

Precise Mass-Market Energy Demand Management Through Stochastic Distributed Computing

Dr. Alex Papalexopoulos, *Fellow, IEEE*, Dr. Jacob Beal, *Member, IEEE*, and Dr. Steven Florek

Abstract—Even though demand response (DR) participation has substantial benefits to the market as a whole, current DR programs suffer from a collection of market, regulatory, infrastructure and technology problems, such as lack of scalability, lack of privacy, imprecision, and nonacceptance by customers. This paper describes how a fundamentally different DR approach, based on service priority tiers for appliances and on stochastic distributed computing, can overcome these problems and be integrated with energy markets. Our approach takes advantage of inexpensive communications technology to estimate the state of home and small-business major electrical appliances and have those appliances respond to power grid state signals within a few seconds. Organizing appliances into service priority tiers allows retail customer power demand to be de-commoditized, making these DR resources a potent force for improving the efficiency of energy markets. This paper describes the proposed methodology, examines how it can be integrated into energy markets, and presents results from mathematical analysis and from simulation of 100 000 devices.

Index Terms—Demand response, distributed computing, locational marginal pricing, power market auctions, stochastic control.

I. INTRODUCTION

IT HAS BEEN well understood that participation of customers in the form of demand response (DR) in wholesale organized energy markets, administered by RTOs/ISOs, helps increase the competition and improve the efficiency of these markets. The benefits from DR participation in wholesale energy markets are substantial:

- RTOs/ISOs can better balance supply and demand, using DR as an alternative to dispatching high-priced generation resources to meet demand, thus reducing market prices and market volatility [1].
- Supply market power can be mitigated by increasing competition and creating additional incentives against generators bidding at high prices [1].
- System reliability and resource adequacy can be enhanced because DR resources can provide fast balancing of the transmission grid in case of loss of generation or other unexpected events [2].

Regulators and policy makers have taken substantial measures over the years to encourage participation of DR in wholesale energy markets. For example, Congress has instituted national policy to enhance DR participation [3] and FERC has is-

sued over the last several years a series of Orders, including Order No. 890, Order No. 890-A, Order No. 719, and Order No. 745 for the same purpose.

In response, several RTOs and ISOs have instituted various types of DR programs. Some of these programs are designed to respond to reliability and emergency conditions, while others are designed to allow wholesale customers, qualifying large retail customers, and aggregators, or curtailment service providers (CSPs) of retail customers to participate directly in the day-ahead and real-time energy markets, certain ancillary service markets, and capacity markets.

Despite these efforts, DR participation has been less than encouraging, particularly for the residential and small-business customers who consume a plurality of all power. Major obstacles include technology, market design, regulation, and poor alignment between the incentives of utilities and CSPs. These obstacles are described in more detail in Section II.

In this paper, we will show how a fundamentally different approach to the DR problem, based on appliance-level service priority tiers and stochastic distributed computing, can overcome these obstacles. Our proposed approach is ideal for self-organization of loosely coordinated devices present in homes and small businesses, micro grids and autonomous systems [4]–[6] and provides fast, accurate, and robust demand shaping for collections of thousands to millions of devices. It also addresses market, incentive, and regulatory obstacles to DR participation: it is non-coercive, voluntary, bottom-up, better preserves customer privacy, and buffers customers from price volatility. In the remainder of this paper, we will present our approach, examine how it can be integrated into energy markets, and present results from analysis and simulation.

II. OBSTACLES TO MASS-MARKET DEMAND RESPONSE PROGRAMS

A mixture of obstacles, including utility incentives, regulation, technology, and market design, have strongly limited participation in DR programs, particularly among the small customers who make up a plurality of electrical power demand.

Due to a mixture of regulation, market structure, and incentives, utilities have until recently emphasized emergency programs, which use high capacity payments to attract customers and in many cases provide more emergency DR than is needed by the RTOs/ISOs. Moving customers to “earlier,” more frequent or price-based triggers, however, has not worked as expected. Moreover, many states prohibit default dynamic pricing for residential customers. As a result, it is difficult to align pricing for generation, load and DR resources; this misalignment can lead to gaming opportunities, cost shifting and perverse incentives.

Manuscript received October 02, 2012; revised February 13, 2013, April 19, 2013; accepted April 29, 2013. Paper no. TSG-00683-2012.

A. Papalexopoulos is with ECCO International, San Francisco, CA 94104 USA (e-mail: alexp@eccointl.com).

J. Beal is with Raytheon BBN Technologies, Cambridge, MA 02138 USA, (e-mail: jakebeal@bbn.com).

S. Florek with ColorPower, Berkeley, CA, USA (e-mail: sflorek@color-power.org).

Color versions of one or more of the figures in this paper are available online at <http://ieeexplore.ieee.org>.

Digital Object Identifier 10.1109/TSG.2013.2263396

The process of regulation has tended to pose another set of problems for deploying DR resources. Revenues from ancillary services markets for regulation and spinning reserve are essential to funding DR, but are precluded by certain reliability councils, such as WECC. In addition, utilities are perceived to have significant advantages in terms of market power, influence with regulators, and ability to shift administrative costs. State-approved IOU DR programs are constrained to reflect multiple policy priorities, which may prevent program costs from being fully recovered in wholesale markets. This is indicative of some mixture of two problems: first, that markets may have structural problems that undervalue DR, and second, that current DR programs depend on incentive structures, load control hardware, or installation channels that are too expensive relative to the true value of DR.

Further, infrastructure and technological barriers still exist which have prevented massive penetration of DR resources into the market. There are also market gaming concerns due to “baseline” issues, arbitrage between zonal and nodal pricing, double counting, selective bidding, etc. With prior approaches, actual curtailment of demand can only be estimated; it cannot be measured. Further, with the instrumentation that is typical of current deployments, these estimates are grossly imprecise and inaccurate. As a result, even though DR resources should be treated exactly comparable with supply resources, RTO/ISO operators cannot fully depend on them at the present time.

Finally, current DR programs are difficult to understand, require sophisticated energy customers, are often perceived as invasive of privacy, and are inherently coercive as they are based on either direct load control or punitive dynamic rate structures. Retail customers, however, prefer simplicity and consistency and these preferences are in conflict with the volatility of the wholesale markets. This is consistent with the experience from other markets. For example, for decades telecommunication services were priced using time-of-use per-minute/per-KB rates, ostensibly to manage peak congestion. After deregulation and the resulting disruptive innovations (including cooperative congestion management, e.g., TCP/IP), the vast majority of customers today opt to purchase telecommunication services as a monthly service based on usage and service tiers, rather than as a commodity. Free market choice does not necessarily imply forced participation in government-supervised centralized auction markets; it is also about choice of market venue participation and different value-exchange models.

Indeed, a significant body of research shows that real-world customers do not act as so-called “rational economic agents” who are continuously pondering their marginal consumption utility vs. the real-time market price [7]. For ongoing expenses (e.g., electricity, internet service, cell phones, rent, etc.) consumers, in general, strongly prefer fixed-cost “level of service” plans over uncertain dynamically determined costs, even when the uncertain pricing is expected to result in some savings [7]. In effect, consumers are trading a slightly higher price for a lower risk and a lower cognitive burden. Thus, for example, broadband internet providers offer consumer service classes based on advertised speed and cell phone providers offer service classes based on coverage, minutes-per-month, and free calling groups.

III. DEMAND MANAGEMENT VIA APPLIANCE SERVICE TIERS AND DISTRIBUTED COMPUTING

Our DR approach, called ColorPower, has previously been presented in [4], [5] and [6], which focus mainly on the algorithmic problem of handling coordinated decision making in a system of thousands to millions of devices. This section reviews the basic architecture presented in those prior papers and extends it by examining how the ColorPower approach to service tiering enables it to cope with or bypass the obstacles discussed in Section II.

ColorPower departs from prior DR approaches by the combination of three key properties:

- Customers are not directly involved in auction markets.
- Customers control DR availability in terms of tiers of service provided by individual appliances.
- The coordination problem of determining which devices should consume power at what times is solved through distributed aggregation and stochastic control.

This combination allows novel approaches to deployment and market integration that address or avoid many of the obstacles discussed in Section II. Some prior work on demand management has addressed some portions of the first two properties (e.g., [8]–[12]), but has suffered from some combination of the following problems: they may require customers to perform load control (which makes DR unreliable and intrusive), expose customers directly to market fluctuations (which has regulatory and adoption problems as noted above), require centralized management of individual loads (which is hard to scale), or tier service only in terms of electrical reliability (which is intrusive because it typically shuts off whole customers or neighborhoods and has regulatory barriers because it is perceived to operate at the expense of socioeconomically disadvantaged customers).

At its root, the ColorPower approach is based on the fact that residential and small-business customers have a great deal of flexibility in their energy needs and are willing and able to adjust levels of demand, *if doing so can be made convenient for them*. This has been demonstrated in a number of studies (e.g., [13], [8], [9]), but problems in customer adoption, scalability, and market integration have previously made it difficult to take advantage of this fact.

Under this approach, the customer marks devices using “colors” that correspond to control plans that provide different levels of *reliability of services provided by the device*. These plans are then executed by an external or internal load controller that communicates indirectly with a central ColorPower controller operated by an ISO/RTO, a utility, aggregator, microgrid operator, or other appropriate authority. The meaning of a “color” is dictated by the class of device. For example, a pool pump may be marked “green” if its load can potentially be shut down at any time or “yellow” if it can only be shut down during peak power events, while an HVAC system might have a tight “green” range of allowed temperature variability, a looser “yellow” range, and loosest “red” range. These “color” markings are then used to organize devices into priority tiers (of which there may be many more than colors), dictating

when devices are available for demand shaping and the order in which classes of device services will be enabled or disabled.

Note that the least invasive devices to control are typically those that autonomously cycle off and on (such as HVACs, water heaters, pool pumps, refrigerators) or those that can gracefully degrade service (e.g., commercial or public lighting, variable speed motors). Serendipitously, these also tend to be larger loads than the devices that customers expect instantaneous non-curtable service from, such as residential lighting, clock radios, and other small electronics.

The ColorPower control algorithm then manages instantaneous demand to meet an externally provided DR goal using a feedback loop. This goal may be supplied by an energy market, as we will discuss in Section V, or by any other policy enacted by an appropriate grid control authority. The feedback loop operates by aggregating demand flexibility information into a global estimate of total prioritized customer flexibility, which is then broadcast back to the device load controllers along with the current demand target, and each local controller uses its local information about color, appliance type, and current state to make a stochastic control decision. With each iteration of aggregation, broadcast, and control, the overall system moves toward the target demand, allowing the system as a whole to rapidly achieve any given target demand and closely track target ramps [4], [5].

An important assumption with this approach is that participating loads report both their sheddable load and *the amount of load currently being curtailed*. This latter requirement permits direct measurement of demand response by sub-metering within the switching mechanism for individual devices,¹ making it unnecessary to estimate baselines from historical whole-premises metering. The magnitude of an individual curtailed load can instead be estimated by the per-device sub-metering, using a model of that device based on measurements taken during recent periods when the device was not curtailed. These individual measurements are then summed across the network by a distributed algorithm for aggregate estimation (e.g., self-organizing spanning trees in [5] and [6]) to produce an estimate of the current amount of sheddable and curtailed demand in each tier.

This distributed approach, by keeping most information local to a device and collecting and redistributing only aggregate summary information, drastically reduces the amount of communication needed for effective control. This allows demand shaping to operate in seconds on systems of millions of devices using low-cost load control hardware. Aggregation also has the beneficial side effect of preserving the privacy of individual customers: their demand information simply becomes part of an overall statistic.

It is important to note that ColorPower is not linked to any particular approach to demand forecasting. Rather, it provides a means of responding to the challenges predicted by forecasts (e.g., peak power events) and a means for correcting for deviations between forecast and reality (e.g., generation variability in wind and solar).

¹The assumption of per-device sub-metering is enabled by the rapidly decreasing cost and increasing availability of the necessary electronics.

ColorPower also enables a drastically different approach to customer incentives for participation. First, notice that since control is not dependent on dynamic pricing, there is a great deal of freedom in the design of customer incentive models. Indeed, there is no reason that ColorPower could not be connected to a dynamic pricing model or a “kickback” model where customers are paid a portion of the market value of the DR, if one of these were determined to be the best way to incentivize customer adoption. As discussed in Section II, however, customers tend to strongly prefer “service level” pricing that makes costs more predictable.

With ColorPower, these sorts of “service level” plans are easy to design: we can simply map the “color” designations of electrical devices to plans. By designating a set of devices with particular colors, a customer is opting into a certain mixture of plans. From a pricing discrimination perspective, what a customer now is choosing is the level of electrical *time-inflexibility* that they wish to pay for. A “more flexible” color mean less certainty of when a device will run (e.g., time when a pool pump runs), or lower quality service delivered by a device (e.g., wider temperature ranges, slower electrical vehicle charging). These types of economic decision-making are eminently compatible with customer desires and economic design, as evidenced by the wide range of quality-of-service contracts offered in other industries: in the case of ColorPower, customer incentives might come in many different forms, depending on distribution channel, market identification, local regulations, and appliance type: for example, there might be a rebate on the power for the appliance, a rebate on the price of the appliance itself, or other incentives such as appliance maintenance discounts.

Thus, we see that with regard to the obstacles to DR programs discussed in Section II, ColorPower offers a potential to cope with or bypass all of them: the separation of incentive model from system operation avoids exposing customers to dynamic price and allows non-coercive participation plans to be formulated intuitively in terms of graceful degradation of the services provided by individual appliances. At the same time, this separation also gives freedom for alignment with a wide range of regulatory and market structures. The distributed control approach lowers infrastructure and hardware requirements, while simultaneously giving much higher predictability about the current availability of DR resources. Finally, these benefits combine to lower the cost of deployment and operation, making DR potentially economically feasible across a much broader range of participants and circumstances.

IV. SCALABLE DEMAND CONTROL WITH COLORPOWER

In this section, we provide a brief review of the algorithmic work on scalable demand shaping presented in [4], [5], and [6], in order to give a foundation for understanding the new results on integration of the ColorPower methodology with energy markets presented in the remainder of the paper.

A. Formal Definition of ColorPower Control Architecture

In a system of devices managed using ColorPower, at any given time each device (or device behavior mode) is either *enabled*, meaning that it can draw power freely, or *disabled*,

Symbol	Definition
color	Customer configuration of quality of service plan for an electrical device
priority tier	Set of devices collected by color and class. Lower numbered tiers are shut off first.
enabled	Device or device mode is allowed to consume power
disabled	Device or device mode is not allowed to consume power
flexible	Device or device mode can switch between enabled and disabled
refractory	Device or device mode cannot switch
EF,DR,EF,DR	States combining Enabled, Disabled, Flexible, and Refractory
$T_{DF,i}$	Fixed rounds of disabled refractory time for tier i
$T_{DV,i}$	Maximum random rounds disabled refractory time for tier i
$T_{EF,i}$	Fixed rounds of enabled refractory time for tier i
$T_{EV,i}$	Maximum random rounds enabled refractory time for tier i
t	A point in time
a	A ColorPower agent
$s(t, a)$	State of demand for agent a at time t
$ X_{i,a} $	Power demand (watts) in state X for color i at agent a
$ s(t) $	State of total power demand (watts) at time t
$ X_i $	Total power demand (watts) in state X for color i
$\hat{s}(t)$	Estimate of $s(t)$
$ \hat{X}_i $	Estimate of $ X_i $
$g(t)$	Goal total enabled demand for time t
$c(t, a)$	Control state for agent a at time t
$p_{off,i,a}$	Probability of a flexible color i device disabling at agent a
$p_{on,i,a}$	Probability of a flexible color i device enabling at agent a
D_i	Demand for i th color and above
b	Boundary color, which mixes disabled and enabled demand
f	target minimum ratio of flexible to refractory demand
$r(t)$	Ratio of the distances from $g(t)$ to the boundary color edges (D_b, D_{b+1})
α	Proportion of goal discrepancy corrected each round
r_c, r_q	Number of rounds for controller to converge/quiesce

Fig. 1. Table of key terms and variables for ColorPower.





Fig. 2. Modified Markov model for ColorPower device state switching, from [4]: states with an “E” are enabled, “D” are disabled, “F” are flexible, and “R” are refractory.

meaning that it has been shut off or placed in a lower power mode². In order to prevent damage to appliances and/or customer annoyance, devices must wait through a *refractory* period after switching between disabled and enabled before they return to being *flexible* and can switch again (the length of this period depends on device type). These combinations give four device states (e.g., **EF** is enabled and flexible), through which each device moves according to a modified Markov model as shown in Fig. 2. Devices move randomly from **EF** to **DR** with probability p_{off} per round and from **DF** to **ER** with probability p_{on} . Devices move from **ER** to **EF** by a randomized timeout of $T_{EF} + U(0, T_{EV})$ seconds, where $U(0, T_{EV})$ is a uniform random draw from the interval $[0, T_{EV}]$; the transition from **DR** to **DF** is the same, substituting T_{DF} and T_{DV} .

The control problem for a ColorPower system is shown by the block diagram in Fig. 3. The system comprises a set of n elec-

Fig. 3. Block diagram of the ColorPower control architecture, from [4]: the demand and priority tier states $s(t, a)$ of each device a at time t are aggregated to produce an estimate $\hat{s}(t)$ of the total demand and priority tier states over all devices. This is broadcast to each device, along with a demand goal $g(t)$, and each device uses local information to set its control state $c(t, a)$ of transition probabilities. Loads then evolve according to this control state, subject to exogenous disturbances.

trical devices, each of which apportions its demand between k priority tiers, where lower numbered tiers are intended to be shut off first (e.g., 1 for “green” pool pumps, 2 for “green” HVAC, 3 for “yellow” pool pumps, etc.), and where each tier has its own time constants (e.g., $T_{EF,2}$ is the fixed portion of the refractory timeout for Tier 2). The state $s(t, a)$ of each device a at time t is the magnitude of power demand in each state for each priority tier (e.g., $|EF_{1,a}|$ is the number of watts of enabled and flexible demand in tier 1 at device a). These values are summed over all reporting devices (some implementations may use sampling for efficiency reasons) using a distributed algorithm and fed to a state estimator to get an overall estimate $\hat{s}(t)$ of the true state $s(t)$ of total demand in each state for each tier (e.g., $|EF_1|$ is the estimated total enabled and flexible demand in tier 1). This estimate is then broadcast to all devices, along with the demand goal $g(t)$ for the next total reduction in enabled demand over all tiers. The controller at each device a uses local information on state, color, and device type to set its control vector $c(t, a)$, defined as the set of transition probabilities $p_{on,i,a}$ and $p_{off,i,a}$ for each tier i . Finally, demands move through their states according to these transition probabilities, subject to exogenous disturbances such as changes in demand due to customer override, changing environmental conditions, imprecision in measurement, etc.

Note that the aggregation and broadcast algorithms must be chosen carefully to ensure that the communication requirements are lightweight enough to allow control rounds a few seconds long on low-cost hardware. The choice of algorithm depends on the network structure: for example, [5] assumed a mesh network, and thus used spanning tree aggregation and gossip-based broadcast, which are fast and efficient in that environment (for details, see [5]). In general, however, the system must use aggregation and broadcast because it is not currently economically feasible to deploy reliable communications for load control devices with enough bandwidth to poll individual state and deliver individual instructions for millions of devices every few seconds. Aggregation reduces bandwidth requirements by com-

²Note that an enabled device may not actually demand any power: for example, a toaster that is not currently in use.

pressing the state of many devices into a single “collective state” message to a central system, and while broadcast reduces bandwidth requirements by sharing a single “collective control” message from the central system with many devices.

B. ColorPower Control Algorithm

The ColorPower control algorithm determines the control vector $c(t, a)$ by a stochastic controller formulated to satisfy four constraints. We present here the constraints shaping the control algorithm and sketch the key engineering decisions in its implementation. For full details and mathematical exposition, see [4].

The four constraints that the controller must satisfy are:

a) Goal Tracking: The total enabled demand in $s(t)$ should track $g(t)$ as closely as possible: i.e., the sum of enabled demand over all tiers i should be equal to the goal. The demand goal $g(t)$ is determined by the utility (or other control authority) based on economic or reliability considerations. For example, it may be calculated by an energy simulator using forecasting market and system information to reduce market clearing prices (LMPs or SMPs) by a certain amount. It may also be calculated to address reliability or grid problems as they arise during various system conditions. This is particularly important as the fraction of variable power generation resources, such as wind or solar, continues to rapidly increase.

Goal tracking is formalized as the equation:

$$g(t) = \sum_i (|EF_i| + |ER_i|) \quad (1)$$

b) Tier Priority: Devices with lower numbered tiers should be shut off before devices with higher numbered tiers. This is formalized as:

$$|EF_i| + |ER_i| = \begin{cases} D_i - D_{i+1} & \text{if } D_i \leq g(t) \\ g(t) - D_{i+1} & \text{if } D_{i+1} \leq g(t) < D_i \\ 0 & \text{otherwise} \end{cases} \quad (2)$$

so that devices are enabled from the highest tier downward, where D_i is the demand for the i th tier and above:

$$D_i = \sum_{j \geq i} (|EF_j| + |ER_j| + |DF_j| + |DR_j|) \quad (3)$$

c) Fairness: At any give time, there is at most one *boundary tier* where the goal dictates that some demand should be enabled and other demand should be disabled. Fairness means that any two devices a and a' in the boundary tier b will have the same statistical expectation of enabled vs. disabled demand:

$$\frac{E[|EF_{b,a}| + |ER_{b,a}|]}{E[|DF_{b,a}| + |DR_{b,a}|]} = \frac{E[|EF_{b,a'}| + |ER_{b,a'}|]}{E[|DF_{b,a'}| + |DR_{b,a'}|]} \quad (4)$$

This means that no device is privileged over others of the same tier, but allows any given device to implement expected likelihood differently based on device type and customer settings. For example, variable lighting may uniformly dim, while

a pool pump chooses whether to shut off, and an air conditioner adjusts its duty cycle.

d) Cycling: Devices in a tier trade off which are enabled and which are disabled such that no device is unfairly burdened by initial bad luck. This is ensured by asserting the constraint:

$$(|EF_b| > 0) \cap (|DF_b| > 0) \Rightarrow (p_{on,a,b} > 0) \cap (p_{off,a,b} > 0) \quad (5)$$

This means that in the boundary tier b , whenever some enabled devices are flexible and some disabled devices are flexible, these devices always have a chance of changing their state. For this last constraint, there is a tradeoff between how quickly devices cycle and how much flexibility is held in reserve for future goal tracking; we balance these with a target ratio f , which specifies a lower bound on the permitted ratio of flexible to refractory power:

$$f \leq \min\left(\frac{|EF_b|}{|ER_b|}, \frac{|DF_b|}{|DR_b|}\right) \quad (6)$$

Since the controller acts indirectly, by manipulating the p_{on} and p_{off} transition probabilities of devices, the only resources available for meeting these constraints is the demand in the flexible states **EF** and **DF** for each tier. When it is not possible to satisfy all four constraints simultaneously, the ColorPower controller prioritizes the constraints in order of their importance. The highest priority is given to customer relations, in the form of the fairness constraint and on any exogenously imposed contractual guarantees on circumstances under which tiers can be disabled (e.g., “Yellow” HVAC devices may not be controlled except in response to peak demand events or grid emergencies). Fairness is implemented by ensuring eventual statistical equivalence in the expected distributions of device demand, though the values of $c(t, a)$ used to implement this may vary by device type and configuration; contractual guarantees are implemented by forcing the p_{off} of currently prohibited tiers to be zero.

The goal tracking, tier priority, and fairness constraints are then handled, in that order, by treating the available demand flexibility as a “budget” and allocating flexibility to each constraint in turn until it is all allocated or reserved for possible future needs. Goal tracking uses a proportional controller to correct a fraction α of the difference between current and goal demand. Tier priority then corrects transient priority inversions by exchanging as much disabled demand as it can in lower tiers for equal amounts of enabled demand in higher tiers. Finally, cycling, lowest priority because it operates at the slowest time-scale, computes the maximal steady-state exchange of enabled and disabled demand in the boundary tier that will not violate (6). These flexibility allocations are then transformed into transition probabilities by normalizing against the current flexible demands.

Each device then samples its distribution and adjusts its state accordingly, and the law of large numbers means that this will result in an aggregate system behavior expected to produce near-optional system behavior; in those cases where behavior deviates, the feedback between aggregate state estimate and local control ensures that it will rapidly be corrected.

The ColorPower control algorithm is fully specified in [4].

C. Effectiveness of the ColorPower Control Algorithm

The convergence and resilience of the ColorPower control algorithm has been evaluated both analytically and empirically via simulation in [4]. The analysis focused on the case of rapid change of demand goal, e.g., following a major equipment failure or renewable generation fluctuation. In this case, the expected number of control rounds r_c for the controller to converge to within ϵ watts of the goal is logarithmic:

$$r_c = \frac{\log \epsilon - \log \Delta}{\log(1 - \alpha)} \quad (7)$$

(where Δ is the initial watts of distance from the goal and α is the proportionality of error that goal tracking attempts to correct each round), provided that a flexible reserve of at least $(1)/(2 + 2f)$ is maintained within any given tier. Thus, for example, the expected time to converge to within 1% of the goal given 10 second rounds and $\alpha = 0.8$ would be 30 seconds—quite fast enough to handle most rapid shedding conditions.

The time for the ColorPower controller to return to a quiescent state, from which another such rapid shift is guaranteed to succeed, is conservatively estimated to be less than:

$$r_q = r_c + r_p + 2(T_{DF,b'} + T_{DV,b'}/2 + T_{EF,b'} + T_{EV,b'}/2) \quad (8)$$

where b and b' are the boundary tiers before and after the goal change respectively, and where r_p is equal to:

$$r_p = \max(T_{DF,b'} + T_{DV,b'}, T_{EF,b} + T_{EV,b}) \quad (9)$$

Simulation studies in [4] confirm these results and demonstrate that the controller scales well on a range of at least 10^3 to 10^6 devices, that it can closely track ramps even when the ramp direction changes frequently, and that control is robust against changing device populations, noise in estimates, and high degrees of heterogeneity between devices.

V. INTEGRATION WITH WHOLESALE ENERGY MARKETS

Having reviewed how ColorPower implements DR, we now turn to the question of how this capability can be integrated with wholesale energy markets. In this section we present a methodology for integrating the DR capacity created by ColorPower with wholesale energy markets. Our objective is to create a price responsive retail demand curve without sending price signals to retail customers or their devices, due to the regulatory and customer preference obstacles discussed in Section II. We will thus first discuss the problems of current DR integration methods, then present a new methodology for integrating DR capacity into wholesale energy markets that is enabled by ColorPower demand priority tiers.

A. Problems With Current Integration Methods

Many believe that smart meter deployment and dynamic pricing will result in a retail auction-based solution where customers will be forced to act as rational economic agents and respond to changing prices. It is further envisioned that when this becomes too burdensome, electrical device manufacturers will create devices with energy market trading intelligence, which will be programmed by customers with price ranges to

control the maximum price at which they operate. This widely held view contradicts what customers have generally been demonstrated to want from services with ongoing expenses: predictable costs based on usage or service quality tiers [7]. Discussion of other key obstacles to dynamic electricity pricing on the retail level may be found in [14]–[17].

Current DR approaches typically treat flexibility as a supply resource: “virtual power plants or proxy generator resources” supplying “negawatts.” In this framework, bids represent increasing levels of inconvenience with corresponding monotonically-increasing costs of paying groups of retail customers to forgo consumption. DR then competes against generators as a source of “supply,” and may participate in the day-ahead market (DAM) including the reliability unit commitment (RUC), the hourly or 5-minute real-time energy market (RTM) and the day-ahead or real-time ancillary services (AS) markets.

While practical for large industrial customers, this creates problems for retail residential customers that severely limit the value of this approach. For example, the management of the data required for scheduling residential customers at the zonal level is taxing since it requires aggregation of numerous very small end-use customers, with frequent migration, i.e., enrollments and de-enrollments in a DR program. Further, using standard historical load distribution factors (LDFs) derived from the EMS state estimator may also be problematic.

There are also serious settlement concerns. Setting aside the complexity of the settlements of the curtailed load portion with aggregators and the rest of the load with the load servicing entity (LSE) and the subsequent bilateral settlement between aggregator and LSE outside of the ISOs settlement process, the determination of actual DR delivery is very problematic. It is expected to be derived from measurement of aggregate meter usage, calculated from a pre-determined, administrative set baseline. The baseline problem, i.e., the challenge of measuring what the customer would have done without the payment, has proven in practice an intractable problem that has resulted in limited participation of active DR and even worse to ultimately unreliable DR resources. Finally, the gaming opportunities that result from the baseline problem can be severe. A comprehensive treatment of this gaming problem is contained in [18]. Verification of performance is also problematic, as this is typically done by comparing against baseline on an aggregate basis, where the source of problems cannot be identified, rather than by summing the difference of baseline versus actual for each customer or device.

There are two additional problems worth noticing that arise from conflicts between prices used in current practices. The first problem relates to the potential double payment for DR that arises if the aggregator or the curtailment service provider (CSP) receives a payment that exceeds the difference between the locational marginal pricing (LMP) and the wholesale price implied in the current retail price. Such double payments increase rates and can incent inefficient DR with costs to the market that exceed its benefits. The second problem is related to the potential arbitrage between a low zonal price used to settle the DAM load schedule and the higher nodal price used to settle the curtailed demand. This problem can be limited in magnitude, but

not eliminated, by administrative rules such as constraining the number of hours that a DR resource can be dispatched and by requiring the presence of physical control devices. These sort of rules, however, are dealing with symptoms rather than the source of the problem and can have the side effect of badly truncating the potential value of active DR participation in the wholesale market.

B. New Proposed Integration Method

The self-identified priority tiers of the ColorPower approach to demand response enable us to propose a new indirect model for retail power participation in energy markets. Since the demand has been differentiated into tiers with a priority order, the demand in each tier may be separately bid into the market. The price for each tier is set based on the cost of supplying DR from that tier, which in turn is linked to the incentives necessary to secure customer participation. The tiers are then ordered by priority to form a monotonically-decreasing demand curve. Demand can thus express preferences before market clearing in the form of lower buy bids for lower priority tiers of demand. This allows aggregated demand to send price signals in the form of the decreasing buy bid curve. Market information thus flows bi-directionally. Curtailment tolerance from customers is signaled by the degree to which they opt into variously incentivized plans. Because customers can self-identify the flexible devices within their total load profile, this effectively creates two tiers of power demand, flexible P_f and inflexible P_i , that can be aggregated and monitored at the utility level. The flexible load can be further partitioned using priority tiers, into P_{f1} , P_{f2} , and so on. Each tier can be separately forecasted by the utility, and bids can be submitted into the market to procure supply for the entire power demand P_t such that:

$$P_t = P_i + (P_{f1} + P_{f2} + \dots + P_{fn}) \quad (10)$$

A ColorPower system can thus be used by a retail utility or an aggregator to bid an aggregate retail demand curve into the ISO markets on behalf of its customers, without requiring customers to interact with real-time prices directly or to submit a price curve for demand response. Customer demand is instead aggregated and represented by a utility procurement trader, who purchases as much power as is deemed rational to purchase, i.e., up until supply prices rise above the levels bid by decreasing priority tiers.

For customers, the colors correspond to contractual constraints on when resources with that color are available for flexibility. For example, a thermostat controlling an HVAC system might have a “green” range for any-time flexibility and a “yellow” range for flexibility during peak-power and emergencies. Demand in tiers whose constraints make them currently unavailable are bid just as unmanaged demand, while bids for potentially flexible demand is selected to maximize profits subject to contractual constraints with customers, just as generators currently bid supply. Designing a mechanism for optimization of such bids is beyond the scope of this paper, but is expected to follow the lines of prior work on reliability pricing such as that described in [11], [12]




Fig. 4. While energy markets have produced fairly smooth supply curves, without mass-market DR, it is impossible to distinguish high- and low-priority energy demand, leading to inefficient demand curves that are flat or nearly flat (a). Using ColorPower, however, can produce a more smoothly decreasing demand curve (b) based on aggregated, differentiated tiers of measured customer demand. (a) Inflexible demand curve. (b) Priority tier demand curve.

The device population then shares the limited available power resource by setting the demand goal $g(t)$ for the ColorPower algorithm to the amount of power purchased. This results in devices degrading the service they provide (e.g., temperature control, circulation of water in a pool) to the level where the value of demand is equal to the marginal cost of power. Thus, any shortages are handled by automatically diffusing small curtailments across the population of devices. This automatic diffusion of inconvenience across the entire population of devices is the cornerstone capability of the proposed approach.

In this way, a relatively small amount of flexible demand can buffer overall power demand volatility by yielding consumption to inflexible components as needed based upon customer-expressed priority, while ColorPower determines which tiers should consume power and fairly distributes power and curtailment to all customer devices in a partially supplied tier.

C. Equalization of Supply and Demand Market Power

ColorPower, integrated as we have proposed, can transform energy markets by eliminating irrational marginal power fulfillment at the retail level. Currently, wholesale energy markets typically feature “hockey stick”-shaped supply bids during shortages, due to exploitation of scarcity pricing by market participants or expense of running peaking generation resources. Both generation and demand benefit from reducing inefficient energy demand fulfillment. Fig. 4(a) shows a scenario based on current markets, with the last few percent of available supply priced ten

times higher than the lowest priced supply—a fairly typical multiplier. Against this is a typical nearly flat demand curve, with a relatively small amount of DR capacity supplied primarily by larger industrial or commercial customers.

A more smoothly decreasing demand curve, as shown in Fig. 4(b), reduces the incentive for strategically withholding supply from markets in order to sell it at scarcity pricing levels. In times of high demand, marginal price is thus set by flexible demand bids, not supply bids. This is expected to severely erode the profitability of energy speculation while lowering prices for the public. This demand curve can also reduce the need for financial hedging, by using physical demand reduction to reduce under-capacity risk. Reductions in hedging costs will also reduce costs across all energy market participants and ideally to the public.

We may also note that exposing individual retail customers directly to auction pricing through dynamic retail pricing gives them significantly less market power compared to the ColorPower aggregate buying approach. In the dominant proposals for dynamic retail pricing, individual customers are presented with a “take-it-or-leave-it” price from their power supplier, and only have the choice to micromanage consumption over their entire premise. Such pricing necessarily places the consumer at a market disadvantage, particularly given the personal cost of monitoring and managing energy in an individual home or small business. The ColorPower model suggests instead a role for a utility (or other organization) to act as an active manager of aggregate demand fulfillment. Incremental costs for an aggregator are much lower due to the economy of scale for one organization to monitor and manage demand delegated to it by large numbers of customers. The utility can then act as an effective negotiator between supply and demand to provide the best overall balance of service levels and cost for their customers, as opposed to being a mere pass-through distributor of a commodity that marks up whatever price is currently being demanded by generation.

D. Numerical Example

We now demonstrate the potential value of ColorPower and our proposed method for integrating it with energy markets by simulating its operation in an example scenario and comparing results to the same scenario using a conventional energy market without mass-market DR.

We will use a “hot summer day” demand scenario in which demand starts at a low level at 10 A.M., climbs to a peak of 100 MW in the mid-afternoon (driven largely by outside temperature affecting HVAC systems), then falls again by 5 P.M.. Fig. 5(a) shows this demand scenario as a stacked graph where demand has been marked with three colors (green, yellow, and red), each of which is further decomposed into three priority tiers (particular color shades).

The tiers and loads for this scenario are constructed from a rough approximation of observed patterns of device power consumption and the consumer flexibility observed in prior studies: from the meta-analysis in [19], it appears that at least 10–12% of peak load in the mass residential market is readily accessible for DR. This rises to 30% or more with appropriate enabling technology, and we expect that the finer granularity of service




Fig. 5. Simulation results with 100 000 independently fluctuating power loads (a): demand is shown as a stacked graph of the demand from each tier, with saturated colors showing enabled demand and faint colors showing disabled demand. After each transition between tiers, there is a brief period where some fraction of demand is refractory and unable to be switched, which is indicated by cross-hatching. The goal is the dashed blue line and the current total power consumption is the solid magenta line that nearly precisely follows it. Event-based control of the same demand priority tiers is shown in (b). (a) ColorPower. (b) Event-based control.

plan enabled by ColorPower could increase the penetration yet further.

We then mapped this base assumption onto a plausible distribution of energy over colors and priority tiers. We begin with an assumption of approximately half the load being due to air conditioners, with each air conditioner split into bands of temperature flexibility assigned into three tiers: bottom red, top yellow, and middle green. Other major load categories comprise pool pumps (top green), flexibility in water heaters (bottom green), lighting (middle yellow and middle red). All other loads are grouped into bottom yellow and top red. The precise values, of course, are conjectural until supported by specific field data, so this experiment must be regarded as a demonstration of capability, rather than a precise quantification of expected benefit.

The markets for our scenario are based on the examples shown in Fig. 4. The supply curve is held fixed throughout the scenario, while the demand curve varies with the current demand. For the non-DR case, the curve is scaled uniformly with total demand; for the ColorPower case, the nine priority tiers are mapped in color order to the nine levels making up the demand curve in Fig. 4, and each set of three same-color priority tiers is scaled together according to the color distribution in the demand scenario (e.g., at 10 A.M. the three red priority tiers contain 3.08 MW at \$150/MWh, 1.71 MW at \$120/MWh, and 1.88 MW at \$90/MWh). The demand is spread over 100 000 heterogeneous device loads (some of which might actually be different service levels within the same device), with initial loads varying by 19× between the largest and smallest devices

(0.1 kW to 1.9 kW, with a mean of 1 kW), and a distribution of device demands over tiers proportional to the initial distribution of aggregate demand over tiers. For simplicity, we assume that the per-device demand and tier demand scale identically over time, and model each load in terms of its mean load across a duty cycle³.

For simplicity of modeling, we treat all power as though dispatched by a real-time market (as enabled by ColorPower). The demand of all 100 000 devices is thus bid together as a set of 9 blocks, with the same price applying to all loads within a given tier. The market only indirectly controls which devices will be enabled, however, through the setting of $g(t)$: which tiers are actually enabled is affected by the distribution of refractory (non-controllable) demand. At each control round, the goal $g(t)$ for the ColorPower algorithm is set by comparing the current supply and demand curves and setting $g(t)$ to the market equilibrium (where the two curves cross). Thus, when demand is low, there will be no curtailment of demand, but as demand rises the low-priority tiers gradually become uneconomical and DR is applied to improve market efficiency. In a real deployment, of course, a similar effective market structure would be assembled from supply and demand bids on various time scales with the assistance of forecasting or other market simulation applications.

The ColorPower control algorithm is simulated with the following parameters: control rounds are 10 seconds long, error in state estimates is $\pm 0.1\%$, and initially all demand is enabled. Every tier has identical parameters: refractory time variables T_{EF} , T_{EV} , T_{DF} , and T_{DV} are all 40 rounds, flexible reserve ratio $f = 1$, and proportional control constant $\alpha = 0.8$.

Fig. 5(a) shows the results of simulating ColorPower control of this scenario under goals set by a real-time energy market. When demand is low, no devices are curtailed; when demand is high, up to 25% of demand is curtailed. As total demand rises, enabled demand hits plateaus where a particular demand tier is gradually curtailed. These plateaus appear whenever the demand curve shifts to require a higher marginal supply price and that supply price is above the bid of the lowest priority enabled tier. As demand continues to rise, that tier is gradually curtailed, until it is entirely consumed and it is once again more valuable to increase supply than to curtail demand. Plateaus similarly appear as demand falls, in regions where supply price has fallen enough that it is once again more valuable to increase demand than to decrease supply.

In this scenario, only “green” power is being curtailed, so with a typical device mixture that curtailment would likely correspond mainly to temporary shutdowns of pool pumps, small drops in the temperature of hot water in water heater tanks, and small rises in temperature in the houses of participating customers as their HVACs relax service constraints, with the amount of rise depending on the customer’s choice of how to configure their temperature/color tradeoff.

The more efficient market enabled by ColorPower priority tiers leads to a total energy cost of \$7522.06. In contrast, the conventional scenario has no ability to price discriminate and

virtually no DR, and thus has a significantly higher total energy cost of \$8563.84. Using ColorPower in this scenario thus gives an increase in market efficiency of 12.2%. Note also that post-event rush-in, a potentially severe problem for both traditional DR and price signal-based control systems, which can severely impact grid and market operations, is managed gracefully by ColorPower’s ongoing feedback control.

The round-by-round control provided by ColorPower is quite precise. The only significant deviations from goal occur during the transition between boundary tiers, when the goal approaches the limit of the current boundary tier faster than refractory demand returns to flexibility, but the algorithm has not yet been given license to control demand from the next tier. Even so, the system never departs by more than 729 kW from the goal and has a mean error of only 41 kW. As for the other constraints: fairness is never violated, while tier priority is always maintained, except for brief periods during the change of boundary tiers, when approximately 200 kW is temporarily “borrowed” from the new boundary tier. Cycling is also effectively suspended during these transitions.

For illustrative comparison, we also simulate a conventional event-based demand controller. In this scenario we assume that demand is grouped into the same priority tiers, but switch entire tiers off or on with a single “event” signal whenever less than 50% of the tier can be served within the current goal. The observed behavior, shown in Fig. 5(b), has almost identical economic value, with an energy cost of \$7516.48. More and larger transients occur, however, due to the coarse granularity of control: the system is within 1 MW of the goal only 77.5% of the time and at its worst is over 4.7 MW off, with a mean error of 624 KW. Goal tracking is thus badly violated and cycling and fairness rendered moot, even though tier priority is always precisely maintained. This might, of course, be improved to give performance more similar to ColorPower by sending event signals to smaller groups of devices and greatly increasing the communication requirements (and therefore the per-device costs of deploying such a system)—for a detailed discussion of this tradeoff, see [5], [6]. Note also that this comparison assumes the same flexibility tiers, while conventional control that switches whole appliances, houses, or neighborhoods, is likely to provide much less availability of flexibility because it is more intrusive.

These results thus indicate that the ColorPower approach, when integrated appropriately with energy markets, should have the capability to significantly improve both market efficiency and grid reliability by flexibly and resiliently shaping demand. The key open question, to be established by field trials in future work, is how deep a penetration of mass-market DR can be enabled by the sub-appliance regulation enabled by ColorPower.

VI. GENERALIZATION BEYOND DEMAND RESPONSE

Although this paper has focused primarily on DR, ColorPower has the potential to offer more general demand management services, which can also help manage other resources such as dispatchable demand and distributed supply.

For example, currently generation resources can be designated as “load-following.” With ColorPower, certain load resources can be designated as “supply-following,” and matched

³This is a safe assumption because the number of loads is large and the randomization of refractory time assures desynchronization of duty cycles.

against generation from intermittent renewable energy sources (RES), such as solar and wind. This can help with balancing and ramping challenges for large-scale RES deployment, instead of the current approach of solely relying on expensive flexible conventional generation, which greatly increases costs and can undermine environmental goals.

Supply-following loads can also be dispatched as opportunistic or emergency power sinks. These can be coupled to real-time telemetry signals, such as substation SCADA demand or wind farm telemetry. For example, a demand tier consisting of “water heaters that raise reservoir temperature to consume excess renewable power” can wait at low or negative buy bids reflecting system conditions under which excess generation needs to be consumed. This type of demand can be used to consume excess power from hydro plants or excessively windy days, with customers paid by generation or transmission to consume excess power and assist the utility in avoiding negative consequences such as contractual penalties or environmentally disruptive dam overspill conditions.

The same approach demonstrated above can be easily extended to include other distributed energy resources (DER) and storage resources, such as solar, combined heat and power (CHP), gas turbine, and other “behind-the-meter” distributed resources that can be modeled as responsive resources. These could be aggregated into economically significant blocks by extending ColorPower to include supply tiers, then bid into markets as supply or demand segments to compete against imported power resources, using economic models like the ones proposed in [20]. If configured as a supply resource, they can bid against other supply resources normally; if configured as a demand resource, they can be interleaved with DR resources to create mixed dispatch orders, e.g., a generation or storage resource might be in a priority tier dispatched before more disruptive emergency DR resources are used.

Storage resources combine characteristics of dispatchable demand and generation: their demand or supply power corresponds to their rate of discharge or charge, while the length of time each resource may be available depends on the amount of energy currently stored. Storage resources can thus be aggregated using ColorPower and bid in at opportunistic prices, like dispatchable demand, to absorb excess supply, and assigned to an appropriately priced priority tier to determine when they discharge. This can be further refined by assigning different levels of storage to different priority tiers, so that, for example, a storage resource might always maintain a small operating reserve to be sold into an ancillary services market.

VII. CONCLUDING REMARKS

The current status of DR programs is not satisfactory due to a combination of obstacles, including utility incentives, market design, regulation, technology, and program design. This paper describes how the ColorPower approach to DR, based on tiers of appliance services and distributed stochastic control, can overcome these obstacles to penetrate the mass-market of residential and small-business consumers. We have shown how the our approach can be integrated with energy markets to achieve

the same economic, reliability and balancing goals as dynamic pricing and traditional DR, with superior potential capabilities and without their drawbacks. We have also simulated the operation of a ColorPower system comprising 100 000 devices integrated with energy supply and demand curves that are used to set the DR goal, demonstrating the capability of the this approach to quickly and resiliently shape demand without exposing the customer to either dynamic pricing or coercive load control. Together, these results show that ColorPower is a good candidate for overcoming the obstacles that have preventing DR from penetrating the mass market and realizing its massive potential value for utilities, customers, and energy markets.

REFERENCES

- [1] “Demand resource compensation in organized wholesale energy markets,” FERC Docket No. RM10-17-000, Order No. 745, Mar. 15, 2011.
- [2] “ERCOT event on February 26, 2008,” National Renewable Energy Lab., Tech. Rep. TP-500-43373, Jul. 2008.
- [3] Energy Policy Act of 2005, Pub. L. No. 109-58, Section 1252(f) 119 stat. 594, 965.
- [4] J. Beal, J. Berliner, and K. Hunter, “Fast precise distributed control for energy demand management,” in *Proc. IEEE SASO*, Sep. 2012.
- [5] V. Ranade and J. Beal, “Distributed control for small customer energy demand management,” in *Proc. IEEE SASO 2010*.
- [6] V. V. Ranade, “Model and control for cooperative energy management,” M.S. thesis, Massachusetts Institute of Technology, Cambridge, MA, USA, Jun. 2010.
- [7] H. Mitomo and T. Otsuka, “Consumers’ preference for flat rates: A case of media access fees,” in *Proc. Int. Telecom. Soc., 17th Biennial Conf.*, 2008.
- [8] D. Hammerstrom *et al.*, “Pacific Northwest Gridwise Testbed Demonstration Projects: Part I. Olympic Peninsula Project,” Pacific Northwest Natl. Laboratory, Richland, WA, USA, Tech. Rep. PNNL-17167, Oct. 2007.
- [9] D. Hammerstrom, T. Oliver, R. Melton, and R. Ambrosio, “Standardization of a hierarchical transactive control system,” in *Proc. Grid Interoperability Conf.*, Nov. 2009.
- [10] S. Widergren, K. Subbarao, D. Chassin, J. Fuller, and R. Pratt, “Residential real-time price response simulation,” in *Proc. IEEE Power Energy Soc. Gen. Meet.*, Jul. 2011, pp. 1–5.
- [11] H.-P. Chao and R. Wilson, “Priority service: Pricing, investment, and market organization,” *Amer. Econ. Rev.*, vol. 77, no. 5, pp. 899–916, Dec. 1987.
- [12] S. Oren, S. Smith, and R. Wilson, “Service design in the electric power industry,” Electric Power Research Inst., Palo Alto, CA, USA, Tech. Rep., 1990, Pricing Strategy Associates, Orinda, CA, USA.
- [13] S. Darby, “The effectiveness of feedback on energy consumption. A Review for Defra of the literature on metering, billing and direct displays,” Environmental Change Institute, University of Oxford, Tech. Rep., 2006.
- [14] S. Borenstein, “Real-time retail electricity prices, theory and practice,” in *Electricity Deregulation, Choices and Challenges*, J. Griffin and S. Puller, Eds. Chicago, IL, USA: Univ. Chicago Press, 2005, pp. 317–357.
- [15] A. Faruqui, “2050: A pricing odyssey,” *Electr. J.*, vol. 19, no. 8, pp. 3–11, Oct. 2006.
- [16] K. Costello, “An observation on real-time pricing: Why practice lags theory,” *Electr. J.*, vol. 17, no. 1, pp. 21–25, Jan.–Feb. 2004.
- [17] E. Hirst, “Price-responsive demand in wholesale markets: Why is so little happening?,” *Electr. J.*, vol. 14, no. 4, pp. 25–37, May 2001.
- [18] F. A. Wolak, J. Bushnell, and B. F. Hobbs, “The California ISO’s proxy demand resource (PDR) proposal,” Market Surveillance Committee of the California ISO May 1, 2009.
- [19] A. Faruqui and J. Palmer, “The discovery of price responsiveness: A survey of experiments involving dynamic pricing of electricity,” *EDI Quart.*, vol. 4, no. 1, Apr. 2012.
- [20] H. Kim and M. Thottan, “A two-stage market model for microgrid power transactions via aggregators,” *Bell Labs Tech. J.*, vol. 16, pp. 101–107, 2011.



of clients such as regulators, governments, utilities, and ISOs. He has designed some of the most complex energy markets in the world. Prior to forming ECCO International he was a Director of the Electric Industry Restructuring Group at the Pacific Gas and Electric Company, San Francisco. He has made substantial contributions in the areas of network grid optimization and pricing, energy market design, and competitive bidding, and implementation of EMS applications and real time control functions. He has published numerous scientific papers in IEEE and other journals. Dr. Papalexopoulos is the 1992 recipient of PG&E's Wall of Fame Award, and the 1996 recipient of IEEE's PES Prize Paper Award. He is also President and CEO of ColorPower, a clean tech startup company focused on research, development and commercialization of demand side management technologies.

Alex D. Papalexopoulos (M'80–SM'85–F'01) received the electrical and engineering Diploma from the National Technical University of Athens, Greece, and the M.S. and Ph.D. degrees in electrical engineering from the Georgia Institute of Technology, Atlanta, GA, USA. He is President and Founder of ECCO International, San Francisco, CA, USA, a specialized energy consulting company which provides consulting and software services on electricity market design and system operations and planning within and outside the U.S. to a wide range



Intelligence Laboratory; he is also a research affiliate of MIT CSAIL and a Science Commons Fellow.

Jacob Beal is a scientist at Raytheon BBN Technologies, Cambridge, MA, USA, and the inventor of ColorPower. His research focuses on the engineering of robust adaptive systems, and particularly on the problems of aggregate-level modeling and control for spatially distributed systems like residential energy demand, robotic swarms, and natural or engineered biological cells. Dr. Beal completed his Ph.D. in 2007 under Prof. Gerald Jay Sussman at the Massachusetts Institute of Technology, Cambridge, MA, USA (MIT) Computer Science and Artificial



tions at Siebel Systems, Omnicom, and Countrywide Home Loans. He has wide-ranging interdisciplinary interests across economics, marketing, physics, algorithms, artificial intelligence, and information systems architecture. After attending the Illinois Institute of Technology, Mr. Florek founded one of the first public-facing Internet Service Providers A*Net in 1992.

Steven Florek is the R&D director of ColorPower, Berkeley, CA, USA. He has consulted across nearly 100 Fortune 500 companies in nearly every industry, delivering breakthrough solutions to organizations to achieve market leadership and delight customers. Prior to ColorPower, Mr. Florek led product design efforts for a GPS ankle bracelet tracking system for parolees, a public venue incident management system featuring WiFi-based locationing, and predictive crime analytics system for law enforcement.