Application of Inductive Monitoring System to Plug Load Anomaly Detection

Christopher Teubert[1], Scott Poll[2]

[1] *NASA USRP Intern, Iowa State University, Ames, IA, 50011, USA*
teubert@gmail.com

[2] *NASA Ames Research Center, Moffett Field, 94035, USA*
scott.poll@nasa.gov

ABSTRACT

NASA Ames Research Center's Sustainability Base is a new 50,000 sq. ft. LEED Platinum office building. Plug loads are expected to account for a significant portion of the overall energy consumption. This is because building design choices have resulted in greatly reduced energy demand from Heating, Ventilation, and Air Conditioning (HVAC) and lighting systems, which are major contributors to energy consumption in traditional buildings. In anticipation of the importance of plug loads in Sustainability Base, a pilot study was conducted to collect data from a variety of plug loads. A number of cases of anomalous or unhealthy behavior were observed including schedule-based rule failures, time-to-standby errors, changed loads, and inter-channel anomalies. These issues prevent effective plug load management; therefore, they are important to promptly identify and correct. The Inductive Monitoring System (IMS) data mining algorithm was chosen to identify errors. This paper details how an automated data analysis program was created, tested and implemented using IMS. This program will be applied to Sustainability Base to maintain effective plug load management system performance, identify malfunctioning equipment, and reduce building energy consumption.

1. INTRODUCTION

Over the past several years there has been tremendous interest in green technologies and sustainable practices within the building industry. As technology improvements have reduced energy consumption from Heating, Ventilation, and Air Conditioning (HVAC) and lighting systems, plug loads constitute larger percentages of a building's total load. Managing plug loads can lead to dramatically reduced building energy consumption (Lobato, Pless, Sheppy, & Torcellini, 2011; Kaneda, Jacobson, & Rumsey, 2010).

In preparation for deploying a plug load management system in Sustainability Base, which was not yet occupied at the time of this investigation, a pilot study was conducted in another office building on the NASA Ames campus (Poll & Teubert, 2012). The system monitored and controlled plug loads through the use of smart power strips, each of which had four channels (receptacles) for devices to be plugged into.

Over the course of the pilot study several issues were observed. Most serious of these were (i) failure of schedule-based plug load management rules to go into effect, (ii) failure of a device to go to low-power or standby mode, (iii) changing a device plugged into a channel, and (iv) inter-channel load relationship anomalies. These issues prevent effective plug load management; therefore, they are important to promptly identify and correct. We describe each of these issues in greater detail in the following paragraphs.

Schedule-based rule failures occur when rules to turn devices off or on at specified times, as commanded by the plug load management system, fail to go into effect. This could happen as a result of loss of communication or faulty hardware. These failures reduce the effectiveness of active plug load management, thereby increasing energy waste.

Time-to-standby failures are when a device fails to enter a low-power mode. This error, which can be symptomatic of a device malfunction, leads to greatly increased energy consumption.

Changed loads refers to a configuration change of the devices plugged into a power strip. Usually this means that a device has either been replaced with a newer model or that a different device has been plugged into that channel. A configuration change such as this is only an issue if the system administrators are not notified of the change. For example, changing loads without updating the associated schedule-based rules could lead to data loss or damage if a computer is inadvertently de-energized.

Inter-channel anomalies refer to a situation where the relationship between two channels is undesirable. One example of such an anomaly would be if a monitor is in active mode while the computer is off. The error could be symptomatic of malfunctioning equipment. In the future, this could also indicate failure of load-sensing control, which turns off devices based on the behavior of a 'master' device. For example, a rule could be created so that when a computer is off, the peripherals (speakers, printer, monitor) would be powered down as well. Load-sensing control was not investigated for the pilot study but it will be in Sustainability Base.

Developing a model-based system to identify the aforementioned anomalies for each channel would be labor intensive and would not scale to a plug load management system for an entire building with hundreds or thousands of loads. Therefore, it was decided to use a data-driven approach to do automated analysis. Data-driven algorithms (Kantardzic, 2011) are capable of analyzing vast amounts of data to pick out unusual or unhealthy behavior and therefore lend themselves nicely to building plug load management at NASA Ames' Sustainability Base.

The Inductive Monitoring System (IMS) tool (Iverson, 2004) was chosen for this application because of its ability to learn healthy behavior without having to create a complex model for each channel. IMS creates a knowledge base of nominal behavior from judiciously chosen training data sets. New plug-load data are then compared to healthy behavior to pick out anomalies. If not addressed, these anomalies could lead to increased power consumption, decreased effectiveness of the plug load management system, or even damage to plug load devices. Once an anomaly is identified, building personnel are automatically notified so that they may address the issue.

The Sustainability Base IMS application uses device power draw data collected by plug load monitoring power strips located in copy rooms, break rooms and at workstations. The power strips measure and transmit power draw once per second to a cloud-based data service which records minimum, mean, and maximum power draw at one minute intervals. The volume of the data (1440 records per device each day) makes it necessary to implement automated analysis.

The main contributions of this research are (i) observation of potential issues with plug load management, (ii) definition of raw and derived plug load parameters that identify different anomalies, (iii) application of Inductive Monitoring System to identify faulty plug load devices or improper usage, and (iv) development of an automated program to process plug load data and notify appropriate personnel of problems that require attention.

2. Inductive Monitoring System

Inductive Monitoring System is a data mining algorithm designed to detect deviation from healthy system behavior. The first step in using IMS is off-line learning, or the establishment of a knowledge-base of healthy behavior. To do this a series of vectors of data previously determined to be healthy, or training data, are fed in one-by-one to the program. K-means clustering (Bradley & Fayyad, 1998) is used to group data into multi-dimensional clusters; different regions of the cluster space may represent different operating modes of the system. If the vector is determined to be close to one of the existing clusters, the cluster is expanded to include it. If the vector is too far from the clusters it becomes the beginning of a new cluster. Parameters are used to control how the clusters are expanded or created; the default IMS parameters were used in this study.

Once the healthy clusters are fully formed, new data sets are then analyzed. Each vector of the testing data is compared with the formed clusters, and the closest cluster is determined for comparison. The composite score is defined as the Euclidean distance between the vector and the closest point on the nearest cluster in multidimensional space. IMS also calculates the contribution, or local score, of each individual parameter to the composite score.

The IMS tool has been used in a number of complex systems. Following the Columbia (STS-107) accident in 2003, IMS was used to analyze the telemetry from four temperature sensors located in each wing of the orbiter. IMS analyzed data from launch/ascent and on-orbit and was able to detect anomalies much earlier than the monitoring systems used in mission control (Iverson, 2004).

An IMS based program has been used by the International Space Station (ISS) flight control team in mission control to monitor operations the Control Moment Gyroscopes (CMGs) and External Thermal Control System (ETCS). This program has successfully identified multiple anomalies in these systems (Iverson, Spirkovska, & Schwabacher, 2010).

2.1. Sustainability Base Plug Load IMS Application

Sustainability Base IMS will be used as an important tool for the building's health management. Output from IMS will allow operators to identify and address unhealthy plug load behavior promptly, thereby increasing the effectiveness of the plug load management system and maintaining high system efficiency. The most critical element in ensuring useful results from the IMS algorithm is the definition of input vectors. Sample training vectors are shown in Table 1. Each row is an input vector whose parameters are defined by the column headers. Combining raw and derived quantities is essential to permit visibility of different failure types. For Sustainability Base three parameters were chosen:

- Raw power draw: This is employed to discover cases of changed power loads. Changed power draw could also be symptomatic of a larger problem, or an example of normal behavior that has not been previously observed.
- Consecutive minutes in the idle mode power range, as applicable: This feature is meant to find cases where a device fails to go to standby or low-power mode. The ranges for idle modes were defined a priori.
- Piecewise function corresponding to the time when the device is drawing power, and zero, otherwise: These values are used by IMS to find cases of schedule-based rule failures, or times when the device is drawing power when it should not.

Power Draw (W)	Idle Time (hrs)	On Time (hrs)
100	0.54	10.33
102	0.56	10.35
160	0.00	10.36
102	0.02	10.38
...	...	...

Table 1. Sample Training Vectors

The power strips transmit power draw measurements every second to a cloud-based data service, which records data at one minute intervals. During mode transitions a device will sometimes spend parts of a minute in different modes, causing the system to record an average power draw value in a range where the device does not typically operate. If included in the training data, these transitory values can prevent the IMS system from detecting anomalies when a changed device draws steady-state power in these intermediate ranges. Recall from the input vector that duration is only captured for the idle mode in the second parameter, the first parameter of the input vector is only raw power draw.

In order to eliminate this phenomenon the raw power draw is filtered to remove intermediate values. Values where the relative difference is greater than 10% are filtered out, where the relative difference is defined as the change in power draw divided by the average, as shown in Eq. (1). The raw power draw of an example load before and after applying the filter is shown in Figures 1a and 1b, respectively.

$$RD = \frac{P_{i+1} - P_i}{(P_{i+1} + P_i)/2} \tag{1}$$

The formed vectors are used as input to the IMS algorithm and the resulting three local scores are used for post processing. Each local score corresponds to the distance, expressed as a percentage, from an input vector parameter to the closest cluster of the healthy training data. A simple filter is applied to remove occasional misleading spikes in the local score that do not correspond to legitimate errors.

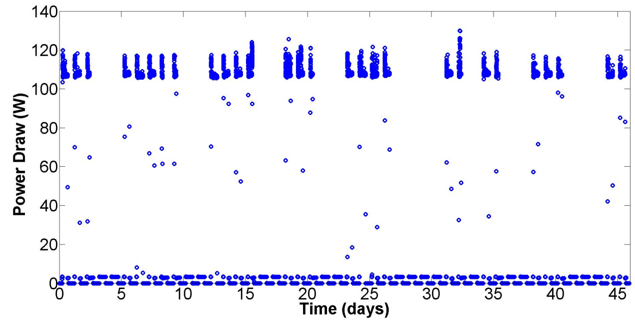

(a) Before Filter

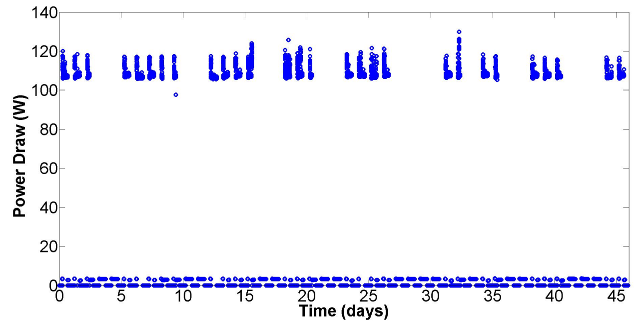

(b) After Filter

Figure 1. Results of Power Draw Filter

An error is indicated when there are more than 5 consecutive minutes in which the resulting changed power load local score (corresponding to first parameter of input vector) is above 3%, time-to-standby local score (second parameter) above 10%, or rules local score (third parameter) is above 5%. Thresholds for individual IMS local score parameters were obtained by observing the noise fluctuations in the three parameters. Adjustments were made manually until the thresholds were at a level where the IMS program reliably filtered out noise while still detecting anomalies.

These thresholds will likely have to be adjusted for the Sustainability Base deployment. Monitoring the plug loads as described above allows the appropriate personnel to be notified of errors so that they may be corrected, thereby preventing power waste, optimizing plug load management system performance, and possibly extending the life of the devices in question.

In the case where the day's local scores are all below 1% (i.e., the system never deviated more than 1% from nominal behavior) the day's behavior is considered healthy and is added to the training vector for processing the next day of data. This allows the program to better define healthy behavior as time goes on, and prevents the system from picking up deviations in power draw that come from gradual normal system behavior changes.

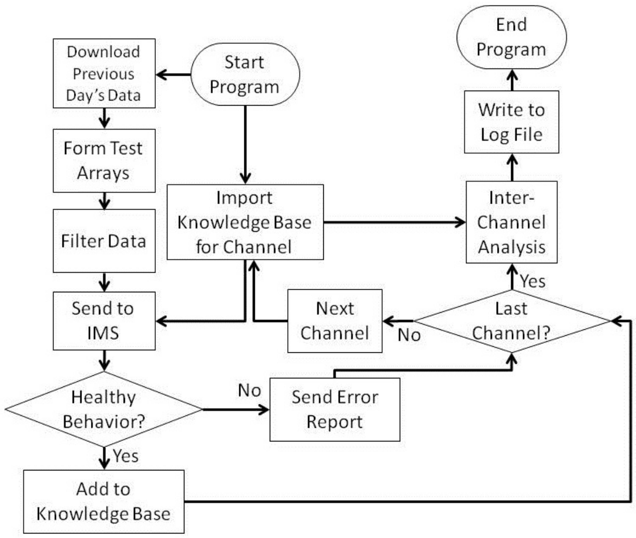

Figure 2. Sustainability Base IMS Program Flow

A separate IMS build may be used to monitor other Sustainability Base subsystems. For example, sensor measurements from the lighting, ground source heat pump, underfloor air distribution, and photovoltaic systems can be used to create additional knowledge bases. The IMS capability will be expanded to incorporate additional systems as they are introduced to the building and as the needs of Sustainability Base change.

2.2. Automated Sustainability Base Plug Load Monitoring System

An automated IMS system analyzing plug load system health (see Figure 2) will be implemented by running the Perl program every morning on a server to analyze the past day's data. It will then find unhealthy behavior and notify building personnel by email as necessary. The program is modular, allowing for future additions as new needs arise.

The first task that the program executes is downloading the previous day's data from the cloud server using the system API. Each channel's plug load information is imported and parsed into the test vectors.

The stored knowledge base of nominal behavior is read from files for each channel. The test vectors and nominal clusters are fed into IMS. An additional test is done to find cases of missing data using the timestamps included in the raw data.

If all local scores are below a certain threshold the vectors are appended to the training vectors for use in future days' analyses. This allows the system to learn so that it may better characterize nominal behavior, thereby both reducing false positives and more accurately recognizing errors. Additionally, allowing the individual to mark the identified false positives,

```
-----------------------------------------
Plug Load Monitoring
Weekly Report
-----------------------------------------

November 8, 2011

ERRORS
High Priority:
- Nodes 1-5: Continued communication error
11/6/11 - 11/13/11
- Workstation Rm 288/Ch3.0, Desktop Computer:
Sustained unusual behavior from 11:36
11/8/11 - 15:22 11/8/11

Low Priority:
- Copy Rm 287/Ch5.0, Shared Copier:
Failure to reach standby mode from
06:00 11/7/11 - 22:00 11/7/11

STATISTICS
Total Energy Use:  3243 kWh
Last Week's Energy Use:  3254 kWh
Energy Difference:      11 kWh

Report generated automatically at 02:33
on November 8, 2011
-----------------------------------------
```

Figure 3. Sample Weekly Report

and having the system then add the marked data to the training vectors could lead to a greater reduction in the frequency of false positives.

The resulting local score vectors are processed using the methods described in Section 2.1. Errors are then separated into three categories based on priority of notification. For high priority errors, such as prolonged communication errors or drastically changed loads, a notification is emailed immediately to the system administrators so that the issue may be resolved. Medium priority errors, such as rule failures, are saved as part of a weekly report emailed to the system administrators. Low priority errors, such as short time-to-standby delay or short-lived communication errors, are saved in a log file located on the server.

A sample weekly report generated using pilot study data is shown in Figure 3. Such a report tells the contact what errors are occurring in the system, where they are occurring, and when they occurred. Errors are sorted by seriousness of the anomaly. The report also includes some statistics on power use and difference from the previous week so that the user may better understand their environmental impact.

Note that we deliberately chose a weekly report that showed several alert types (communication failure, failure to reach standby mode, etc.). Consequently, the energy difference in Figure 3 does not represent the energy savings from a week that employed rules versus a week that did not. In fact, in this report the week of November 8 and the previous week had no rules in place, hence the similar energy consumption. Nor does it represent the difference in energy consumption between a week in which the identified anomalies were remediated and a week where the anomalies were not remediated. Assessment of remediation of identified anomalies was beyond the scope of this study. As the system is deployed in Sustainability Base, each identified anomaly and potential remediation will be individually assessed so as to not adversely affect system operations in the event that the identified anomaly is insignificant.

This program was found not to require a large amount of time or processing ability. However, significant data storage is necessary to cache plug load data, training vectors, and program logs. Each channel requires on average 126 kB of storage each day for the plug load data at one minute resolution. System cost and impact can be reduced by employing a multi-purpose server, running additional programs for other functions, rather than a dedicated server.

3. Results

Findings of applying IMS to pilot study plug loads are discussed in this section. During the pilot study examples of five types of errors were found: changed loads, time-to-standby errors, rule failures, inter-channel anomalies, and missing data. For missing data it was not necessary to use IMS so it was analyzed using a simple anomaly detection scheme.

These errors are discovered without defining a priori what types of errors to expect or any model attributes such as rule times, time to transition to standby mode, or normal power draw. The only information that is needed is a sample of healthy data and a range of idle mode power draw for each channel, as applicable. The IMS program is then able to find unusual behavior by comparing new data to the knowledge base learned from the sample data. IMS is capable of discovering errors that have never been seen before in the system. Additional investigation of such cases can determine whether the new behavior is detrimental or insignificant to system operation.

Examples of schedule-based rule failures, time-to-standby errors, changed loads and inter-channel anomalies are provided in the following subsections. In each case the raw data are plotted together with an indication of the points that have been flagged by IMS as being abnormal. The dominant IMS local score output is also presented in a corresponding plot as percentage relative error, which is a normalized distance from the relevant test vector parameter to the nearest cluster in the knowledge base.

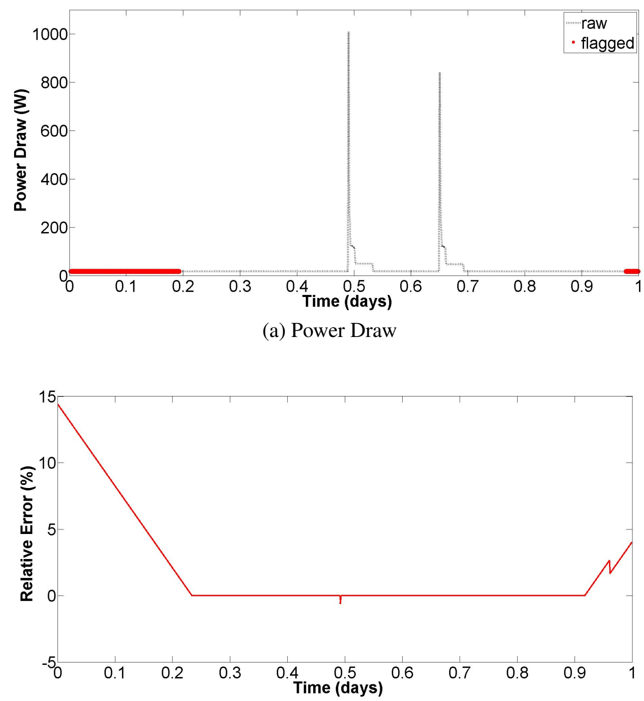

(a) Power Draw

(b) On Time (Third Vector Parameter) Local Score

Figure 4. Case of Rule Failure

3.1. Schedule-Based Rule Failures

The first error revealed during the plug load management trial was the occasional failure for schedule-based rules to go into effect. This is likely because of loss of communication or a malfunctioning channel.

An example of this was observed with a printer. Schedule-based rules were set up to turn off the printer between the hours of 10PM and 6AM to conserve power, but failed to go into effect because of communication issues. Figure 4a shows the power draw of the printer. The black dotted line is the raw power draw for the printer and the red points have been flagged by IMS for schedule-based rule failures; note that because of the threshold applied to the IMS local score output, not all points from 10PM (0.92) to 6AM (0.25) are highlighted. The local score over the same time period in Figure 4b increases the longer it has been since the device was supposed to be powered down.

The Sustainability Base IMS found when rule failures occurred for all cases tested. Fixing this type of error would ensure that the plug load management system eliminates power consumption during non-business hours, when the device should be turned off, and maintain the system effectiveness.

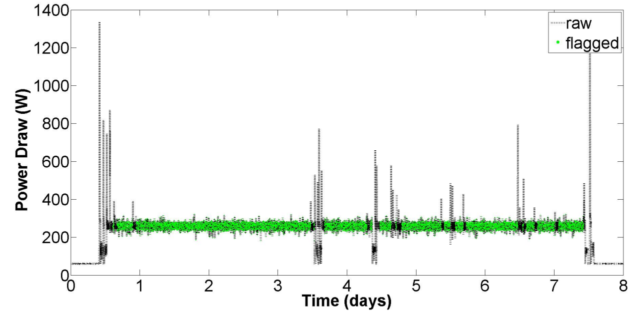

(a) Power Draw

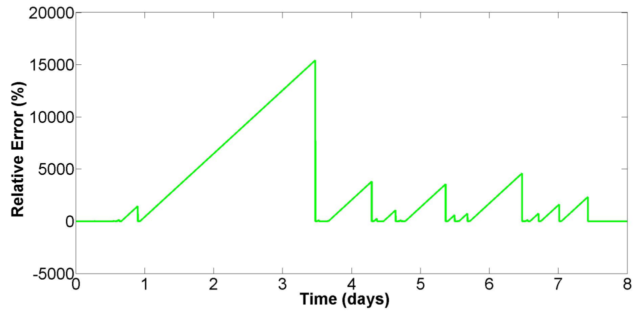

(b) Idle Time (Second Vector Parameter) Local Score

Figure 5. Case of Time-to-Standby Error

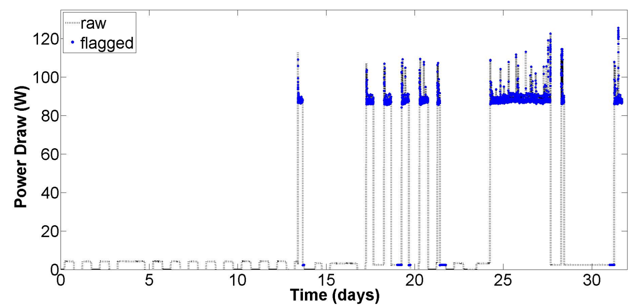

(a) Power Draw

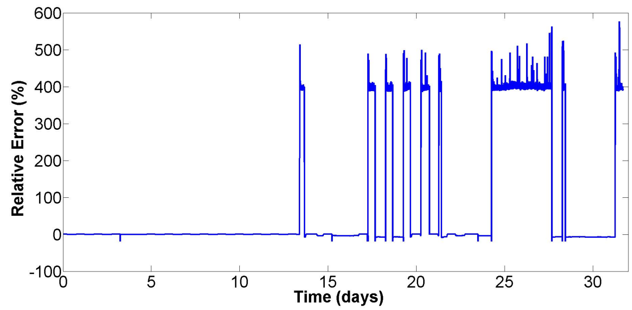

(b) Power Draw (First Vector Parameter) Local Score

Figure 6. Case of Changed Load

3.2. Time-to-Standby Errors

Time-to-standby errors were observed in a couple of cases. They can occur as a result of device malfunctions and often require equipment maintenance or possibly replacement.

The power draw for a malfunctioning copier from the pilot study can be seen in Figure 5a. The black dotted line is the raw power draw for the copier and the green points correspond to instances where IMS has found time-to-standby errors. Note that the device fails to enter low-power mode for several days. These cases correspond to an increased local score during the same time periods as seen in Figure 5b.

IMS was able to reliably pick out when time-to-standby errors occurred in all cases tested. Using built-in low power functionality was found in the pilot study to be one of the most effective methods of reducing power draw. Therefore, it is important to identify cases of this anomaly and correct them.

3.3. Changed Loads

During the pilot study there were occasions when occupants changed the devices connected to the power strip. They did this in order to replace or change the location of their devices. This can result in data loss or damage if a computer with particular shutdown procedures is plugged into a controlled outlet instead of an uncontrolled outlet.

IMS revealed several cases where a device was swapped with another device. The power draw for one such channel can be seen in Figure 6a. This channel originally had a set of speakers (days 1-13), but they were replaced with a computer (days 13-31). The data points that IMS has picked to be the changed load have been marked with blue points. Figure 6b is the local score of that channel during the time period; the output prior to day 13 is near zero, meaning that IMS had seen similar data before.

IMS was able to pick out that the active mode of the computer was a changed load, but the phantom load of the computer was too close to the active mode of the speakers and therefore was not detected. Similarly, the only case that IMS was unable to detect was when a computer was replaced with speakers. The active power draw of the speakers matched the phantom load of the computer previously on that channel. This points to the need of active configuration management. It is likely that additional derived quantities could be used to find the changed load in this case. This is discussed further in the Conclusion Section.

3.4. Inter-Channel Anomalies

We also did a preliminary investigation of detecting inter-channel anomalies. These anomalies could indicate that a device is malfunctioning or that load-sensing control rules have failed to go into effect.

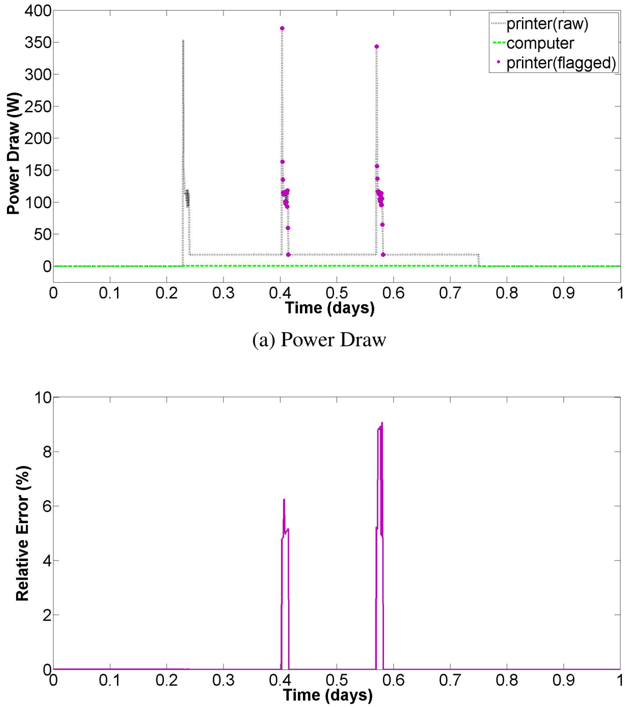

(a) Power Draw

(b) Global Score

Figure 7. Case of Inter-Channel Anomaly

An example of this can be seen in Figure 7a, which shows the power draw of a printer. In this example the printer (the black dotted line) has three large spikes in its power draw. The first spike occurs when the device is re-energized, while the second and third are from print jobs. In this case the printer was receiving print jobs while the associated computer (the green dashed line) was off. The IMS algorithm saw this as unusual and flagged the points noted in magenta.

Figure 7b shows the IMS global score during that same time period. Note that IMS was able to distinguish between the typical start-up load and the anomalous print jobs. The time stamp in the third parameter of the IMS vector allowed IMS to distinguish between the start-up power spike, which occurred at the same time every day, and potentially anomalous behavior.

Note that the data for these channels considered individually are normal, it is the correlation between the channels that has changed relative to the knowledge base. Detecting these inter-channel errors can uncover abnormal device usage or behavior. For the example shown here the behavior actually reflects the fact that this was a shared printer which received a print job from another computer, but it was presented to IMS in such a way as to make it a test case for inter-channel anomalies.

4. Conclusion

These results from the pilot study have proven the effectiveness of IMS for plug load health monitoring. Sustainability Base IMS successfully detected schedule-based rule, time-to-standby, changed load and inter-channel errors in the system. Such a system is expected to be an effective aid in preventing energy waste, improving plug load management system effectiveness, and avoiding system damage. The IMS is currently being deployed to Sustainability Base. Some fine tuning is expected to strike the right balance between flagging irrelevant issues and missing relevant ones.

The Sustainability Base IMS system will provide support to facility managers and occupants to identify usage anomalies. For this study it was not directly tied into the plug load management system, which employed only schedule-based rules to change the on/off state of the channels. The commercial plug load management system was not able to identify the anomalies noted by the IMS system and so was unaware of operational faults that negatively impact energy usage and system usability.

The results from the pilot study (Poll & Teubert, 2012) show that proper setup of device power management settings lead to significant energy savings. The IMS can be employed to find cases of incorrect setup or malfunctioning equipment. Additional analysis would be required to accurately estimate the magnitude of this savings.

Additional research is planned to extend the Sustainability Base IMS to detect other types of errors. The programs were created in a modular fashion to allow for such expansions. These could include additional derived quantities in order to better pick out device changes where the power draw closely matches that of the previous device, or to detect new errors as the needs of the system change.

Another future addition to the Sustainability Base IMS program is the ability to do real time analysis. This will require a fast and reliable method of accessing plug load data. Enabling such a system will increase its effectiveness by notifying system administrators and occupants sooner of potential problems and will allow for some real-time system health statistics to be displayed for Sustainability Base.

Acknowledgment

We thank the personnel of the Intelligent Systems Division who participated in the plug load pilot study. We extend our gratitude to Bryan Matthews for his help with the IMS algorithm. Finally, we thank members of the Diagnostics and Prognostics Group in the Intelligent Systems Division for helpful discussions and feedback.

REFERENCES

Bradley, P. S., & Fayyad, U. M. (1998). Refining initial points for k-means clustering. In *Proceedings of the fifteenth international conference on machine learning* (pp. 91–99).

Iverson, D. L. (2004). Inductive system health monitoring. In *Proceedings of the 2004 international conference on artificial intelligence.* CSREA Press.

Iverson, D. L., Spirkovska, L., & Schwabacher, M. (2010). General purpose data-driven online system health monitoring with applications to space operations. In *Proceedings of the fifty-third annual isa powid symposium.* Research Triangle Park, NC: International Society of Automation.

Kaneda, D., Jacobson, B., & Rumsey, P. (2010). Plug load reduction: The next big hurdle for net zero energy building design. In *Proceedings of 2010 aceee summer study on energy efficiency in buildings.* Washington, D.C.: American Council for an Energy Efficient Economy.

Kantardzic, M. (2011). *Data mining: Concepts, models, methods, and algorithms*. Hoboken, New Jersey: John Wiley and Sons, Inc.

Lobato, C., Pless, S., Sheppy, M., & Torcellini, P. (2011). *Reducing plug and process loads for a large scale, low energy office building: Nrel's research support facility* (Tech. Rep. No. NREL/CP-5500-49002). National Renewable Energy Laboratory.

Poll, S., & Teubert, C. (2012). Pilot study of a plug load management system: Preparing for sustainability base. In *Proceedings of 2012 ieee green technologies conference.* Institute of Electrical and Electronic Engineers, Inc.