Steamrolling Residential Electric Demand

Ashutosh Dhekne, Sridhar Duggirala, Noah Garfinkle, and Thomas Reese

Abstract—This project implements intelligent scheduling of 'smart' residential appliances as a demonstration of demand management optimized for increasing efficient incorporation of renewable energy sources. We discuss an innovative approach where the grid can perform trial runs to know the responses of a group of houses to the available renewable supply. We then present detailed results for various simulations we have performed as a proof of concept.

Keywords—Cyber-Physical Systems, Residential Electric Demand, Demand Management, Renewable Energy, Smart Home

1 Introduction

Residential energy consumption comprises over 37% of the total electricity production in the United States[1]. This electric consumption is often controlled directly by consumers, causing diurnal demand peaks on typical energy load profiles. These peaks, however, mean that demand often does not fit the preferred supply patterns of electricity producers. This difference between supply and demand forces the electric grid to rely on less efficient peaking generators during times of high demand. As these times of peak demands are outside the control of producer, the fluctuations would lead to wastage of natural resources. This project attempts to manage the demand of the end users by exploiting the flexibility in the household appliances such as Heating, Ventilating, and Air Conditioning (HVAC) systems. For exploiting such flexibility, the presence of a 'smart controller' which schedules household devices is assumed. Such homes are referred as 'smart homes'.

This paper begins by briefly outlining the challenges, opportunities and significance of this problem. Sections 2, 3, and 4 lay out the objectives, necessary background and the design methodology of our proposed solution. Sections 5 and 6 examine the proposed solution for a number of simulated test runs. Section 7 concludes with what the authors consider to be the most important implications of this work and the significant future work for the topic to achieve even more merit and popularity is presented in Section 8.

1.1 Problem Statement

Traditionally, electricity has been available only from non-renewable sources. It was produced from natural materials that stored an enormous amount of chemical, potential or nuclear energy. However, with increasing amounts of renewable energy being harnessed these days, the issues related to electricity supply and demand have undergone radical change. As renewable sources of energy such as wind and solar energy are freely available, using such sources to the maximum extent is in the best interest of the environment. This paper leverages the renewable sources of energy to match the peak demands and smoothen the load on non-renewable sources. Hence, the desired objective is given in Equations 1 and 2.

$$\frac{Demand}{Supply} \rightarrow 1 \tag{1}$$

and

$$\frac{PeakDemand}{AverageDemand} \rightarrow 1 \tag{2}$$

• *A. Dhekne, S. Duggirala and T. Reese are graduate students in the Department of Computer Science at The University of Illinois at Urbana-Champaign.*

• *N. Garfinkle is a graduate student in the Department of Civil and Environmental Engineering at The University of Illinois at Urbana-Champaign.*

Manuscript received December 19, 2014.

1. http://www.c2es.org/technology/overview/electricity

1.2 Objectives

Reducing peak loads, smoothing the usage curve, and minimizing the consumption of non-renewable energy sources through improved utilization of renewable energy sources is seen as a way for both power providers and consumers to benefit. Typical peak loads are served by peaking generators that are expensive to operate and cause more burden on the environment. Serving a flattened load would reduce the utilization of peak generators, thus leading to monetary, and environmental savings. These savings can be passed on to consumers via rate reductions.

There are several requirements that are to be satisfied by the system that implements such a solution. First, the new system should not introduce any instabilities into the power grid which might cause brown outs or damage to sensitive equipment. Second, along with maintaining stability and efficiency, such system should also protect the privacy of the end-users. Hence, the ideal solution should not transmit any additional information (if possible reduce it) than that is transmitted currently to the grid.

2 Background

2.1 Residential Energy Consumption

If grid power utilization was perfectly flat throughout the day, an energy provider could use a cheap source of constant output such as nuclear or coal for their entire energy portfolio. There would be no additional demand to meet, power generation could hit a peak efficiency, and fear of grid overload would be non-existent. In reality, however, these stable sources must be augmented with variable sources, often in the form of natural gas.

During an average day, power use on a grid follows a diurnal curve. This is due to the habits of individuals as they run through their daily routine. These two peaks in power usage cause excess energy consumption and local variability on the grid. Using variable sources, while being both expensive and inefficient, is necessary to power today's energy demands.

In addition to traditional sources, renewable sources such as wind and solar add further variability to the grid. While this power is essentially free to generate, demand does not necessarily follow the curve of its highly variable availability.

2.2 Market and Environmental Analysis

More than 50% of a home's energy use can be attributed to a few high energy appliances, including: HVAC, refrigeration, water heaters, and clothes washers and dryers. Currently these devices are manually operated and hence the diurnal peak is a reflection of the working habits of the general population. These devices, when observed at a macro level lead to inherent inefficiencies present in the grid. Such devices are referred to as 'dumb devices'.

Of these appliances the majority could be operated outside of direct manual triggers. Their operations could thus be delayed to run on a different schedule. Also, since many of these appliances progress through various operational stages with the possibility of leeway between each stage, they could lend themselves some freedom of when to operate.

Appropriate exploitation of this notion allows the majority of a household's energy consumption to be shifted to a time when energy is not in peak demand. Doing so would lead to higher efficiencies and cheaper energy costs. We call devices that are capable of performing/operating such a schedule as 'smart devices'. In this paper, we further explore this idea and show that, at scale, adjusting many homes' usage using such smart devices would help flatten the energy demand curve.

2.3 Power Sources

Although the current sources of power available to consumers have already been introduced and discussed, it is worth very briefly formalizing the terminology and implications of each.

Base generation consists of power which is derived from highly concentrated, cheap sources. However, base sources such as coal power plants are difficult to dynamically control, and operate most efficiently if allowed to output a relatively fixed supply of power over time. Such facilities tend to be large and

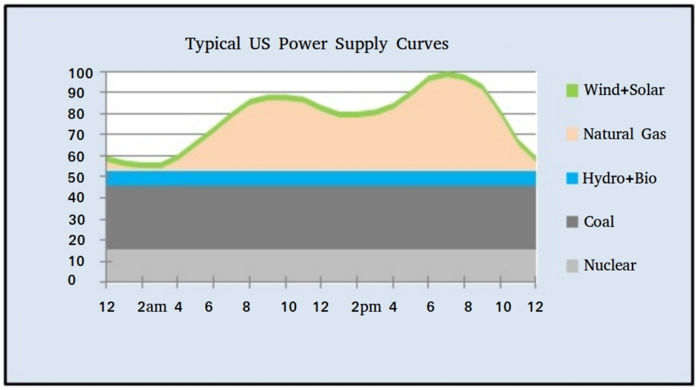

Fig. 1. Illustration of demand variability of various power sources [2]

expensive to scale, requiring years to alter base capacity.

Unlike base sources, peaking power sources are very flexible; capable of being brought on and off at will, often utilizing automated control techniques. These sources of energy such as natural gas tend to be more expensive than base sources, but easier to scale over time.

Renewable power sources exploit abundant, freely available natural sources of energy such as wind and solar. Unfortunately, these sources are highly variable, and entirely outside of the control of power utilities. Additionally, production facilities tend to be decentralized, with frequent additional and scaled capacity added. Such variability in supply and capacity will benefit from control schemes designed to better utilize power as it becomes available.

Figure 1 illustrates the average daily variability of several sources of power. This figure is particularly helpful for readers seeking to visualize the flat supply of base sources, large flexibility of peaking generators such as natural gas turbines, and the high variability of renewable sources.

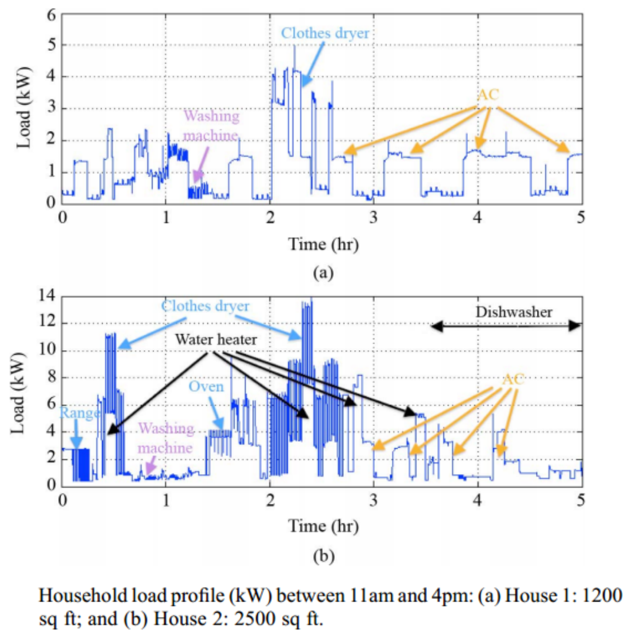

Fig. 2. Daily household load due to various major appliances [4]

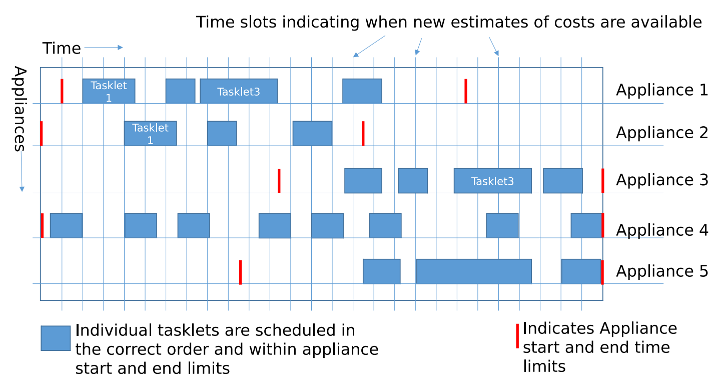

Fig. 3. Illustration of the tasks and tasklets in an appliance

3 Methodology: Inside the Home

3.1 Appliances, Tasks and Tasklets

A typical home has many appliances including lamps, TV, refrigerator, dish washer, washing machine, heater, air conditioners, microwave ovens etc. Some of these appliances have a set number of stages that they go through. For example, a washing machine goes through soaking, washing, spinning, and drying stages as shown in Figure 2. When a washing machine is scheduled, it will go through these stages in sequence. We call such a sequence, a task. Thus the act of starting the washing machine once is one task. The various stages the appliance undergoes during a task are called tasklets. Figure 3 explains the relationship between tasks and tasklets in an appliance.

We allow a task to specify its start and end-by time. We also allow a start-by time for each tasklet. This effectively allows us to model all possible bounds on the tasks and individual tasklets. The scheduler then takes data from all appliances and proceeds to create the viable schedule for this house. It also receives a 24 hour prediction for the availability of renewable and non-renewable power.

The system is able to take into account both smart and dumb houses. A dumb house will continue operating as normal by demanding the needed power. Within a smart home, any smart device is assumed to provide the house its task requirements and dumb devices are treated as single tasks with an immediate start time.

This paper assumes that the prediction of the renewable and non-renewable sources of energy is compiled by the grid and is communicated (through the Internet) to the controllers in every house. Each house, in return communicates the schedule's aggregate demand back to the grid.

3.2 Manager

Household appliances with scheduling capabilities can be linked together with a commercially available router designed for connected homes, such as the Securifi Almond+. Existing communication protocols like the Zigbee or the Z-Wave protocols can be used for this communication. After receiving requests to schedule tasks from appliances, the router then uses a scheduling algorithm such as the brute force implementation shown in Algorithm 1. It can also fetch the grid's prediction of renewable and non-renewable power supply from the Internet.

Algorithms 1 and 2 describe how the brute force scheduler operates. Algorithm 1 demonstrates the simple nature of a brute force scheduler, while algorithm 2 describes how each possible combination of task parameters can be generated by the scheduler as new tasks are submitted.

The subset of feasible solutions is generated using the $generatePossibleSchedules$, which considers every possible combination of all feasible values for each task in order to create the comprehensive list which the brute force algorithm acts upon.

The scheduler may be triggered by *a*) submission of new task requests; *b*) receipt of a new power estimate from the grid; *c*) power consumption from dumb appliances; or *d*) interruption of currently scheduled tasks. Once a schedule is set, the manager is responsible for triggering scheduled tasks at the appropriate time. Tasks which are currently in progress when a new schedule is to be generated will continue with the existing schedule for the tasklet which has already begun (tasklets are not interrupted). The scheduler is however, free to reschedule future tasklets within the allowable constraints of the overall task.

While the scheduler can be adapted to meet various priorities which might be shared between the consumer and the power provider, the current scheduler seeks to minimize the amount of non-renewable power consumed. This has two desirable properties. First, minimizing non-renewable power consumption ensures that available renewable power is utilized as much as possible within the scheduling constraints. Second, rewarding only each unit minimization of nonrenewable power helps to preserve stability – discussed in section 6 – because a large peak in renewable power availability can only shift the schedule proportional to the amount saved at that individual time step.

The scheduler must also seek to minimize the likelihood of the house drawing too much current at any time step, tripping the circuit breaker. This consideration is particularly important because the start-up tasklet of many high-energy-consumption appliances (such as washing machines) may have a huge spike constrained to a small time period as motors start.

3.3 Smart Meter

Smart metering is increasingly common for new residential construction. Such meters – capable of recording high resolution power consumption profiles for each residence – already possess most of the measurement and communication capabilities necessary for coordinating schedules between the grid and consumers. Two additional communications are required for the current implementation of this project. First, the meter must be capable of connect-

Algorithm 1 Scheduler

```
procedure SCHEDULER(ctaskList, startTime, cpowerEstimate, fuseCapacity)    ▷ Compute the best possible schedule
    possibleSchedules ← generatePossibleSchedules(taskList, startTime)
    bestSchedule ← None
    bestScheduleScore ← High
    for schedule in possibleSchedules do
        score ← evaluateSchedule(schedule, powerEstimate, fuseCapacity)
        if score < bestScheduleScore then
            bestSchedule ← schedule
            bestScheduleScore ← score
    return bestSchedule
```

Algorithm 2 GeneratePossibleSchedules

```
procedure GENERATEPOSSIBLESCHEDULES(ctaskList, startTime)    ▷ Produce a list of feasible schedules
    for task in taskList do
        taskScores ← []
        lastAllowableStartTime ← task.maxEndTime − taskLength    ▷ Set the possible start times for the task
        allowableStartTimes ← [tfortinrange(task.minStartTime, lastAllowableStartTime]    ▷ Create a nested list of tasklet gaps
        allowableTaskletGaps ← []
        for tasklet in task.tasklets do
            taskletGap ← [gforginrange(0, tasklet.maxGap)]
            taskletGaps.append(taskletGap)
            combinedTaskletGaps ← list(it.products(*taskletGaps))
        for startTime in allowableStartTimes do
            for taskletGap in combinedTaskletGaps do
                generateAllowableTaskParameters
    combinedAllowableTaskSchedules ← list(it.product(*allAllowableTaskParameters)
    return combinedAllowableTaskSchedules
```

ing with the home manager, a task easily accomplished with many existing smart meters. Second, the meter must be capable of securely receiving power estimates from the power grid, if the manager lacks this functionality.

4 Methodology: System

4.1 Single Home

We currently employ a brute force scheduler in every house to come up with the most efficient schedule for the individual house. This schedule takes all the possible flexibility available from every appliance into account and then computes all possible combinations. According to the input predictions of the renewable and non-renewable supply, the scheduler will try to maximize the use of renewable power.

4.2 Multiple Homes

Each house follows an identical scheduling rule set. The grid decides whether or not to provide renewable energy to any given house. This simplifies the problem of finding the global optimum by offloading the work to the grid instead of using a distributed algorithm.

4.3 The Grid

The grid has an estimate of the available renewable energy for the entire day. It selects a subset of houses and sends them a modified version of this estimate (we call this a trial run). It then receives a quotation about the demand from the individual houses. The grid is free to perform numerous modifications to the estimate as well as the subset of houses to send this estimate to, and perform more trial runs. Finally the grid picks the best trial run and submits that to the houses for real scheduling.

Since the chosen estimate is one of those it already had a trial run for, the grid can predict the demand precisely and take efficiency related decisions before the new schedules kick in. Such actions could involve making modifications to peaking or base generators allowing the power producers to run more efficiently.

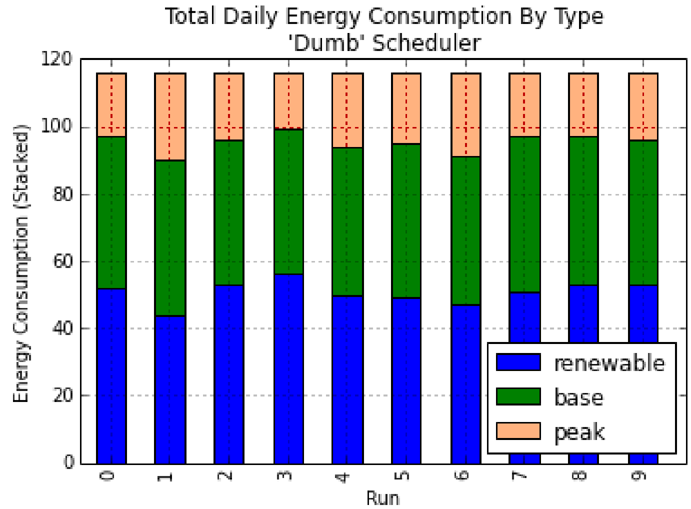

Fig. 4. Energy consumption for 10 randomly generated scenarios with a 'dumb' scheduler

5 Demand Management

The primary advantage of smart scheduling of demand in order to utilize renewable sources of energy, should be improved efficiencies for the power supplier. Such efficiencies can be described first by flattening the peak demand (Equation 2) and then by meeting as much of the demand as possible using renewable sources (Equation 1).

An adequately flattened load would be seen to utilize little peaking power (compare the peak component in Figure 4 with Figure 5).

Because peak power is generally more expensive than base power, which will quickly become more expensive than renewable power as renewable power sources reach scale, we would expect the increased utilization of renewable power (and associated offsetting of peak supplies) to decrease costs for suppliers. These efficiency benefits can then be passed on to consumers.

As an example, to comprehend the impact, Figure 6 shows the difference between grid costs for 500 randomized simulations when the appliances are allowed to have a large scheduling flexibility (smart) versus when the appliances are allowed no scheduling time flexibility (dumb). For this argument's sake, we consider the cost of renewables to be 1 unit, non-renewable base generators to be 2 units and non-renewable peaking generators to be 3 units.

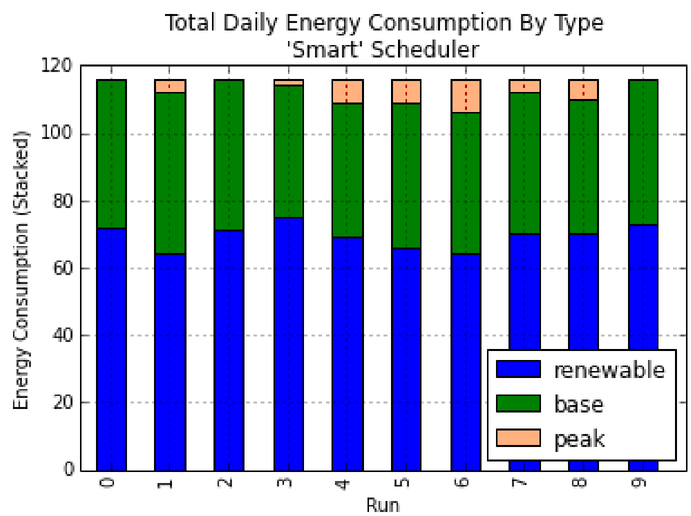

Fig. 5. Energy consumption for 10 randomly generated scenarios with a 'smart' scheduler

An innovative approach in our paper is to communicate the imminent aggregate demand back to the grid. This allows the grid to make smart decisions about tuning the schedules of the peaking generators. Essentially this gives some foreknowledge to the grid allowing it to make very efficient decisions. It can decrease the overhead margins considerably and turn off peaking generators when they are no longer needed. This is a big leap from the standard diurnal average estimation available today.

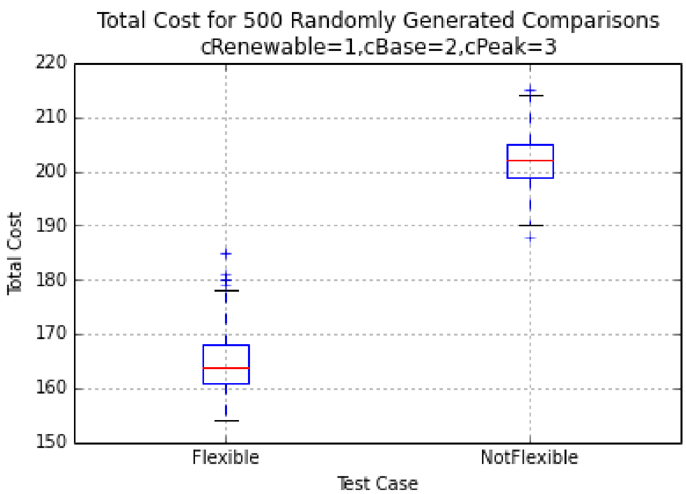

Fig. 6. Monte Carlo analysis of cost for 500 randomly generated simulations with 'smart' and 'dumb' appliances

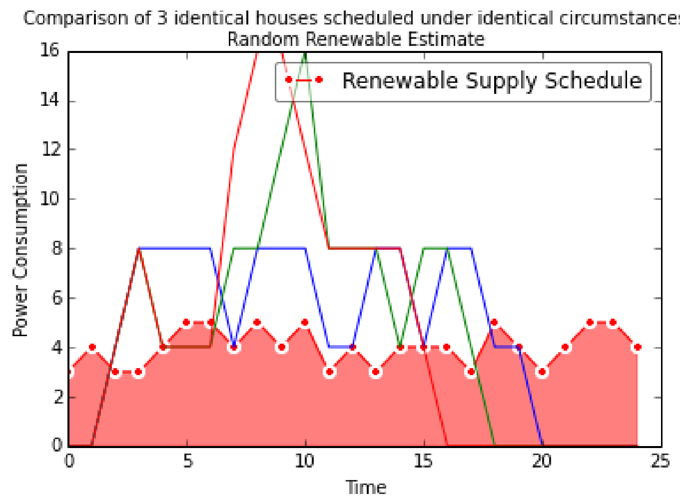

Fig. 7. Stability analysis for a scenario with multiple homes and a highly variable supply of renewable energy

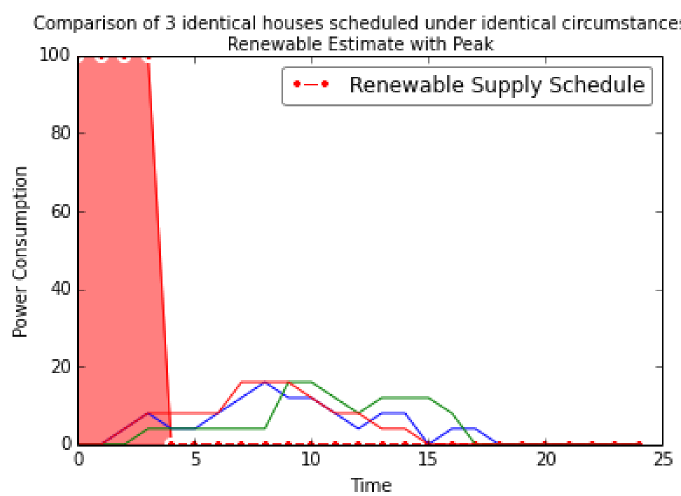

Fig. 8. Stability analysis for a scenario with multiple homes and a large peak of renewable energy availability

6 STABILITY

Stability is an important property that should be satisfied by the system. In this scenario, preserving stability would require that the schedules of majority of the houses should remain unaltered in the presence of a sudden rise (or sudden drop) in the supply of renewable source.

In this section we illustrate the performance of the proposed algorithm with the following variations in supply of renewable energy: a) peak supply for a short duration of time; b) moderate supply for a short duration of time; c) moderate supply for random durations of time; and d) no supply for the entire duration of the day.

Through handling of these cases we argue that the system we proposed is a stable one. Further, the proposed solution would not lead to adverse scenarios such as failure of the grid. These simulations are shown in Figure 7-9 demonstrate these cases for a multi-home scenario. Use of the algorithm will never make the grid more susceptible to problems than what it already is through use of traditional systems. In many cases, the algorithm makes the grid more efficient.

Since the algorithm gets the entire day's renewable estimates from the grid, it can perform cost benefit analysis for the whole period of flexibility keeping both Equation 1 and Equation 2 in perspective. The algorithm also always respects the house's total power capacity and the requirements of individual equipments. It therefore never chooses a schedule that can blow the fuse.

Finally, since the grid is ultimately responsible for sending out the appropriate final renewable availability estimates, the grid can just choose another trial run if it feels a particular run is reaching too close to the limits.

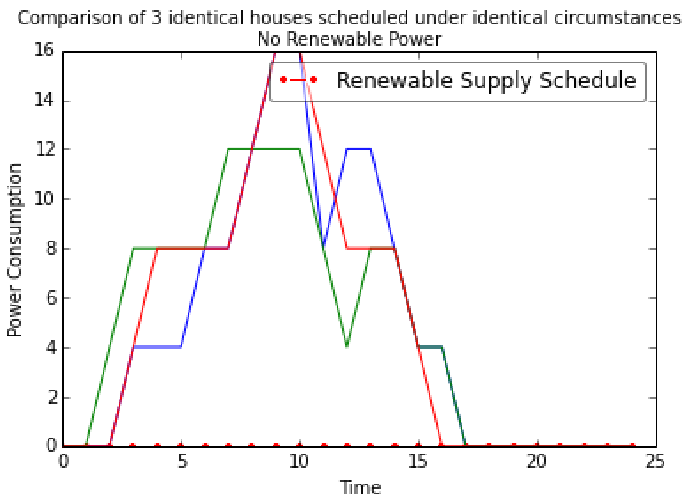

Fig. 9. Stability analysis for a scenario with multiple homes and no availability of renewable power

7 Conclusions

In the experiments and empirical analyses performed thus far, it is seen that commercially available technologies make it feasible for power suppliers, consumers and appliances to coordinate residential energy demand in a way which improves overall system performance, lower costs, and benefits the environment. Additionally, it is shown that fundamental properties of the algorithm can be structured such that no new instabilities are introduced to the power grid, and little additional information flow is required from consumers to producers.

Especially for cases with available sources of renewable energy and flexible, cyclic high power devices with exploitable patterns of consumption, it is seen that both consumers and producers can benefit from improved system efficiencies.

8 Future Work

A primary objective of future work should be the implementation of a more efficient task scheduler. The current brute-force approach will rapidly prove infeasible when scheduling large number of appliances, scheduling for long time periods or performing multiple scheduling passes while coordinating with the grid.

The overall scheduling methodology can also benefit from consideration of heuristics, which can be implemented in several ways. First, the power grid could establish patterns of power consumption, and provide incentives for individual homes to subscribe to a particular pattern for its own consumption. This methodology could replace trial runs of the scheduler, allowing the grid and households to cooperate more closely without having to share additional information. Additionally, heuristics can be used to reduce the complexity of schedulers by taking advantage of known features and demand classes such as pattern driven, cyclic high power devices.

Finally, this project will benefit greatly from physical prototyping. While smart appliances are less accessible currently, smart plugs are commercially available and may be able to automate some appliances. Additionally, the success of Internet of Things technologies such as Nest thermostats have demonstrated the ease by which HVAC systems – one of the largest energy consumers in residential buildings – can be controlled. The routers capable of acting as central managers and smart meters are both commercially available and increasingly found in households around the nation.

Appendix A
Source Code

All source code for this project can be found at https://github.com/Gibolt/SmartGrid. Where appropriate, comments within the source code seek to describe the proper implementation of each class and algorithm. All source code is offered without warranty, and is subject to change at any time.

Appendix B
Software Diagram

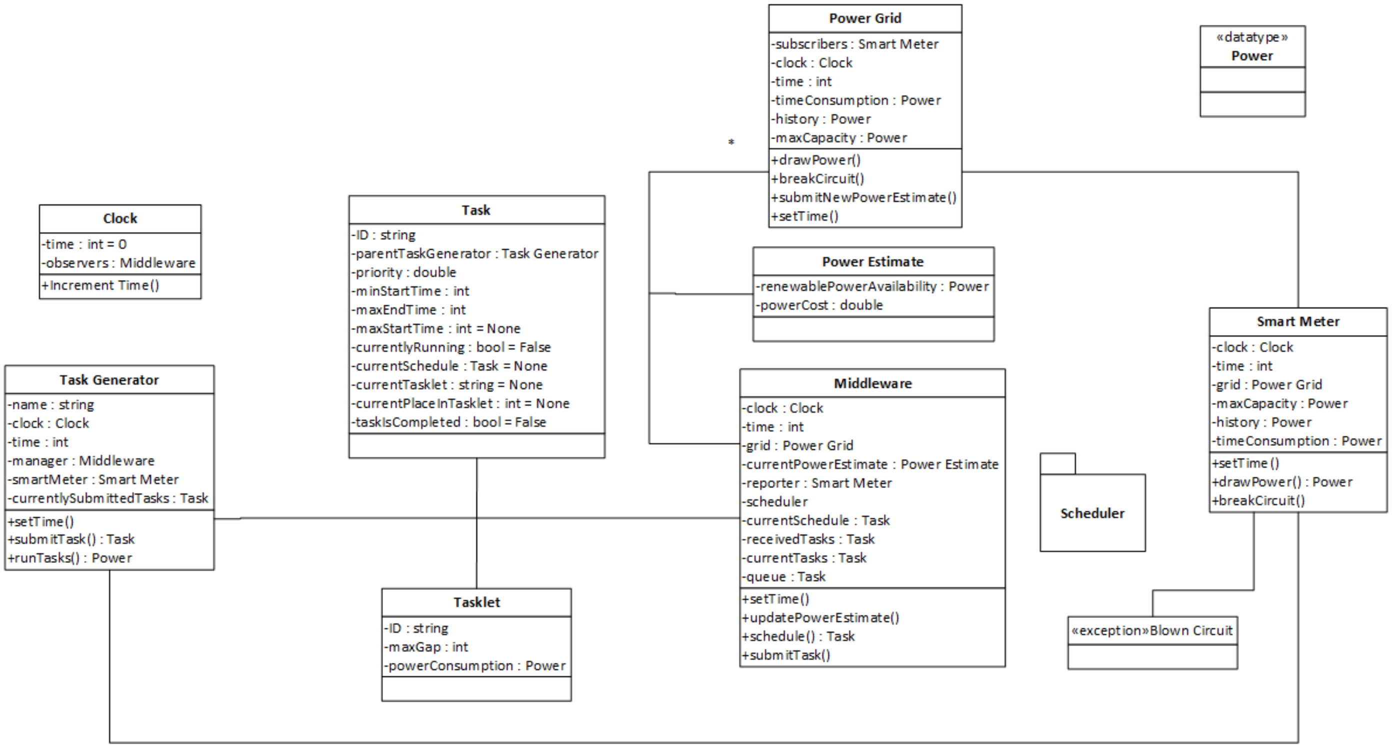

Fig. 10. UML diagram of the classes created for this project

Acknowledgments

The authors would like to thank Professors Marco Caccamo and Sibin Mohan of the University of Illinois at Urbana-Champaign. The authors would also like to thank their classmates who tolerated multiple discussion sessions about this project.

References

[1] Center for Climate and Energy Solutions, *Electricity Overview* http://www.c2es.org/technology/overview/electricity

[2] Dan Mantena, *Fuel Diversity in U.S. Electricity Markets* http://theenergycollective.com/sites/theenergycollective.com /files/imagepicker/151581/Wind Solar Needs NG Graph2.jpg

[3] A.H. Mohsenian-Rad, V.W.S. Wong, J. Jatskevich, R. Schober, A. Leon-Garcia, "Autonomous Demand-Side Management Based on Game-Theoretic Energy Consumption Scheduling for the Future Smart Grid," in IEEE Transactions on Smart Grid, vol. 1, 2010, pp.320-331

[4] M. Pipattanasomporn, M. Kuzlu, S. Rahman, Y. Teklu, "Load Profiles of Selected Major Household Appliances and Their Demand Response Opportunities," in IEEE Transactions on Smart Grid, vol. 5, 2014, pp.742-750