Policy-Driven Distributed and Collaborative Demand Response in Multi-Domain Commercial Buildings

Archan Misra
Advanced Technology Solutions
Telcordia Technologies
1 Telcordia Drive, Piscataway, NJ USA

Henning Schulzrinne
Department of Computer Science
Columbia University
1214 Amsterdam Avenue, New York, NY, USA

ABSTRACT

Enabling a sophisticated Demand Response (DR) framework, whereby individual consumers adapt their electricity consumption in response to price variations, is a major objective of the emerging Smart Grid. We first point out why the current model, of EMS-based centralized control of a static repository of high load appliances, is inappropriate for supporting DR in future commercial buildings and campuses, where the consuming appliances are controlled by multiple users. To enable DR in such multi-domain environments, we envision a more collaborative and autonomous model, where a large set of heterogeneous smart electrical devices autonomously self-organize and negotiate their collective DR. Enabling this vision requires the development of new networking primitives for autonomic, hierarchical overlay formation, new energy profiles that can represent aggregate characteristics of groups of devices and new hierarchical distributed optimization techniques.

Categories and Subject Descriptors

J.7 [**Computer Applications**]: Command and Control

General Terms

Design, Management

1. INTRODUCTION

The Smart Grid seeks to embed greater intelligence in the generation, distribution and consumption of electricity by developing an information infrastructure that supports:

- Dynamic pricing of electrical power that considers the mismatch between energy production and demand, and that integrates both conventional and non-conventional energy producers (such as the micro-grid [1]).
- Dynamic adaptation of power consumption by individual consumers, who adjust the operation of "smart appliances" in response to such pricing signals (e.g., by dynamically setting thermostat thresholds or through off-peak scheduling of pool filtration systems).

While much of the initial pilots and media attention has centered on the installation of Smart Meters in residential environments, we argue that supporting DR in commercial building environments (e.g., in a multi-office highrise or a university campus consisting of multiple departments and buildings) is an equally important goal, with a unique set of challenges. Indeed, [2] reveals that mid-sized commercial buildings consume approx. 20% of the total electricity generated in the US.

Currently, commercial buildings and campuses usually have a facilities manager manually control the power consumption of various networked devices. In response to load control signals that occur over timescales of hours or days, the manager uses an industry-grade EMS (Energy Management System), such as Honeywell EBI™or Carrier i-Vu™, to send load control signals, carried over messaging standards such as LonWorks and BACNet, to various devices, such as elevators or HVAC systems. Projecting the emerging mechanisms of DR, one might argue that supporting automated DR in such environments is simply a matter of augmenting the EMS to accept pricing signals from the electrical utility and then use pre-defined policies to automatically control the high-load devices–e.g., shut down a set of elevators if the price exceeds 40 cents/kWh.

We argue, however, that this form of policy-driven control of individual devices is an inadequate model for the majority of commercial building environments and that several peer-to-peer and hierarchical agent-based network interaction capabilities are needed to make DR a practical reality. We first observe that devices, such as washing machines, water coolers, are progressively becoming 'smarter' and capable of dynamical adjustment of power consumption [3]. Also, note that studies [4] show that (see Figure 1) the 'heavy hitters' (e.g., HVAC) only consume 36% of the energy in a typical commercial building, with a significant percentage of power being consumed by more moderate-load devices, such as computers and lights. Accordingly, to maximize the benefits of DR, it is essential to extend the ability to exert price-driven adjustment of power consumption from the presently small set of high-load devices (such as elevators and HVAC units) to eventually encompass a wider spectrum of medium-load appliances, such as window ACs, water chillers and light controllers. For such 'pervasive' DR adaptation, the current model of EMS-controlled adaptation suffers from two key drawbacks:

- **Centralized Control**: Currently, the EMS gateway is a centralized repository of the power consumption characteristics of all devices and occupant policies (e.g., ice maker is run at 11 pm unless pricing exceeds a threshold) and has direct control privileges over all these devices. However, in reality, commercial buildings have multiple inhabitants (e.g., multiple companies occupying different floors of an office highrise), who effectively define different *domains* of administrative control. For example, the building itself may be owned by a real estate company (which controls the elevator banks) but the smart appliances within the building may be controlled by multiple different organizations or entities, with their own preferences (for example, company A may wish to control the blade servers in its data center or the water coolers in its cafeteria). More precisely, smart appliances in commercial building or campus environments will be associated with a *multiplicity of control domains*, with each domain likely defining DR policies for appliances that it controls. The policy-driven adaptation model of DR must thus be capable of respecting these different scopes of domain control, while ensuring that the appliances collectively adhere to certain overall constraints (e.g., the cumulative consumption of the building stays below 20 kW).

- **Static Device Configuration:** Currently, the EMS exercises control over a relatively small and static set of devices (e.g., elevators and HVAC units) that are manually configured into the EMS device inventory. Such a manual approach to inventory management will, however, become unacceptably onerous when one tries to adapt a much wider range of devices and smart appliances, such as copiers, printers and water coolers! It should be clear that what is required is a set of standardized mechanisms for devices to be automatically discovered (with bare minimal human intervention) and 'co-opted' into the DR infrastructure.

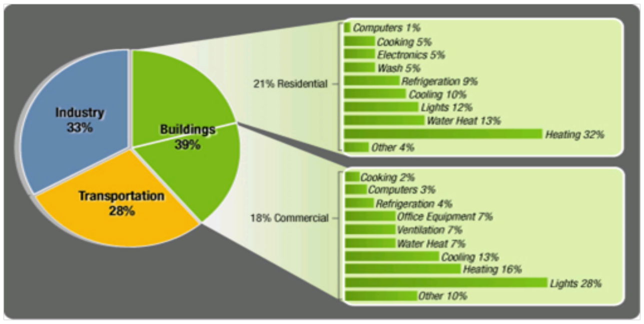

Figure 1: Breakup of Energy Consumption in Residential and Commercial Buildings (US), reproduced from [4]

In this paper, we shall suggest a model of Demand Response Adaptation and Control via Hierarchical, Distributed, Multi-Domain Autonomy (DRACHMA) that appears to be well suited to tackle these challenges. Section 2 will explain the basic DRACHMA model of operation and identify several key functional capabilities required to support such an autonomous model of device discovery and distributed Demand Response. Subsequently, Section 3 will provide a brief outline of the key technical problems and suggest some early thoughts on relevant promising techniques.

2. THE DRACHMA PARADIGM

The DRACHMA paradigm proposes that wide-spread DR in commercial building environments is achieved by having electrical devices communicate among themselves to organize autonomously into appropriate organizational control groups (possibly hierarchical), and negotiate among themselves the most appropriate form of collective DR adaptation. For our purposes, we define a *domain as either a collection of sub-domains or a collection of smart appliances/devices whose DR behavior is specified by a specific administrative entity.* The key issue here is that the contract with the utility (power grid) may have been negotiated by a different domain or operational entity (the university utility office or the office building owner), whereas the policies governing DR should be expressible by finer-grained domains (such as the chemistry department or the individual organizations renting in the commercial building). For example, a bio-chemistry laboratory unit should be able to easily specify that, on a particular Saturday, it is more important to keep the chiller unit active throughout the day (to ensure that a particular experiment can continue uninterrupted) and instead turn off all printers/copiers, while a travel agency office on a different floor should be able to preferentially increase the thermostat by 2°C, instead of turning off printers and water cooling units.

Figure 2 shows a high-level view of the DRACHMA model of pervasive, decentralized DR. Each of the "ovals" represents a different domain (a logical entity that expresses policies for smart devices within its control)–the challenge is to coordinate DR adaptation across all devices and domains within (and eventually, beyond) the building.

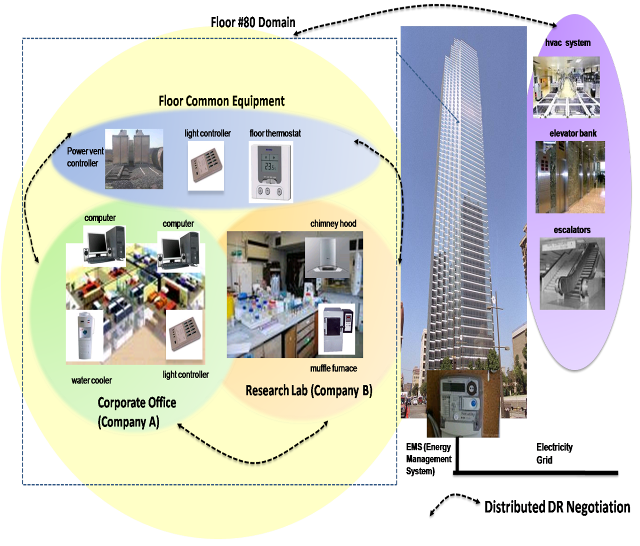

Figure 2: Decentralized DR in DRACHMA. Devices organize themselves into (possibly hierarchical) domains); domains collaboratively negotiate their collective adaptation. For example, the 80^{th} floor domain includes 3 different sub-domains.

A DRACHMA-based operational architecture must thus possess several key capabilities:

1. **Secure, Easy, Technology-Independent, Auto-configuration of individual domains:** Smart appliances should be equipped with a relatively simple set of protocols that allow close to plug-n-play formation of individual autonomous domains. For example, office workers should be able to easily buy off-the-shelf office equipment (e.g., printers, water-coolers) and specify energy adaptation policies for a domain that comprises both such customer-procured equipment, as well as pre-installed premises equipment (e.g., lighting controllers for intra-office corridors). Such autoconfiguration must, of course, conform to appropriate requirements on *security*, so that a neighboring office cannot directly dictate policies for the printers in my office and support *privacy*, so that a neighbor cannot observe how often I run my dishwasher. It is important to distinguish our problem of logical domain formation from the underlying challenges of enabling localized communication among a distributed set of smart appliances. While wireless networking standards (such as ZigBee-based Smart Energy Profile [7]) enable self-configuring, extended-area wireless connectivity among a physical neighborhood of such smart devices, our focus is on enabling subsets of such devices to form logical groupings, on the basis of organizational boundaries, that are independent of specific PHY/LINK layer technologies. For example, a bank's domain (and adaptation policy) may be defined across multiple branches or may span multiple floors of the same building. The auto-configuration mechanisms should also be easy and require only minimal human intervention–e.g., without requiring consumers to enter a common password or PIN on multiple devices to ensure that these devices become part of a single domain.

2. **Robust cross-domain coordination:** Besides enabling a group of devices to autonomously self-organize into a domain, the framework must also be capable of ensuring rapid and easy configuration of domain hierarchies. An open question relates to determining the likely structural nature and granularity of such domain hierarchies (e.g., is peer-to-peer better than hierarchical?) and ensuring that such hierarchies can accommodate proposed novel forms of energy consumers, e.g., allow the load adjustment policies for employee PHEVs to be different from those associated with the server farm of the employees' office.

3. **Distributed DR adaptation:** DRACHMA empowers individual domains (or devices) to establish the DR preferences for their personal devices, e.g., an office can indicate the relative priorities in the operation of its cooling units. However, if the DR signal indicates a cumulative bound on the power consumption by all offices in the building or indicates a price/kW that depends on the cumulative load, all of the individual domains must orchestrate their adaptation, in a decentralized but likely hierarchical fashion, to meet this bound. This approach allows, for example, DR by a hierarchy of domain controllers (representing domains at different granularities), with electrical devices attaching to the leaves of the hierarchy. Ideally, this should be accomplished by a single adaptation framework (and associated signaling protocols) that scales both in terms of numbers of devices being controlled and the hierarchical depth of the autonomous domains. The adaptation mechanism must be capable of adjustable autonomy, allowing individual users to override apparently optimal settings–for example, allowing a consumer to say "I need the dishwashers in the cafeteria on NOW", even though it may be cheaper to defer this load to off-peak hours.

2.1 Alternatives to the DRACHMA Model

One alternative to DRACHMA's quasi-decentralized DR model is to retain centralized control but automate device registration. In this approach, devices within a building (or campus) would auto-discover the identity of a centralized EMS and 'register' themselves with that EMS, which also stores DR policies for different organizations (entered via a Web-based interface). The EMS could then respond to dynamic pricing signals by first centrally determining the best set of device adaptations that maximize some global utility and then issuing appropriate signals to each individual device. While such centralization simplifies the adaptation framework, it makes it harder for individuals to dynamically alter adaptation policies for specific sets of devices. Centralized control also raises potential privacy concerns, as now each department of each organization must expose its consumption preferences directly to the third party (building operator) EMS. In contrast, DRACHMA enables a vision of *"consumption virtualization"*, where adaptation decisions are appropriately de-coupled from the physical topology of the electrical network. For example, a distributed data center or server farm, running across two offices located in cities A and B, respectively, may be viewed as a geographically dispersed domain with advanced load balancing policies–e.g., if pricing is more expensive in city A than in city B, then the company's (domain) adaptation policy may be to simply shut down the servers in office A and migrate the entire computational load to office B (see [6] for an early investigation of this idea). Another alternative is to make DR completely *decentralized and device-driven*–in this vision (e.g., [8]), each appliance/device comes with its built-in adaptation logic that responds to pricing signals. While certainly potentially compelling in the long run, this approach assumes that DR is performed solely through pricing signals, as any form of cumulative load control would require *collective adaptation* by groups of devices. Moreover, this approach could increase the configuration burden on individuals, who must now set their preferred adaptation parameters for potentially hundreds of devices.

3. KEY TECHNICAL INNOVATIONS

We now describe three areas where innovative research is required to make the DRACHMA vision a reality.

3.1 Autonomic Domain Formation

The problem of autonomic, distributed domain formation is similar to peer-to-peer overlays in the Internet. We will require protocols that, possibly via application-layer multicast or publish-subscribe mechanisms, enable devices to discover other domain members over a distributed network and then perform multi-party notification and negotiation. For example, the dishwasher should be able to inform other overlay members that the cafeteria personnel have pushed the "wash now" button. Moreover, the domain nodes should be able to select a 'leader' (for the purposes of negotiating DR with other domains) in a secure and resilient manner, that is robust to the dynamic addition or deletion of domain member nodes. A variety of well-understood IP-based protocols, such as DNS-SD [9], gossiping protocols and SIP-based P2P signaling [10] may prove useful for building such a scalable signaling framework, providing features such as message redirection and group notification.

3.2 Composition of Energy Profiles

DRACHMA envisages 'appliance domains' (i.e., a collection of appliances) to be the atomic unit of negotiation for DR adaptation, with a domain leader (controller) representing both the consumption profile and DR policies of the entire domain. The domain leader needs to expose only the *aggregate* consumption characteristics of all its constituent devices. For example, the leader of 'company A's domain' need not externalize exactly how many water coolers or window airconditioners it has, nor exactly how its apportions its aggregate load among these different appliances. To implement this approach, we need to understand how aggregate consumption characteristics of a group of potentially dissimilar devices (with significant diversity in both consumption and control mechanisms) can be derived and expressed from the load profiles of individual appliances. For example, HVAC units have different recovery time transients [5] and deadband ranges of $\geq 5^\circ$F, implicitly defining the granularity at which their consumption may be adapted. Likewise, high-intensity discharge (HID) lamp systems typically require a minimum burn time of 10 hours (to avoid reduced lamp life) before they can be turned completely off, once activated. Accordingly, we must develop a taxonomy of power consumption profiles (with attributes such as peak load, minimum activation duration. etc.) for both individual, and groups of, devices.

3.3 Distributed, Hierarchical Utility Maximization

The distributed negotiation of DR adaptation among domains (potentially in a hierarchical fashion) is a very complex problem, because different domains (organizations) may have different preferences (sensitivities to prices) and objectives (e.g., maximizing comfort vs. reducing energy bills). One formal approach to modeling this problem is via an economics-based utility maximization framework. Formally, let $\{A_1, A_2, \ldots, A_N\}$ denote the set of units whose consumption pattern can be controlled, with $P_i(A_i, C_i(t))$ denoting the power consumed by device A_i at time t when its control parameter assumes the value $C_i(t)$. One can then define a utility function (reflecting features such as occupant comfort) associated with individual, or group of units (e.g., number of operational elevators in the elevator bank), as $U_X(\{(A_{x1}, C_{x1}(t)), (A_{x2}, C_{x2}(t)),\})$ for a set of devices $X = \{A_{x1}, \ldots, A_{xL}\}$. Finally, the framework incorporates a set of K constraints (denoted by the set $\Gamma = \{\gamma_1, \ldots, \gamma_K\}$), related to either device operational characteristics (e.g., an HID controller must be on for at least 10 minutes) or occupant objectives. DR can then be viewed as the optimization:

$$\max \sum_{x \in X} U_x(A_{x1}, C_{x1}(t)) - \sum_i P_i(A_i, C_i(t)) * \lambda_t \qquad (1)$$

$$\text{such that } \gamma_k = \text{true } \forall \gamma_k \in \Gamma,$$

where λ_t is the price per kW (indicated in the DR price signal) at time t. The basic idea is that the best operating point is one which maximizes the utility at the minimum possible cost. While distributed Network Utility Maximization (NUM) techniques have been studied extensively for non-hierarchical communication networks, DRACHMA requires the development of a utility maximization approach where utilities are expressed, and pricing feedback is received, at different domain granularities. Also, the framework must embrace the gradual emergence of micro-grids comprising local energy storage systems (e.g., batteries) and production sources (e.g., solar cells), which are often characterized by stochastically varying production rates (e.g., solar cell efficiency drops dramatically on a cloudy day). In such scenarios, DR adaptation requires advances to distributed optimal control-based approaches (e.g., [11]) to maximize the overall utility in a stochastic sense, under uncertainties in production and demand profiles, over a longer time horizon. Moreover, devices and domains must also be capable of providing future predictions of demand, to enable better store-vs-consume decisions on the micro-grid.

4. REFERENCES

[1] R. Lasseter, R.H and P. Piagi (2004), MicroGrid: A Conceptual Solution, PESC 2004.

[2] Environmental Information Administration, EIA Annual Energy Outlook, June 2008, available at: http://www.eia.doe.gov/oiaf/aeo/pdf/0383(2008).pdf

[3] D. Hammerstrom, et al, Pacific Northwest GridWise$^{\text{TM}}$Testbed Demonstration Projects, Part II. GridFriendly$^{\text{TM}}$Appliance Project, PNNL Technical Report 17079, October 2007.

[4] A. Chen, Working Toward the Very Low Energy Consumption Building of the Future, available at: http://newscenter.lbl.gov/feature-stories/2009/06/02/working-toward-the-very-low-energy-consumption-building-of-the-future/

[5] M. Gupta, S. S. Intille, and K. Larson, Adding GPS-control to traditional thermostats: An exploration of potential energy savings and design challenges," Pervasive, September 2009.

[6] A. Qureshi, R. Weber, H. Balakrishnan, J. Guttag and B. Maggs, Cutting the Electric Bill for Internet-Scale Systems, Proceedings of ACM SIGCOMM, September 2009.

[7] The ZigBee Alliance, ZigBee Smart Energy Application Profile, available at: http://zigbee.org/en/spec_download/zigbee_downloads.asp.

[8] B. Nordman, The Case Against the Smart Grid, I4energy Seminar Series - CITRIS – CITRIS / UCB, October 2009.

[9] S. Srinivasan, A. Moghadam, and H. Schulzrinne, BonAHA: service discovery framework for mobile Ad-Hoc applications, IEEE CCNC, January 2009.

[10] L. Li, Y. Ji, T. Ma and L. Gu, Locality-Aware Peer-to-Peer SIP, IEEE ICPADS, December 2008.

[11] S. Eswaran, A. Misra and T. La Porta, Control-theoretic Optimization of Utility over Mission Lifetimes in Multi-hop Wireless Networks, IEEE SECON, June 2009.