A Test of Commercially Available Products for Estimating End Uses from Smart Meter Data

M. Baker,G Hicks, S. Rodriguez
SBW Consulting
Bellevue, Washington
mbaker@sbwConsulting.com

M. Fuller
Evaluation Measurement & Verification
Pacific Gas and Electric Company
San Francisco, California

***Abstract*—A test was conducted to determine whether one-minute interval measurements of a home's total electrical use could be disaggregated into standardized end uses, without having to provide information about each home, e.g., demographics, appliance and equipment inventory, or building characteristics. One-minute total energy use measurements were obtained for 160 homes. In addition, direct end use measurement systems were installed in six of these homes.**

Three vendors of commercially available disaggregation products were hired to estimate end uses for all 160 homes. The identity of the directly measured homes was not disclosed to these vendors. The vendors were asked to disaggregate the one-minute total use data into 14 end uses for a period of five months starting from December 2014. During this period, disaggregation estimates from all vendors were found to have large estimation errors for almost all of the end uses, sometimes failing to even identify an end use while other times significantly overestimating energy consumption.

Estimates from one vendor matched or came close to matching the end use measurements for some important end uses such as air conditioning, cooking, lighting, and refrigeration, but had large errors for other end uses. While vendor estimates were most accurate for end uses aggregated over the five-month analysis period, accuracy markedly degraded for individual months or individual homes.

Keywords—residential; end use loads; accuracy testing;

I. INTRODUCTION

Many utilities now have the capability of collecting from all of their customers, via so-called smart meters, measurements of electrical use for intervals as short as one-minute. Pacific Gas and Electric (PG&E) has deployed such meters. The data available from these meters has many possible applications. One of which is to provide customers with information that may help them operate their homes more efficiently. In particular, PG&E wanted to determine whether the one-minute interval measurements of total electrical energy supplied to a home could be accurately disaggregated into end uses, without using any other information about each customer, e.g., demographics, appliance and equipment inventory, or building characteristics. If this proved possible, PG&E would be able to routinely provide their customers with information about how much they were spending on end uses such as refrigeration or space cooling. PG&E sponsored the test described in this paper to determine whether any of the companies that offer smart meter data analysis services (referred to in the paper as vendors) could provide accurate end use estimates given only smart metering data collected from PG&E residential customers.

One-minute total energy use measurements were obtained for 160 homes. Out of the 160 participating homes, disaggregated end use measurements were also directly obtained for six test homes (also referred to in this paper as test sites). Approximately 60 devices or circuits were monitored in each of the test sites. Three vendors were selected to test their ability to accurately estimate end uses for a subset of participating homes. The meters serving each of the homes participating in this study were modified so that they could continue to record normally for billing purposes, but could also provide measurements of electric energy use for each minute. Homes participating in the test were assigned an arbitrary identifier that masked their identity. All of the test sites were included in the homes provided to each vendor but they were not identified. The vendors were not provided any information about the characteristics of any of the PG&E customer's homes. Data was provided on 87 homes to two of the vendors and on 85 homes to the third. There was some intentional overlap in the homes assigned to each vendor. These data were used by the vendors in estimating 14 standardized electrical end uses for each of their assigned homes. Vendor end use estimates were compared at the hourly level to direct measurements for the test sites.

II. CHARACTERISTICS OF TEST SITES

PG&E requested volunteers for this test from its employees and contractors. Owners of 37 homes expressed interest in the test. The home had to be their primary residence and be served by a PG&E electric meter. If the home had gas end uses, they also had to be served by a PG&E meter. Each homeowner provided information about the home's location and the characteristics of the structure, occupants and energy using equipment. All but nine of these homes were eliminated for the following reasons: Solar electric installations on the house, combined space and water heating systems, plans to remodel, and antiquated wiring.

Device/circuit inventories were completed for nine homes. This included a complete inventory of all electric and gas powered devices served by the PG&E meters. Electric devices were included in the inventory if they were permanently connected to a circuit, i.e., hard-wired, or the homeowner

reported that the device was always plugged into the same outlet. We ignored devices that were occasionally plugged into an outlet or were often plugged into different outlets. All electric circuits in each home were traced so that each device could be associated with a specific circuit. These inventories were analyzed and PG&E selected six of the nine inventoried homes for the test based on many technical and cost factors.

Two of the test homes were near San Francisco and the other four were in our near Stockton, California. Stockton is hotter in the summer season. It has approximately five times as many cooling degree days as San Francisco. The homes were of similar size, with floor area ranging from 2,100 to 3,000 square feet. Five out of six of the homes were more than 15 years old. One was built in 2007. Five out of six homes had 2 adult occupants. Of these, only two had occupants that were children (one each). One home had five adult occupants and no children. All homes had gas space heating equipment, either central or wall units. Four of the homes had central air conditioning. Two of the homes only had electric cooking equipment. The other four had a mix of gas stove tops and electric ovens. None of the homes had a gas oven. Four homes had one refrigerator. One of them had two refrigerators and one home had four. All homes heated water with a single gas fired water heater. All homes had a single electric clothes washer and a single electric clothes dryer, except one home which had a gas heated clothes dryer. Two of the homes had a pump for a swimming pool or spa. Three homes had a charger for an electric vehicle. Three homes had one television and three had two televisions.

III. Definition of Electric End Uses

During the design of this study, we searched for but were unable to find any national or international standards for classifying energy using equipment by end use. However, clear definitions were needed to guide the vendor estimates of end use and our measurement of end uses in the test homes. 14 end uses were defined: Space Heater, Air Conditioner, Domestic Water Heater, Pool/Spa Heater, Lighting, Refrigerator/Freezer, Cooker, Clothes Dryer, Clothes Washer, Dish Washer, Electric Vehicle, Spa/Pool Pump, Other Pump, to determine the least cost and most reliable plan for installing power measurement equipment and the associated equipment for wireless communication within the house and with our remote data collection systems. The plan included intentional redundancy. For example, we measured (where feasible) the main feed to each electric breaker panel and we measured each of the circuits controlled by those breakers. This allowed for a comparison of the panel total use to the sum of the use on each of the breakers. Conventional true power measurement devices were used that require installation of split-core current transformers (CTs) around one leg of single phase circuits or both legs of two phase circuits (such as those serving clothes dryers). Corresponding potential transformers (PTs) were also connected to obtain the voltage(s) for each circuit.

The circuit and device inventory allowed us to identify the end use of each electric circuit. In many cases a circuit served more than one end use. A plan was developed for disaggregating each of these circuits. In some cases, we were able to install wireless plug load power monitors on devices or plug strips serving groups of devices that served to disaggregate the power use into the target end uses. In other cases, most often involving circuits that serve lighting fixtures, we used the plug load power monitors to measure all of the non-lighting devices on the circuit and then defined a virtual data collection channel for lighting to be the difference between power measured at the breaker and the sum of non-lighting plug load measurements. This "virtual" channel technique was also used to derive other non-lighting end uses for certain circuits.

It was not possible to completely disaggregate some circuits. Any circuit that served more than two "hard-wired" end uses was classified based on the end use which had the largest rated power draw. These could have been separately measured but it would have required re-wiring a portion of the home, which was not feasible for this test. In some cases, more than one end use was present in a single device, e.g., an outdoor fountain or a bathroom exhaust fan that had integrated lighting. These devices would have to be disassembled in order to separately measure the end uses, which was also beyond the scope of this test.

Although it is not possible to quantify the degree of misclassification in the measured end uses, we believe that such misclassification introduces only a small error when we compare the measured end uses to the vendor's estimates.

IV. Collection of End Use Measurements

One-minute interval measurements of electric and gas end uses were obtained from the test homes for a full year, including the period used for comparison with the vendor estimates of electric end uses. Approximately 600 separate measurements, either gas or electric use, were obtained. All measurement points within each home were wirelessly connected to a multi-protocol (ZWave and ZigBee) controller and Other.

V. Installation of Power Meters

We analyzed the circuit and device inventory along with the physical layout of electric panels and devices in each home which communicated with our primary remote storage via a cellular connection. The gateway device polled each power measurement once a minute and recorded the energy used since the last polling. These recordings were pushed to our remote database once each minute.

All collected data was tested to determine whether it fell within expected ranges and for other indications of error conditions such as repeated identical values. Check sum comparisons, e.g., total power feeding a panel compared to the sum of the power to each of its breakers, were particularly useful in diagnosing problems with the measurement system. We also routinely looked for measurement points that failed to function.

VI. Vendor Estimates of Electric Energy End Uses

PG&E sought firms who had market-ready products which were capable of disaggregating total one-minute electric use into separate end uses for residential customers. PG&E

selected three vendors for this test and compensated them for their services. Each vendor utilized proprietary algorithms to develop estimates for each of the end uses defined by PG&E. Two vendors were asked to prepare these estimates for 87 homes; the third was asked to estimate end uses for 85 homes. The test sites were included among the homes sent to each vendor, but the vendors were not told which homes had end use measurement systems. The vendors delivered one-minute estimates for each home's end uses. These were delivered a few days after the end of each month during the test period.

VII. Accuracy of Vendor End Use Estimates

The accuracy of the vendor estimates of electric end uses was tested using the direct measurements of these end uses in the six test homes. Both the vendor estimates of end use and our measurements of end use were aggregated to the hourly level prior to making any comparison. The hourly aggregation was required because it was not possible to synchronize the exact boundaries of each minute between the total use data supplied to the vendors and the measurements that we obtained for each end use. Further, the precision of our end use measurements, especially for small loads could result in one or more minutes of zero use for some end uses followed by a value for a minute that was the sum of use for a number of preceding minutes. This would not lead to a meaningful comparison with the vendor estimates. Finally, PG&E did not have any interest in the accuracy of the vendor estimates at the one-minute level and felt that a test at the hourly level would be sufficient for any possible application of the vendor products.

To ensure the vendor comparisons were based on the quality of the disaggregation and not the quality of the data, we dropped from the comparison any hour where the sum of the one-minute use provided to the vendors was more than 10% different from the sum of our measured use. There are at least a couple reasons for discrepancies. The first reason is that the one-minute readings from the modified PG&E meters were not complete. There were also instances where the PG&E meter use was zero for more than an hour, which seems unlikely. In addition, there were errors in our measurements, including omissions and invalid data. At times some of the equipment went offline and data during those periods was lost. Some pieces of equipment sometimes reported erroneous values (which were zeroed out in the dataset).

For the test period, the total use provided to the vendors was within 10% of our measured use for 65% of the hours. These hours were used in the comparison of measured use to vendor estimates which is shown in Fig. 1.

PG&E may use the data obtained in this test in evaluating the accuracy of other vendor products. Therefore, the kWh scale for the panels in Fig. 1 and Fig. 2 have not been included. The panels in each figure use the same scale and thus the height of each bar can be compared across panels. The number that appears above each bar is the percent of the measured end use.

The Green vendor came closest to estimating most of the significant end uses. However, that vendor was 630% high in its estimate of the Space Heater end use. None of the homes used electricity for their primary heat source, so the Space Heater end use was associated with the electric fans in the heating equipment that distribute the heat throughout the home.

Although, the test period was in the winter and early spring, there was some cooling and the Green vendor estimated 90% of the measured Air Conditioner use in this period, substantially more accurate than the other two vendors. In addition, the Green vendor estimated 100% of the Cooker and Spa / Pool Pump end uses. The other two vendors estimated between 50 and 70% of some end uses, but in general were further from the measured use, except for Clothes Dryers where the other two vendors were somewhat more accurate. All vendors were about as accurate for the Electric Vehicle end use, but all missed the measured use by 50% or more.

Fig. 2 is a summary of reported pool pump energy use by site. Two of the 6 sites had pool pumps. These charts demonstrate that accurate reporting in aggregate did not mean that reporting was accurate on a site by site basis. The Green vendor reported the correct total, and the Blue vendor was within 30% of the total. Both Green and Blue vendors reported a false positive at Site 6 which had a fountain pump. Fountain pumps may appear similar for short periods, but fountain pumps operate continuously and not on a schedule like pool pumps. At Sites 2 and 4 the reported energy use was 30% to 60% less than the measured total.

VIII. Conclusion

We find some evidence that commercially-available disaggregation products can identify some end uses accurately from one-minute total home use. The Blue Vendor was able to identify 12 of 13 specified end uses while Red and Green were able to Identify 8 of 13 (Fig. 1). Vendor accuracy varied for different end uses when averaged across 6 sites. False positives of energy use are another concern. The Blue and Green vendors both identified a pool pump at site 6 where there was none (Fig. 2). Being able to identify many end uses is ideal, but the results must be accurate on a site by site basis. Further development of disaggregation algorithms is needed before they are sufficiently accurate to provide customers with accurate estimate of how much they spend on most end uses.

This data set will be used by PG&E as a litmus test for disaggregation vendors that are being considered for customer add on services. The data set will not be completely released so that the test is still a useful evaluation metric.

Fig. 1. All end uses - percent of measured for three vendor estimates (all homes - December 2014 thru April 2015).

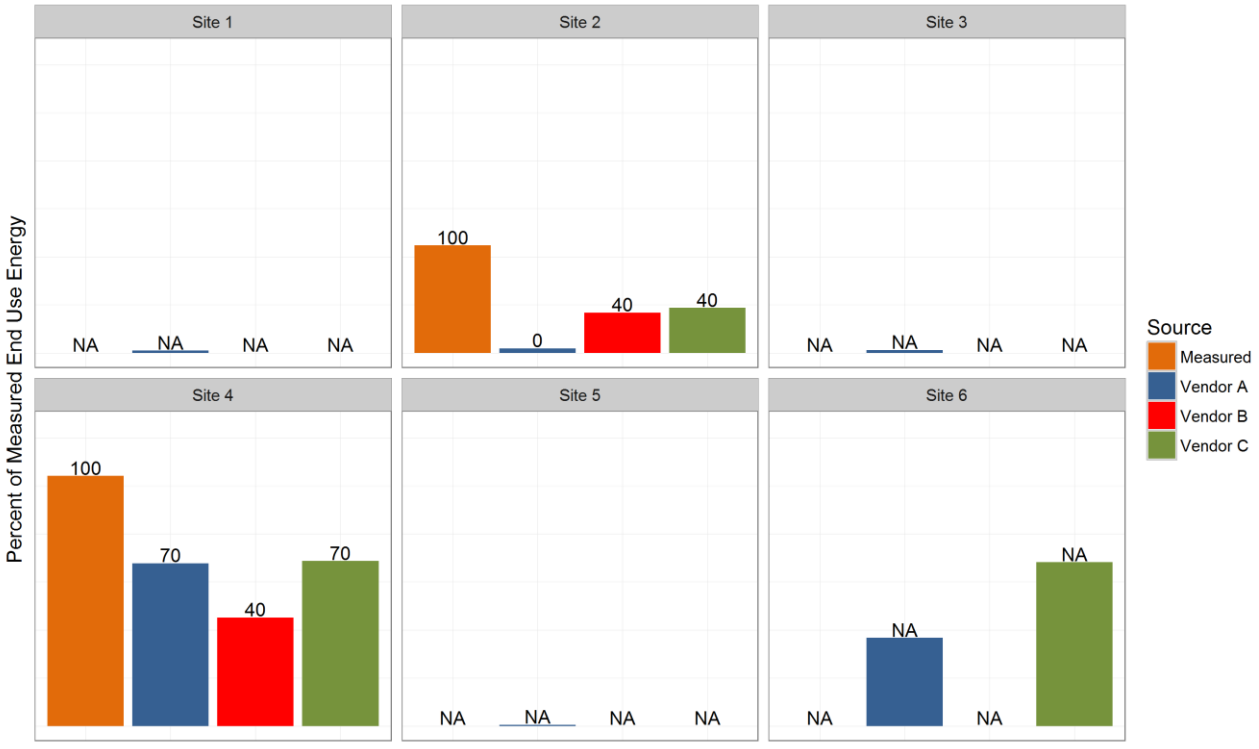

Fig. 2. Spa / Pool Pump - percent of measured by site for three vendor estimates (December 2014 thru April 2015).