Defining CPS Challenges in a Sustainable Electricity Grid

Jay Taneja, Randy Katz, and David Culler
Computer Science Division
University of California, Berkeley
Berkeley, California 94720
{*taneja,katz,culler*}*@cs.berkeley.edu*

***Abstract*—Cyber-Physical Systems (CPS) are characterized as complex distributed systems exhibiting substantial uncertainty due to interactions with the physical world. Today's electric grids are often described as CPS because a portfolio of distributed supplies must be dispatched in real-time to match uncontrolled, uncertain demand while adhering to constraints imposed by the intervening transmission and distribution network. With the increased control complexity required by deep penetration of fluctuating renewable supplies, the grid becomes more profoundly a CPS and needs to be addressed *as a system*. In this evolving CPS, a large fraction of supply is under-actuated, a substantial portion of demand needs to become dispatchable, interactions among distributed elements are no longer unidirectional, and operating requirements of elements are more dynamic. To more sharply define these CPS challenges, we obtain a yearlong, detailed measurement of the real-time blend of supplies on the primary California grid dispatched to meet current demand and then scale the solar and wind assets, preserving uncontrolled weather effects, to a level of penetration associated with California's 2050 GHG targets. In this representation of a future sustainable grid, we assess the impact of demand shaping, storage, and agility on the reconstituted supply portfolio, characterize resulting duration curves and ramping, and investigate the distributed control and management regime. We articulate new operational and market opportunities and challenges that may materialize from intermittent periods of abundance and scarcity in the overall energy network. We find that in a sustainable grid, lulls in renewable production during winter are more critical than peaks in demand during summer, capacity for load shifting and energy storage are more valuable as renewables penetration increases, and that grid balancing requires integrated management of supply and demand resources.**

***Keywords*-electricity; cyber-physical systems; smart grid; renewable energy**

I. INTRODUCTION

Modern electric grids serve as a canonical example of a cyber-physical system (CPS); they comprise a complex distributed system in which a portfolio of electric power generation resources must be managed dynamically to meet an uncontrolled time-varying demand. This demand must be met while adhering to constraints imposed by the transmission network, generator ramp rate capabilities, and emissions limits, with sufficient reserve to handle faults and failures. The primary control loop is typically realized through a system operator solving an iterative unit commitment problem; based on a prediction of load, an assignment of generation capacity is made on an hour-by-hour basis through a day-ahead auction that accounts for transmission limits and losses and maintains a certain amount of reserve. The matching of generation to load is refined through hour-ahead and 5-minute-ahead markets based on recently observed demand. The feedback from load to generation is ultimately manifested through power quality observations, i.e., frequency fluctuations and voltage deviations resulting from any mismatch. Generators respond to mismatches by engaging more or less of the reserve. Commercially, there are widespread efforts to utilize information technology to improve the efficiency and effectiveness of this CPS through so-called "smart meters" that monitor loads and report on 15-minute intervals, synchrophasors that observe power quality at intermediate points in the transmission grid, and delivery of pricing signals to trigger a response from demand. These efforts begin to introduce information planes to augment the physical planes of classic electric grids [24].

However, the integration of large amounts of fluctuating renewable resources, such as wind and solar, make the control of future electric grids fundamentally more challenging. In addition to uncertainty in demand, supplies are no longer completely dispatchable and often do not replicate the slow ramp rates and high inertia that characterize traditional generation. Some argue that to achieve deep penetration of renewable resources, the operational model of the grid must be turned around, changing the paradigm from load-following supplies to supply-following loads [18]. In contrast to the utility-centric grid that arose from the industrial revolution, this new model has been called a consumer-centric grid, with CPS elements of networking and control more prominent. Indeed, a recent study by the California Council on Science and Technology on how to meet California's GHG emissions target of 80% below 1990 levels by 2050 concluded that the largest leap in realizing this goal is technology to maintain the balance between load and available supply, termed "Zero-Emissions Load Balancing" [9]. The report notes even if only the state-legislated minimum of 33% renewable energy is met, firming these fluctuating renewables with fossil fuel-based generation, often intermittent and peaker plants, would alone exceed the entire 2050 emissions allotment, notwithstanding emissions from the remainder of electricity and transportation fuels.

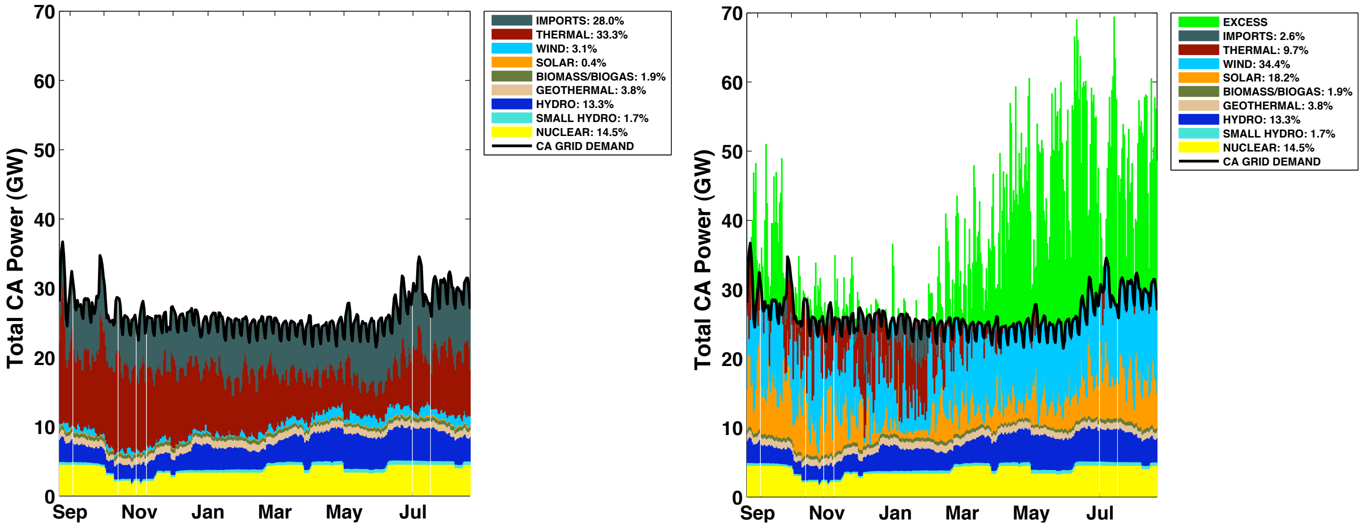

(a) Present-Day Grid

(b) Grid Scaled to 60% Renewables

Figure 1. California electricity blend for a year, August 2010-August 2011. Note that biomass and biogas are combined.

	Present-Day Grid			*Scaled Scenario - 60% Renewables*		
Generation Type	**Capacity / Peak (GW)**	**Total Energy (%)**	**Capacity / Load Factors (%)**	**Capacity / Peak (GW)**	**Total Energy (%)**	**Capacity / Load Factors (%)**
Renewables						
Geothermal	2.600 / 1.095	3.8%	38.7% / 92.0%	2.600 / 1.095	3.8%	38.7% / 92.0%
Biomass/Biogas	1.145 / 0.616	1.9%	43.5% / 80.9%	1.145 / 0.616	1.9%	43.5% / 80.9%
Small Hydro	1.380 / 0.646	1.7%	31.7% / 67.8%	1.380 / 0.646	1.7%	31.7% / 67.8%
Wind	2.812 / 2.470	3.1%	29.1% / 33.2%	57.116 / 22.995	34.4%	15.8% / 39.3%
Solar	0.403 / 0.457	0.4%	28.7% / 25.3%	29.792 / 30.636	18.2%	16.1% / 15.7%
Non-Renewables						
Nuclear	4.456 / 4.581	14.6%	86.0% / 83.6%	4.456 / 4.581	14.5%	86.0% / 83.6%
Hydro	12.574 / 6.286	13.3%	27.7% / 55.5%	12.574 / 6.286	13.3%	27.7% / 55.5%
Imports	N/A / 11.055	28.0%	N/A / 66.6%	N/A / 9.291	2.6%	N/A / 7.2%
Thermal	44.339 / 27.014	33.3%	19.7% / 32.4%	44.339 / 19.528	9.7%	5.7% / 13.0%
Total	**80.764 / 47.128**	**100.0%**	**32.6% / 55.8%**	**130.882 / 47.128**	**100.0%**	**20.1% / 55.8%**

Table I

SUMMARY OF GENERATION IN CALIFORNIA - PRESENT-DAY AND SCALED TO 60% RENEWABLES.

As a small step toward meeting this challenge, we seek to frame how CPS methods can be brought to bear on the problem of maintaining the dynamic match between supply and demand in a grid with a deep penetration of renewables. The techniques we rely on are common to many CPS problems: pervasive monitoring, modeling, and mitigation employing a rich information plane and distributed intelligence. Specifically, we provide an empirical characterization of the dynamics of a sustainable grid in order to identify key future challenges in this type of CPS and how they differ from those currently being addressed in the grid.

The empirical basis for this study is a recently released trace of all categories of generation assets under management by the California Independent System Operator (CA ISO). From this, we perform a scaling study to articulate the dynamics on the scale of the entire state if a much larger array of renewable assets were deployed, subject to the same uncontrolled environmental variations. We then characterize the future duration curves that will drive operations and capital investment. From these curves, we assess the impact of hypothetical electrical storage capacity, demand management strategies to effect supply-following, and coordinated management of the generation portfolio, and lay out the CPS challenges inherent in each area. We also see opportunities for energy-intensive but agile industrial processes.

We recognize that the specific mix and dynamics of available renewable resources differs substantially around the world, so particular observations are not expected to be universal, but the methodology of obtaining them should be broadly applicable. Also, we conduct a lumped study at the scale of the entire state unconstrained by the particular transmission infrastructure currently in place. The equivalent study to capture these constraints would need to be conducted at a regional level with finer granularity. Furthermore, we have not sought to forecast changes in the demand profile due to population growth, economic activity, or efficiency measures, as is typical of policy-focused studies, which also tend to utilize statistical characterizations of the supply and demand, rather than attending to their temporal dynamics. Finally, we do not impose any particular economic framework or constraints upon the study. We ask, "what would the grid be like with a deep penetration of renewables today?" Hopefully, understanding that will give rise to the investigation of economic mechanisms effective in managing a grid like that, despite substantial differences from the highly-constrained mechanisms in place today.

II. THE CALIFORNIA ELECTRICITY GRID

Partial deregulation of the California electricity grid in the mid-1990s required the sale of some power generation stations by the small handful of electric utilities to a number of independent power producers. Balancing authority over 80% of the electricity grid of California was granted to a single entity, the California Independent System Operator (CA ISO). CA ISO is an independent, non-profit corporation that monitors, coordinates, and controls the electrical power system in the state, managing electricity flow across the 25,865-mile transmission network and operating wholesale power markets [2]. Recently, CA ISO released hourly supply data for its ten different types of generation sources starting from August, 2010 [11].

Figure 1(a) shows a yearlong breakdown of the sources of electricity consumed in the CA ISO operating region, from August, 2010 to August, 2011. Overall demand, defined here as the sum of these generation sources, is also plotted; we do not consider the slight discrepancies in demand and supply that are often addressed with grid ancillary services. Electricity demand (and, thereby, production) varies on multiple timescales: *daily* with peaks in the late afternoon and nadirs in the middle of the night, *weekly* with weekends on average 9.6% lower than weekdays, and *seasonally* with winter load on average 15.8% lower than air-conditioner-driven summer load. Though the magnitudes of these variations are unique to California, the challenge of matching highly-variable electricity demand with a portfolio of generation resources is common to all electrical grid operators.

A crucial driver to the creation of the existing blend of generation resources is the state energy policy that governs the economics of supply deployment – chiefly, California has mandated a renewables portfolio standard (RPS) target, whereby 33% of the state's generated electricity (in energy, not capacity) must come from renewable sources by the year 2020 [4]. This builds upon a previous RPS mandate of 20% of generation by 2010 [3].

The left columns of Table I summarize the characteristics of these sources over the year based on the hourly data. The capacity is based on ratings provided by the California Energy Commission (CEC) for each of the 1007 generators under CA ISO management with the exception of Imports, which reflects power purchased from other operators in the Western U.S. interconnect (WECC) [10]. Capacity factor is the ratio of mean delivered power to rated power, and load factor is the ratio of mean delivered power to peak power.

Of these generation sources, the renewables proportion consists of wind, solar, geothermal, biomass, biogas, and small hydroelectric (only facilities less than 30 MW) generation. Additionally, renewables also comprise an unpublished proportion of the imported energy, which itself comprises 28% of the total generation during the year under study. The imported energy comes from the U.S. Southwest (primarily coal and nuclear, with some solar) as well as the U.S. Pacific Northwest (mainly coal and hydroelectricity). Thermal is a mix of mostly natural gas combined-cycle and single-cycle plants providing base, intermittent, and peaker capacity.

This simple characterization demonstrates that the overall capacity factor of the thermal plants is lower than that of the renewables. This contrasts with common wisdom that three times as much solar or wind capacity must be deployed as compared to fossil fuel capacity. The critical difference is the degree of control over when that capacity is dispatched.

The problem that receives the most attention in management of the CA grid today is the yearly peak demand driven largely by the spike in summer afternoon air conditioning load. The capital investment to meet this summer peak is enormous. The highest 10% of demand (in excess of 42 GW) persists for only 41 hours a year (0.47%) and the last 20% for only 183 hours (2.1%). The generation data show that the problem of planning for load-following is even worse; an additional 33 GW of capacity are never called into use [10].

A. Temporal Variations

Several facets of Figure 1(a) reveal features of the dynamics of the grid resources. First, the nuclear generation used as a stable baseload resource in California comes from two facilities, each with two operating units. This small number of units statewide results in clearly quantized power states due to unit maintenance or repair. Second, hydroelectric sources produce more electricity in the summer, coinciding with melting of the snowpack in the Sierra Nevada mountain range. Third, much of the day-to-day variation in overall demand is met by thermal generation, which primarily includes facilities fueled by natural gas, but also some fueled by petroleum coke and coal; though many of these facilities can provide consistent baseload power, their operation suggests they are being dispatched as load-following and peaking generation to cope with demand variability. Barely visible in the figure, solar and wind, despite recently being the fastest growing sources in the generation blend, comprise only 3.5% of the total generated energy.

Figure 2 shows five days of generation in summer and winter. Note the clear difference in daily peak power – the full summer and winter averages are 35.6 GW and 29.1 GW, respectively – but also note the relative similarity of the nightly minima – the summer and winter averages are 22.6 GW and 20.2 GW. Among the resources used to cope with this variability are the aforementioned thermal generation, but also large hydroelectric generation, which can provide

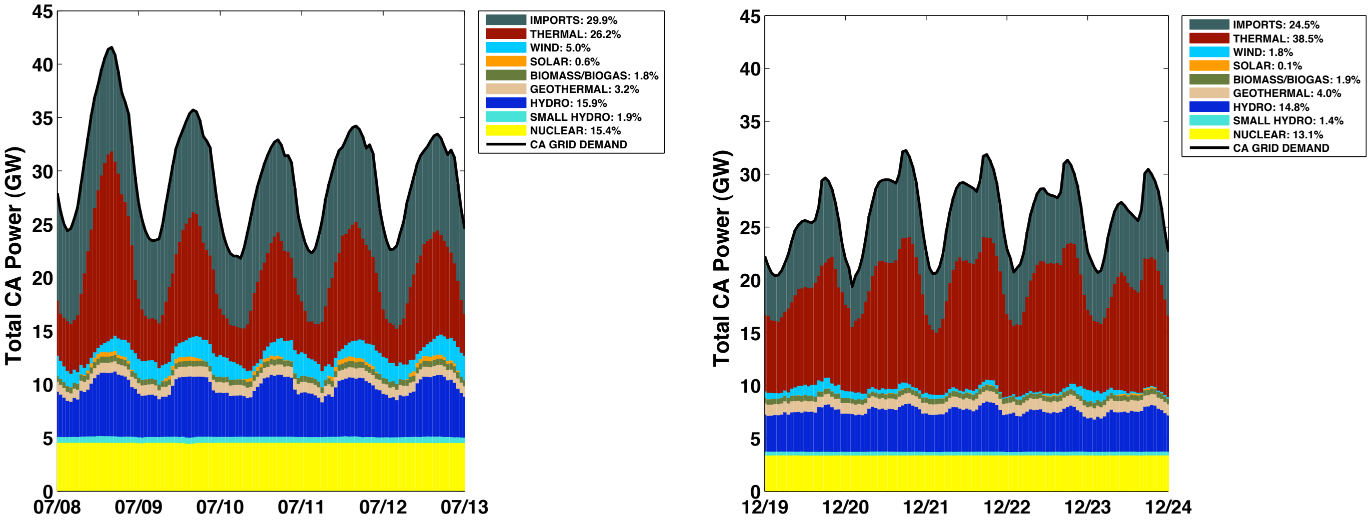

(a) Summer

(b) Winter

Figure 2. Five days of present-day California electricity blend in the (a) summer and (b) winter.

extra power during peak hours through scheduled generation and can be used to match supply and demand. However, we also see the characteristic patterns of the non-dispatchable renewable sources, solar and wind. Since much of this study is concerned with these generation sources, we examine them more thoroughly.

B. Solar and Wind Power in California

At this time, there are 403 MW of utility-operated solar capacity within the CA ISO operating region. This does not include residential or commercial rooftop solar photovoltaic (PV) deployments, of which there is currently over 1 GW of capacity in the state [13]. The energy generated by these panels instead offsets electric demand within the local distribution tree using grid-connected inverters.

Of the solar capacity in the CA ISO region[1], over 400 MW are from parabolic trough solar thermal facilities in southern California, with the balance coming from two solar PV sites. Additional facilities in planning and construction phases will employ advanced solar technologies, including concentrating solar thermal and PV.

All of the 2.8 GW of wind power farms in California are onshore, consist of low (<150m) wind turbines, and are located in 8 of California's 58 counties with roughly half the capacity in each of Northern and Southern California. Additionally, the advent of new technologies, such as high-altitude and offshore wind turbines, will enable massive potential wind resources.

Looking more closely at the temporal variations in these non-dispatchable sources, there are important seasonal and daily patterns. Solar generation peaks during the daytime but it varies by season, as its power profile is dictated by orientation and tilt relative to Solar Normal. Its gross features are very predictable, but occlusions such as clouds cause rapid ramps. Most of the solar production in California is from solar-thermal generators, so ramps are damped and generation can be delayed somewhat into the evening after the sun has gone down. Wind, on the other hand, provides power at more and varying hours of the day, but is less predictable and tends to have larger ramps in generation [28].

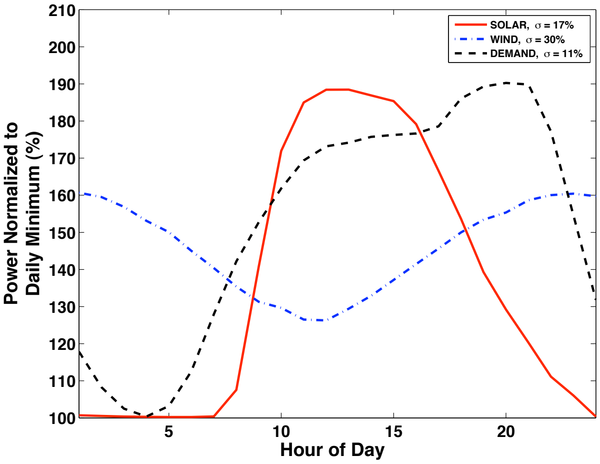

Figure 3. Mean of solar generation, wind generation, and demand normalized to daily minimum over the hours of the day for a year. The mean standard deviation for all hours of each curve shows relative variability.

Figure 3 compares the mean total demand, solar generation, and wind generation for the hours of the day over the year. All of the data are normalized to represent a percentage of the daily range of each category, independent of scale. The shoulder and tail to the night in solar generation shows the impact of solar thermal generation. Wind power does tend to be higher at night and lower during the day; some have argued that the combination of wind and solar provides a good match to demand [15]. Still, the evening peak in demand represents a challenge. Furthermore, wind is less

---

[1] Due to data gathered from two sources, solar peak generation exceeds rated capacity. This does not alter our observations.

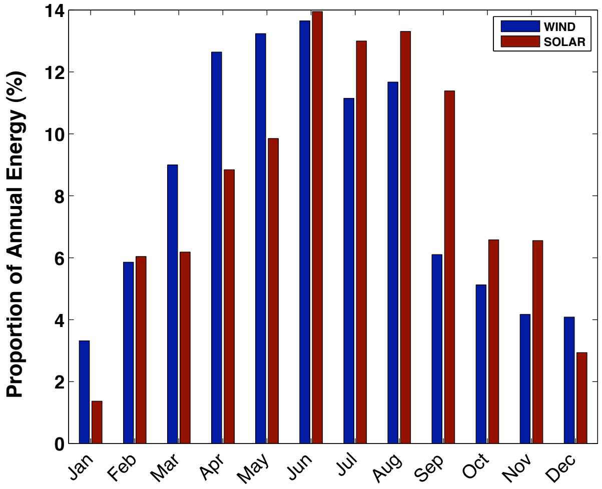

Figure 4. Monthly proportion of annual energy contribution for solar and wind generation sources.

stable – the mean standard deviation across the hours of the normalized wind generation is 30%, slightly less than double the corresponding value of 17% for solar generation.

Seasonally, solar production peaks in the summer when the sun is at its most intense. Studies on the U.S. wind resource indicate that its winter production potential is double its summer potential [21]. However, the data presented in Figure 4 for California's wind and solar production show the opposite: far more wind energy was produced in the summer. As this analysis deals with only a year of data and is subject to the specific installations in the state, care should be exercised in drawing conclusions, nonetheless it underscores the importance of understanding the profile of local renewable resources. As these sources become a more significant proportion of the grid mixture, seasonal variability of wind and solar becomes of increasing importance.

Having established an understanding of how a large grid operates today, we continue by modeling and examining possible future generation blends and their inherent challenges. By scaling the elements of the grid blend differently for a decarbonized grid, we add to the body of work that lays out scenarios with renewables comprising 33% [12], [17] and beyond [9]. Creating arbitrary yet temporally representative mixes of generation sources, this study allows data-driven evaluation of the ability of the electricity grid to meet our changing goals, enables assessment of conventional wisdom in grid planning, and frames a wide array of critical CPS challenges for evolving electricity grids.

III. TOWARDS A SUSTAINABLE GRID

Starting from the relatively fine-grained monitoring of present-day supplies (hourly by type) over a considerable period of time (a year), we can scale various aspects in order to characterize how the CPS challenges change as we approach a sustainable grid. This complements the many policy-oriented studies which operate on statistical models of energy availability and usage by focusing on the temporal dynamics that are of essence in developing a viable control regime with deep penetration of fluctuating renewables.

A. *Scaling Methodology*

To model a grid on the scale of the CA ISO region with a large fraction of renewable energy in the blend, we take our yearlong hourly time series and scale the non-dispatchable components, solar and wind, each by a constant factor. Then we reduce the primarily fossil fuel-driven components, imports and thermal, until the demand – the sum of the original blend – is met. Any power produced beyond present-day demand, i.e., after thermal and imports are fully displaced, is identified as *excess*. All other components of the energy blend are left unchanged. Many other scaling rules are possible, each with advantages and disadvantages, so before further analysis we acknowledge our basic assumptions.

Proportional scaling of renewables in this manner preserves the temporal variations of the existing wind and solar installations at hourly through yearly timescales. However, these dynamics are dictated partly by the current geographic distributions of resources and the particular design goals of those deployments, such as technology selection and orientation. For solar, deployment throughout a north-south-oriented state like California results in a distribution of lengths of day. We expect geographic diversity of the future solar resource to be similar, as much potential prime solar deployment area remains in the lightly populated desert areas of Southern California. Though, it may be optimized more for winter production rather than summer peaks. For wind, geographic diversity is important for reducing overall variability of the resource [16]. Existing wind diversity arises from deployments split equally between Northern and Southern California, which is likely to remain.

The choice of which resources are displaced by increases in renewables is a policy decision, equivalent to the loading order. It is influenced by policy goals, contracts, and economic factors. To meet GHG targets, we seek to displace high-carbon fossil fuel resources, i.e., thermal generation. Currently, load-following is accomplished by a mix of thermal, import, and hydro generation; the imports may have a mix of renewables. Furthermore, the very notion of imports to (or exports from) the grid suggests that it exists in a larger context which may be engaged in its own GHG optimization. Taking the view that the CA ISO region is sufficiently large that it is reasonable to try to achieve energy balance within it, we first displace imports and then thermal generation, but recognize that in reality it would be desirable to have a more sophisticated interplay, say to flatten the duration curve of the remaining thermal generation.

We also maintain the existing demand profile. Though a litany of factors will alter demand – including population growth, changes in per capita electricity consumption, and the emergence of electric vehicles – estimates of the mag-

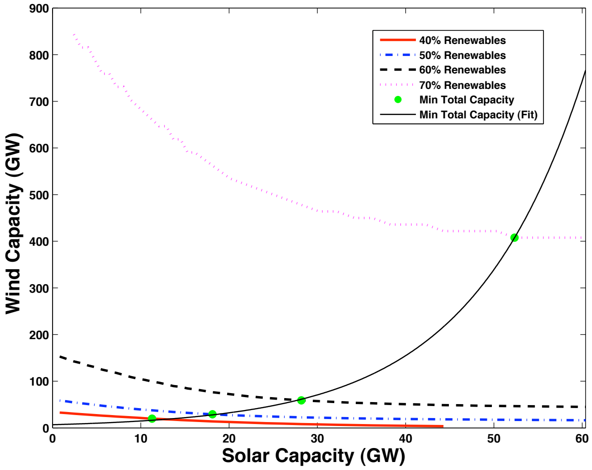

Figure 5. Mix between wind and solar at four renewables penetration levels. Curve is fit to points of minimum total installed renewables capacity.

nitude of change vary significantly. Additionally, since we make no projection on the time horizon of these changes and we only seek to elucidate issues that may arise in a future grid, we avoid demand scaling and focus on interactions of demand and supply dynamics.

Finally, in increasing the penetration of renewables, there are many possible blends that produce a specific renewables penetration level measured on an energy basis for the full year, as seen in Figure 5. The approach we take is to use the particular combination of solar and wind that minimizes the total installed capacity of these resources. This point is highlighted for each level of renewables penetration, and an exponential curve is fit to those points:

$$W_c = 6.7078 * e^{0.0785 * S_c} [GW], \tag{1}$$

where W_c and S_c represent wind and solar capacities in gigawatts (GW). In this study, we choose to select a renewables threshold of 60%, because this ensures that in addition to present-day baseload resources (hydroelectricity and nuclear), roughly 10% of energy is supplied by fossil fuel resources, similar to other studies [9]. Further, as the renewables penetration increases to levels beyond 60%, the installed capacity needed increases far more quickly as a result of temporal variations. The threshold of 60% represents a balance that emphasizes the challenges of variability without overwhelming other resources on the grid.

B. *Characterization of a Grid with 60% Renewables*

Figure 1(b) shows the temporal dynamics of a "scaled" grid supply blend over the year following the methodology above, with renewables scaled to 60% and a solar-wind mixture that minimizes the overall renewables capacity according to Figure 5. The resulting capacity and load factors are summarized in the right columns of Table I.

Looking at the resulting electricity breakdown over the year, we see the large effect of seasonal variations on the availability of renewables. The critical constraint is no longer summer peak demand, but rather the winter lulls in renewable generation. For much of the year, a significant amount of energy is produced in excess of demand. The new challenge in summer is to utilize the surplus of renewable energy, which could be sloughed, exported, or used to enable new energy-agile practices in industries using intermittent, seasonal, inexpensive electricity. The detail on five days of summer in Figure 6(a) shows that intra-day discrepancies between renewables generation and the present-day demand curve are no longer relevant, as there is excess production almost all day, with infrequent occurrences of fossil fuel-based generation needed to meet demand.

On the other hand, there are months of winter where thermal and even imported generation are necessary to meet demand, such as in Figure 6(b). This is a result of irregular solar availability and high variance in wind production coupled with reduced hydroelectric resources. Section IV-A provides further analysis of the effects this seasonal pattern has on fossil fuel generation.

IV. Opportunities for Cyber-Physical Systems

A grid with deep renewables, as represented by the dynamics in the scaled supply blend above, presents a family of CPS challenges and opportunities that go far beyond those in the grid in operation today. At core, the challenges arise from the shift from primarily modulated, dispatchable supply to primarily uncontrolled, non-dispatchable supply, and the resultant opportunities from the increased intelligence and communication needed to allow the energy network to function *as a system*. Here we explore some of the dominant CPS thrusts relative to the temporal dynamics of our year in a grid with deep renewables: the coordinated management of the entire portfolio, the potential to modulate (or dispatch) demand, the utilization of storage resources, and grid-driven demand reduction. Ultimately, all of these aspects need to come together in a manner that addresses the additional level of fidelity associated with transmission constraints, plant dynamics, demand adjustment mechanisms, and markets.

A. *Cooperative Portfolio Management*

A sustainable grid presents a family of CPS challenges associated with coordinating heterogeneous, distributed processes to better manage increasingly critical dispatchable resources. For instance, consider the duration curves in Figure 7 – the present-day thermal curve indicates the traditional need for critical peak demand response. A large fraction of dispatchable capacity is utilized only a tiny fraction of the time. By curtailing or shifting load away from these critical times, capital investment in generation can be greatly reduced. However, the impact of this volatile minority of the portfolio is amortized over the large body of supplies that are essentially providing baseload and seasonal adjustment.

In a sustainable grid, this dispatchable resource takes on a new and critical role of providing firming to renewables, which, recall from Figure 3, are more variable and less

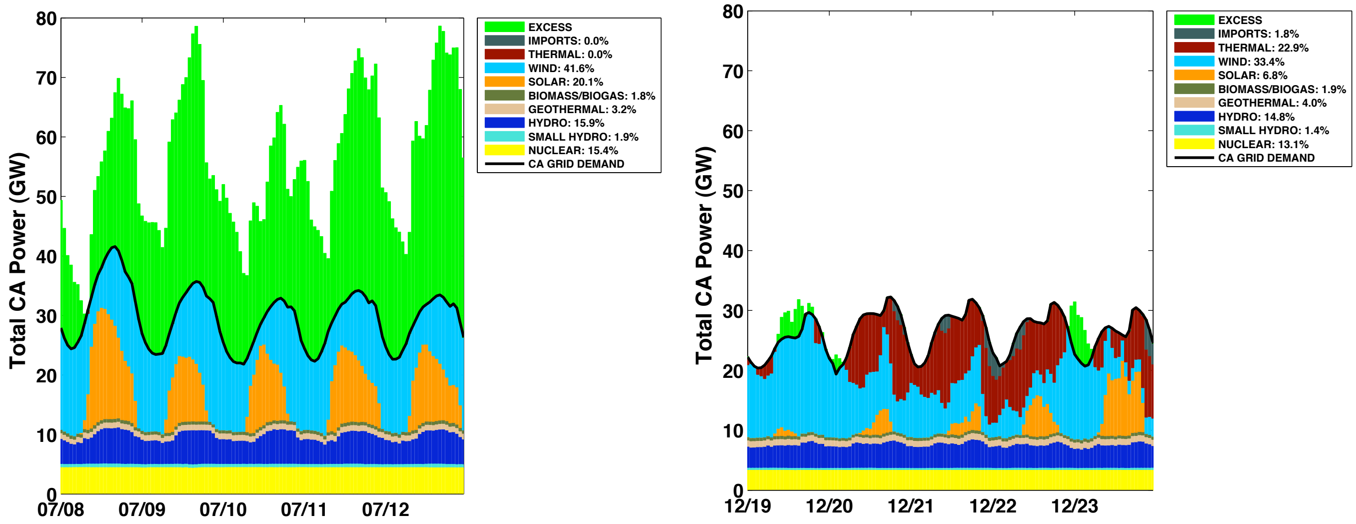

(a) Summer (b) Winter

Figure 6. Five days of California electricity blend scaled to 60% renewables in the (a) summer and (b) winter.

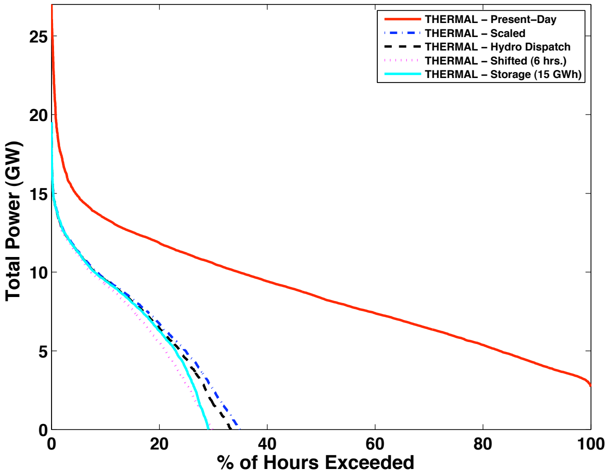

Figure 7. Comparison of duration curves for thermal generation before and after scaling, as well as after supply or demand modifications.

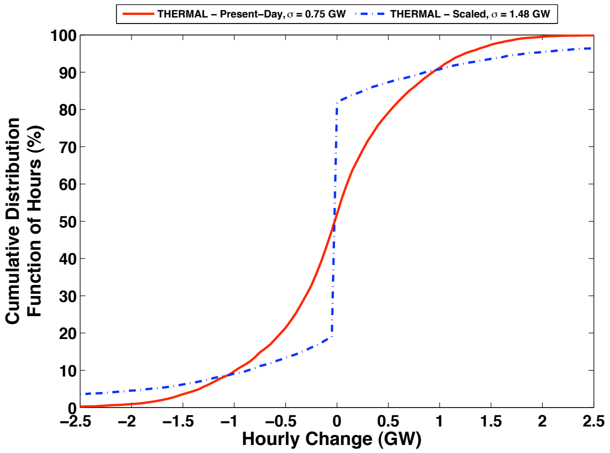

Figure 8. Comparison of distributions of hourly change for thermal generation before and after scaling. Increased variability in renewables results in less often but higher magnitude dispatch of fossil fuel resources.

predictable than grid demand. The nature of this can be seen in the thermal curve for the 60% renewable "scaled" blend in Figure 7. Most of the time, these resources are unused, but the reduction in peak, from 27 GW to 19 GW, is less than the reduction in overall usage. Further, the upper 25% of resources is used a tiny fraction of the time. More sophisticated coordination algorithms would use storage, load shifting, and curtailment to eliminate this sharp peak, even if overall usage of these resources increases marginally. The control regime of these dispatchable resources is likely to differ markedly from today, because it is not enough to only modulate reserve to preserve power quality, instead ramping some resources all the way off and back on becomes essential. However, the frequency of these transitions needs to be minimized for each individual plant.

Yet another constraint on this new CPS is the agility of dispatchable generation; a distribution of the hourly changes experienced by thermal generation is in Figure 8. Though the most common "change" in the scaled grid is zero, the changes have a longer-tailed distribution than those in the present-day grid. We cannot ascertain whether increased changes are feasible, as CA ISO would not release data on allowable ramp rates of the generation fleet. Nonetheless, we emphasize that the dispatchability and agility that fossil fuels provide are critical to operating a grid composed primarily of renewables. This presents a new role for fossil fuels: rather than providing baseload power and seasonal adjustments, fossil fuels are a precious storage resource that can be used only once, providing much-needed firming to cope with deep uncertainty in both supply and demand.

B. *Supply-Following Loads*

Existing grid architectures are predicated on the notion that operators can control generation supplies in order to match demand (load-following), but that while demand is predictable, the ability to shape it (its *elasticity*) is limited.

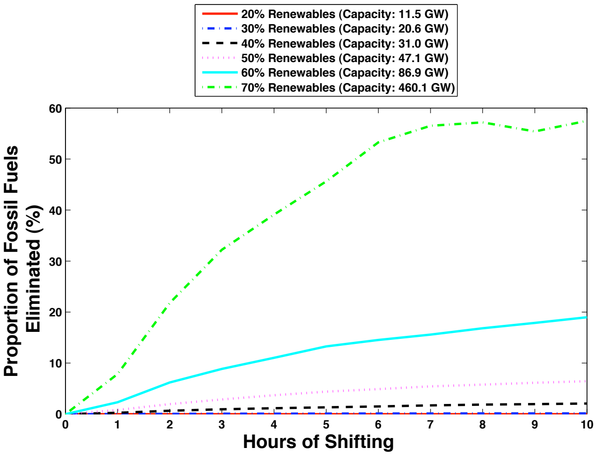

Figure 9. The effect of load shifting on the fossil fuels in the generation blend. Note that as renewables penetration increases, the absolute percentage of fossil fuels in the blend decreases.

Deep penetration of fluctuating renewables presents the dual challenge that the fraction of supply resources that are controllable is reduced and those resources are needed to firm renewables in addition to following demand. With increased communication and intelligence, it is particularly attractive to shift part of the burden of maintaining the match onto the demand side, i.e., to create a class of *supply-following loads*. A body of research work has explored such mechanisms for a wide variety of loads [5], [6], [8], [19], [22], [28] and substantial attention has been devoted to signaling protocols to achieve certain kinds of demand response [1]. Here we seek to quantify the potential gain of such approaches and to characterize the trade-offs in duration, magnitude, and agility that can guide their development.

Conceptually, a supply-following load is a mechanism that can shift demand away from periods of energy deficit to periods of energy surplus. For example, in Figure 6(b), shifting would allow loads adjacent to periods of excess to advance or delay operation in order to make use of renewable generation, thereby reducing imported and thermal generation. Periods with more excess have more potential for shifting, and increasing the time horizon allows more loads to shift away from fossil fuel-based generation.

To characterize the potential gain of such shifting as a function of the length of the shifting time horizon, we run a fairly aggressive shifting algorithm over the scaled trace. A key observation is that the opportunities for shifting occur near the crossover points of excess and deficit. First, a list of all possible shifting opportunities is constructed, containing "candidate" hours where fossil generation is used that are within the shifting time horizon of a "target" hour with excess generation. This list may contain multiple target options for a candidate, as well as multiple candidates vying for a target. Beginning with candidates that have to shift the longest number of hours, target hours are matched to candidates by identifying candidates that only have one potential target hour. The lesser of the candidate and target amount is shifted to displace the excess; if the excess is eliminated, the target is removed. Ties are broken by favoring delaying loads over advancing loads. This process is iterated until there are no more (target, candidate) pairs.

This method is optimistic because (1) decisions to shift loads are made *post hoc* given full knowledge of future loads and (2) demands and loads are advanced or delayed with no change in energy consumed. It provides a reasonable upper bound on the potential for load shifting capabilities. It also serves as a target for more realistic real-time algorithms that utilize prediction of demand and supply conditions. The second assumption reflects the difficulty in making realistic claims about the shiftability of a diverse class of loads in a top-down analysis. A more detailed shifting model, say for shifting of thermostatic loads by precooling, might incorporate a penalty for shifting.

To assess the overall potential for load shifting to displace fossil fuel generation, we conducted experiments to vary the hours of shifting and renewables penetration limits; blends are computed using the method described in Section III-A. Figure 1(b) reveals the fundamental limits of shifting: it has limited impact in summer when renewable supply is abundant, and in winter when it is scarce. It is most effective in the shoulder seasons where there are intermittent periods of excess alternated with periods of fossil fuel generation. Figure 9 shows the gain throughout the full year over a range of renewables penetration and shifting durations. The value of shifting increases strongly with the penetration – the number of target hours where excess energy is present increases faster than the decrease in the number of hours where fossil fuels are used for generation. Since there is more excess to utilize and less fossil to displace, it will be much more important in future more sustainable grids than today. However, the overall impact of shifting is fundamentally limited by seasonal effects.

The most exciting CPS aspects of supply-following loads lie in going beyond the simplistic algorithms modeled here. Rather than merely minimize total fossil generation, it is essential to utilize shifting to flatten the duration curve for those supplies. By targeting peak hours of fossil generation, we arrive at a novel variant of peak shaving. In addition, the agility associated with shifting can be used to take advantage of those opportunities of excess.

C. Use of Storage

Many issues addressed by supply-following loads can also be addressed by energy storage. Currently, 128 GW of storage capacity are already deployed throughout the world (about 6 GW in California), but over 99.9% of it is pumped hydroelectric storage [14], [26], where water is pumped from lower- to higher-altitude reservoirs when electricity is cheap and later released through generation turbines when it is expensive. Though other, more dynamic storage media such as compressed air caverns, flywheels, and novel battery chemistries are in development [27], most require additional maturation to be viable at scale.

Putting aside these technological challenges, opportunities for CPS in energy storage lie in modeling and managing storage to reduce variability and maximize utilization in the unpredictable and dynamic landscape of a renewables-heavy grid. Though storage management is conceptually similar to that of supply-following loads, there are key differences. Energy storage allows shaping of supply rather than demand, and is substitutable for any type of generation – unlike demand, where loads are unique, any unit of supplied power is equally useful. However, to provide this capability, storage bears a roundtrip efficiency factor, known to be 75% for pumped hydro [23]. Nonetheless, the ability to shift supply becomes increasingly valuable in a blend with deep renewables. To demonstrate this, we model a rudimentary storage resource that charges whenever excess is available and discharges without loss in place of imported and thermal generation whenever capacity is available. Figure 7 shows the duration curve of the thermal resource for a scaled grid that employs 15 GWh of this greedy peak shifting storage; the storage reduces the hours where thermal generation is needed, but does little to address peak hours. More refined storage controllers may target the peak and flatten the demand curve, and is an open area for CPS research.

Additionally, coordination of disparate generation resources, including those traditionally used for baseload power, becomes especially valuable in this regime. As another example of the value of systems co-optimization, we examine the large hydroelectric resource, consisting of facilities with capacity over 30 MW. Looking at the daily patterns in hydroelectric production in Figure 6, it is clear that some load-following is occurring daily to match supply to overall demand and that the magnitude of this dispatch varies by season – it peaks in the summer and reduces thereafter. However, if the objective of this dispatch were altered from load-following to firming shortfalls of supply relative to demand, this generation could provide substantial zero-emissions load balancing capacity. Using a simple omniscient algorithm that reassigns dispatch from this resource to first remove imported and then thermal generation, we attempt to quantify the utility of this resource. Figure 7 compares this method to the 60% scaled blend with no demand modification, the same blend with up to 6 hours of demand shifting, and the blend with 15 GWh of storage. Though the magnitude of improvement is less than shifting, this type of strategy can be extremely valuable. Interestingly, such scheduling of this resource has only modest impact on the sharp peak. In those rare hours where there is very little solar and wind, current reservoir management results in little available hydro resource. An open question is whether by predicting those correlated lulls and provisioning this or other storage resources accordingly, the peak could be eliminated. Otherwise, demand curtailment is the final resort.

D. *Load Curtailment and Energy Efficiency*

In a renewables-heavy grid, the critical periods change from summer peaks in energy demand to winter lulls in renewable energy generation. This elicits a concurrent change in strategies for load curtailment, situations in which grid operators target the hours that need the most expensive generation by selectively interrupting non-essential loads. In our analysis, these critical hours appear in the left peak of the duration curves in Figure 7 – this is when high-cost, very low-utilization generation resources are operated. When this demand cannot be shifted elsewhere on the duration curve using other mechanisms, load curtailment can avoid these costly peaks. For example, in the scaled grid profiled above, the peak 1% of hours over the year for thermal generation each consume at least 14 GW and up to 19.5 GW – over 28% of the capacity. Though this long tail is well understood for electricity demand, the difficulty in predicting and responding to peak events in a more dynamic electricity system is a critical challenge for the CPS community.

Additionally, since peaks in the thermal resource on a renewables-heavy grid are no longer coupled to demand, reduction of energy consumption at any hour of day becomes more valuable. Examples of this may include decreases in nighttime lighting loads or improvements in HVAC setback schedules. Though energy efficiency has long been a CPS goal [7], [20], [25], shortages in renewable energy that can constrain generation resources make it ever more important. Further, more energy-efficient demand improves the performance of strategies like load shifting and storage because there are more targets and less candidates for shifting.

V. CONCLUSION

Modern electric grids are an essential class of cyber-physical system. While they were originally designed in an era when energy was plentiful and communication was almost non-existent, they are being transformed in an age of zero-emissions production and pervasive information. They present archetypical elements of a CPS: the need to orchestrate a heterogenous collection of distributed resources in the presence of deep uncertainty, where the main elements of the system are actuated by forces outside the domain of control. As we move toward sustainable grids, the uncertainty and underactuation is not just in the energy demand, but also in the generation of renewable energy. This setting defines a completely new set of CPS challenges.

In this study, we follow a methodology reflecting a mantra of 3Ms: Monitor, Model, and Mitigate. Utilizing a recently released, large-scale, and detailed monitoring of CA ISO generation resources, we are able to represent and characterize the temporal dynamics of the grid today at the scale of the state of California. Based on this, we build a model of what the grid would look like with sufficient renewable sources to provide 60% of the electricity consumed. We see that this presents different challenges from those dominating the discussion today. Rather than critical peak demand driven by cooling loads, the summer has copious energy availability, but the winter has deep lulls. Fuels take on a new role of providing precious storage resources; possessing tremendous energy density and duration, they can be dispatched once to fill the gaps in zero-emissions renewable supplies.

In this setting, we have explored abstractly several of the critical avenues for mitigating the supply-demand matching problem. These include the introduction of supply-following loads, energy storage resources, demand curtailment, and integrated coordination of diverse generation resources. Each of these brings substantial benefits and each has limitations. Some are at the fine scale of rapid ramp rates and coordinated response, others are at the coarse scale of seasonal variations and total reservoir capacity. To tackle the challenges presented by a sustainable energy grid, all will need to be employed as a CPS in a coordinated fashion.

This study is only a very preliminary step, but it makes a significant stride by examining temporal dynamics of a large-scale energy network, rather than only manipulating statistical characterizations of network elements. In this manner, more of the nature of the CPS challenges is revealed. However, we look only at large-scale interactions and balance at hourly timescales. This analysis needs replication at finer granularity that exposes specific constraints of the transmission network, individual plants, and loads. Furthermore, the algorithms presented are simplistic and merely characterize the potential and limitations of a few opportunities. In every respect, these mechanisms and their associated protocols need to be developed far more fully, both to formalize a theory of operation in a sustainable grid and to invent analysis, prediction, and algorithmic techniques that achieve these bounds, as well as technological alternatives that remove such barriers as seasonal limits and ramp rates. Together, these CPS efforts will move us a step closer to the design of a sustainable energy network.

ACKNOWLEDGMENTS

This work was supported by the National Science Foundation under grants CPS-0932209 and CPS-0931843.

REFERENCES

[1] Automated Demand Response. http://www.auto-dr.com/.
[2] California Independent System Operator. http://www.caiso.com/.
[3] California Senate Bill 107 - 2006.
[4] California Senate Bill 2 - 2011.
[5] Dynamic Demand. http://www.dynamicdemand.co.uk/.
[6] B. Aksanli, J. Venkatesh, L. Zhang, and T. Rosing. Utilizing Green Energy Prediction to Schedule Mixed Batch And Service Jobs in Data Centers. In *Proceedings of the 4th ACM Workshop on Power-Aware Computing and Systems (HotPower '11)*, 2011.
[7] A. Aswani, N. Master, J. Taneja, D. Culler, and C. Tomlin. Reducing Transient and Steady State Electricity Consumption in HVAC Using Learning-Based Model-Predictive Control. *Proceedings of the IEEE*, 100(1):240 –253, 2012.
[8] A. Brooks, E. Lu, D. Reicher, C. Spirakis, and B. Weihl. Demand dispatch: Using real-time control of demand to help balance generation and load. *Power and Energy Magazine, IEEE*, 8(3):20–29, May/June 2010.
[9] California Council on Science and Technology. California's Energy Future – The View to 2050. http://ccst.us/publications/2011/2011energy.php, 2011.
[10] California Energy Commission. Power Plant Database. http://energyalmanac.ca.gov/powerplants/Power_Plants.xls.
[11] California ISO. Daily Renewables Watch. http://www.caiso.com/green/renewableswatch.html.
[12] California ISO. Summary of Preliminary Results of 33% Renewable Integration Study, 2011.
[13] California Solar Initiative. California Solar Statistics. http://www.californiasolarstatistics.ca.gov.
[14] Electric Power Research Institute. Electricity Energy Storage Technology Options, 2010.
[15] G. Hoste, M. Dvorak, and M. Jacobson. Matching Hourly and Peak Demand by Combining Different Renewable Energy Sources: A Case Study for California in 2020, 2009.
[16] International Energy Agency. Variability of Wind Power and Other Renewables - Management Options and Strategies, 2005.
[17] KEMA, Inc. Research Evaluation of Wind Generation, Solar Generation, and Storage Impact on the California Grid, 2010.
[18] A. Krioukov, S. Alspaugh, P. Mohan, S. Dawson-Haggerty, D. E. Culler, and R. H. Katz. Design and Evaluation of an Energy Agile Computing Cluster. Technical Report UCB/EECS-2012-13, EECS Department, University of California, Berkeley, 2012.
[19] A. Krioukov, P. Mohan, S. Alspaugh, L. Keys, D. Culler, and R. Katz. NapSAC: Design and Implementation of a Power-Proportional Web Cluster. In *Proceedings of the First ACM SIGCOMM Workshop on Green Networking*, 2010.
[20] J. Lu, T. Sookoor, V. Srinivasan, G. Gao, B. Holben, J. Stankovic, E. Field, and K. Whitehouse. The Smart Thermostat: Using Occupancy Sensors to Save Energy in Homes. In *Proceedings of the 8th ACM Conference on Embedded Networked Sensor Systems (SenSys)*, 2010.
[21] X. Lu, M. B. McElroy, and J. Kiviluoma. Global Potential for Wind-generated Electricity. *PNAS*, 2009.
[22] T. J. Lui, W. Stirling, and H. O. Marcy. Get smart. *Power and Energy Magazine, IEEE*, 8(3):66–78, May/June 2010.
[23] D. MacKay. *Sustainable Energy – Without the Hot Air*. http://www.withouthotair.com, 2009.
[24] MIT Energy Initiative. The Future of the Electric Grid. http://web.mit.edu/mitei/research/studies/documents/electric-grid-2011/Electric_Grid_Full_Report.pdf, 2011.
[25] L. Parolini, B. Sinopoli, B. Krogh, and Z. Wang. A Cyber-Physical Systems Approach to Data Center Modeling and Control for Energy Efficiency. *Proceedings of the IEEE*, 100(1):254 –268, 2012.
[26] Richard K. Miller & Associates. Utility Energy Storage – Pumped Hydroelectric Storage. http://www.rkma.com/utilityenergystorageSAMPLE.pdf.
[27] B. Roberts. Capturing grid power. *Power and Energy Magazine, IEEE*, 7(4):32–41, July/August 2009.
[28] J. Taneja, D. Culler, and P. Dutta. Towards Cooperative Grids: Sensor/Actuator Networks for Renewables Integration. In *Proceedings of the 1st IEEE Int'l Conference on Smart Grid Communications*, 2010.