# A Commercial Disaggregation System for Residential and Light Commercial Buildings

Hal Alles, Bradley A. Zenger
Verlitics LLC
Portland, OR
hal_alles@msn.com, baz@pivotal-investments.com

*Abstract*— **The cost effective disaggregation of electrical loads is a technology offering great promise, but viable, commercial products have been elusive. Verlitics is demonstrating a disaggregation technology that offers the fastest way to a commercially viable product, while providing a robust, extensible technology platform. The technology has demonstrated the ability to deliver multiple value streams, including energy, equipment and occupancy monitoring resulting in compelling financial payback opportunities. Verlitics is currently demonstrating commercially viable technology in both 50/60Hz and split-phase and 2 or 3-phase applications. Verlitics believes its technology will ultimately identify 90 percent of targeted devices with accuracy of 90 percent in 90 percent of buildings in the target market.**

**Keywords— NILM, disaggregation, continuous commissioning**

## I. INTRODUCTION

Verlitics as a private and venture funded start up company had the goal of developing a commercially viable disaggregation system for the residential and light commercial markets. Verlitics and partner company market studies indicated the full cost of the product and service to the customer needed to provide a return on investment of less than 18 months. Verlitics determined that a specialized power monitor was needed and not currently available, and that the installed price required is less than $500. In addition, a disaggregation service that provided energy savings information needed to be priced at no more than about $100 per year for residential applications, but could be higher for light commercial customers. Within these parameters, it was determined there were several million serviceable customers in the US, sufficient to make a viable business, provided that existing efficient distribution channels could be used.

The MIT technology development work in the period 1984 - 1995 [1, 2] demonstrated transition and state based disaggregation for a small set of devices based on power measurements at a rate of few samples per second. The authors commented that there was little additional useful information using higher sample rates. However, Verlitics determined that, in the context of present day appliances and digital processing, there is significantly more information available if the sampling rate matches the line frequency. Each AC cycle is a natural quantization of energy usage. Many appliances use SCR's and/or triacs to control large loads, and these devices are naturally synchronized to each AC cycle. Verlitics, therefore, chose to use the moderate sampling rate of 60 samples per second or 50 samples per second as appropriate. The basic algorithm uses a few seconds of "on-transition" data to identify generic Devices, and maintains the states of these Devices until an off-transition can be matched to an on-Device. Devices that use continuously variable energy are separated and tracked through a separate process. The identified Device Cycles spanning multiple days are further analyzed for Cycle patterns and correlations to the Cycles of other Devices. This information is used to correct errors in individual Cycles and to identify Appliances (dishwasher, clothes dryer, clothes washer, furnace, etc) composed of multiple generic Devices and/or find a pattern of Device Cycles that identify single Device Appliances (coffee maker, crock pot, electric stove, etc).

## II. VERLITICS MONITOR

The Verlitics monitor uses standard integrated circuits supplied in suitable variations from multiple suppliers. The power meter IC is designed for use in standard utility power meters and encodes two channels of data. Since the target market includes light commercial (typically 3-phase), two power meter IC's are used so that up to 4 channels can be measured and reported.

Each channel requires a current transformer around the service leg supply wire and a direct connection to the leg. This typically requires a splice inside the breaker box, so a licensed electrician is required for a typical installation.

The power meter IC has an internal sampling and compute rate of ~4000 samples per second, depending on its clock frequency. The AC cycle sync is a phase locked loop that generates an integer number of clock cycles per AC cycle such that a fixed number of equally spaced samples are taken per AC cycle by the power meter IC (typically 67 samples). Clock synchronization significantly reduces aliasing effects, especially for devices controlled by SCRs and/or triacs.

For each AC cycle, the microprocessor reads from the power meter IC four, parameters (16-bit resolution) for each channel. These parameters are:

- Voltage: 0.1 volt resolution
- Current: 10 mA resolution
- Power: 1 watt resolution
- Q-reactive power: 1 var resolution

With this resolution, the 16-bit values are able to encode the range of values produced by a 200-amp service.

## III. DISAGGREGATION OVERVIEW

Fig 1 illustrates the turn-on power transition for six common household appliances. Stable operating conditions are typically achieved within 0.2 to 5 seconds. Fig 2 illustrates six additional appliances. The Q-reactive power transitions are typically as distinctive and may be similar or quite different. The differences in these characteristic traces are the foundation of Verlitics disaggregation.

The Verlitics disaggregation process has data structures for “Instance”, “Device”, and “Appliance”. An Instance is a complete on to off cycle of a Device or Appliance. The data structure includes the start time, duration, energy used while on, and multiple other parameters that are used for characterization and identification. A normal Instance has a distinct on-transition and an off-transition, which is recognized by a step change in power completed in three or fewer few AC cycles. A periodic Instance represents a Device that has a continuously variable load (period < 1 second). Its on-transition occurs when the variations begin and its off-transition occurs when the variations end. A “slow instance” captures the otherwise un-captured energy used between an off-transition and the next on-transition.

A Device is a structure for collecting and tracking similar Instances associated with the same device. At the start of processing, there are no Device structures. As each Instance is created, it is compared to the existing Devices. If it is matched, it is linked to that Device. If there is no match, the Instance becomes the seed Instance for a new Device.

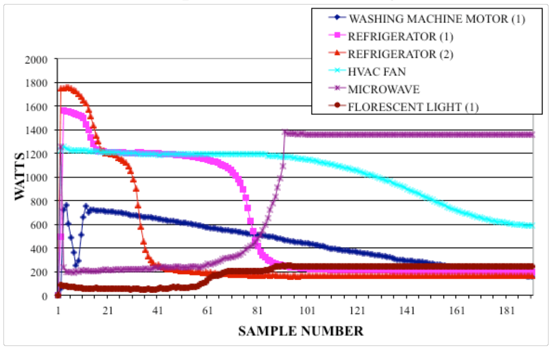

Fig. 1. Six examples of startup power traces for typical household appliances

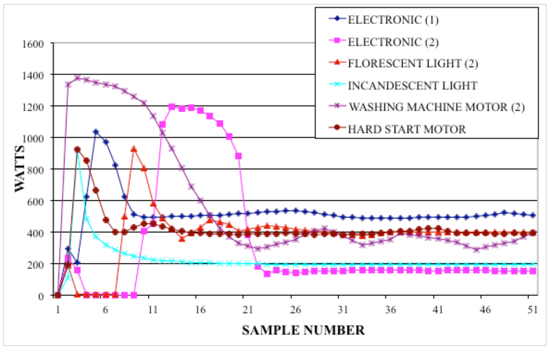

Fig. 2. Six additional examples of household appliance startup traces

An Appliance is what a typical customer might recognize and care about. Some Devices are also Appliances. For example, an incandescent light generates a clean on-off transition and has no other active components, so it is both a Device and an Appliance. However, many Appliances are composed of several elemental Devices. For example, a dishwasher has at least a heater and a motor, which create multiple Instances during one cleaning Cycle. Heaters are very common Devices, and many kitchen appliances have similar heaters that are typically controlled by a thermostat or timer. Based on the pattern of Instances of the same heater Device (same power usage) Appliances such as a coffeemaker, crock-pot, and toaster can be separately identified.

A Cycle is the representation of the time that an Appliance completes a functional cycle of operation.

Samples are processed through the Verlitics disaggregation algorithms one at a time sequentially, repeated for each service leg (2 or 3). Extensive field trials highlighted the identification of many periodic Devices, which when on, have signals that combine with normal on-transitions, potentially masking the identity of other Devices. The “Preprocessor” significantly reduces these masking signals using adaptive non-liner filter techniques described later in this paper. The Preprocessor also creates and matches periodic Instances to Devices, and detects normal on- and off-transitions.

When 2-phase and 3-phase Devices make a transition, they occur on multiple legs at essentially the same time. A correlation calculation is used to ensure the transitions are similar, and therefore made by the same Device. Instances are created separately for each service leg and linked through a multiphase Instance and Device. Single-phase and all combinations of multiphase-Instances and Devices are kept separately (3 groups for 2-phase, 7 groups for 3-phase). This reduces the number of choices when matching off-transitions to the appropriate on-Instance.

Typically a normal on-transition (positive) stabilizes after 0.5 to 5 seconds. The start up shape is then matched to an existing Device or a new Device is created. The shape is first converted to a “Feature Description” which is used to determine its “Generic Type”, and then it is compared to existing Devices of the same Generic Type. There are fewer than 50 Generic Types including Devices like heaters, lights, refrigerator, microwave, etc. There are also catchall Generic Types such as noise, peaked, and unknown. The energy used by each on-Instance is accumulated each sample time. Likewise, the total energy and on time of all the Instances of a Device are accumulated in the Device.

The audit function ensures that the total power, used by all of the on-Instances on each leg, does not sum to more than the total power of that leg. It is generally easier to detect a Device turning on than turning off, so an excess of on-power can develop. The audit function uses a variety of techniques to select the on-Instance that is turned off.

The set of Instances and Devices produced by this process are independent entities and may have some errors, typically caused by two or more transitions occurring within the 0.2 – 5 second period after the start of a transition, or mismatching a negative transition to the wrong on-instance. A second pass analysis process ("Macro Processor") looks for macro behaviors of the Instances and Devices to correct errors and identify Appliances. Typically data blocks spanning 1 – 30 days are used for this analysis.

The Macro Processor is Appliance oriented and contains Appliance specific knowledge. There is one routine for each major class of Appliance such as refrigerators, microwaves, clothes dryers, coffeemakers, A/C-heat pump, etc. These routines search for expected component Devices making up the Appliance and use time correlation to associate these component Devices. For example many refrigerators run a defrost cycle about once each day. During a defrost cycle a ~700W heater turns on at the same time as a refrigerator Cycle and lasts several minutes. Using 5 – 10 days of data, the heater Device can reliably be found, and associated (along with its energy) with the refrigerator Appliance. This heater typically adds about 10% to the refrigerator total energy.

The Macro Processor also looks for changes in specific parameters over time to generate alerts for pending failures and required maintenance. These include such things as: furnace filter replacement, refrigerator servicing; or the identification of a pending relay failure.

While the Macro Processor requires data spanning several days, it can be operated incrementally after a few days of data are accumulated, providing timely alerts such as "The stove is on" at 9pm or "The A/C is no longer working".

## IV. FEATURE DESCRIPTION OF TURN-ON TRANSITIONS

Fig 3 is a graphical representation of the start up transition of a typical refrigerator. It also illustrates the features used to describe the transition. Fig 4 is an overlay of traces from 11 instances of this Device. The traces are consistent and repeatable. The first feature in Fig 3 is a positive step where the value is 1749 Watts at sample 2. The next feature is a maximum negative slope where the value is 1433 Watts at sample 15. The remaining features include a minimum slope, another maximum negative slope, and a checkpoint value of 206 Watts at sample 50. Checkpoint values are spaced 50 samples apart. This set of feature triplets is used to determine the Generic Type and match this Instance to a Device.

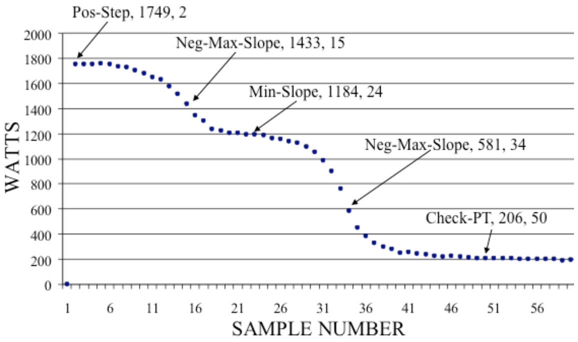

Fig. 3. A typical startup power trace for a refrigerator marked with descriptive feature

Different refrigerators may differ in the values and locations of these features, but they have the same set and sequence of features. The Generic Type identification is made by identifying a specific sequence, which includes a positive step, maximum negative slope, minimum slope, and maximum negative slope. The values and times of each feature must also be within specified, broad ranges.

With a Generic Type identification of "refrigerator", the set of Devices of Generic Type: refrigerator, are evaluated for best match to the captured Instance. The matching process returns an error value based on the weighted comparisons of feature types, values and when they occurred. If the smallest error is below a Generic Type Device-dependent threshold, then the Instance is associated with the Device. If there is no match, a new Device is created with the Instance as the seed. When a Device is linked to the 5th, 10th, 25th, 50th, or 100th Instance, each Instance is pair wise compared with every other Instance to find the seed Instance that produces the lowest total error.

Field trial data from a building with seven refrigerators were successfully disaggregated. There were four different refrigerator types, and four of the refrigerators were of the same type, three of these were served by one leg, and the other leg served one. There were five Devices created, with one Device linking to all the Instances generated by the three identical refrigerators on the same service leg. The process determined the number of identical refrigerators by occasionally finding three instances on at the same time, and often finding two instances on at the same time.

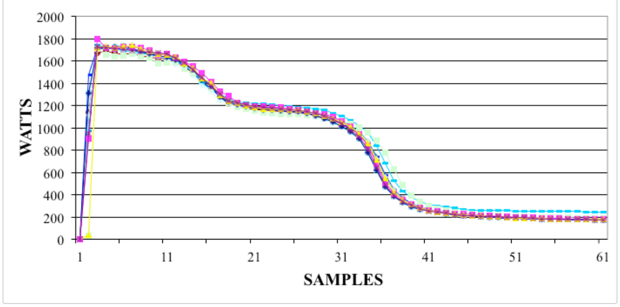

Fig. 4. Eleven overlaid power traces for a refrigerator illustrating the repeatability of the signature

## V. REMOVING MASKING SIGNALS

As a result of extensive field-testing of the technology, it was recognized that in many buildings there are electrical devices connected to the electrical service that, through their operation, introduced electrical noise that could mask the detection of other devices. The Preprocessor, discussed previously, removes the masking signals produced by these devices. The Preprocessor recognizes when a masking signal is present, characterizes it, and selects the most effective removal technique. Four different removal techniques are used to remove masking signals of the types: periodic, chaotic, rectangular, and pulse.

Fig 5 illustrates an example of a periodic masking signal, produced by a piston air compressor, combined with a normal trace and the results of the removal process. After the removal process, the remaining trace can be reliably matched to the appropriate Device.

Fig 6 illustrates an example of a chaotic masking signal produced by the switching power supply of a computer combined with a normal positive transition. After filtering, the changes caused by filtering are constrained by the masking signal limits through an iterative process. This helps prevent overly aggressive filtering of sharp features such as peaks and steps necessary for identification of Instances. Chaotic masking signals are more difficult to remove than periodic signals, but the Preprocessor significantly improves the ability to properly match Instances to the proper Device from all of the samples.

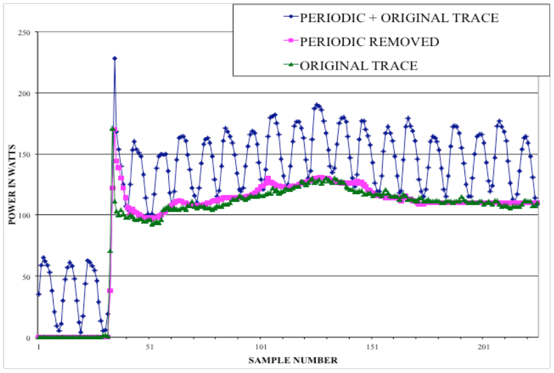

Fig. 5. Example of removing a periodic masking signal from a power trace

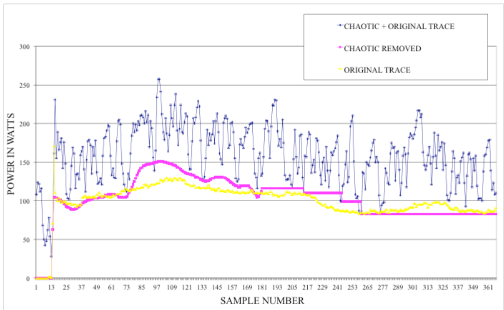

Fig. 6. Example of removing a chaotic signal from a power trace

## VI. CONTINUOUS COMMISSIONING

Verlitics' disaggregation technology is able to monitor changes in the behavior of Appliances with motors through changes in monitored electrical parameters that can be linked to changes in the operation of the Appliance. Field trials have validated the value of this capability. Such changes in Appliance behavior usually translate to a loss of efficiency and/or the need for preventative maintenance or in some cases prevention of Appliance failure. The approach employed monitors the values of one key feature of the start up transition. If the feature changes by more than a threshold amount over a period of weeks or months, an alert is generated.

Practical examples include refrigerators and air conditioners, which become less efficient as coolant is lost. This loss of coolant can be detected by the Verlitics technology and an alert for the user generated. Through this same approach airflow reduction in systems has been detected and identified. The pending failure of startup relays has also been demonstrated. These Verlitics' benefits demonstrate significant value, beyond the exclusive energy efficiency monitoring of many other disaggregation technologies. The market for equipment monitoring is well established and could benefit from this unique application of disaggregation technology.

## VII. FIELD TRIALS

Verlitics collected over 250,000 hours of data from more than 50 different buildings (mostly residential). This data set has been used to verify its disaggregation technology through multiple validation processes both in the laboratory and in the field. The initial Verlitics process that creates Instances and Devices typically captures > 90% of the energy used by appliances (that make transitions) for all buildings, and more than 95% for most. Over 10 Appliance specific routines for the Macro Processor have been developed and verified. These Appliances typically use 50% to 90% of the total energy used in these buildings. About 10 more Appliance specific routines are needed to achieve the 90% identified for 90% of the buildings.

## VIII. CONCLUSION

Several additional trials are currently in process both domestically and internationally while more Appliance specific routines are being developed. The set of trial experiences to date indicate that the 90%-90%-90% goal is achievable by Verlitics disaggregation in the near term.

## REFERENCES

[1] U.S Pat. No. 4,858,141 issued Aug. 15, 1989 to Hart, et al.
[2] U.S Pat. No. 5,483,153 issued Jan. 9, 1996 to Leeb, et al.