Lowering energy demand during peak hours

By Johannes Blomquist

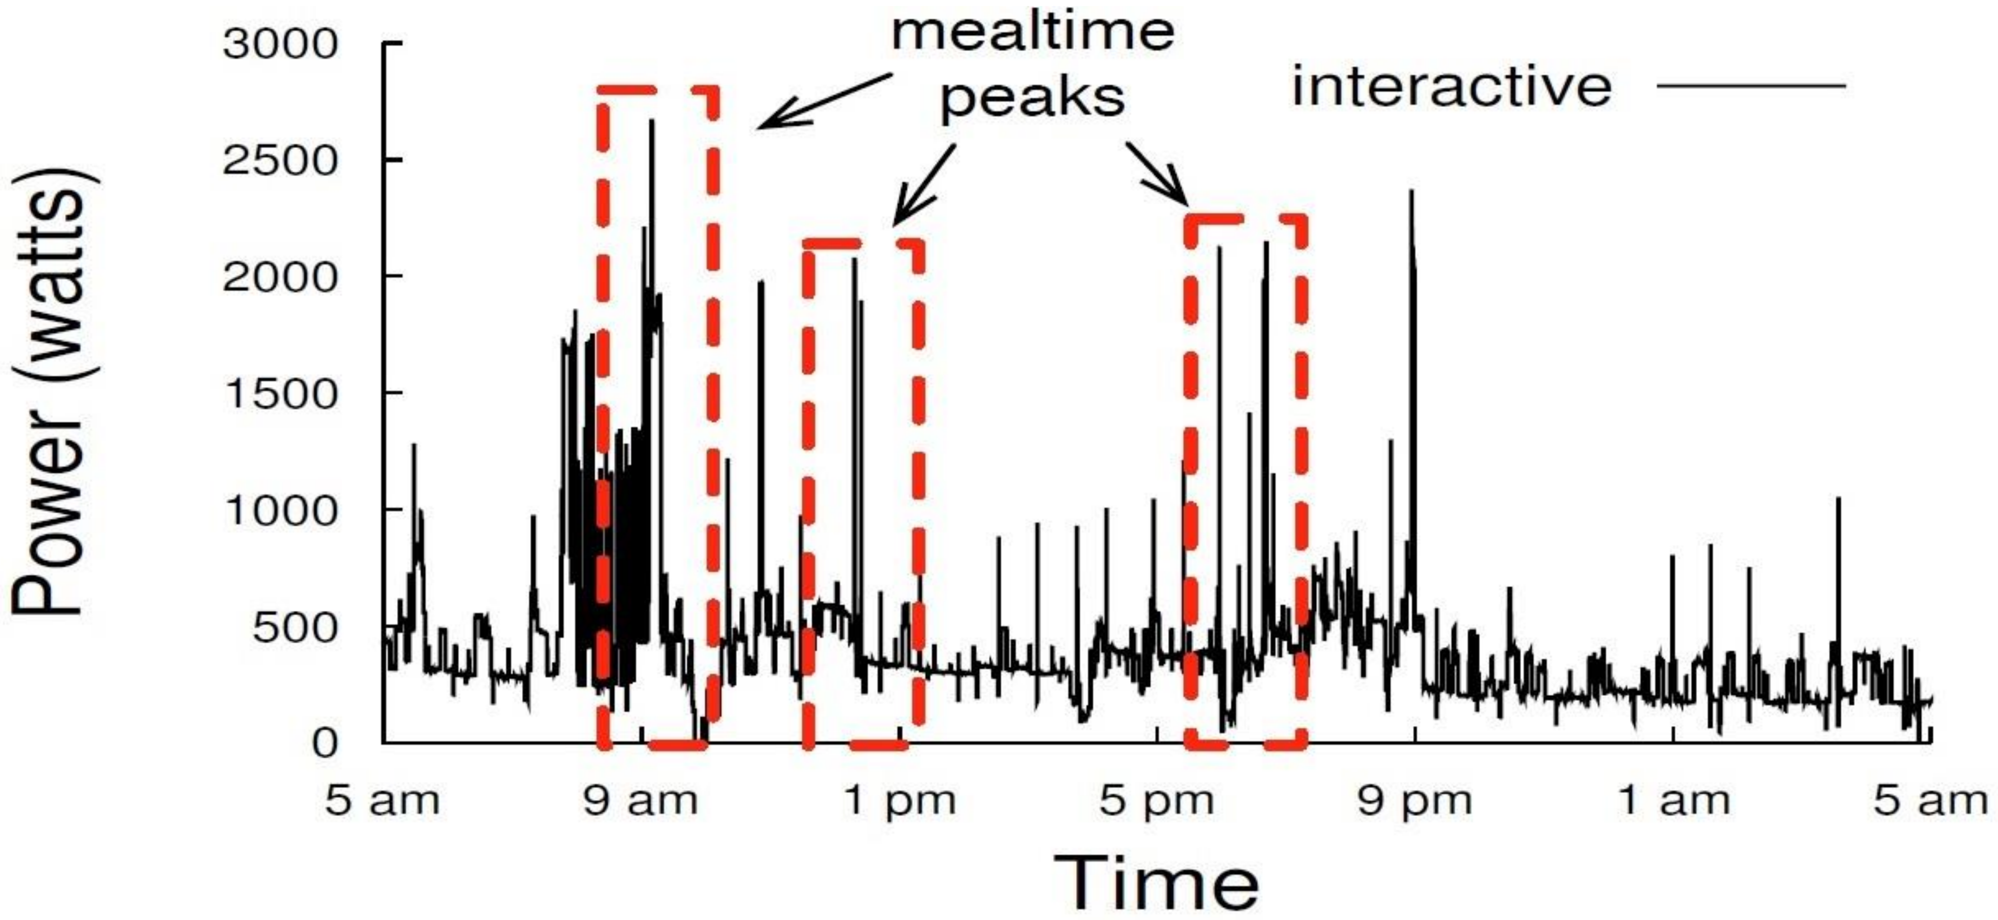

Problems:

- *People have very similair daily routines*
- *Supplying the peak energy is very expensive*

Motivation

- Clients might get charged for their maximum peak demand
- Providing clients financial incentive to take some action
  - Residential and commercial buildings account for over 75% of electricity consumption in the US
- Benefits both client and power supplier

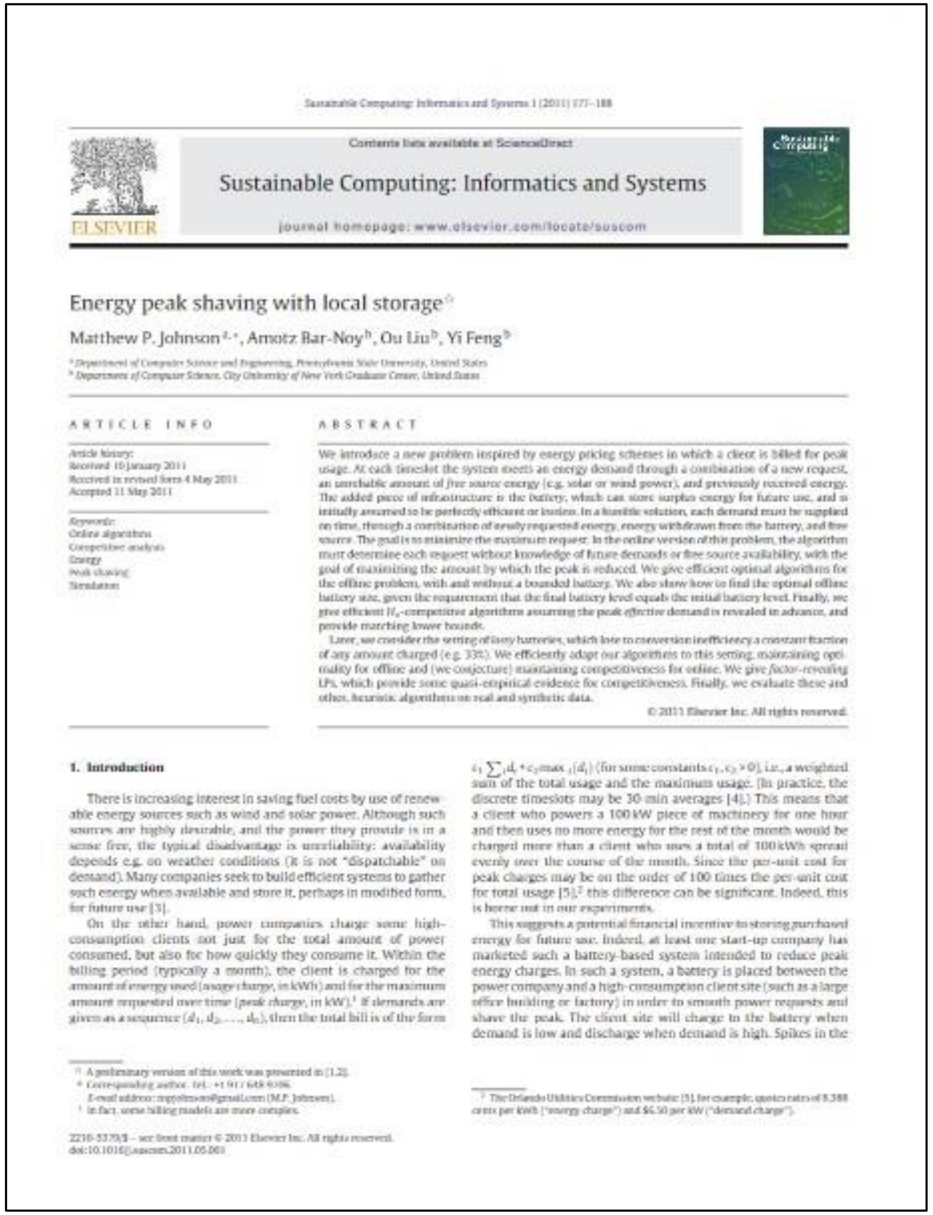

Sustainable Computing: Informatics and Systems 1 (2011) 177–188

Contents lists available at ScienceDirect

Sustainable Computing: Informatics and Systems

journal homepage: www.elsevier.com/locate/suscom

Energy peak shaving with local storage

Matthew P. Johnson[a,*], Amotz Bar-Noy[b], Ou Liu[b], Yi Feng[b]

[a] Department of Computer Science and Engineering, Pennsylvania State University, United States
[b] Department of Computer Science, City University of New York Graduate Center, United States

ARTICLE INFO

Article history:
Received 10 January 2011
Received in revised form 4 May 2011
Accepted 11 May 2011

Keywords:
Online algorithms
Competitive analysis
Energy
Peak shaving
Simulation

ABSTRACT

We introduce a new problem inspired by energy pricing schemes in which a client is billed for peak usage. At each timeslot the system meets an energy demand through a combination of a new request, an unreliable amount of free source energy (e.g. solar or wind power), and previously received energy. The added piece of infrastructure is the battery, which can store surplus energy for future use, and is initially assumed to be perfectly efficient or lossless. In a feasible solution, each demand must be supplied on time, through a combination of newly requested energy, energy withdrawn from the battery, and free source. The goal is to minimize the maximum request. In the online version of this problem, the algorithm must determine each request without knowledge of future demands or free source availability, with the goal of maximizing the amount by which the peak is reduced. We give efficient optimal algorithms for the offline problem, with and without a bounded battery. We also show how to find the optimal offline battery size, given the requirement that the final battery level equals the initial battery level. Finally, we give efficient H_n-competitive algorithms assuming the peak effective demand is revealed in advance, and provide matching lower bounds.

Later, we consider the setting of lossy batteries, which lose to conversion inefficiency a constant fraction of any amount charged (e.g. 33%). We efficiently adapt our algorithms to this setting, maintaining optimality for offline and (we conjecture) maintaining competitiveness for online. We give *factor-revealing* LPs, which provide some quasi-empirical evidence for competitiveness. Finally, we evaluate these and other, heuristic algorithms on real and synthetic data.

© 2011 Elsevier Inc. All rights reserved.

1. Introduction

There is increasing interest in saving fuel costs by use of renewable energy sources such as wind and solar power. Although such sources are highly desirable, and the power they provide is in a sense free, the typical disadvantage is unreliability: availability depends e.g. on weather conditions (it is not "dispatchable" on demand). Many companies seek to build efficient systems to gather such energy when available and store it, perhaps in modified form, for future use [3].

On the other hand, power companies charge some high-consumption clients not just for the total amount of power consumed, but also for how quickly they consume it. Within the billing period (typically a month), the client is charged for the amount of energy used (*usage charge*, in kWh) and for the maximum amount requested over time (*peak charge*, in kW).[1] If demands are given as a sequence $(d_1, d_2, \ldots, d_n)$, then the total bill is of the form $c_1 \sum_i d_i + c_2 \max_i(d_i)$ (for some constants $c_1, c_2 > 0$), i.e., a weighted sum of the total usage and the maximum usage. (In practice, the discrete timeslots may be 30-min averages [4].) This means that a client who powers a 100 kW piece of machinery for one hour and then uses no more energy for the rest of the month would be charged more than a client who uses a total of 100 kWh spread evenly over the course of the month. Since the per-unit cost for peak charges may be on the order of 100 times the per-unit cost for total usage [5],[2] this difference can be significant. Indeed, this is borne out in our experiments.

This suggests a potential financial incentive to storing purchased energy for future use. Indeed, at least one start-up company has marketed such a battery-based system intended to reduce peak energy charges. In such a system, a battery is placed between the power company and a high-consumption client site (such as a large office building or factory) in order to smooth power requests and shave the peak. The client site will charge to the battery when demand is low and discharge when demand is high. Spikes in the

☆ A preliminary version of this work was presented in [1,2].
* Corresponding author. Tel.: +1 917 648 9096.
E-mail address: mpjohnson@gmail.com (M.P. Johnson).
[1] In fact, some billing models are more complex.
[2] The Orlando Utilities Commission website [5], for example, quotes rates of 8.388 cents per kWh ("energy charge") and $6.50 per kW ("demand charge").

2210-5379/$ – see front matter © 2011 Elsevier Inc. All rights reserved.
doi:10.1016/j.suscom.2011.05.001

2012 IEEE International Conference on Pervasive Computing and Communications, Lugano (19-23 March 2012)

SmartCap: Flattening Peak Electricity Demand in Smart Homes

Sean Barker, Aditya Mishra, David Irwin, Prashant Shenoy, and Jeannie Albrecht[†]
University of Massachusetts Amherst
[†] *Williams College*

***Abstract*—Flattening household electricity demand reduces generation costs, since costs are disproportionately affected by peak demands. While the vast majority of household electrical loads are interactive and have little scheduling flexibility (TVs, microwaves, etc.), a substantial fraction of home energy use derives from background loads with some, albeit limited, flexibility. Examples of such devices include A/Cs, refrigerators, and dehumidifiers. In this paper, we study the extent to which a home is able to transparently flatten its electricity demand by scheduling only background loads with such flexibility. We propose a Least Slack First (LSF) scheduling algorithm for household loads, inspired by the well-known Earliest Deadline First algorithm. We then integrate the algorithm into SmartCap, a system we have built for monitoring and controlling electric loads in homes. To evaluate LSF, we collected power data at outlets, panels, and switches from a real home for 82 days. We use this data to drive simulations, as well as experiment with a real testbed implementation that uses similar background loads as our home. Our results indicate that LSF is most useful during peak usage periods that exhibit "peaky" behavior, where power deviates frequently and significantly from the average. For example, LSF decreases the average deviation from the mean power by over 20% across all 4-hour periods where the deviation is at least 400 watts.**

I. Introduction

Recent studies indicate that residential and commercial buildings account for over 75% of electricity consumption in the United States [2]. As a result, designing new "green" buildings and retrofitting existing buildings with green technologies has become both an important research challenge and societal need. In the residential sector, many techniques exist to reduce either a home's energy footprint or its energy bill. For instance, smart buildings may use motion sensors to track occupants and opportunistically disconnect *loads*[1] in empty rooms [11]. Alternatively, consumers may participate in automated demand response programs increasingly offered by electric utilities, which automatically turn off home appliances when the demand for electricity is high [10]. These intelligent load management schemes reduce a home's energy footprint and its bill by automatically disconnecting loads from power when necessary or convenient. This paper focuses on an intelligent load management scheme for flattening household electricity usage or demand.

Flattening demand implies reducing the difference between the peaks and troughs in a home's electricity usage, thereby creating a flatter usage pattern that lessens the deviation from the average usage. Demand flattening has the potential to benefit residential consumers as the electric grid becomes smarter and more efficient, since peak demands have a disproportionate affect on grid capital and operational costs, including transmission, generation, and fuel costs. For instance, demand flattening significantly reduces transmission and distribution losses, which account for nearly half (47%) of residential energy consumption [3], since these losses are proportional to the square of current.

To incentivize demand flattening, utilities are transitioning from flat pricing models to variable time-of-use or peak-load models [4], [5], [8], [17]. Since the marginal cost to generate electricity rises as demand increases, utilities are beginning to add surcharges to bills based on a consumer's peak usage. For example, a utility may determine the bill, in part, based on a customer's largest half-hour of electricity demand within a day, regardless of the total day's energy consumption. The new electricity pricing models provide consumers strong incentives to regulate not only their total energy consumption, but also their consumption profile. In particular, these new pricing models incentivize customers to lower their peak consumption by flattening their usage.

Unfortunately, while conceptually simple—to control its demand, a home need only decide when to disconnect its loads—intelligent load management has proven difficult to implement in practice. One reason is that disconnecting loads requires active consumer involvement during peak periods, such as turning off unnecessary lights, programming a thermostat, or postponing washing clothes. Prior studies have shown that compelling consumers to change their household routines is challenging [9]. While providing occupants real-time feedback of their power consumption may initially incentivize them to reduce their usage, once the novelty wears off occupants typically revert to their previous habits. Even for consumers that wish to actively manage their load, choosing which loads to disconnect and when is a complex decision that must be continuously re-evaluated based on information that is constantly changing. To address the problem, we have designed *SmartCap*, a system for automatically monitoring and controlling household loads.

As a key step in SmartCap's design, this paper studies the extent to which homes are able to flatten their home electricity demand without affecting home occupants or requiring their active involvement. We explore the impact of modifying *background* electrical loads that are completely transparent to home occupants and have no impact on their perceived comfort. While the vast majority of electrical loads in homes are *interactive* and have little scheduling flexibility

[1] We use the term load throughout the paper to refer to any appliance or device in the home that draws electricity.

978-1-4673-0258-6/12/$31.00 ©2012 IEEE 67

Sustainable Computing: Informatics and Systems 1 (2011) 177–188

Contents lists available at ScienceDirect

Sustainable Computing: Informatics and Systems

journal homepage: www.elsevier.com/locate/suscom

Energy peak shaving with local storage☆

Matthew P. Johnson[a,*], Amotz Bar-Noy[b], Ou Liu[b], Yi Feng[b]

[a] Department of Computer Science and Engineering, Pennsylvania State University, United States
[b] Department of Computer Science, City University of New York Graduate Center, United States

ARTICLE INFO

Article history:
Received 10 January 2011
Received in revised form 4 May 2011
Accepted 11 May 2011

Keywords:
Online algorithms
Competitive analysis
Energy
Peak shaving
Simulation

ABSTRACT

We introduce a new problem inspired by energy pricing schemes in which a client is billed for peak usage. At each timeslot the system meets an energy demand through a combination of a new request, an unreliable amount of *free source* energy (e.g. solar or wind power), and previously received energy. The added piece of infrastructure is the battery, which can store surplus energy for future use, and is initially assumed to be perfectly efficient or lossless. In a feasible solution, each demand must be supplied on time, through a combination of newly requested energy, energy withdrawn from the battery, and free source. The goal is to minimize the maximum request. In the online version of this problem, the algorithm must determine each request without knowledge of future demands or free source availability, with the goal of maximizing the amount by which the peak is reduced. We give efficient optimal algorithms for the offline problem, with and without a bounded battery. We also show how to find the optimal offline battery size, given the requirement that the final battery level equals the initial battery level. Finally, we give efficient H_n-competitive algorithms assuming the peak *effective* demand is revealed in advance, and provide matching lower bounds.

Later, we consider the setting of lossy batteries, which lose to conversion inefficiency a constant fraction of any amount charged (e.g. 33%). We efficiently adapt our algorithms to this setting, maintaining optimality for offline and (we conjecture) maintaining competitiveness for online. We give *factor-revealing* LPs, which provide some quasi-empirical evidence for competitiveness. Finally, we evaluate these and other, heuristic algorithms on real and synthetic data.

© 2011 Elsevier Inc. All rights reserved.

1. Introduction

There is increasing interest in saving fuel costs by use of renewable energy sources such as wind and solar power. Although such sources are highly desirable, and the power they provide is in a sense free, the typical disadvantage is unreliability: availability depends e.g. on weather conditions (it is not "dispatchable" on demand). Many companies seek to build efficient systems to gather such energy when available and store it, perhaps in modified form, for future use [3].

On the other hand, power companies charge some high-consumption clients not just for the total amount of power consumed, but also for how quickly they consume it. Within the billing period (typically a month), the client is charged for the amount of energy used (*usage charge*, in kWh) and for the maximum amount requested over time (*peak charge*, in kW).[1] If demands are given as a sequence $(d_1, d_2, \ldots, d_n)$, then the total bill is of the form $c_1 \sum_i d_i + c_2 \max_i \{d_i\}$ (for some constants $c_1, c_2 > 0$), i.e., a weighted sum of the total usage and the maximum usage. (In practice, the discrete timeslots may be 30 min averages [4].) This means that a client who powers a 100 kW piece of machinery for one hour and then uses no more energy for the rest of the month would be charged more than a client who uses a total of 100 kWh spread evenly over the course of the month. Since the per-unit cost for peak charges may be on the order of 100 times the per-unit cost for total usage [5],[2] this difference can be significant. Indeed, this is borne out in our experiments.

This suggests a potential financial incentive to storing purchased energy for future use. Indeed, at least one start-up company has marketed such a battery-based system intended to reduce peak energy charges. In such a system, a battery is placed between the power company and a high-consumption client site (such as a large office building or factory) in order to smooth power requests and shave the peak. The client site will charge to the battery when demand is low and discharge when demand is high. Spikes in the

☆ A preliminary version of this work was presented in [1,2].
* Corresponding author. Tel.: +1 917 648 9096.
E-mail address: mpjohnson@gmail.com (M.P. Johnson).
[1] In fact, some billing models are more complex.

[2] The Orlando Utilities Commission website [5], for example, quotes rates of 8.388 cents per kWh ("energy charge") and $6.50 per kW ("demand charge").

2210-5379/$ – see front matter © 2011 Elsevier Inc. All rights reserved.
doi:10.1016/j.suscom.2011.05.001

2012 IEEE International Conference on Pervasive Computing and Communications, Lugano (19-23 March 2012)

SmartCap: Flattening Peak Electricity Demand in Smart Homes

Sean Barker, Aditya Mishra, David Irwin, Prashant Shenoy, and Jeannie Albrecht[†]
University of Massachusetts Amherst
[†]*Williams College*

***Abstract*—Flattening household electricity demand reduces generation costs, since costs are disproportionately affected by peak demands. While the vast majority of household electrical loads are interactive and have little scheduling flexibility (TVs, microwaves, etc.), a substantial fraction of home energy use derives from background loads with some, albeit limited, flexibility. Examples of such devices include A/Cs, refrigerators, and dehumidifiers. In this paper, we study the extent to which a home is able to transparently flatten its electricity demand by scheduling only background loads with such flexibility. We propose a Least Slack First (LSF) scheduling algorithm for household loads, inspired by the well-known Earliest Deadline First algorithm. We then integrate the algorithm into SmartCap, a system we have built for monitoring and controlling electric loads in homes. To evaluate LSF, we collected power data at outlets, panels, and switches from a real home for 82 days. We use this data to drive simulations, as well as experiment with a real testbed implementation that uses similar background loads as our home. Our results indicate that LSF is most useful during peak usage periods that exhibit "peaky" behavior, where power deviates frequently and significantly from the average. For example, LSF decreases the average deviation from the mean power by over 20% across all 4-hour periods where the deviation is at least 400 watts.**

I. Introduction

Recent studies indicate that residential and commercial buildings account for over 75% of electricity consumption in the United States [2]. As a result, designing new "green" buildings and retrofitting existing buildings with green technologies has become both an important research challenge and societal need. In the residential sector, many techniques exist to reduce either a home's energy footprint or its energy bill. For instance, smart buildings may use motion sensors to track occupants and opportunistically disconnect loads[1] in empty rooms [11]. Alternatively, consumers may participate in automated demand response programs increasingly offered by electric utilities, which automatically turn off home appliances when the demand for electricity is high [10]. These intelligent load management schemes reduce a home's energy footprint and its bill by automatically disconnecting loads from power when necessary or convenient. This paper focuses on an intelligent load management scheme for flattening household electricity usage or demand.

Flattening demand implies reducing the difference between the peaks and troughs in a home's electricity usage, thereby creating a flatter usage pattern that lessens the deviation from the average usage. Demand flattening has the potential to benefit residential consumers as the electric grid becomes smarter and more efficient, since peak demands have a disproportionate affect on grid capital and operational costs, including transmission, generation, and fuel costs. For instance, demand flattening significantly reduces transmission and distribution losses, which account for nearly half (47%) of residential energy consumption [3], since these losses are proportional to the square of current.

To incentivize demand flattening, utilities are transitioning from flat pricing models to variable time-of-use or peak-load models [4], [5], [8], [17]. Since the marginal cost to generate electricity rises as demand increases, utilities are beginning to add surcharges to bills based on a consumer's peak usage. For example, a utility may determine the bill, in part, based on a customer's largest half-hour of electricity demand within a day, regardless of the total day's energy consumption. The new electricity pricing models provide consumers strong incentives to regulate not only their total energy consumption, but also their consumption profile. In particular, these new pricing models incentivize customers to lower their peak consumption by flattening their usage.

Unfortunately, while conceptually simple—to control its demand, a home need only decide when to disconnect its loads—intelligent load management has proven difficult to implement in practice. One reason is that disconnecting loads requires active consumer involvement during peak periods, such as turning off unnecessary lights, programming a thermostat, or postponing washing clothes. Prior studies have shown that compelling consumers to change their household routines is challenging [9]. While providing occupants real-time feedback of their power consumption may initially incentivize them to reduce their usage, once the novelty wears off occupants typically revert to their previous habits. Even for consumers that wish to actively manage their load, choosing which loads to disconnect and when is a complex decision that must be continuously re-evaluated based on information that is constantly changing. To address the problem, we have designed *SmartCap*, a system for automatically monitoring and controlling household loads.

As a key step in SmartCap's design, this paper studies the extent to which homes are able to flatten their home electricity demand without affecting home occupants or requiring their active involvement. We explore the impact of modifying *background* electrical loads that are completely transparent to home occupants and have no impact on their perceived comfort. While the vast majority of electrical loads in homes are *interactive* and have little scheduling flexibility

[1] We use the term load throughout the paper to refer to any appliance or device in the home that draws electricity.

978-1-4673-0258-6/12/$31.00 ©2012 IEEE

Energy peak shaving with local storage

MATTHEW P. JOHNSON, AMOTZ BAR-NOY, OU LIU, YI FENG

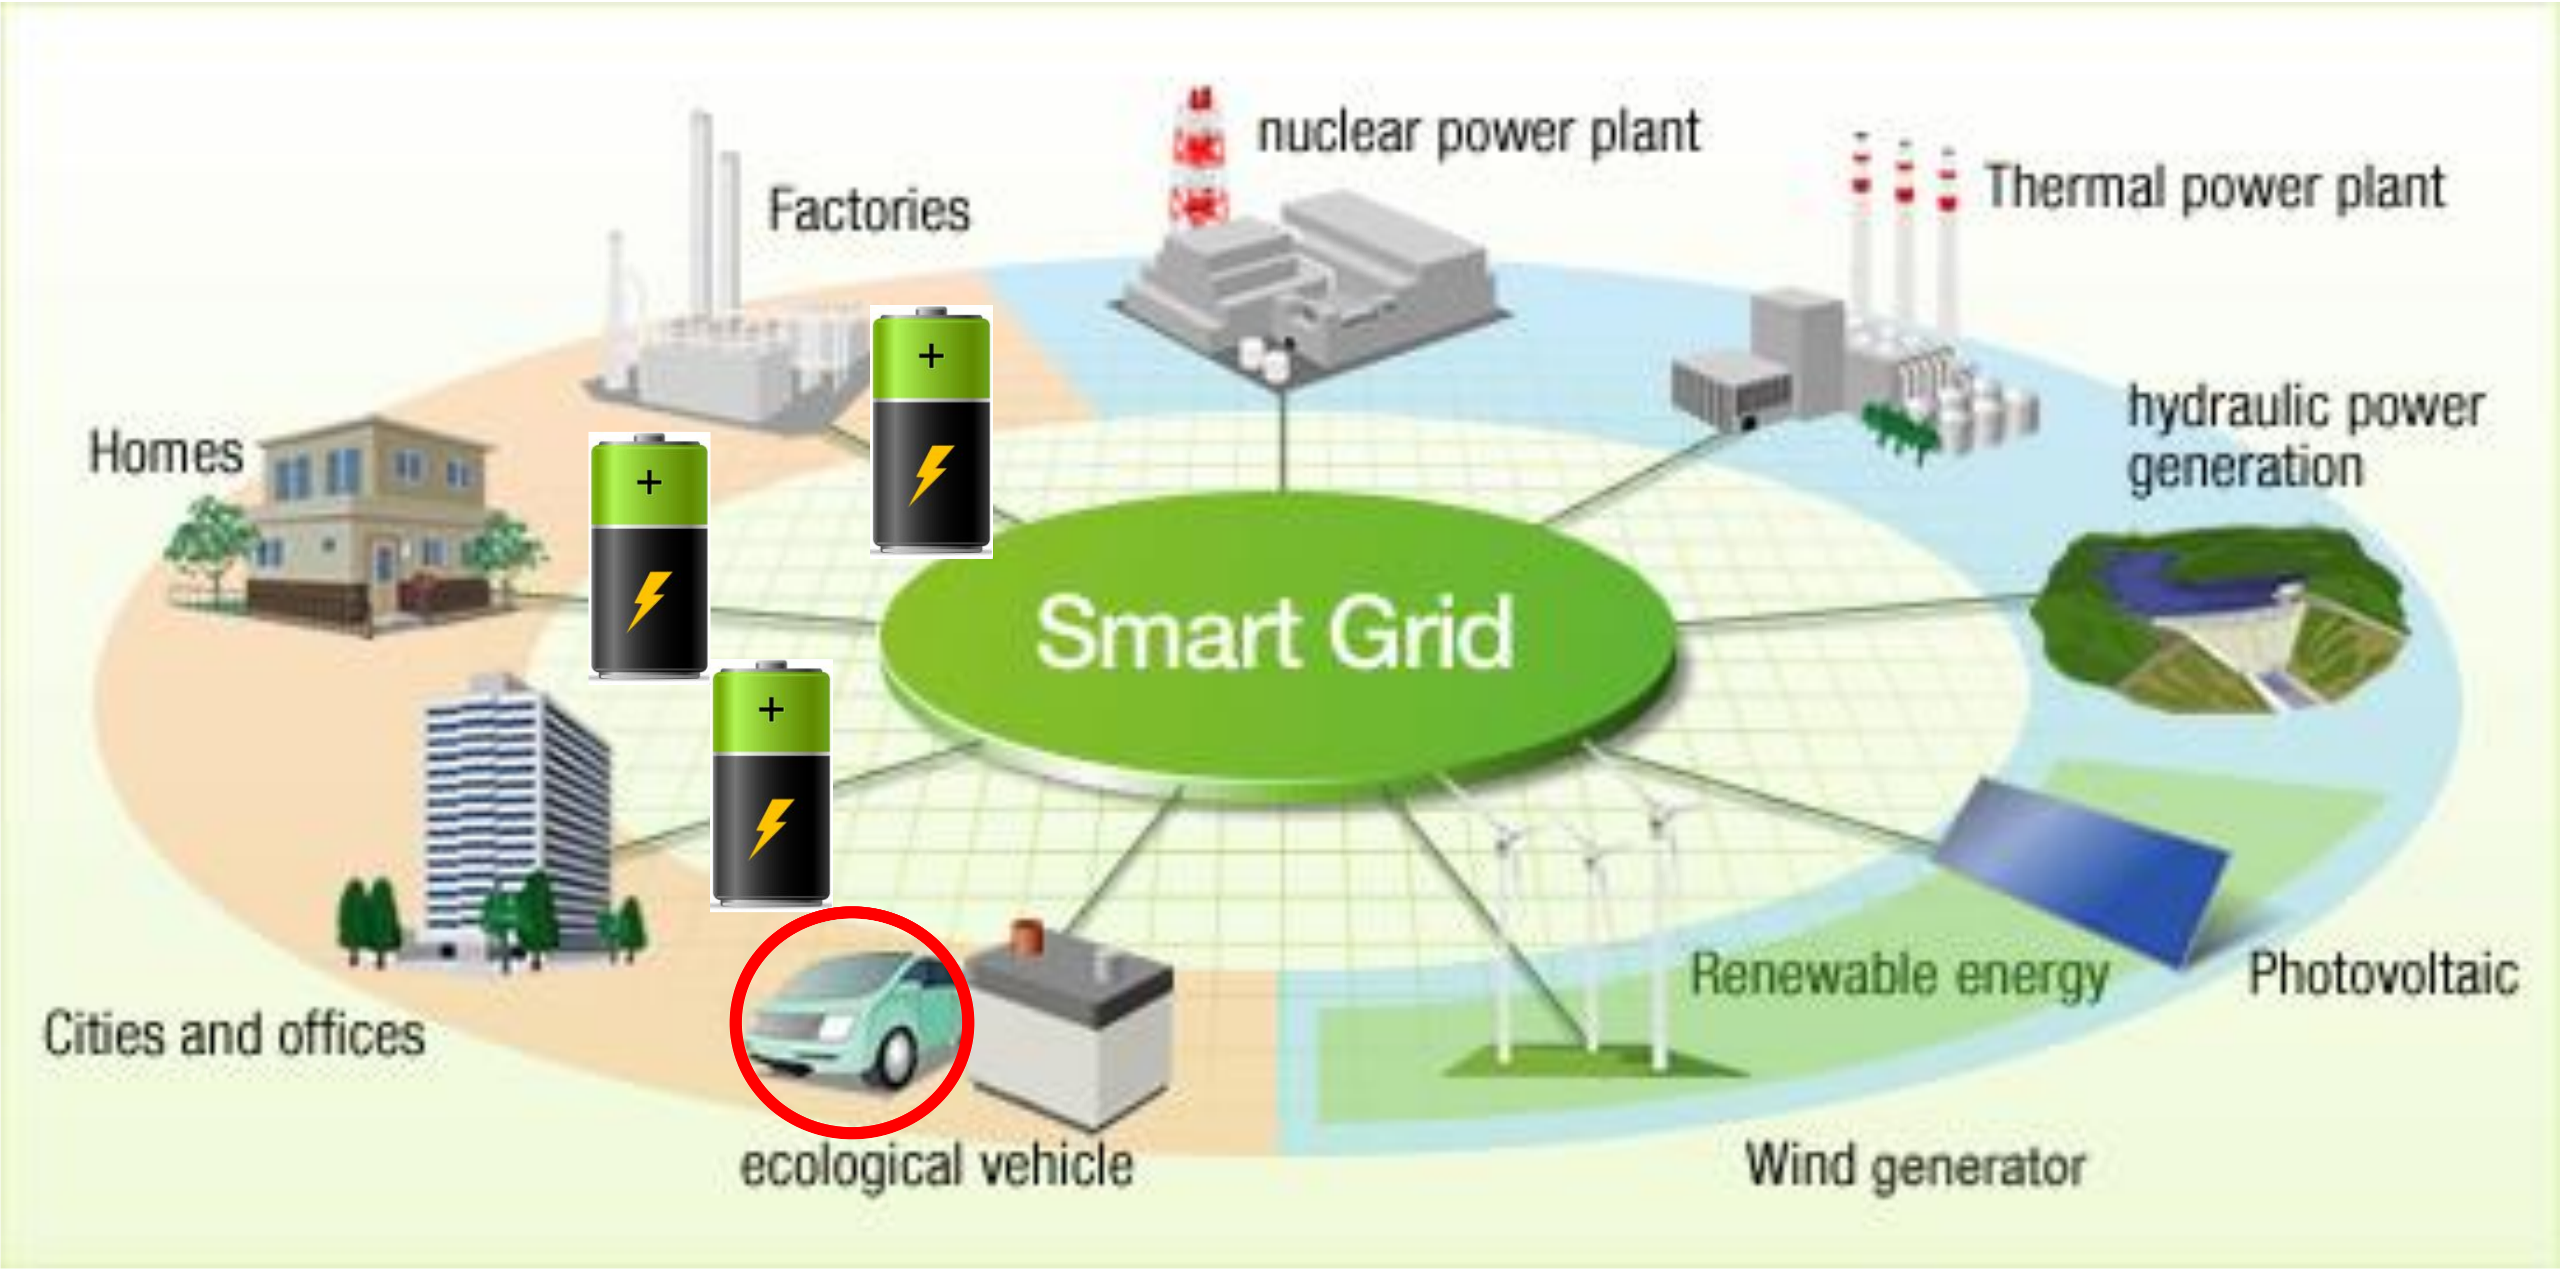
nuclear power plant
Thermal power plant
Factories
Homes
hydraulic power generation
Smart Grid
Cities and offices
Renewable energy
Photovoltaic
ecological vehicle
Wind generator

Difficulties:

- Energy charge and discharge algorithms

Energy charge and discharge algorithms

- System modes (threshold):
  - Request exactly the demand
  - Request more then demand and charge the battery
  - Request less then demand and discharge the battery
- Overflow and Underflow
  - Underflow is not allowed in the system
- Online and Offline algorithms

Energy charge and discharge algorithms

Alg.	Battery	Online	Threshold T_i	Running time
1.b	Bounded	No	$\hat{\mu}(1, n)$	$O(n^2 \log n)$
2.a	Bounded	Yes	$D - \frac{D-\hat{\mu}(1,i)}{H_n}$	$O(n^2 \log n)$
2.b	bounded	Yes	$D - \frac{D-\mu(s_i,i)}{H_{(n-s_i+1)}}$	$O(n \log n)$

Difficulties:

- Energy charge and discharge algorithms
- Battery size

Battery size

- A large part of the initial system cost is the battery’s capacity
- A completely flat request curve is possible with a large battery
- Knowledge of maximum demand is needed
- Goal is to find the smallest battery that can achieve the optimal peak request

Batte

- A larg ... ity
- A cor ... tery
- Know
- Goal ... timal peak

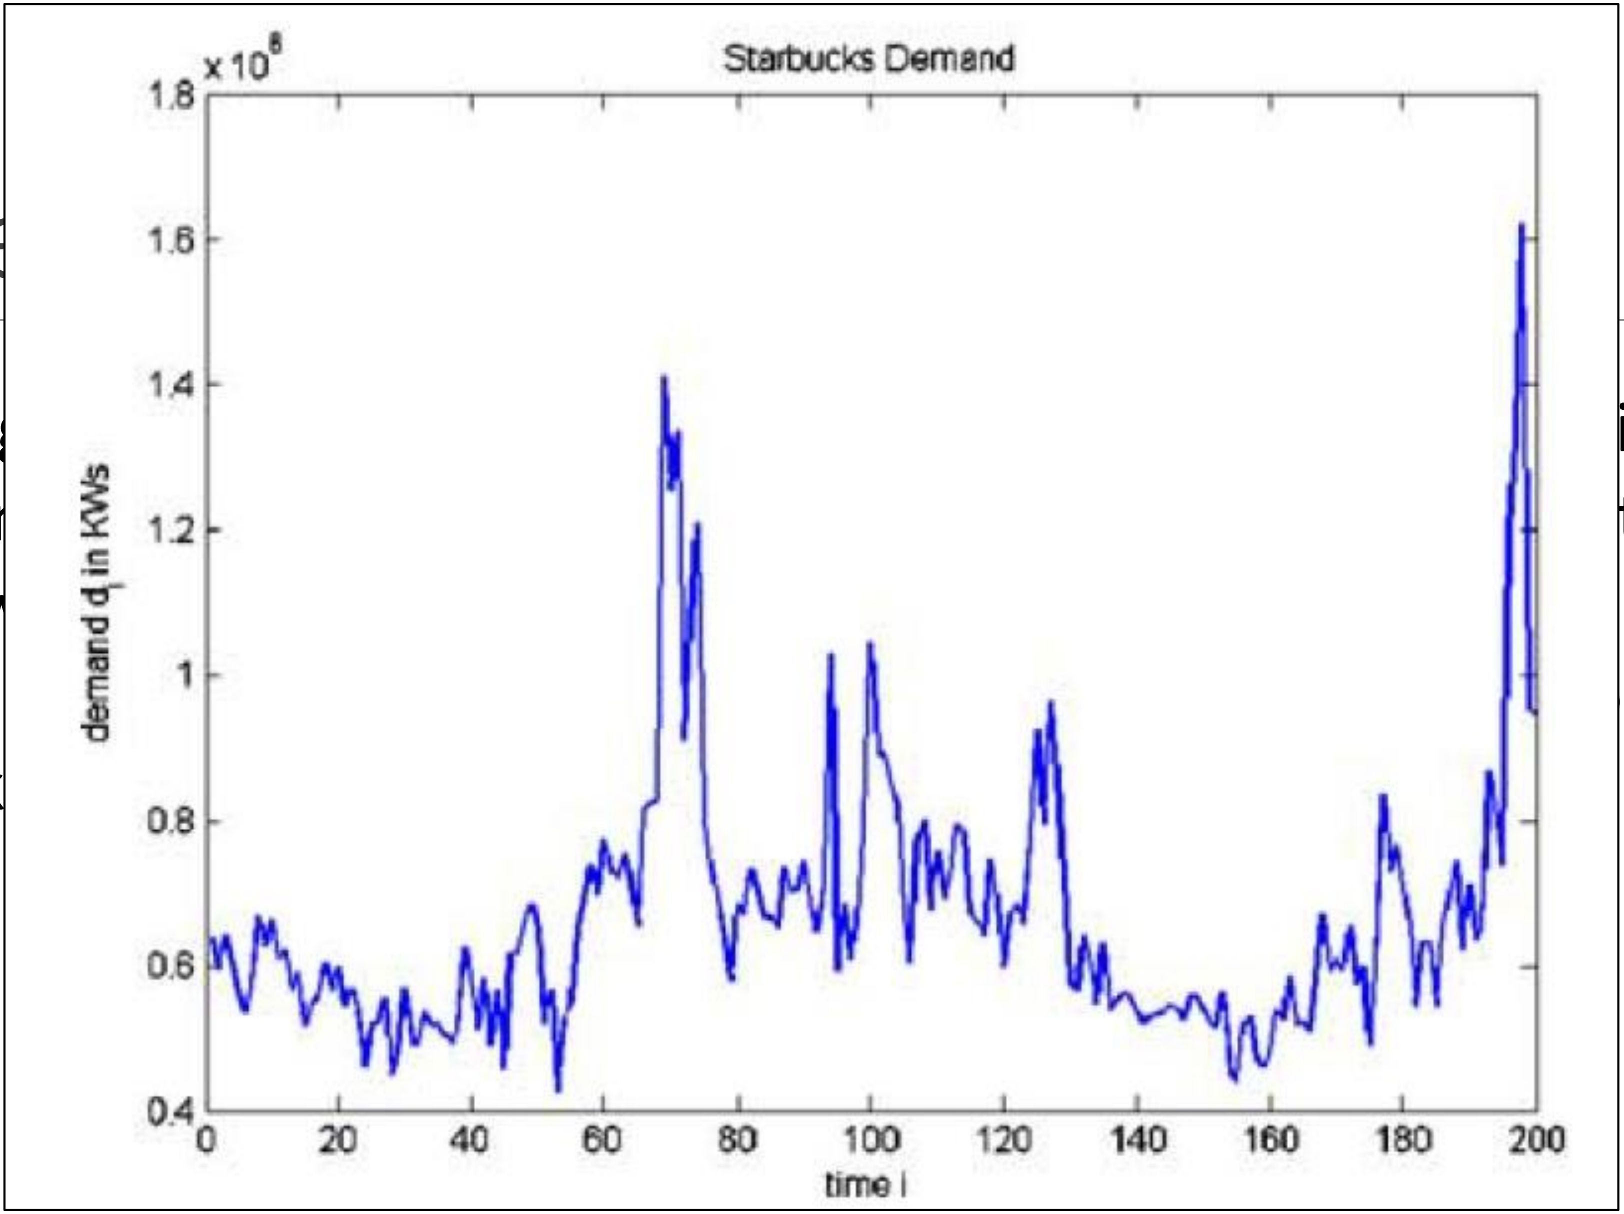

Batte

- A larg ity
- A cor tery
- Know
- Goal timal peak

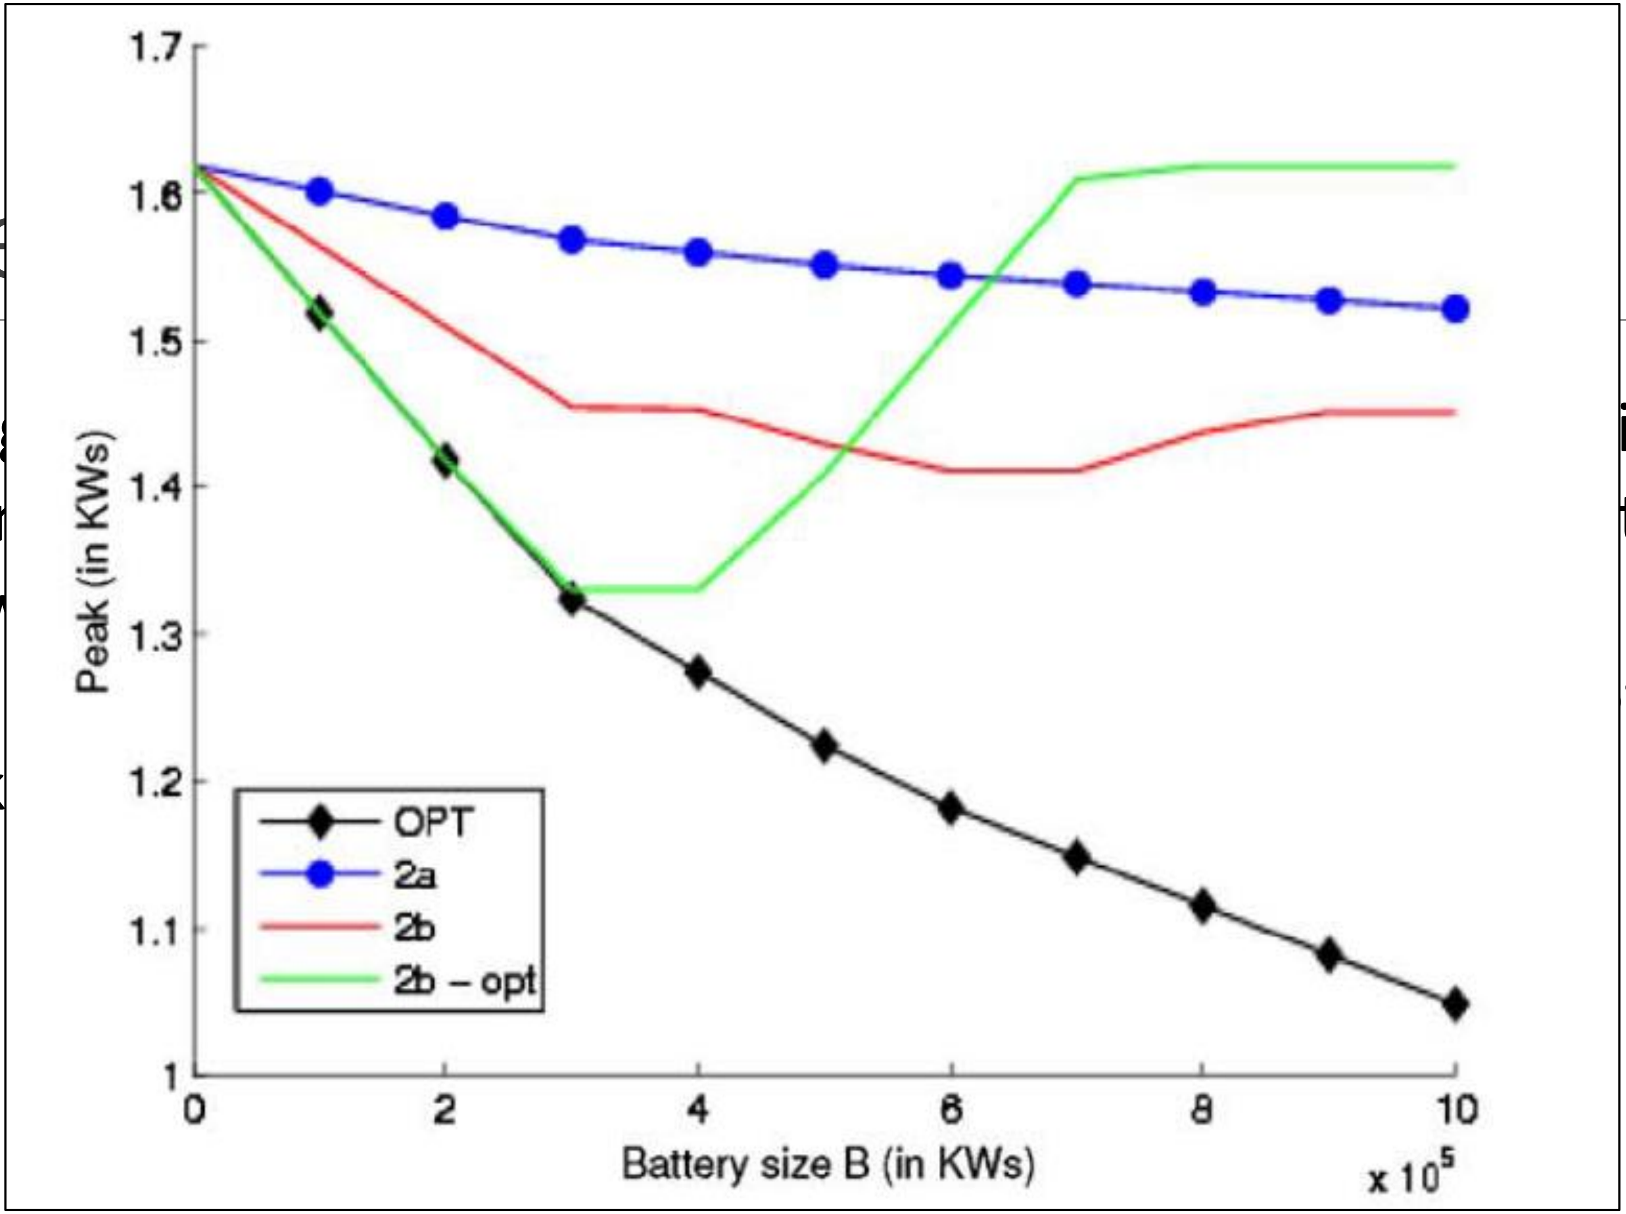

Difficulties:

- Energy charge and discharge algorithms
- Battery size
- Lossy batteries

Lossy batteries

- AC/DC transforming suffers some energy losses
- The bigger the battery, the greater the losses

Lossy b

- AC/DC
- The big

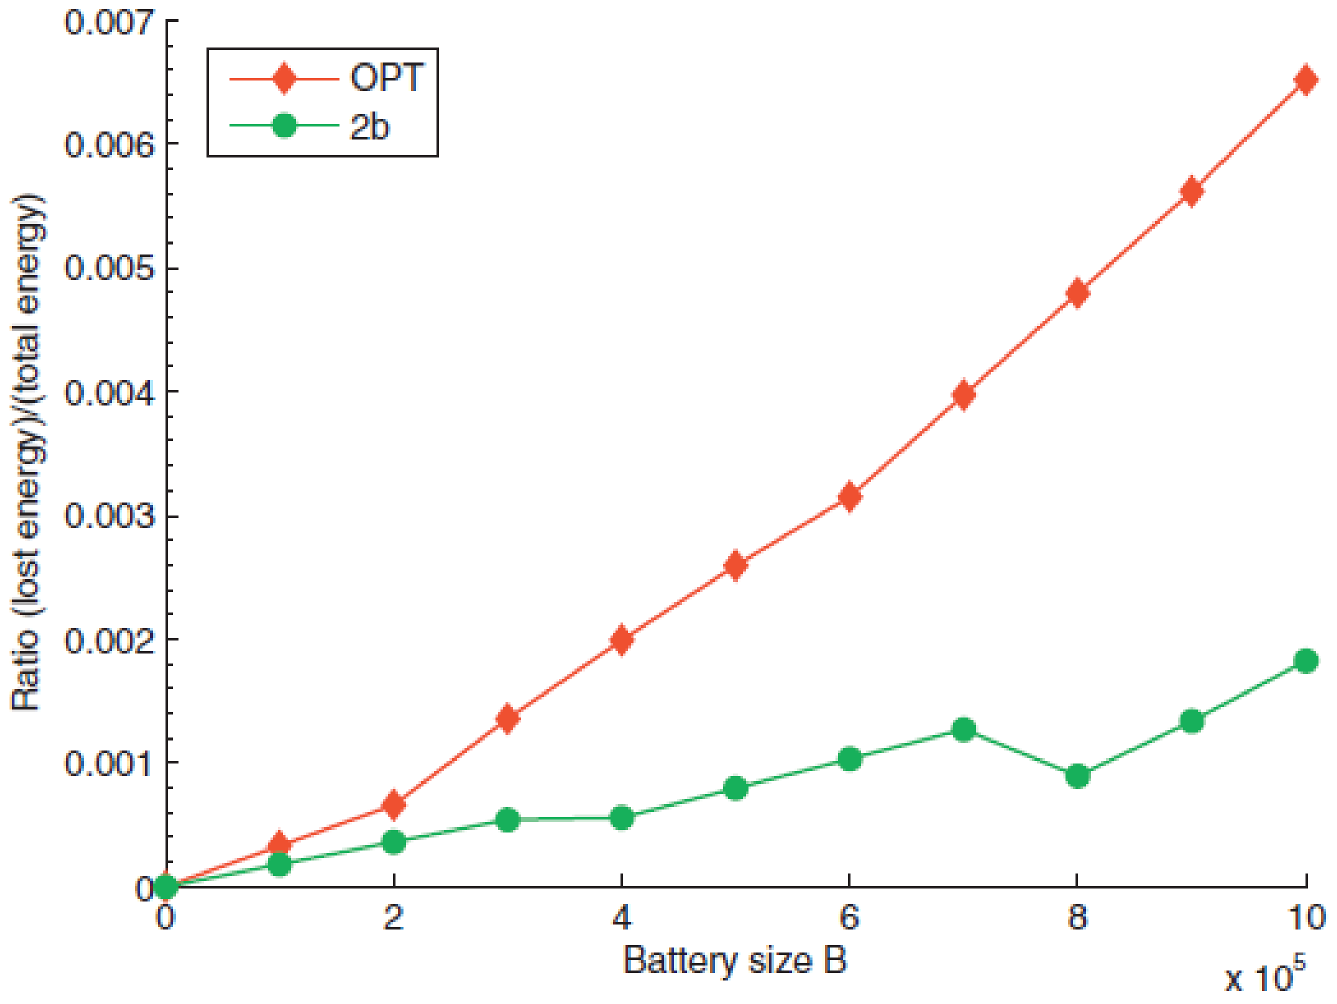

(b) Varying B; $L = 0.33$.

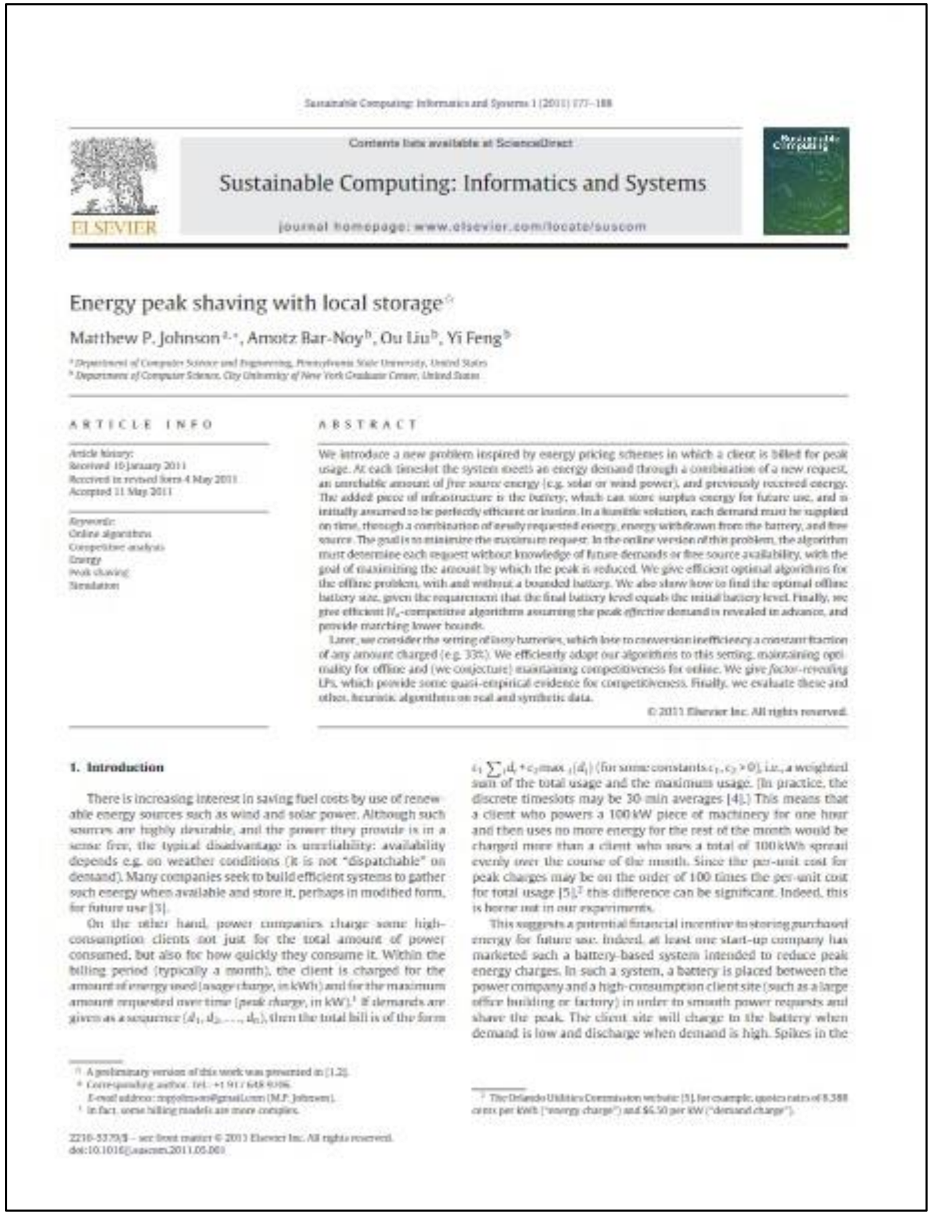

Sustainable Computing: Informatics and Systems 1 (2011) 177–188

Contents lists available at ScienceDirect

Sustainable Computing: Informatics and Systems

journal homepage: www.elsevier.com/locate/suscom

Energy peak shaving with local storage☆

Matthew P. Johnson[a,*], Amotz Bar-Noy[b], Ou Liu[b], Yi Feng[b]

[a] Department of Computer Science and Engineering, Pennsylvania State University, United States
[b] Department of Computer Science, City University of New York Graduate Center, United States

ARTICLE INFO

Article history:
Received 10 January 2011
Received in revised form 4 May 2011
Accepted 11 May 2011

Keywords:
Online algorithms
Competitive analysis
Energy
Peak shaving
Simulation

ABSTRACT

We introduce a new problem inspired by energy pricing schemes in which a client is billed for peak usage. At each timeslot the system meets an energy demand through a combination of a new request, an unreliable amount of *free source* energy (e.g. solar or wind power), and previously received energy. The added piece of infrastructure is the battery, which can store surplus energy for future use, and is initially assumed to be perfectly efficient or lossless. In a feasible solution, each demand must be supplied on time, through a combination of newly requested energy, energy withdrawn from the battery, and free source. The goal is to minimize the maximum request. In the online version of this problem, the algorithm must determine each request without knowledge of future demands or free source availability, with the goal of maximizing the amount by which the peak is reduced. We give efficient optimal algorithms for the offline problem, with and without a bounded battery. We also show how to find the optimal offline battery size, given the requirement that the final battery level equals the initial battery level. Finally, we give efficient H_n-competitive algorithms assuming the peak *effective* demand is revealed in advance, and provide matching lower bounds.

Later, we consider the setting of lossy batteries, which lose to conversion inefficiency a constant fraction of any amount charged (e.g. 33%). We efficiently adapt our algorithms to this setting, maintaining optimality for offline and (we conjecture) maintaining competitiveness for online. We give *factor-revealing* LPs, which provide some quasi-empirical evidence for competitiveness. Finally, we evaluate these and other, heuristic algorithms on real and synthetic data.

© 2011 Elsevier Inc. All rights reserved.

1. Introduction

There is increasing interest in saving fuel costs by use of renewable energy sources such as wind and solar power. Although such sources are highly desirable, and the power they provide is in a sense free, the typical disadvantage is unreliability: availability depends e.g. on weather conditions (it is not "dispatchable" on demand). Many companies seek to build efficient systems to gather such energy when available and store it, perhaps in modified form, for future use [3].

On the other hand, power companies charge some high-consumption clients not just for the total amount of power consumed, but also for how quickly they consume it. Within the billing period (typically a month), the client is charged for the amount of energy used (*usage charge*, in kWh) and for the maximum amount requested over time (*peak charge*, in kW).[1] If demands are given as a sequence $(d_1, d_2, \ldots, d_n)$, then the total bill is of the form $c_1 \sum_i d_i + c_2 \max_i \{d_i\}$ (for some constants $c_1, c_2 > 0$), i.e., a weighted sum of the total usage and the maximum usage. (In practice, the discrete timeslots may be 30-min averages [4].) This means that a client who powers a 100 kW piece of machinery for one hour and then uses no more energy for the rest of the month would be charged more than a client who uses a total of 100 kWh spread evenly over the course of the month. Since the per-unit cost for peak charges may be on the order of 100 times the per-unit cost for total usage [5],[2] this difference can be significant. Indeed, this is borne out in our experiments.

This suggests a potential financial incentive to storing purchased energy for future use. Indeed, at least one start-up company has marketed such a battery-based system intended to reduce peak energy charges. In such a system, a battery is placed between the power company and a high-consumption client site (such as a large office building or factory) in order to smooth power requests and shave the peak. The client site will charge to the battery when demand is low and discharge when demand is high. Spikes in the

☆ A preliminary version of this work was presented in [1,2].
* Corresponding author. Tel.: +1 917 648 9396.
E-mail address: mpjohnson@gmail.com (M.P. Johnson).
[1] In fact, some billing models are more complex.
[2] The Orlando Utilities Commission website [5], for example, quotes rates of 8.388 cents per kWh ("energy charge") and $6.50 per kW ("demand charge").

2210-5379/$ – see front matter © 2011 Elsevier Inc. All rights reserved.
doi:10.1016/j.suscom.2011.05.001

2012 IEEE International Conference on Pervasive Computing and Communications, Lugano (19-23 March 2012)

SmartCap: Flattening Peak Electricity Demand in Smart Homes

Sean Barker, Aditya Mishra, David Irwin, Prashant Shenoy, and Jeannie Albrecht†
University of Massachusetts Amherst
†*Williams College*

***Abstract*—Flattening household electricity demand reduces generation costs, since costs are disproportionately affected by peak demands. While the vast majority of household electrical loads are interactive and have little scheduling flexibility (TVs, microwaves, etc.), a substantial fraction of home energy use derives from background loads with some, albeit limited, flexibility. Examples of such devices include A/Cs, refrigerators, and dehumidifiers. In this paper, we study the extent to which a home is able to transparently flatten its electricity demand by scheduling only background loads with such flexibility. We propose a Least Slack First (LSF) scheduling algorithm for household loads, inspired by the well-known Earliest Deadline First algorithm. We then integrate the algorithm into SmartCap, a system we have built for monitoring and controlling electric loads in homes. To evaluate LSF, we collected power data at outlets, panels, and switches from a real home for 82 days. We use this data to drive simulations, as well as experiment with a real testbed implementation that uses similar background loads as our home. Our results indicate that LSF is most useful during peak usage periods that exhibit "peaky" behavior, where power deviates frequently and significantly from the average. For example, LSF decreases the average deviation from the mean power by over 20% across all 4-hour periods where the deviation is at least 400 watts.**

I. Introduction

Recent studies indicate that residential and commercial buildings account for over 75% of electricity consumption in the United States [2]. As a result, designing new "green" buildings and retrofitting existing buildings with green technologies has become both an important research challenge and societal need. In the residential sector, many techniques exist to reduce either a home's energy footprint or its energy bill. For instance, smart buildings may use motion sensors to track occupants and opportunistically disconnect *loads*[1] in empty rooms [11]. Alternatively, consumers may participate in automated demand response programs increasingly offered by electric utilities, which automatically turn off home appliances when the demand for electricity is high [10]. These intelligent load management schemes reduce a home's energy footprint and its bill by automatically disconnecting loads from power when necessary or convenient. This paper focuses on an intelligent load management scheme for flattening household electricity usage or demand.

Flattening demand implies reducing the difference between the peaks and troughs in a home's electricity usage, thereby creating a flatter usage pattern that lessens the deviation from the average usage. Demand flattening has the potential to benefit residential consumers as the electric grid becomes smarter and more efficient, since peak demands have a disproportionate affect on grid capital and operational costs, including transmission, generation, and fuel costs. For instance, demand flattening significantly reduces transmission and distribution losses, which account for nearly half (47%) of residential energy consumption [3], since these losses are proportional to the square of current.

To incentivize demand flattening, utilities are transitioning from flat pricing models to variable time-of-use or peak-load models [4], [5], [8], [17]. Since the marginal cost to generate electricity rises as demand increases, utilities are beginning to add surcharges to bills based on a consumer's peak usage. For example, a utility may determine the bill, in part, based on a customer's largest half-hour of electricity demand within a day, regardless of the total day's energy consumption. The new electricity pricing models provide consumers strong incentives to regulate not only their total energy consumption, but also their consumption profile. In particular, these new pricing models incentivize customers to lower their peak consumption by flattening their usage.

Unfortunately, while conceptually simple—to control its demand, a home need only decide when to disconnect its loads—intelligent load management has proven difficult to implement in practice. One reason is that disconnecting loads requires active consumer involvement during peak periods, such as turning off unnecessary lights, programming a thermostat, or postponing washing clothes. Prior studies have shown that compelling consumers to change their household routines is challenging [9]. While providing occupants real-time feedback of their power consumption may initially incentivize them to reduce their usage, once the novelty wears off occupants typically revert to their previous habits. Even for consumers that wish to actively manage their load, choosing which loads to disconnect and when is a complex decision that must be continuously re-evaluated based on information that is constantly changing. To address the problem, we have designed *SmartCap*, a system for automatically monitoring and controlling household loads.

As a key step in SmartCap's design, this paper studies the extent to which homes are able to flatten their home electricity demand without affecting home occupants or requiring their active involvement. We explore the impact of modifying *background* electrical loads that are completely transparent to home occupants and have no impact on their perceived comfort. While the vast majority of electrical loads in homes are *interactive* and have little scheduling flexibility

[1] We use the term load throughout the paper to refer to any appliance or device in the home that draws electricity.

978-1-4673-0258-6/12/$31.00 ©2012 IEEE

Sustainable Computing: Informatics and Systems 1 (2011) 177–188

Contents lists available at ScienceDirect

Sustainable Computing: Informatics and Systems

journal homepage: www.elsevier.com/locate/suscom

Energy peak shaving with local storage[☆]

Matthew P. Johnson[a,*], Amotz Bar-Noy[b], Ou Liu[b], Yi Feng[b]

[a] Department of Computer Science and Engineering, Pennsylvania State University, United States
[b] Department of Computer Science, City University of New York Graduate Center, United States

ARTICLE INFO

Article history:
Received 10 January 2011
Received in revised form 4 May 2011
Accepted 11 May 2011

Keywords:
Online algorithms
Competitive analysis
Energy
Peak shaving
Simulation

ABSTRACT

We introduce a new problem inspired by energy pricing schemes in which a client is billed for peak usage. At each timeslot the system meets an energy demand through a combination of a new request, an unreliable amount of *free source* energy (e.g. solar or wind power), and previously received energy. The added piece of infrastructure is the battery, which can store surplus energy for future use, and is initially assumed to be perfectly efficient or lossless. In a feasible solution, each demand must be supplied on time, through a combination of newly requested energy, energy withdrawn from the battery, and free source. The goal is to minimize the maximum request. In the online version of this problem, the algorithm must determine each request without knowledge of future demands or free source availability, with the goal of maximizing the amount by which the peak is reduced. We give efficient optimal algorithms for the offline problem, with and without a bounded battery. We also show how to find the optimal offline battery size, given the requirement that the final battery level equals the initial battery level. Finally, we give efficient H_n-competitive algorithms assuming the peak effective demand is revealed in advance, and provide matching lower bounds.

Later, we consider the setting of lossy batteries, which lose to conversion inefficiency a constant fraction of any amount charged (e.g. 33%). We efficiently adapt our algorithms to this setting, maintaining optimality for offline and (we conjecture) maintaining competitiveness for online. We give *factor-revealing LPs*, which provide some quasi-empirical evidence for competitiveness. Finally, we evaluate these and other, heuristic algorithms on real and synthetic data.

© 2011 Elsevier Inc. All rights reserved.

1. Introduction

There is increasing interest in saving fuel costs by use of renewable energy sources such as wind and solar power. Although such sources are highly desirable, and the power they provide is in a sense free, the typical disadvantage is unreliability: availability depends e.g. on weather conditions (it is not "dispatchable" on demand). Many companies seek to build efficient systems to gather such energy when available and store it, perhaps in modified form, for future use [3].

On the other hand, power companies charge some high-consumption clients not just for the total amount of power consumed, but also for how quickly they consume it. Within the billing period (typically a month), the client is charged for the amount of energy used (*usage charge*, in kWh) and for the maximum amount requested over time (*peak charge*, in kW).[1] If demands are given as a sequence $(d_1, d_2, \ldots, d_n)$, then the total bill is of the form $c_1 \sum_i d_i + c_2 \max_i \{d_i\}$ (for some constants $c_1, c_2 > 0$), i.e., a weighted sum of the total usage and the maximum usage. (In practice, the discrete timeslots may be 30 min averages [4].) This means that a client who powers a 100 kW piece of machinery for one hour and then uses no more energy for the rest of the month would be charged more than a client who uses a total of 100 kWh spread evenly over the course of the month. Since the per-unit cost for peak charges may be on the order of 100 times the per-unit cost for total usage [5],[2] this difference can be significant. Indeed, this is borne out in our experiments.

This suggests a potential financial incentive to storing purchased energy for future use. Indeed, at least one start-up company has marketed such a battery-based system intended to reduce peak energy charges. In such a system, a battery is placed between the power company and a high-consumption client site (such as a large office building or factory) in order to smooth power requests and shave the peak. The client site will charge to the battery when demand is low and discharge when demand is high. Spikes in the

[☆] A preliminary version of this work was presented in [1,2].
[*] Corresponding author. Tel.: +1 917 648 9396.
E-mail address: mpjohnson@gmail.com (M.P. Johnson).
[1] In fact, some billing models are more complex.

[2] The Orlando Utilities Commission website [5], for example, quotes rates of 8.388 cents per kWh ("energy charge") and $6.50 per kW ("demand charge").

2210-5379/$ – see front matter © 2011 Elsevier Inc. All rights reserved.
doi:10.1016/j.suscom.2011.05.001

2012 IEEE International Conference on Pervasive Computing and Communications, Lugano (19-23 March 2012)

SmartCap: Flattening Peak Electricity Demand in Smart Homes

Sean Barker, Aditya Mishra, David Irwin, Prashant Shenoy, and Jeannie Albrecht[†]
University of Massachusetts Amherst
[†] *Williams College*

Abstract—Flattening household electricity demand reduces generation costs, since costs are disproportionately affected by peak demands. While the vast majority of household electrical loads are interactive and have little scheduling flexibility (TVs, microwaves, etc.), a substantial fraction of home energy use derives from background loads with some, albeit limited, flexibility. Examples of such devices include A/Cs, refrigerators, and dehumidifiers. In this paper, we study the extent to which a home is able to transparently flatten its electricity demand by scheduling only background loads with such flexibility. We propose a Least Slack First (LSF) scheduling algorithm for household loads, inspired by the well-known Earliest Deadline First algorithm. We then integrate the algorithm into SmartCap, a system we have built for monitoring and controlling electric loads in homes. To evaluate LSF, we collected power data at outlets, panels, and switches from a real home for 82 days. We use this data to drive simulations, as well as experiment with a real testbed implementation that uses similar background loads as our home. Our results indicate that LSF is most useful during peak usage periods that exhibit "peaky" behavior, where power deviates frequently and significantly from the average. For example, LSF decreases the average deviation from the mean power by over 20% across all 4-hour periods where the deviation is at least 400 watts.

I. Introduction

Recent studies indicate that residential and commercial buildings account for over 75% of electricity consumption in the United States [2]. As a result, designing new "green" buildings and retrofitting existing buildings with green technologies has become both an important research challenge and societal need. In the residential sector, many techniques exist to reduce either a home's energy footprint or its energy bill. For instance, smart buildings may use motion sensors to track occupants and opportunistically disconnect *loads*[1] in empty rooms [11]. Alternatively, consumers may participate in automated demand response programs increasingly offered by electric utilities, which automatically turn off home appliances when the demand for electricity is high [10]. These intelligent load management schemes reduce a home's energy footprint and its bill by automatically disconnecting loads from power when necessary or convenient. This paper focuses on an intelligent load management scheme for flattening household electricity usage or demand.

Flattening demand implies reducing the difference between the peaks and troughs in a home's electricity usage, thereby creating a flatter usage pattern that lessens the deviation from the average usage. Demand flattening has the potential to benefit residential consumers as the electric grid becomes smarter and more efficient, since peak demands have a disproportionate affect on grid capital and operational costs, including transmission, generation, and fuel costs. For instance, demand flattening significantly reduces transmission and distribution losses, which account for nearly half (47%) of residential energy consumption [3], since these losses are proportional to the square of current.

To incentivize demand flattening, utilities are transitioning from flat pricing models to variable time-of-use or peak-load models [4], [5], [8], [17]. Since the marginal cost to generate electricity rises as demand increases, utilities are beginning to add surcharges to bills based on a consumer's peak usage. For example, a utility may determine the bill, in part, based on a customer's largest half-hour of electricity demand within a day, regardless of the total day's energy consumption. The new electricity pricing models provide consumers strong incentives to regulate not only their total energy consumption, but also their consumption profile. In particular, these new pricing models incentivize customers to lower their peak consumption by flattening their usage.

Unfortunately, while conceptually simple—to control its demand, a home need only decide when to disconnect its loads—intelligent load management has proven difficult to implement in practice. One reason is that disconnecting loads requires active consumer involvement during peak periods, such as turning off unnecessary lights, programming a thermostat, or postponing washing clothes. Prior studies have shown that compelling consumers to change their household routines is challenging [9]. While providing occupants real-time feedback of their power consumption may initially incentivize them to reduce their usage, once the novelty wears off occupants typically revert to their previous habits. Even for consumers that wish to actively manage their load, choosing which loads to disconnect and when is a complex decision that must be continuously re-evaluated based on information that is constantly changing. To address the problem, we have designed *SmartCap*, a system for automatically monitoring and controlling household loads.

As a key step in SmartCap's design, this paper studies the extent to which homes are able to flatten their home electricity demand without affecting home occupants or requiring their active involvement. We explore the impact of modifying *background* electrical loads that are completely transparent to home occupants and have no impact on their perceived comfort. While the vast majority of electrical loads in homes are *interactive* and have little scheduling flexibility

[1] We use the term load throughout the paper to refer to any appliance or device in the home that draws electricity.

978-1-4673-0258-6/12/$31.00 ©2012 IEEE

SmartCap: Flattening Peak Electricity Demand in Smart Homes

SEAN BARKER, ADITYA MISHRA, DAVID IRWIN, PRASHANT SHENOY, JEANNIE ALBRECHT

SmartCap, A system that:

- Gathers information from:
    - Sensors
    - Energy prices
    - Grid information
    - Household consumption data
- Goal: Schedule devices intelligently in order to lower peak demand without affecting house occupants

Background loads

- Example of background loads:
  - Refrigerator
  - Freezer
  - HRV (Heat Recovery Ventilation)
  - Dehumidifier
  - A/C
  - Battery chargers
- Accounts for 59% of total energy consumption in this household
- Interactive loads such as lights, TV, PC, Microwave, etc is not controlled by SmartCap

LSF – Least Slack First

- Adopted from the EDF policy used in Real-time systems
- Responsible for turning devices on and off
- Includes a threshold to prevent to many devices scheduled at the same time
- Devices have min and max guard band boundary
- Devices close to max boundary get high priority

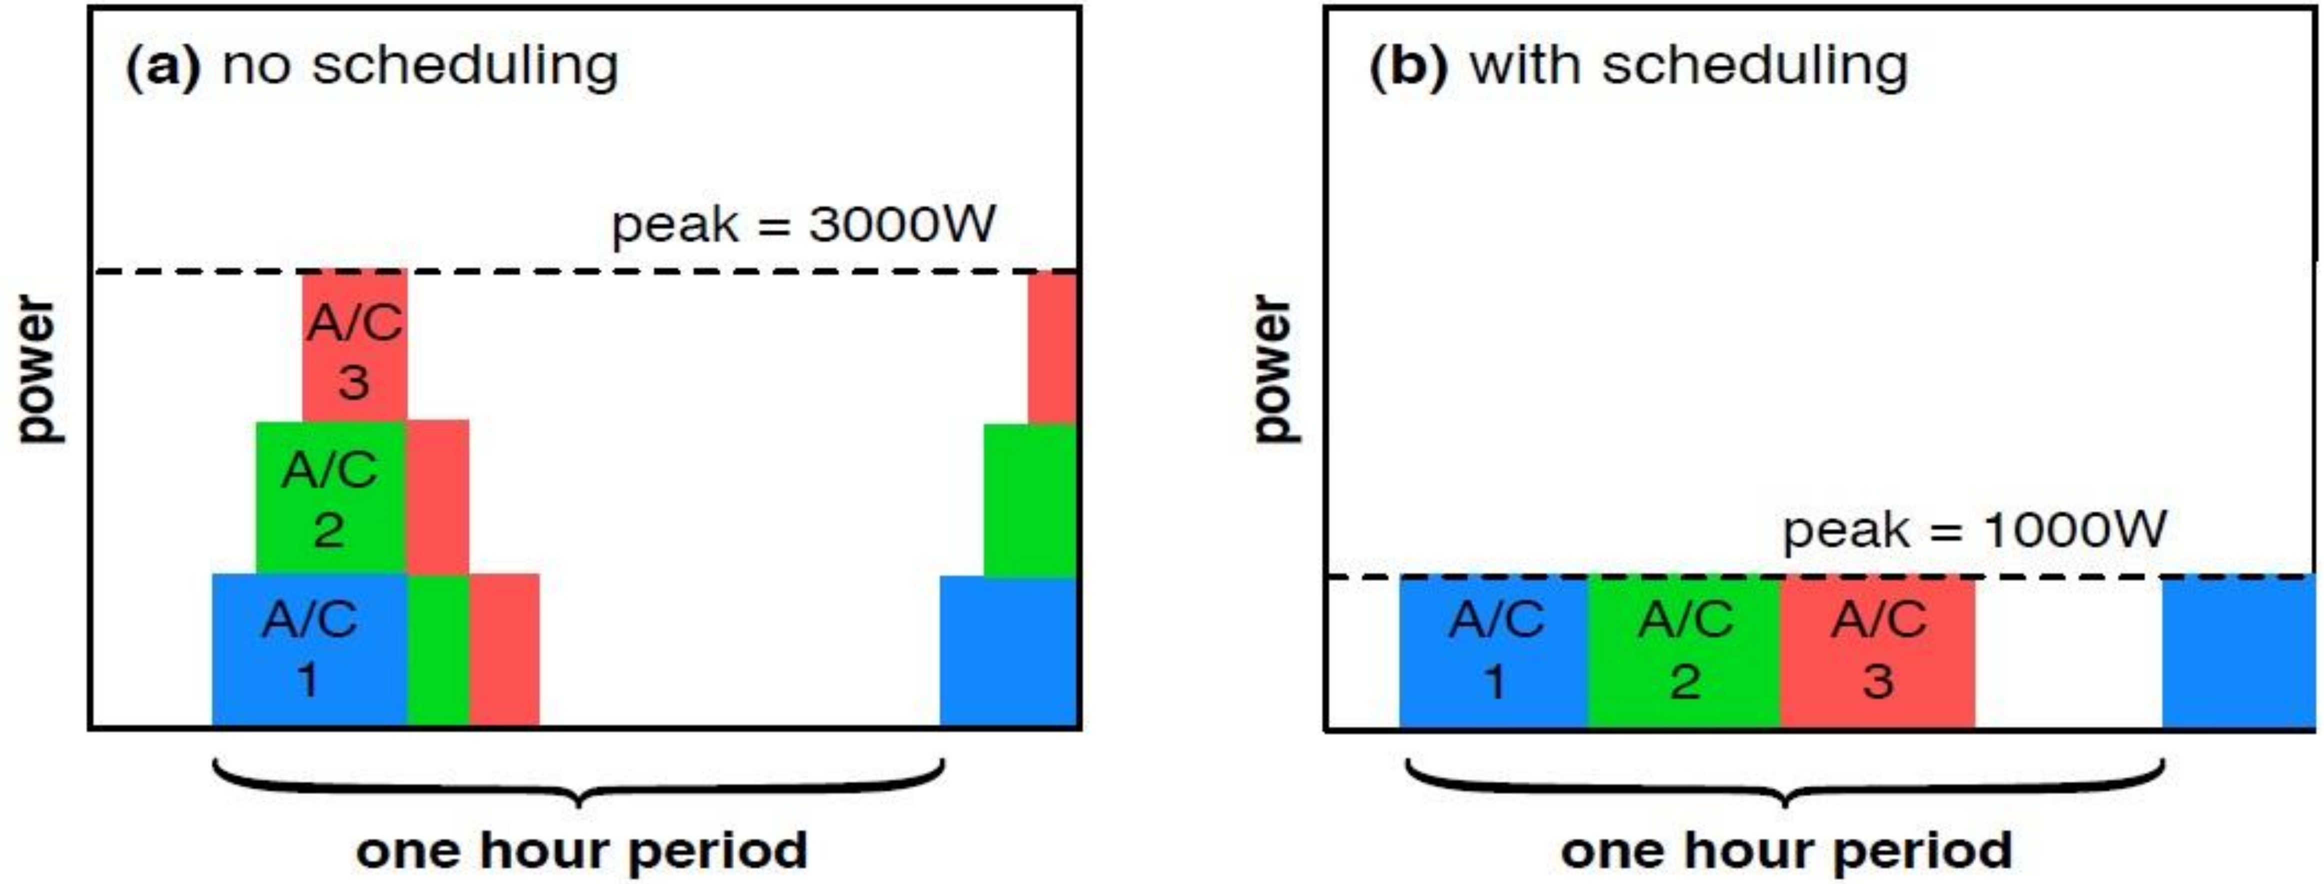
(a) no scheduling
peak = 3000W
power
A/C 3
A/C 2
A/C 1
one hour period
(b) with scheduling
peak = 1000W
power
A/C 1
A/C 2
A/C 3
one hour period

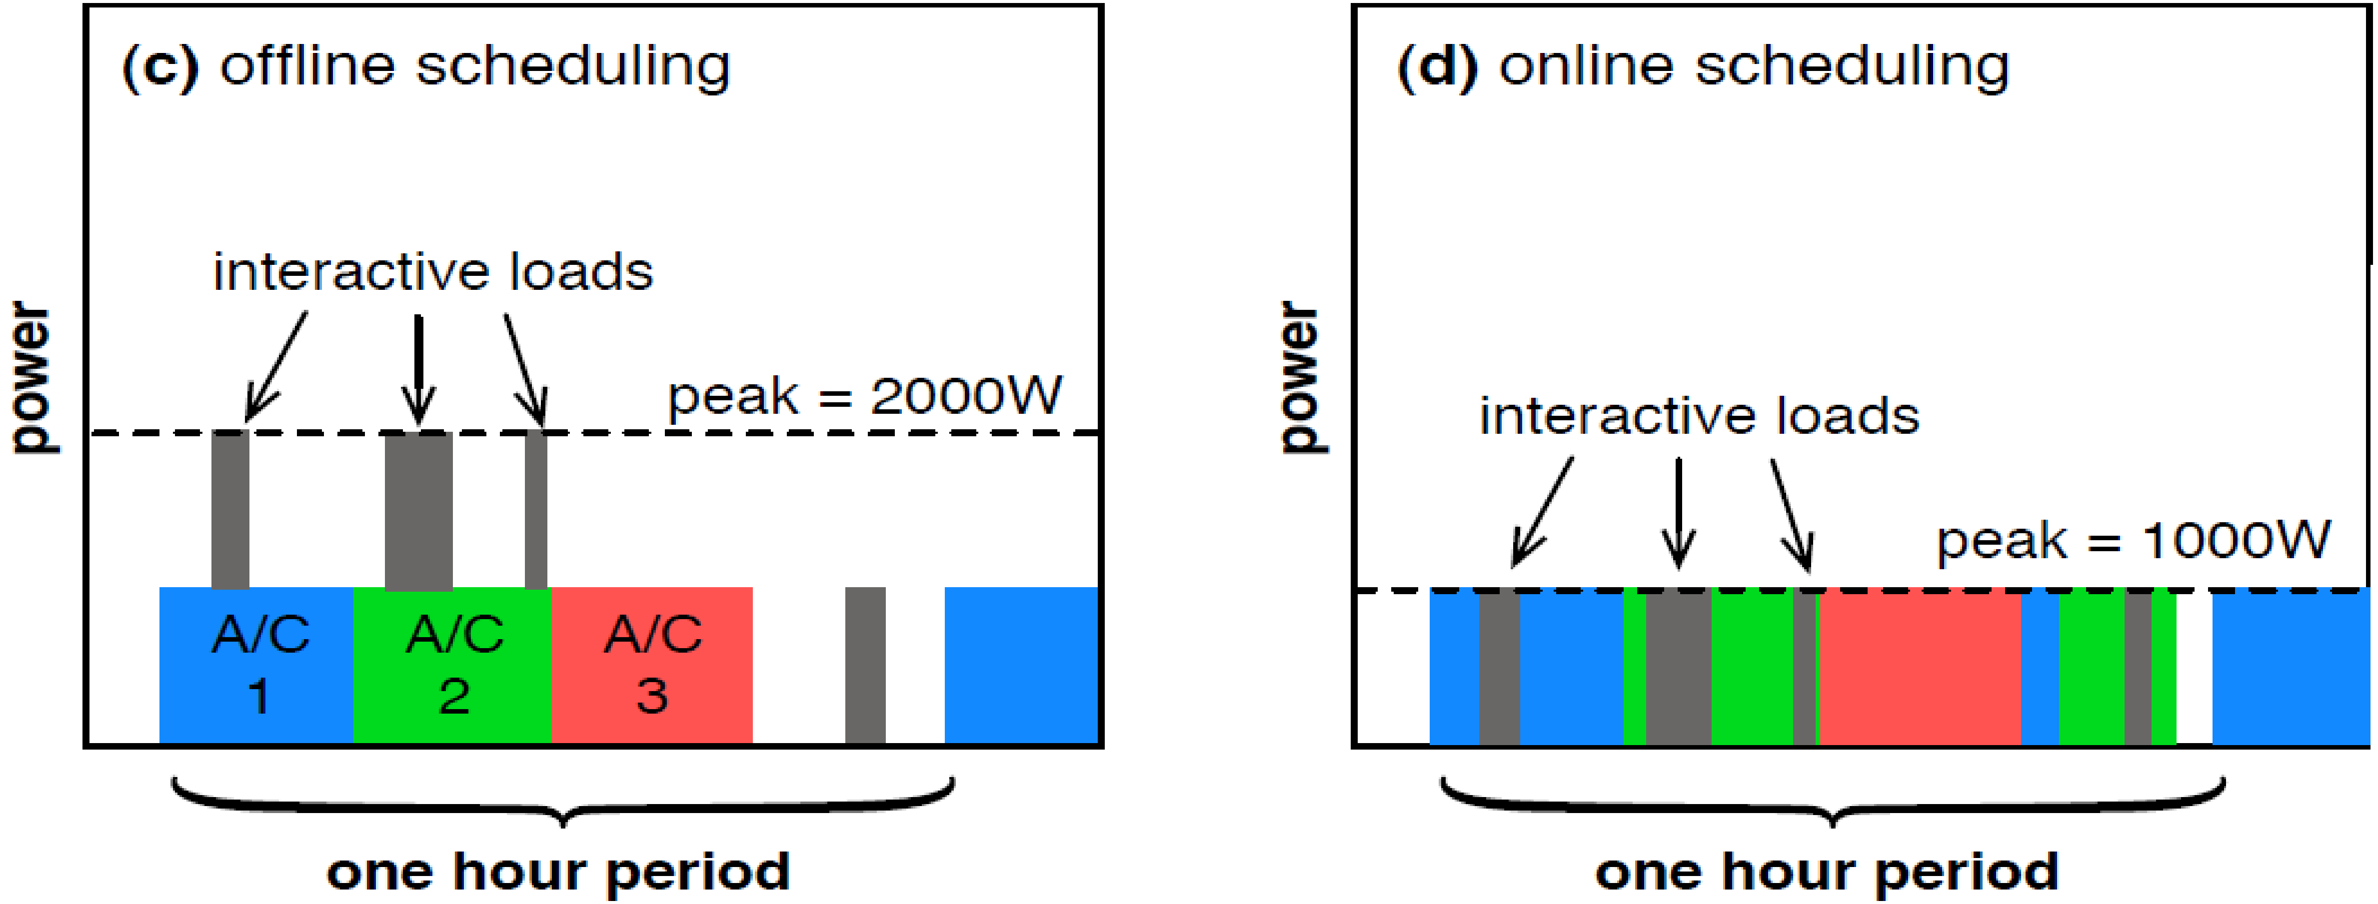
(c) offline scheduling
power
interactive loads
peak = 2000W
A/C 1
A/C 2
A/C 3
one hour period
(d) online scheduling
power
interactive loads
peak = 1000W
one hour period

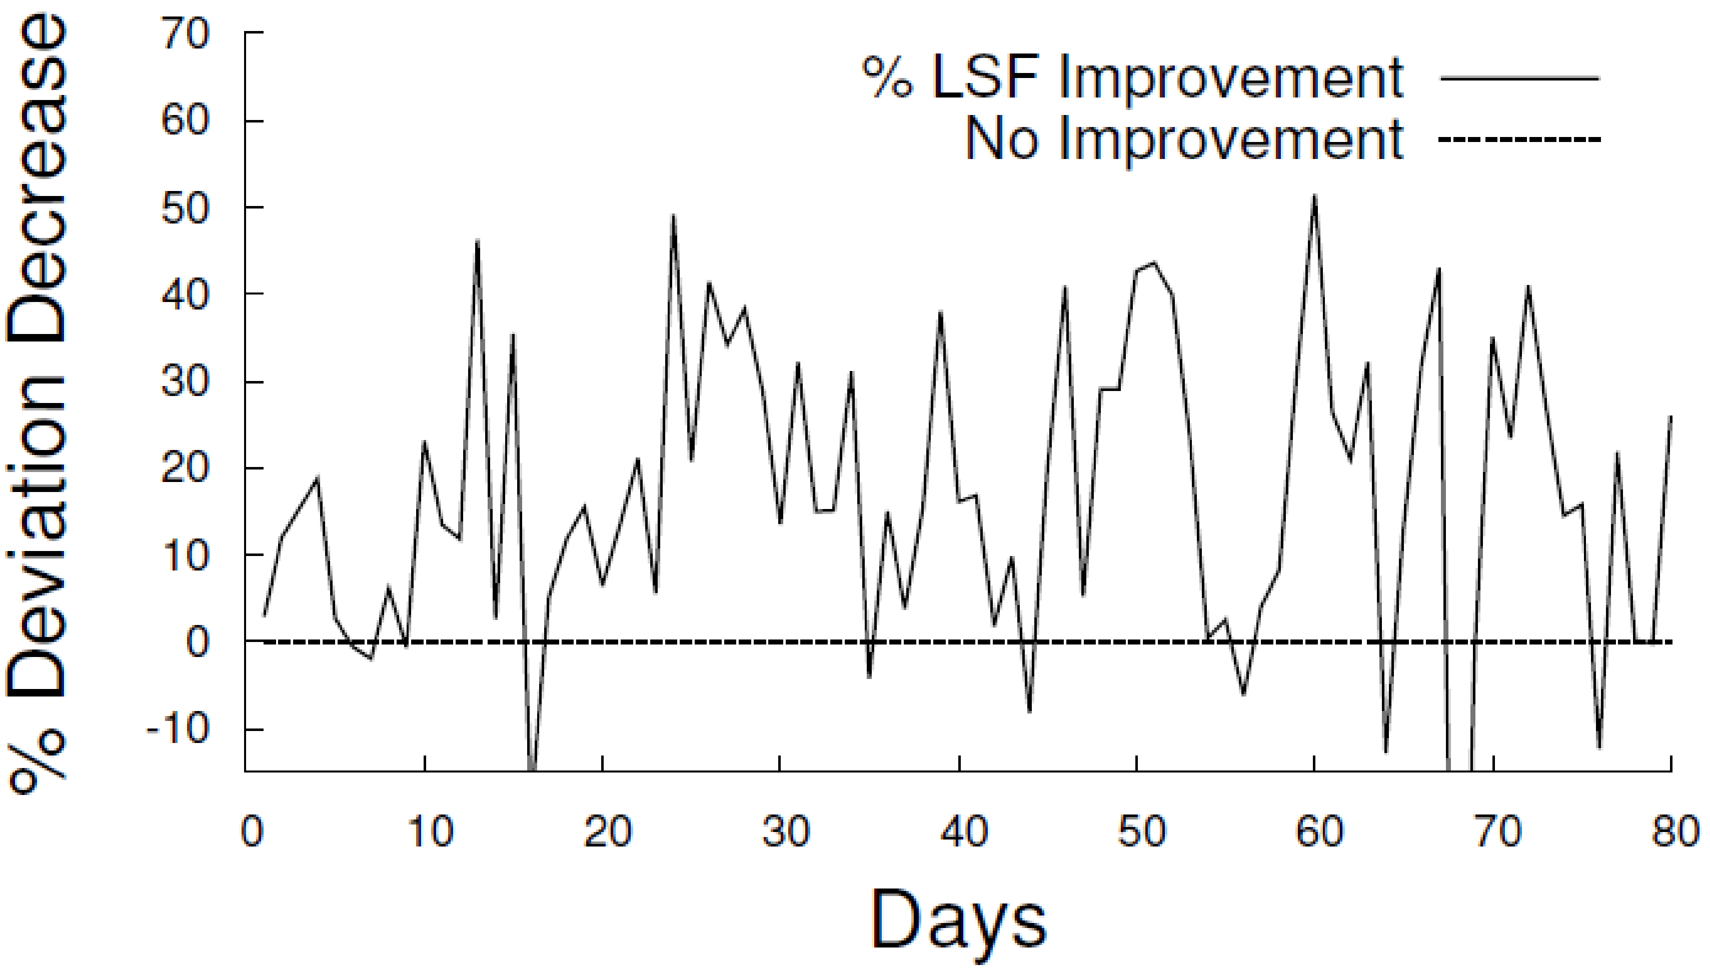

- Profile flattened on over 91% of the days
- Resulting in 16 % flatter profile on average.

Earth Hour city of Guelph, Kanada 2013

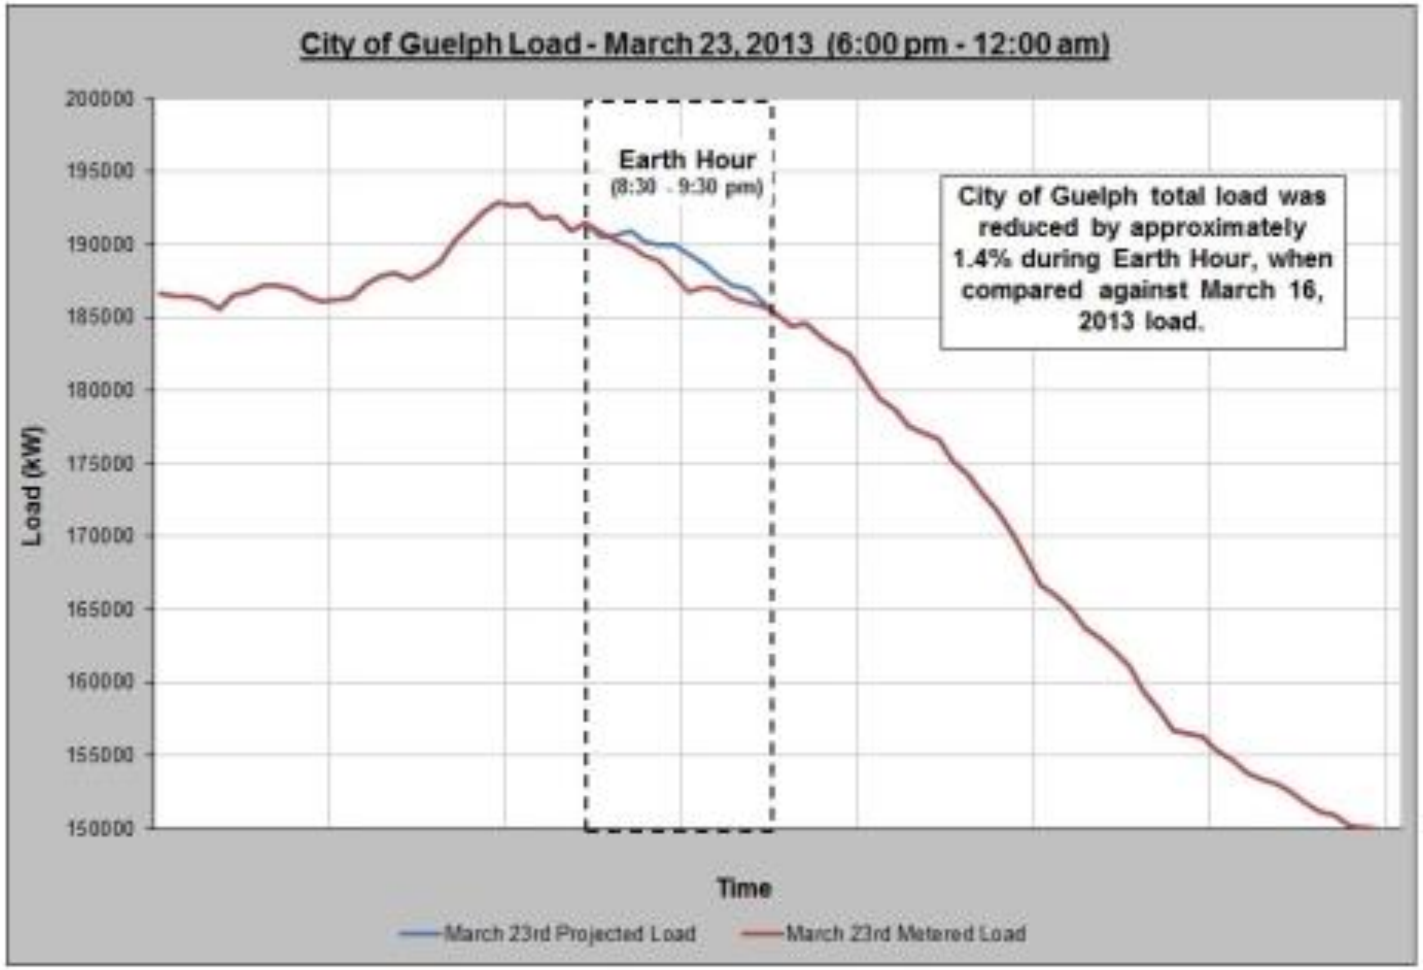

- 2,6 MWh drop was noticed

Thank you!

Any questions?