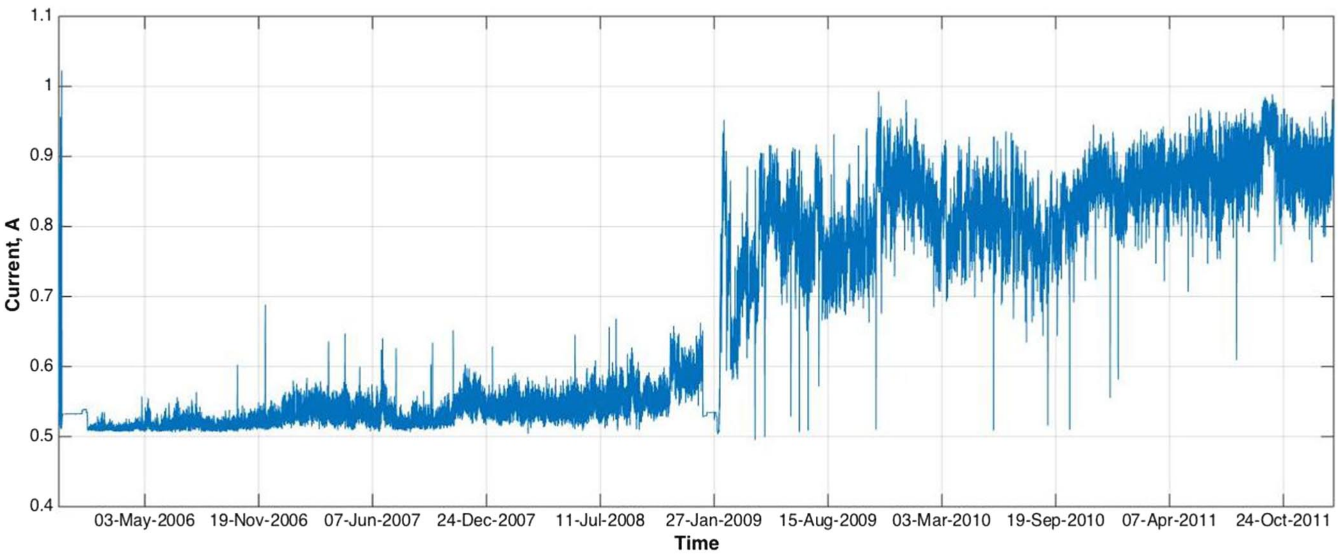

Fig. 2 Nominally performing solid-state power amplifier from an Inmarsat satellite, SSPA-INMARSAT-1.

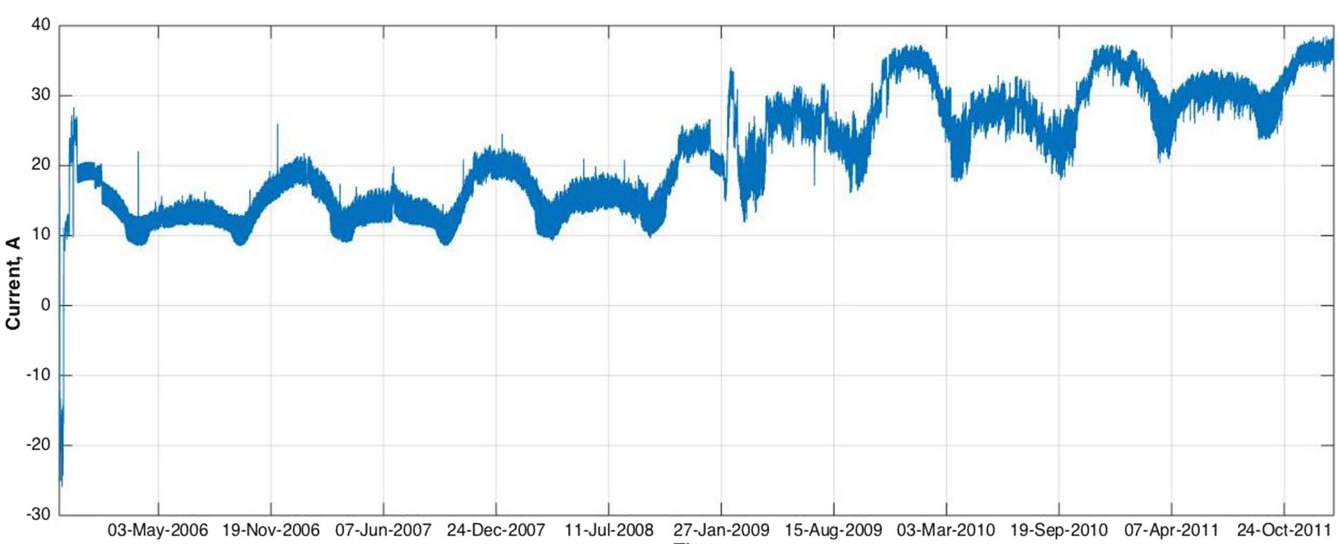

Fig. 3 Nominally performing thermistor in the amplifier payload from an Inmarsat satellite, THERM-INMARSAT-3.

Table 2 Summary of space environment datasets used in the analysis

Types of data	Energetic particle data	Geomagnetic indices	Solar environment
Measurement/index	>2 MeV electron flux, >10 MeV proton flux (daily fluences)	Kp index	Solar wind speeds and densities
Source	GOES	World data center for Geomagnetism (Kyoto)	ACE
Time range	1991–2015	1991–2015	1998–2015
Space environment hazard	Deep dielectric charging; SEUs, TID	Noise or anomalies in power and communications systems	Particle event enhancement

onboard, and then they are averaged and telemetered to the ground with minute or hour sampling resolution. We received minor frame (less than 1 min) sampling resolution for only one satellite, and so we analyze the minute and hourly sampled data in this work.

E. Space Environment Data

We analyze the results of the algorithms with respect to the space environment measurements that are relevant in GEO. Both telemetry and space environment datasets are necessary to quantify space on the satellites it is important to have datasets from the same time period and as spatially colocated as possible [28,29]. We use the Geostationary Operational Environmental Satellites (GOES) >2 MeV electron flux (daily fluence), GOES >10 MeV proton flux (daily fluence), the Kp index (measure of the geomagnetic activity), sunspot daily averages from the World Data Center in Brussels, and solar wind data from the Advanced Composition Explorer (ACE). We have compared telemetry events to other notable space weather events, such as coronal mass ejections (CMEs), interplanetary CMEs, and meteor showers. For all space weather datasets, we obtain data over the entire span of the GEO ComSat telemetry (from 1991 to 2015), except the data from ACE, which became operational in 1998. Table 2 contains a summary of the datasets used.

1. *Energetic Particle*

GOES is operated by the U.S. National Environmental Satellite, Data, and Information Service, providing Earth severe storm tracking and meteorological data. In addition, GOES supports space weather forecasting (see National Oceanic and Atmospheric Association, NOAA, Space Weather Prediction Center**), using the solar x-ray imager and the space environment monitor (SEM), which provides nearly continuous measurements of the energetic particle and magnetic environments in GEO. Because GOES is in GEO, GOES can provide spatially relevant data to the GEO ComSats (not the case for other space weather datasets discussed). The SEM contains an energetic particle sensor (EPS)/high-energy proton and alpha detector used in this analysis, which measures the energetic particle flux (the number of particles through a unit area per unit time). We use the GOES EPS >2 MeV electron flux channel data (daily fluence). Additionally, we use the GOES EPS P4 proton flux channel, which measures protons between 15–40 MeV [30].

2. *Geomagnetic Indices*

The Geomagnetic Equatorial Dst Data Service is hosted by the World Data Center for Geomagnetism in Kyoto, Japan (World Data Center, Data Analysis Center for Geomagnetism and Space Magnetism††). Several types of geomagnetic indices (Dst, Kp, Ap) are calculated at the center that are then verified and archived for public access. This database was used for acquiring values of level 2 Dst and Kp for the duration of the ComSat telemetry for determining dates for severe geomagnetic space weather events between 1991 and 2015.

3. *Solar Environment*

For the solar environment information, we use daily sunspot numbers from the World Data Center Royal Observatory of Belgium in Brussels to compare events with the strength of the solar cycle (for which the sunspot number is a proxy) (World Data Center, Royal Observatory of Belgium‡‡). ACE makes measurements of the solar wind, including the bulk speed, proton density, and ion temperature. ACE has provided operational data since January 1998 and is located at the first Lagrange point (L1), approximately 1.5×10^6 km from the Earth, always observing local dayside. A benefit of ACE's stationary location is that, in combination with the solar wind speeds, one can calculate the time at which the solar wind carrying energetic particles should impact Earth (NOAA§§). The solar wind speeds also provide evidence for magnetopause compression. If the solar wind speeds, and therefore pressure, are high (600–800 km/s), then the magnetopause is likely to compress, placing GEO ComSats outside of the magnetosphere where they are relatively unshielded from the space weather environment (ACE, Science Center¶¶).

F. Known Satellite Anomaly Lists

In addition to space environment measurements, we also examine dates and types of reported satellite anomalies in or near GEO at the same time as events detected in the spacecraft telemetry. Although there is much interest in the creation and maintenance of a robust, open-access, centralized repository of satellite anomalies, it does not exist [31,32]. There exist a handful of anomaly lists maintained by individuals and commercial vendors. The Solar-Terrestrial Physics Division of the NOAA National Geophysical Data Center database contains more than 5000 entries through 1993, but is no longer active, following Dr. Joe Allen's retirement [33]. Satellite News Digest (SND), an industry website, collects information and analyses about satellites, cataloging information for a variety of purposes (Satellite News Digest***). SND's assembled information is largely available to industry members through a subscription, which allows access to SND's anomaly records, maintained by Peter C. Klanowski. Since 1994 (with an English version starting in 1997), SND has maintained an archive of publicly available satellite failure information. SpaceTrak, supported by Seradata, is a professionally maintained satellite information database that contains satellite and payload anomaly information from 1957 through today.††† The information is available through a subscription. Information on the anomaly lists used in this analysis can be found in Table 3.

**Data available online at http://www.swpc.noaa.gov [retrieved 12 June 2017].
††Data available online at http://wdc.kugi.kyoto-u.ac.jp/dstdir/index.html [retrieved 12 June 2017].
‡‡Data available online at http://www.sidc.be/silso [retrieved 12 June 2017].
§§Data available online at http://www.swpc.noaa.gov [retrieved 12 June 2017].
¶¶Data available online at http://www.srl.caltech.edu/ACE/ASC/level2/new/intro.html [retrieved 12 June 2017].
***Data available online at http://www.sat-nd.com/info/about.php [retrieved 05 June 2017].
†††Data available online at http://seradata.com/products/spacetrak [retrieved 05 June 2017].

Downloaded by 73.149.22.231 on May 18, 2018 | http://arc.aiaa.org | DOI: 10.2514/1.I010587

Table 3 Anomaly lists used in this analysis for comparison to ComSat telemetry events

Anomaly list	Description
NOAA National Geophysical Data Center[a]	Maintained by Dr. Joe Allen through individual relationships with operators and vendors, more than 5000 entries up to 1993
Satellite News Digest[b]	Maintained by Peter Klanowski, assembles anomalies from industry starting in 1994
SpaceTrak[c]	Professionally maintained database of satellite information from 1957 to the present, maintained by Seradata

[a]Data available online at https://www.ngdc.noaa.gov/stp/satellite/anomaly/doc/5jsumm.txt [retrieved 05 June 2017].
[b]Data available online at http://www.sat-nd.com/failures [retrieved 05 June 2017].
[c]Data available online at http://seradata.com/products/spacetrak [retrieved 05 June 2017].

Downloaded by 73.149.22.231 on May 18, 2018 | http://arc.aiaa.org | DOI: 10.2514/1.I010587

III. Algorithm Descriptions

Algorithms developed for this analysis include a method for detecting transient events (or “spikes”, jumps, drops) and for detecting change points in the telemetry. Figure 4 shows an example of both types of events. Each event detected has an event date and an event score, which is a metric for the magnitude of the event relative to what is normal in the telemetry. We use a statistical technique consisting of sliding median windows (a variant on Tukey’s method [34]) to find transients and change points. We also considered other statistical methods, such as Bonferroni and Durbin–Watson, but determined Tukey’s methods fit our requirements; no domain knowledge is required, and the algorithms do not contain or impose any component or satellite-specific parameters or thresholds. There is no assumption about the underlying distribution of the telemetry, and no training data are required. Sections III.A and III.B describe the algorithm parameters, scoring metrics, and examples with telemetry for both transient and change point detection algorithms, respectively. A summary description of both algorithms is given in Table 4.

To determine the appropriate statistical representation of a window of data (central tendency), we test the data with a one-sample Komolgorov–Smirnov test [35]. The test rejects the null hypothesis at the 5% significance level that the data are normally distributed (the test result states that the data distribution is not Gaussian). Figure 5 shows the empirical cumulative distribution function (CDF) compared to the standard normal CDF. If the data were normally distributed, they should align with the standard normal CDF. Telemetry streams from other components were also tested and yielded the same result, rejecting the null hypothesis that the data are normally distributed. Therefore, for both algorithms, nonparametric statistical parameters are used; the median is chosen as the statistical metric instead of the mean. The median is more robust to outliers, which is advantageous if the outliers are what we are interested in detecting.

A. Transient Event Detection

Transient event detection is based on the Tukey method, relying on statistics of the dataset to identify deviants [34,36]. In our algorithm, the data are segmented into bins (or windows) that allow for the definition of normal to change over the dataset. The window is shifted along the dataset, comparing each telemetry data point to the rest of the data points in the window. This approach is similar to that used by Hewlett-Packard in their online anomaly detection in data center management, which also uses a variant of the Tukey method [36].

We choose a constant segmentation scheme (same duration window sizes) of seven days. One orbit of a GEO satellite is one day, and so seven days mitigates effects of the orbit. The telemetry sampling rate is hourly, which is 168 data points per window. The one-week window is a tunable parameter: the routine can accept an integer number of data points as the window size (making it dependent on the data resolution). We have analyzed the data with different window durations and have found no impact on the results.

Each telemetry data point x_i is compared to the local median X. For each data point, the number of standard deviations from the local median is recorded. The number of standard deviations is the event score S, which is calculated as follows:

$$S = \frac{|x_i - \text{median}(X)|}{\sigma} \tag{1}$$

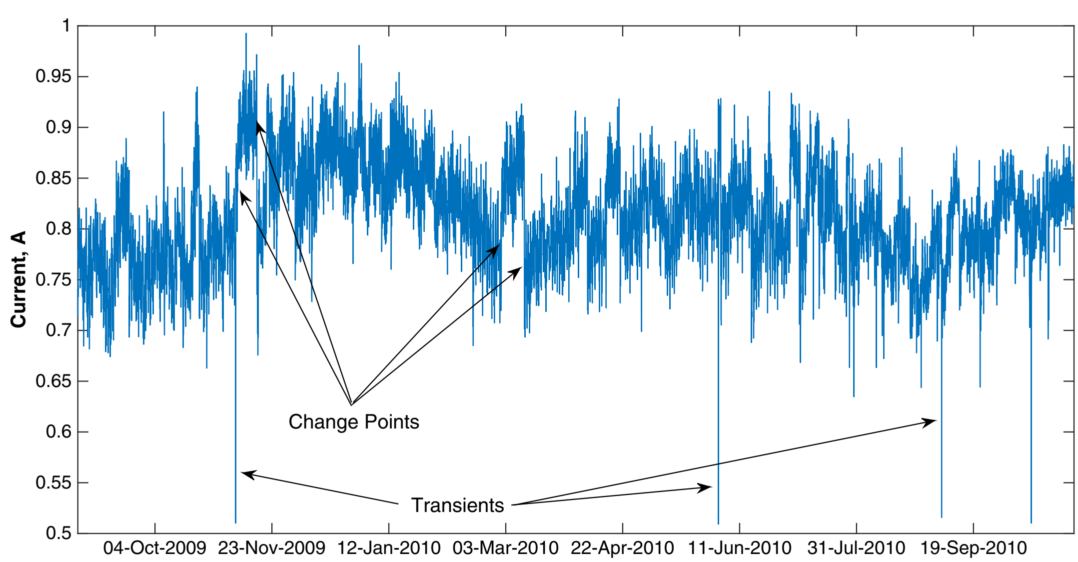

Fig. 4 Plot of solid-state power amplifier telemetry, SSPA-INTELSAT-3. The large transient events and change points are marked as examples of events detected by the two algorithms.

Table 4 High-level description of transient and change point event detection algorithms

Algorithm description	Transient event detection	Change point event detection
Approach	Use single sliding window	Use two adjacent windows, slide windows one data point at a time
Detection	Each telemetry data point is compared to the local median	Compare median between two adjacent windows
Event scoring	Number of standard deviations telemetry data point is from local window median	Percent change in median between two adjacent bins

Selecting the events that have the highest event scores (those that "deviate the most") or selecting the events with scores higher than a threshold lets a user choose how many events to keep.

In Fig. 6, the transient events detected by the algorithm are marked with dashed purple lines on an example of SSPA telemetry, SSPA-INTELSAT-4. The event dates are plotted with standard deviation (in purple, right y axis) and overlaid with the raw telemetry (in blue, left y axis) for greater than three standard deviations from the "local median" (99.73% of the values lie within three standard deviations of the bin median). The transient event identification method detects all noticeable spikes, performing as intended.

Figure 7 shows another example of amplifier telemetry with transient events detected (SSPA-INMARSAT-2). Figure 7 shows the entire dataset in the top panel (2006–2011) and a zoomed-in region from December 2007 to October 2008 in the bottom panel. There are more events detected in the first half of the data (210 spikes detected before January 2009) than in the second half of the data (115 spikes detected after January 2009). This is because the standard deviation in the data is small (0.02276) before a change between December 2008 and February 2009. The data become noisier after 2009 and contain larger changes in median. This change is seen across all Inmarsat datasets examined and may be indicative of a systemwide internal and/or external effect for those satellites.

B. Change Point Event Detection

Change point detection is used in industries including finance, medicine, and sociology. Techniques for change point detection, such as those for climate change detection, genetic time series analysis, and intrusion detection in computer networks, can be found in the literature [36–39].

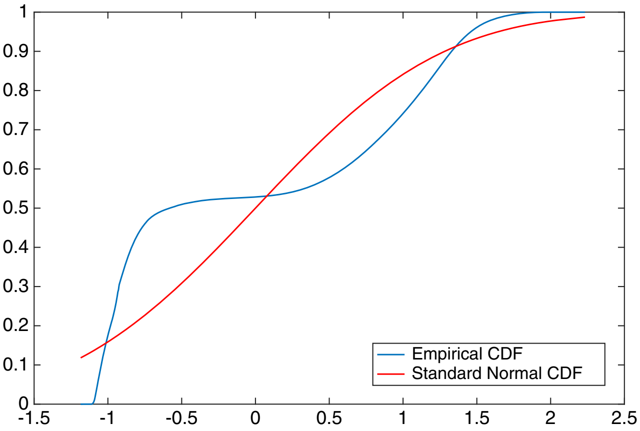

Fig. 5 Empirical CDF compared to the standard normal (Gaussian) CDF for a one-sample Kolmogorov–Smirnov test on an example telemetry stream.

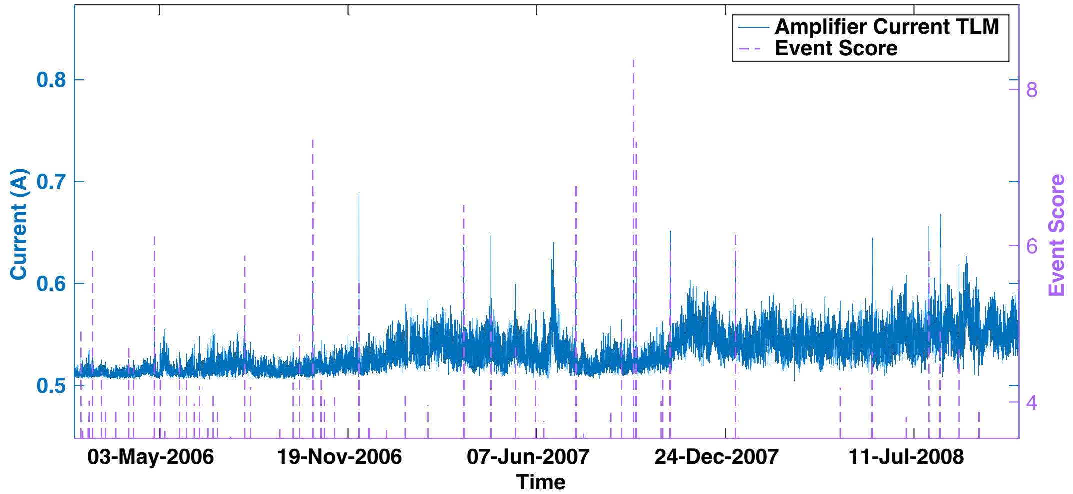

Fig. 6 Transient event detection in SSPA-INTELSAT-4 telemetry stream (blue). Events detected are marked (dashed purple lines), the height of which are measures of the events' deviation from local normal.

Downloaded by 73.149.22.231 on May 18, 2018 | http://arc.aiaa.org | DOI: 10.2514/1.I010587

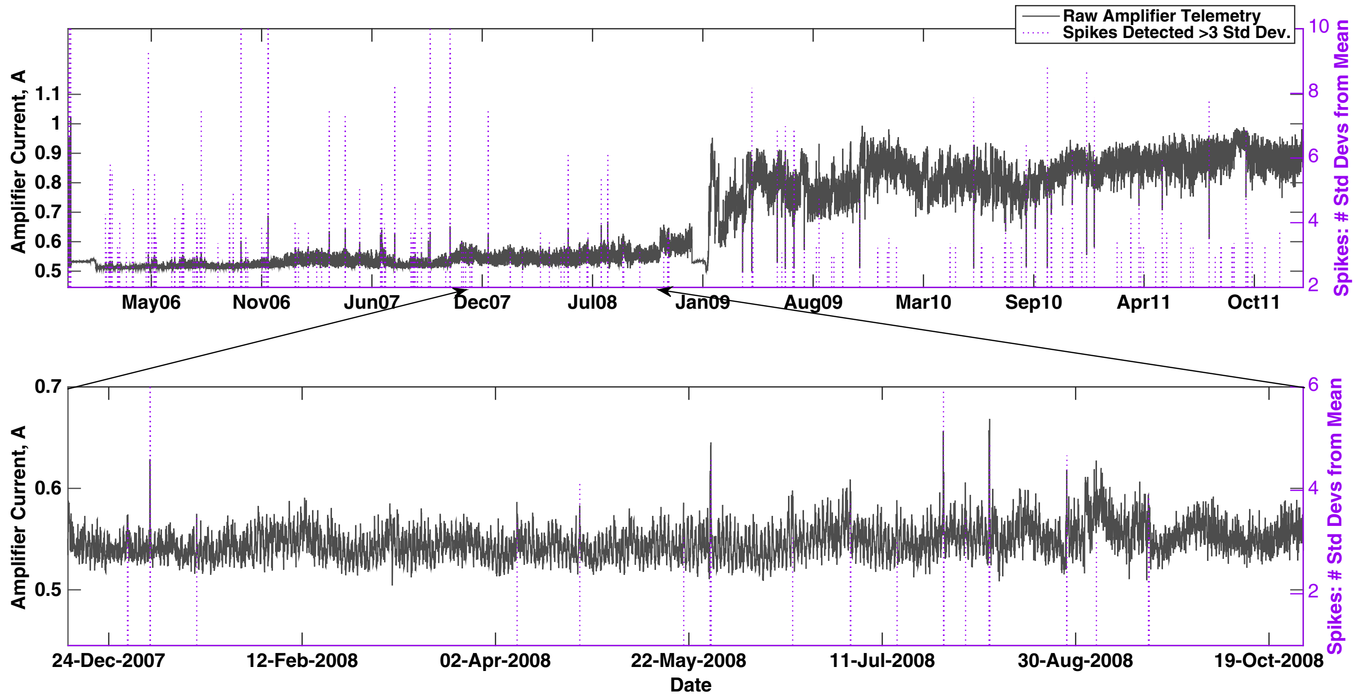

Fig. 7 Transient event detection in SSPA-INMARSAT-2 telemetry. The bottom plot is a zoomed in on December 2007 to October 2008 demonstrating the algorithm's ability to identify spikes from the local median.

Downloaded by 73.149.22.231 on May 18, 2018 | http://arc.aiaa.org | DOI: 10.2514/1.I010587

A common way to detect the change points and determine trends in time series data are by representing the data using piecewise linear approximations (PLAs) or piecewise linear representations [40]. We use PLAs and approximate the data in bins (or windows). The median statistics for each bin are compared. Bin size selection, or segmentation, is an active area of research [41,42] and will be explored in future development of the telemetry event identification algorithms. We currently do not use dynamic bin sizing. We choose to keep the window size the same and optimize its range in time using a moving window technique [43].

The change point detection algorithm employs an optimization scheme for finding the events, or "change dates." The data are initially segmented into seven-day windows (or "bins"), and the weekly bin statistics (median, standard deviation, etc.) are computed. The rationale for the seven-day window is the same as for transient event detection; a GEO orbit is one day, and so a seven-day window mitigates some of the variation during an orbit. For bins that show large changes in median between adjacent bins, an optimization routine is employed using a combination of two moving windows, one data point at a time (telemetry is in hourly resolution) to maximize the difference in the median between two adjacent windows in the local time frame of the originally large change in median. When the change in median is maximized locally, the date of the "event" is recorded as the boundary between the two moving windows.

Figure 8 shows an example of a section of amplifier telemetry from SSPA-INTELSAT-5 after median change detection and optimization. The event score is determined by the magnitude of the change in median at the event date, which is calculated as follows:

$$S = \frac{|\text{median}(X_1) - \text{median}(X_2)|}{|\text{median}(X_1)|} \times 100 \tag{2}$$

The change point detection algorithm reports the largest changes in median using the event score S and numbers the events by their event score in descending order. We select the number of changes to be reported (e.g., 30 changes detected in Fig. 8). The event date with the largest change in median is "#1"; the second largest is "#2", etc. As a result, when compiling all events together, the output is a list of median change event dates ordered by significance, from which one can then select a minimum threshold after evaluating the entire dataset (i.e., we do not want to restrict an amplifier, for example, to only detecting 15 events if many other amplifiers or components have 20 or more events detected, and we do not want to have the algorithm identify 15 event dates from a second amplifier if only 10 have notable changes in median).

The large change occurring in January 2009 is, indeed, detected, identified, ranked #2 for the amplifier. The sharp drop in current at the very beginning of the telemetry (during commissioning) is ranked #1. In Fig. 8, the slight change seen by eye between the #9 ranked event and the #12 ranked event does not have a large-enough change in median to be ranked in the top 30 events selected. For example, if the number of events to detect is increased to 40, the change between the #9 ranked event and the #12 ranked event is detected and ranked #39. The transient detection method did identify the spike between the #9 ranked event and the #12 ranked event.

IV. Results and Event Analysis

We are looking for events that are detected on the same dates from components on the same satellite and event dates common to multiple satellites, which may be indicative of a system-level event and an environment-level event, respectively. We find certain event dates when many if not all of the components of a particular subsystem or satellite have an event detected. To try to explain the common event dates, we compare the events to known spacecraft activity, space weather, and anomaly data sets.

Downloaded by 73.149.22.231 on May 18, 2018 | http://arc.aiaa.org | DOI: 10.2514/1.010587

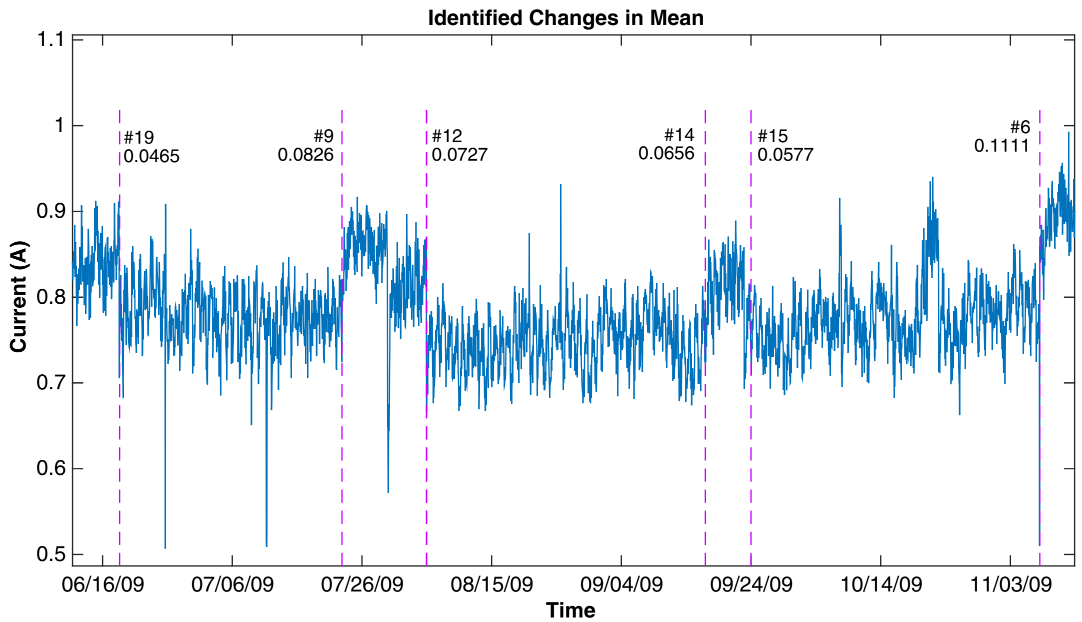

Fig. 8 Median change detection algorithm applied to amplifier telemetry from SSPA-INTELSAT-5 (blue) zoomed into a region (June 2009 to December 2009) with the detected median changes (dashed purple) using weekly bins.

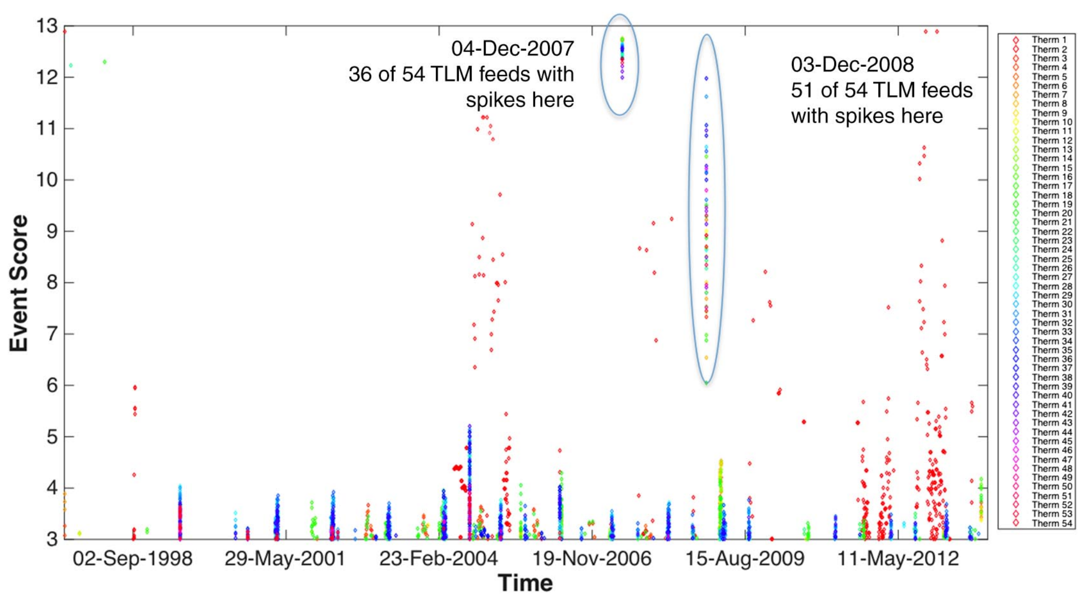

Fig. 9 Transient events detected in 54 thermistor telemetry (TLM) streams in one satellite (INTELSAT-A) from 1998 to 2012. Each thermistor is a unique color.

A. Compiling Events from Individual Types of Components

When comparing the events detected for a type of component on a satellite, we find clusters of events at dates that seem unlikely to be random. For example, for a set of 54 thermistor telemetry streams from one satellite, INTELSAT-A, we have plotted the date and event score for each telemetry stream (each thermistor with a different color). Figure 9 shows that there are dates where there are large transient events across many, if not all, of the telemetry files. Large events (dates when many telemetry streams have high event scores) in December 2007 and December 2008 are annotated in Fig. 9. This suggests that there is some relationship between the transients in the telemetry and system-level events.

THERM-INTELSAT-1 and THERM-INTELSAT-2 in Fig. 9 (shown in red and red-orange) do not show clustering with the other thermistor events. The dispersion exhibited by the event dates from THERM-INTELSAT-1 and THERM-INTELSAT-2 could be due to the physical locations of the thermistors with respect to the others. THERM-INTELSAT-1 and THERM-INTELSAT-2 are from the propulsion system, located near the thrusters, whereas all the other thermistors analyzed are from the high-power amplifier payload.

There appears to be a rough periodicity of events with event scores between 3 and 4. This periodicity is likely explained by the eclipse seasons of the satellite. Figure 10 shows the same thermistors plotted in Fig. 9 with the eclipse seasons highlighted. The light-green periods are spring season eclipses, and the light-red periods are fall season eclipses. The events appear to be related to moving in and out of eclipse. It should be noted that the largest events (highest event scores) and thermistors exhibiting dispersion do not appear to be associated with the eclipse seasons.

B. Compiling Events from Components on One Satellite

We find, when compiling the event dates from all components from one satellite, that there are common event dates, regardless of component type. The occurrence of these events indicates possible system-level events. Binned in days, we have summed the event scores from all the events

Downloaded by 73.149.22.231 on May 18, 2018 | http://arc.aiaa.org | DOI: 10.2514/1.I010587

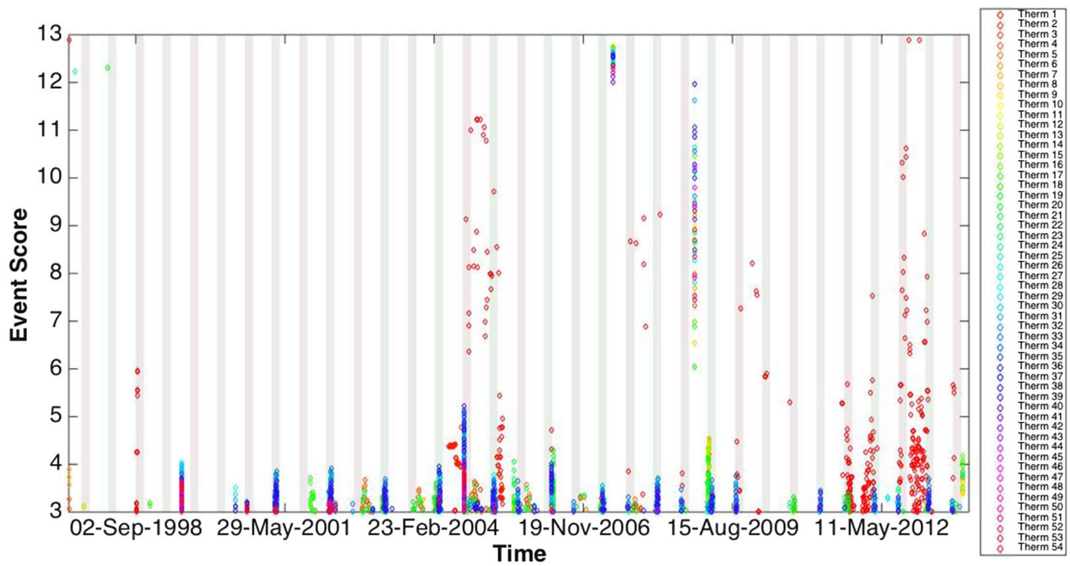

Fig. 10 Events detected in INTELSAT-A thermistor telemetry annotated with eclipse seasons. The light green and red shaded regions are spring and fall eclipse seasons, respectively.

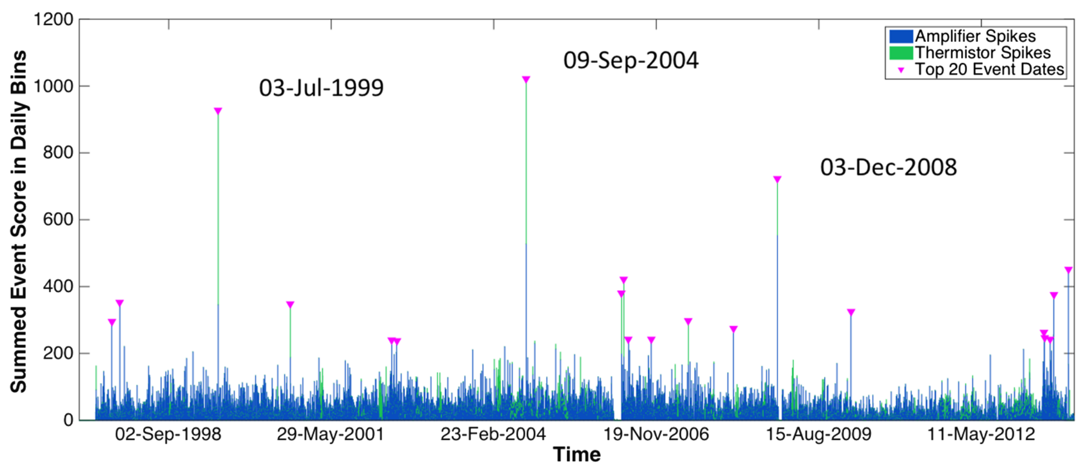

Fig. 11 Daily summed event scores for 185 telemetry streams from one satellite, INTELSAT-A. The amplifier events and thermistor events are summed to show the coinciding top events.

detected in amplifier and thermistor telemetry that we obtained from INTELSAT-A (185 telemetry streams). Figure 11 shows the summed event scores by date. The amplifier events are in blue, and the thermistor events are in green. The top events (dates with the highest summed event scores on that day) are marked with upside-down purple triangles.

For the top event dates, we have compared the date to the space environment, known spacecraft operations, and anomaly lists. Figure 12 shows the events and event scores plotted with space weather metrics identified in Sec. II.E. The events are from an analysis of 199 telemetry streams from INTELSAT-B. Table 5 summarizes the findings of the top events. For each satellite that had a maneuver during its lifetime, the maneuver is detected as one of the top 5 events for each of the satellites. This is a significant finding because this information could be useful for groups interested in a satellite's activity, such as for a space situational awareness application.

For four of the Intelsat satellites, we have SEU dates and times. There are two instances where the event detected in the telemetry occurs within a couple of days of an SEU (see Table 5). We find no statistically significant relationships between the SEUs and the events detected by the algorithms.

C. Compiling Events from Multiple Satellites

The same event detection algorithms were deployed on the other satellites in the dataset. We find that there are dates where the top events from one satellite coincide with a top event from another satellite (INTELSAT-A and INTELSAT-E). From one satellite, we see a top event date of 3 December 2008 (as shown in Fig. 9). In another satellite, we see a top event data of 4 December 2008. The algorithms currently sum event scores by day, summing the event scores for the 12:00 a.m. to 11:59 p.m. time period, allowing for quick examination of data when over a decade of telemetry is analyzed. Upon closer investigation, the events from INTELSAT-A and INTELSAT-E occur at 3 December 2008 23:25:59 and 4 December 2008 04:59:59, respectively. (Note that this analysis is using hourly telemetry, and so the times listed here are not exact to the digits reported. The time reported in this analysis is the time that the telemetry was sampled that hour.) Telemetry from each satellite spans over a decade, and so it is unlikely that these events are unrelated. It could be indicative of an environment-level event.

We have examined the locations and separation of these two satellites on orbit with respect to morning/daylight sector passing the terminator. Because of the proprietary nature of the telemetry and the spacecraft from which the telemetry originates, we cannot comment on the precise

Downloaded by 73.149.22.231 on May 18, 2018 | http://arc.aiaa.org | DOI: 10.2514/1.I010587

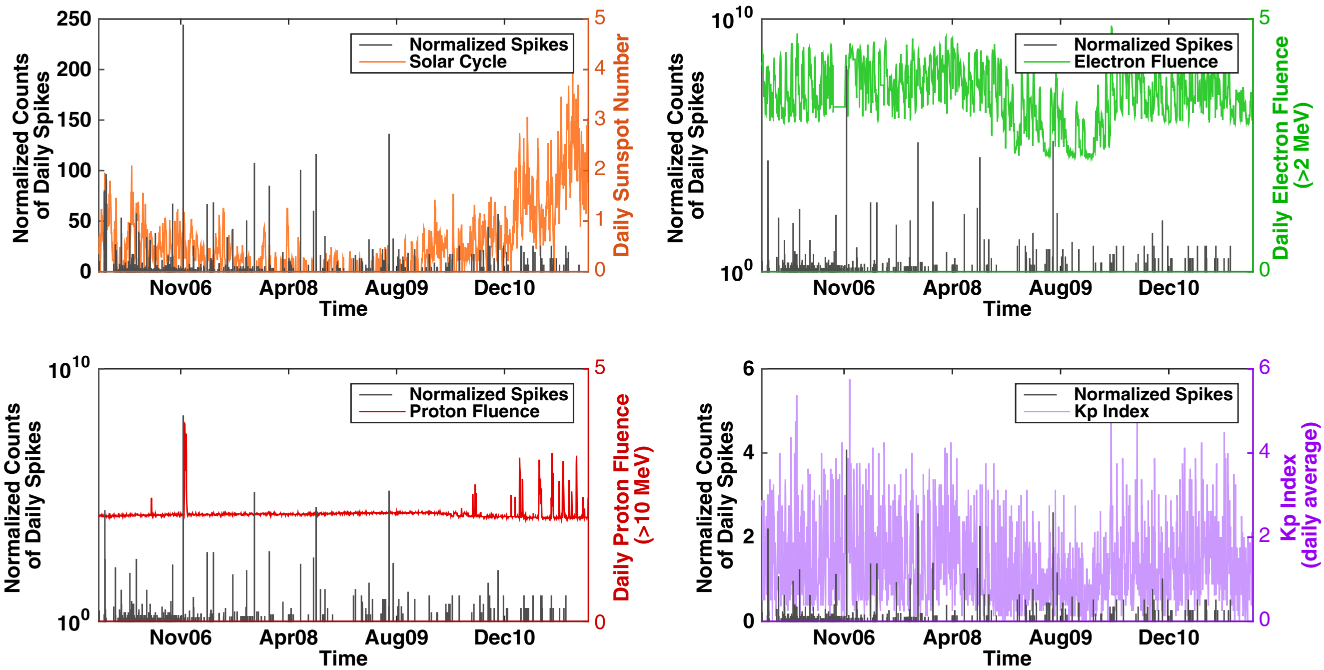

Fig. 12 Plots of daily summed event scores for 199 telemetry streams from one satellite, INTELSAT-B, compared to space weather metrics: sunspot number, daily electron fluence, daily proton fluence, and Kp index.

locations of the spacecraft, but INTELSAT-A is located to the east of INTELSAT-E. INTELSAT-A is the first to experience the event, and then about 5 h later, INTELSAT-E experiences the major event (having moved eastward toward INTELSAT-A's initial position). The time spacing between the locations is roughly 2 h. The top events are recorded within 5 h of each other. It seems very unlikely that these events are unrelated, and there is likely an environmental reason for the largest events for both satellites to occur within a few hours of one another. Further analysis is required to determine the cause of this potential environmental event.

We compare the high-energy electron fluence accumulated before the largest events. Build up of high-energy electrons can lead to charging of dielectric materials, potentially causing catastrophic discharges [44,45]. Comparing the fluence accumulated 1, 7, 10, 14, and 21 days before the large-event dates to a random Monte Carlo sampling of days, we find no statistically significant relationship between accumulated fluence and the events detected by the algorithms [45].

Ten out of the top 20 events for INTELSAT-A and 11 out of the top 20 events for INTELSAT-E detected occur when there is a sharp increase in solar wind speed. This could be due to such events as the passing of a fast CME from a coronal hole, or a slower CME being overtaken by a faster CME. Future work includes a more detailed investigation of the timing of the solar wind changes compared to the event times, and Monte Carlo experiments as a part of a more rigorous study to interpret the findings.

D. Algorithm Performance–Computational Resources

Because the transient detection algorithm involves comparing each individual data point to the local median, the execution time is expected to be linearly dependent on the size of the telemetry file. Running the algorithm on our dataset confirms this expectation of a linear relationship between computational time and the number of data points. For example, for an Inmarsat telemetry file with hourly resolution (6.14 years, 1,284,144 data points), the algorithm takes 0.0721 s. For reference, the total time to run the transient detection algorithm is 0.572 s on 18 telemetry files.

Running the change point detection algorithm on our dataset shows a linear, nearly constant relationship between execution time and number of telemetry data points. For example, for an Inmarsat telemetry file with hourly resolution (6.14 years, 1,284,144 data points), the algorithm takes 0.576 s and the average time is 0.553 s for all telemetry files. For reference, the total time to run the change point detection algorithm on 18 telemetry files is 9.950 s. The transient and change point detection algorithms were run on a MacBook Pro with OSX version 10.9.5, 2.5 GHz processor, 16 GB memory, and MatLab version 8.4.0.150421 (R2014b).

Table 5 Top event date analysis for one satellite, INTELSAT-B[a]

Event date	Space environment	Spacecraft operations	Other reported anomalies
3 Dec. 2008	None, quiet	Transponder SEU, 7 Dec.	Thruster anomaly with GOES-12
9 Sept. 2004	Fast solar wind from coronal hole arrived at Earth, 6–7 Sept. 2004	None	Thaicom outages 12 Sept. 2004
2 July 1999	Moderate $Kp = 4$	None	Echostar IV fuel system anomalies, July 1999, ABRIXAS failure, of onboard batteries
28 Oct. 2013	Handful of powerful CMEs starting 25 Oct., several associated X-class flares	None	Unknown
4 Nov. 1997	None, quiet	Maneuver in progress	Unknown
28 Feb. 2010	28 Feb. 2010 large CME (not Earth-directed)	Maneuver in progress, Transponder SEU, 27 Feb.	AMC-16 further degradation of solar arrays early March 2010

[a]The event dates are compared to the space weather environment, spacecraft operations, and reported spacecraft anomalies.

Table 6 List of top event dates from all satellites analyzed from 1991 to 2015

3 Dec. 2008	4 June 2007
9 Sept. 2004	16 Sept. 2000
2 July 1999	7 March 2008
28 Oct. 2013	1 June 2013
4 Nov. 1997	18 Oct. 2006
28 Feb. 2010	8 July 2013

V. Conclusions

Change point detection and transient detection algorithms have been developed that identify deviations from normal, avoiding component- or satellite-specific conditioning. Although the metrics are well established, these first-step methods are not currently used to evaluate spacecraft health using telemetry. Analyzing housekeeping telemetry from 32 geostationary Earth orbit (GEO) ComSats from 1991 to 2015, we find transient and change point events that occur in many or all of the telemetry files, indicating potential system-level effects on certain dates. Spacecraft maneuvers are detected in five out of five cases. Events from thermistor telemetry show periodicity with eclipse season. At the environment level, a top event date is found that occurs in two satellites: 4 December 2007. The space environment has been considered as well as available operational information and anomaly reports from other spacecraft in GEO. The present findings have been shared with the satellite operators, who were interested in the results and are using the approach to improve anomaly management in their systems. Government agencies were also interested in these findings, and requests for additional information have been supported. Further analysis with the events detected is required to determine if a relationship exists. The top event dates from all satellites analyzed from 1991 to 2015 are:

The authors are actively engaged in pursuing more GEO ComSat telemetry and spacecraft operational information for testing and validation and to continue algorithm development. They would like data from components other than the amplifier system with accompanying operational procedures, command logs, and reported anomalies. It would be expected to find that certain components are more or less sensitive to external environment effects or to systemwide effects, either by their location within the spacecraft or their design and function. Therefore, it is planned to obtain larger datasets that are representative of the entire spacecraft.

In addition, the authors are interested in telemetry from orbits other than GEO. They are in the process of getting access to telemetry from the Van Allen Probes, Lunar Reconnaissance Orbiter, and U.S. Air Force Research Laboratory spacecraft. These spacecraft are selected because each is equipped with dedicated space environment monitoring technology, which will allow for better validation of algorithms. Events that are detected in the telemetry can be directly compared to the local space weather environment as detected onboard. The data providers have also indicated that they would be able to supply command logs and other operational information.

Using the change point and transient event detection algorithms, larger window sizes will be investigated, which may allow detection of events that are anomalous in more of a seasonal context. Adding other detection routines is also being explored, such as changes in the slope or noise envelope (or variance) in the telemetry. The detection algorithms are retrospective, "batch" algorithms, looking back at previous windows of data and analyzing them. For eventual use in operations, it is planned to enable real-time (or near-real-time) detection, moving the algorithms "online". There are several ways that these algorithms could be transitioned into real-time (or near-real-time) use. One of the simplest ways would be to have the seven-day window be a "lookback" window instead of having the data point in the middle of a window. This would work for transient event detection. For change-point detection, one could shorten the first window to perhaps a couple of hours or so, depending on the response time required and the tolerance to false positives.

The authors are also looking into integrating learning algorithms for use after a certain period of time on orbit. Although they do not want to rely on it for early operations or for changing or new environments, learning could be beneficial once in a nominal orbit. The training does not have to be in advance or real time. They are looking into techniques that, after a certain amount of elapsed mission time, can use learning by modeling previously seen telemetry. The response to hazards using the initially proposed algorithms can be fast, but the learning can be slow, supplementing algorithms if there are nominal operating periods. This will require increasing the computational speed of the algorithms as well.

Acknowledgments

The authors would like to thank Inmarsat and Intelsat for contributing the telemetry and technical guidance. Inmarsat and Intelsat are leading the way for scientific research using telemetry by allowing MIT to have access to the data, demonstrating a commitment to academic research to improve space operations. We greatly appreciate the interaction with company staff to help interpret data when needed. The Inmarsat and Intelsat satellite telemetry is proprietary and thus cannot be distributed. The GOES particle data are publicly available. The authors would like to acknowledge the U.S. Air Force Office of Sponsored Research grant FA9550-13-1-0099 and NASA for funding this work through NASA Space Technology and Research Fellowship grant NNX16AM74H.

References

[1] "State of the Satellite Industry," Satellite Industry Association, Sept. 2016, http://www.sia.org/wp-content/uploads/2017/03/SSIR-2016-update.pdf [retrieved 30 May 2017].

[2] Schrijver, C. J., "Socio-Economic Hazards and Impacts of Space Weather: The Important Range Between Mild and Extreme," *Space Weather*, Vol. 13, No. 9, Sept. 2015, pp. 524–528.
doi:10.1002/2015SW001252

[3] "Severe Space Weather Events: Understanding Societal and Economic Impacts: A Workshop Report," National Academies Press, Washington, D.C., 2008, https://www.nap.edu/catalog/12507/severe-space-weather-events-understanding-societal-and-economic-impacts-a.
doi:10.17226/12507

[4] Hastings, D., and Garrett, H., *Spacecraft-Environment Interactions*, Cambridge Atmospheric and Space Science Series, Cambridge Univ. Press, Cambridge, England, U.K., 1996, pp. 142–207, Chap. 6.
doi:10.1017/CBO9780511525032

[5] Fennell, J., Koons, H., Roeder, J., and Blake, J., "Spacecraft Charging: Observations and Relationship to Satellite Anomalies," TR-2001(8570)-5, Aerospace Corp., El Segundo, CA, 2001.

[6] Baker, D., "The Occurrence of Operational Anomalies in Spacecraft and Their Relationship to Space Weather," *Institute of Electrical and Electronics Engineers Transactions on Plasma Science*, Vol. 28, No. 6, 2000, pp. 2007–2016.
doi:10:1108/27.902228

Downloaded by 73.149.22.231 on May 18, 2018 | http://arc.aiaa.org | DOI: 10.2514/1.010587

[7] Miau, J.-J., and Holdaway, R., "The History and Continuing Quest to Reduce the Costs of Spacecraft Ground Stations and Operations," *Reducing the Cost of Spacecraft Ground Systems and Operations*, Vol. 3, Space Technology Proceedings, Springer, The Netherlands, 2000, pp. 1–7. doi:10.1007/978-94-015-9395-3

[8] Roddy, D., "The Space Segment," *Satellite Communications*, 4th ed., McGraw–Hill, New York, 2001, Chap. 7.

[9] Larson, W. J., and Wertz, J. R., (eds.), *Space Mission Analysis and Design*, 3rd ed., Microcosm Press, El Segundo, CA, 1992, p. 330 doi:10.1007/978-94-011-2692-2

[10] Fillery, N., and Stanton, D., "Telemetry, Command, Data Handling and Processing," *Spacecraft Systems Engineering*, edited by P. Fortescue, G. Swinerd, and J. Stark, Wiley, Chichester, England, U.K., Aug. 2011, pp. 439–466, Chap. 13. doi:10.1002/9781119971009

[11] de Lacy, H., and Jones, A., "Shrinking Silicon Feature Sizes: Consequences for Reliability," *Proceedings of the Conference for Military and Space Electronics*, Portsmouth, England, U.K., 2008, https://shop.micross.com/pdf/Micross_Technical_Presentation-Effect_of_Die_Shrinkage.pdf [retrieved 30 May 2017].

[12] Morgan, P., "Fault Protection Techniques in JPL Spacecraft," *Proceedings of the 1st International Forum on Integrated System Health Engineering and Management in Aerospace*, Jet Propulsion Lab., National Aeronautics and Space Administration, Pasadena, CA, Sept. 2006.

[13] Wander, A., and Förstner, R., "Innovative Fault Detection, Isolation and Recovery On-Board Spacecraft: Study and Implementation Using Cognitive Automation," *Proceedings of the 2013 Conference on Control and Fault-Tolerant Systems*, IEEE, Washington, D.C., Oct. 2013, pp. 336–341. doi:10.1109/SysTol.2013.6693950

[14] Tipaldi, M., and Bruenjes, B., "Spacecraft Health Monitoring and Management," *Proceedings of the 2014 IEEE Metrology for Aerospace*, IEEE Publ., Piscataway, NJ, May 2014, pp. 68–72. doi:10.1109/metroaerospace.2014.6865896

[15] Henry, D., Simani, S., and Patton, R. J., "Fault Tolerant Flight Control: Fault Detection and Diagnosis for Aeronautic and Aerospace Missions," *Fault Tolerant Flight Control*, Springer–Verlag, Berlin, 2010, pp. 91–128. doi:10.1007/978-3-642-11690-2_3

[16] Marzat, J., Piet-Lahanier, H., Damongeot, F., and Walter, E., "Model-Based Fault Diagnosis for Aerospace Systems: A Survey," *Proceedings of the Institution of Mechanical Engineers, Part G: Journal of Aerospace Engineering*, Vol. 226, No. 10, Jan. 2012, pp. 1329–1360. doi:10.1177/0954410011421717

[17] Guiotto, A., Martelli, A., and Paccagnini, C., "SMART-FDIR: Use of Artificial Intelligence in the Implementation of a Satellite FDIR," *Proceedings of Data Systems in Aerospace (ESA SP)*, Vol. 532, No. 71, SAGE Publ., Prague, June 2003.

[18] Meitinger, C., and Schulte, A., "Human-UAV Co-Operation Based on Artificial Cognition," *Lecture Notes in Computer Science*, Vol. 5639, Springer–Verlag, Heidelberg, 2009, pp. 91–100. doi:10.1007/978-3-642-02728-4

[19] Bowman, C., and DeSieno, D., "Detecting, Classifying, and Tracking Abnormal Data in a Data Stream," U.S. Patent 8306931, B1, Nov. 2012.

[20] Benedict, B. L., "Investing in Satellite Life Extension—Fleet Planning Options for Spacecraft Owner/Operators," *AIAA SPACE 2014 Conference and Exposition*, AIAA Paper 2014-4445, Aug. 2014. doi:10.2514/6.2014-4445

[21] Chandola, V., Banerjee, A., and Kumar, V., "Anomaly Detection: A Survey," *Association for Computing Machinery*, Vol. 41, No. 3, July 2009, Paper 15. doi:10.1145/1541880.1541882

[22] Ratanamahatana, C. A., Lin, J., Gunopulos, D., Keogh, E., Vlachos, M., and Das, G., "Mining Time Series Data," *Data Mining and Knowledge Discovery Handbook: A Complete Guide for Practitioners and Researchers*, edited by O. Maimon, and L. Rokach, Springer, Medford, MA, 2005, pp. 1069–1103. doi:10.1007/0-387-25465-x_51

[23] Esling, P., and Agon, C., "Time-Series Data Mining," *Association for Computing Machinery*, Vol. 45, No. 1, Nov. 2012, Paper 12. doi:10.1145/2379776.2379788

[24] Sun, P., Chawla, S., and Arunasalam, B., "Mining for Outliers in Sequential Databases," *Proceedings of the 2006 Society for Industrial and Applied Mathematics International Conference on Data Mining*, Soc. for Industrial and Applied Mathematics (SIAM), Philadelphia, PA, April 2006, pp. 94–105. doi:10.1137/1.9781611972764.9

[25] Warrender, C., Forrest, S., and Pearlmutter, B., "Detecting Intrusions Using System Calls: Alternate Data Models," *Proceedings of the 1999 IEEE Symposium on Security and Privacy*, IEEE Publ., Piscataway, NJ, 1999, pp. 133–145. doi:10.1109/secpri.1999.766910

[26] Salvador, S., and Chan, P., "Learning States and Rules for Detecting Anomalies in Time Series," *Applied Intelligence*, Vol. 23, No. 3, Dec. 2005, pp. 241–255. doi:10.1007/s10489-005-4610-3

[27] Weigend, A. S., Mangeas, M., and Srivastava, A. N., "Nonlinear Gated Experts for Time-Series: Discovering Regimes and Avoiding Overfitting," *International Journal of Neural Systems*, Vol. 6, No. 4, Dec. 1995, pp. 373–399. doi:10.1142/S0129065795000251

[28] Baker, D. N., "How to Cope with Space Weather," *Science*, Vol. 297, No. 5586, Aug. 2002, pp. 1486–1487. doi:10.1126/science.1074956

[29] Tretkoff, E., "Space Weather and Satellite Engineering: An Interview with Michael Bodeau," *Space Weather*, Vol. 8, No. 3, March 2010, pp. 1–3. doi:10.1029/2010SW000584

[30] *GOES N Databook*, Boeing Satellite Development Center, El Segundo, CA, Nov. 2009, Sec. 5.

[31] Galvan, D., Hemenway, B., Welser, W., and Baiocchi, D., *Satellite Anomalies: Benefits of a Centralized Anomaly Database and Methods for Securely Sharing Information Among Satellite Operators*, RAND Corp., Santa Monica, CA, 2014.

[32] O'Brien, P., Mazur, J., and Guild, T., "Recommendations for Contents of Anomaly Database for Correlation with Space Weather Phenomena," Rept. TOR-2011(3903)-5, Aerospace Corp., El Segundo, CA, 2011.

[33] Allen, J., and Denig, W., "*Satellite Anomalies*," National Oceanic and Atmospheric Administration, National Geophysical Data Center, Washington, D.C., 1993, https://www.ngdc.noaa.gov/stp/satellite/anomaly/doc/5jsumm.txt [retrieved 05 June 2017].

[34] David, F. N., and Tukey, J. W., "Exploratory Data Analysis," *Biometrics*, Vol. 33, No 4, Dec. 1977, p. 768. doi:10.2307/2529486

[35] Massey, F. J., "The Kolmogorov–Smirnov Test for Goodness of Fit," *Journal of the American Statistical Association*, Vol. 46, No. 253, March 1951, pp. 68–78. doi:10.1080/01621459.1951.10500769

[36] Wang, C., Viswanathan, K., Choudur, L., Talwar, V., Satterfield, W., and Schwan, K., "Statistical Techniques for Online Anomaly Detection in Data Centers," *Proceedings of the 12th IFIP/IEEE International Symposium on Integrated Network Management and Workshops*, IEEE, Washington, D.C., Jan. 2011, p. 385. doi:10.1109/inm.2011.5990537

[37] Reeves, J., Jien, C., Wang, X., Lund, R., and Qi, L., "A Review and Comparison of Changepoint Detection Techniques for Climate Data," *Journal of Applied Meteorology and Climatology*, Vol. 46, No. 6, June 2007, pp. 900–915. doi:10.1175/JAM2493.1

[38] Wang, Y., Wu, C., Ji, Z., Wang, B., and Liang, Y., "Non-Parametric Change-Point Method for Differential Gene Expression Detection," *Public Library of Science*, Vol. 6, No. 5, May 2011, pp. 1–16. doi:10.1371/journal.pone.0020060

[39] Yamanishi, K., Takeuchi, J., Williams, G., and Milne, P., "On-Line Unsupervised Outlier Detection Using Finite Mixtures with Discounting Learning Algorithms," *Data Mining and Knowledge Discovery*, Vol. 8, No. 3, May 2004, pp. 275–300. doi:10.1023/B:DAMI.0000023676.72185.7c

Downloaded by 73.149.22.231 on May 18, 2018 | http://arc.aiaa.org | DOI: 10.2514/1.I010587

[40] Shatkay, H., and Zdonik, S. B., “Approximate Queries and Representations for Large Data Sequences,” *Proceedings of the 12th International Conference on Data Engineering*, IEEE Computer Soc. Press, Washington, D.C., Jan. 1996, p. 536.
doi:10.1109/icde.1996.492204

[41] Keogh, E., Chu, S., Hart, D., and Pazzani, M., “Segmenting Time Series: A Survey and Novel Approach,” *Data Mining in Time Series Databases*, Series in Machine Perception and Artificial Intelligence, Vol. 57, World Scientific, Singapore, June 2004, pp. 1–22.
doi:10.1142/SMPAI

[42] Keogh, E., Pazzani, M., Agrawal, R., and Stolorz, P., “An Enhanced Representation of Time Series Which Allows Fast and Accurate Classification, Clustering and Relevance Feedback,” *Proceedings of the 4th International Conference on Knowledge Discovery and Data Mining*, AAAI Press, New York, Jan. 1998, pp. 239–243.

[43] Kifer, D., Ben-David, S., and Gehrke, J., “Detecting Change in Data Streams,” *Proceedings of the 2004 Very Large Data Bases Conference*, Endowment Inc., Berlin, Jan. 2004, pp. 180–191.
doi:10.1016/b978-012088469-8/50019-x

[44] Bodeau, M., “High Energy Electron Climatology that Supports Deep Charging Risk Assessment in GEO,” *48th AIAA Aerospace Sciences Meeting*, AIAA Paper 2010-1608, Jan. 2010.
doi:10.2514/6.2010-1608

[45] Lohmeyer, W., Carlton, A., Wong, F., Bodeau, M., Kennedy, A., and Cahoy, K., “Response of Geostationary Communications Satellite Solid-State Power Amplifiers to High-Energy Electron Fluence,” *Space Weather: The International Journal of Research & Applications*, Vol. 13, No. 5, May 2015, pp. 298–315.
doi:10.1002/2014SW001147

M. J. Kochenderfer
Associate Editor

Downloaded by 73.149.22.231 on May 18, 2018 | http://arc.aiaa.org | DOI: 10.2514/1.I010587