# A Framework for Enabling Energy Aware Facilities

The goal of this research is to identify ways to inexpensively provide detailed information about energy consumption in buildings and facilitate conservation. By relying on aggregate data, homeowners and facility managers are blind to the contribution of individual appliances and activities to the overall numbers. And while there is adequate evidence that providing real-time appliance-specific data allows users to achieve significant energy savings, the available solutions in the market are either inadequately granular or prohibitively expensive.

We have been developing signal processing, machine learning, and data fusion techniques to extract actionable information from whole-building power meters and other available sensors. The main research objectives are: (a) to create a framework for obtaining disaggregated, appliance-specific feedback about electricity consumption in a building by extracting high-value information from low-cost data sources; and (b) to investigate and develop data mining and machine learning algorithms for making use of appliance-specific electricity data, in order to provide users with recommendations on how to optimize their energy consumption and understand the effects of their energy-related decisions.

A series of residential buildings in Pittsburgh, PA are serving as a test-bed for evaluating and validating our approach. We have partnered with Blueroof Technologies, a non-profit corporation located in McKeesport, PA that researches, develops and provides affordable senior citizen housing with integrated sensor networks and building automation systems. This project also involves collaboration from researchers in Robert Bosch LLC, a leading global provider of consumer goods and building technology, who provides additional technical research assistance and expertise, while ensuring that the results of this project can have a strong potential for broad impact and success in the market.

Although we were awarded this grant late in 2009, already the first generation of our prototype system is able to correctly identify most of the appliances in the residential units where it has been installed, and can give estimates of the consumption of some of these appliances that are accurate to within a 15% difference, although the focus has been on two-state appliances with simple start-up transients (e.g., lights, small compressors and motors and resistive loads). We have also developed novel, low-cost sensing technologies to improve the process of training the system within a home.

The current research is focused on obtaining long (one week or more) and representative (with appliances found in most US homes) fully labeled datasets that can be used to evaluate the effectiveness of the solutions developed so far. We are also investigating novel feature extraction methods that can allow load signatures to generalize across appliance types. Lastly, part of our team is developing techniques to extend the existing framework to appliances that exhibit longer start-up transients and/or multiple states.