Data Centers as Dispatchable Loads to Harness Stranded Power

Kibaek Kim, Fan Yang, *Member, IEEE*, Victor M. Zavala, *Member, IEEE*, and Andrew A. Chien *Fellow, IEEE*

***Abstract*—We analyze how both traditional data center integration and dispatchable load integration affect grid dynamics and efficiency. Our analysis uses parallel optimization solvers, 1000's of renewable generation scenarios, and running thousands of simulations. These enable us to perform sophisticated analysis, including rigorous grid performance modeling. For example, we evaluate scenarios of optimal placement of dispatchable loads in a realistic power grid system and assesses performance in the face of myriad wind and load scenarios.**

Our analysis reveals that significant spillage and stranded power exists in current power grids. Adding data centers with collocating renewable generation (wind farms) can be harmful to RPS goals, increases both stranded power and thremal generation. In contrast, adding data centers that are dispatchable by the power grid both reduces stranded power and improves both grid cost and stability, even at high RPS. In short, dispatchable loads can enable grid scaling to high RPS. These dispatchable computing loads achieve duty factors of 60-80%

***Index Terms*—renewable power, green computing, power grid, energy markets, renewable portfolio standard, cloud computing**

I. Introduction

Over the past two decades, a growing consensus has emerged that the earth's climate is warming at a significant rate with anthropogenic carbon an important contributor [1], [2]. In response, there are growing array of efforts worldwide to reduce the amount of carbon being released into the atmosphere [3], [4]. One area of particular interest is increasing carbon emissions due to use of information and computing technologies (ICT), which were recently estimated at 2% of global emissions [5] and are among the most rapidly growing. In fact, recent estimates suggest that by 2020, ICT will account for 4% of carbon emissions [5]–[7].

Recently, the rise of cloud computing has raised concerns about carbon emissions from data centers [8]. These concerns have spawned research on how to increase data-center energy efficiency [9], [10] and exploit renewable power to supply data-center loads [11]–[14]. A recent strategy pursued by several "hyperscaler" internet companies has been the purchase of wind-power offsets as part of "long-term purchase" contracts [15]. Another well-studied research topic is the optimization of data-center site selection based on cost and on exploitation of renewable power [16], [17]. To our knowledge, all such studies consider benefits and costs from the perspective of the cloud computing operator, who seeks to maximize revenue, reduce total-cost-of-ownership (TCO), and maintain high data-center availability. In this paper we take an alternative approach and consider the impact of the addition of new data centers and collocated renewables on the resilience, efficiency, and flexibility of the power grid.

Ambitious "renewable portfolio standards" (RPS) goals for renewable power as a fraction of overall power are currently being adopted. U.S. states across the Midwest included in the Mid-continent Independent System Operator (MISO) system have adopted standards ranging from 25% (2015) in Illinois to $25 \sim 31\%$ (2025) in Minnesota. California has already reached a 20% renewable mix in 2010 and is on track to reach its 33% target for 2020 [18] for wind and solar power. In September 2015, California adopted an RPS goal of 50% renewable by 2030 [19]. Other state goals include 50% by 2030 in New York, and 10 GW by 2025 in Texas.

Obama's "Clean Power Plan," issued August 2015, calls for a national 32% reduction in electric power carbon emissions by 2030, with renewable power as a critical element. And, the U.S. Department of Energy released a landmark report, "Wind Vision 2015," that describes how the United States can achieve a 35% RPS for wind alone by 2050, a big jump from a combined solar and wind RPS of 5.2% in 2014 [20]. In practical terms, this means that regions with relatively more wind resources, such as Texas, can achieve an RPS goal of over 50% by 2050. These ambitious and transformative goals pose serious challenges for the power grid, including its ability to achieve "merit order" and fairness while maintaining system efficiency and resiliency. In particular, the fluctuation of renewable wind and solar power generation can create periods with large amounts of stranded power.

Motivated by the dual goals of achieving high RPS in the power grid and supporting large-scale computing, we address the following questions:

1) What impact will the growth of data-center installations have on the future power grid?
2) Should renewables be collocated next to data centers?
3) Can we simultaneously enable scalable computing and greater renewable penetration?

To explore these questions, we developed a computational framework that uses a detailed power grid system model and cutting-edge, parallel optimization solvers. We use these capabilities to explore a range of scenarios and characterize the impact of increasing RPS levels and data-center deployments on key power grid metrics such as system cost, stranded power, and data-center duty factors. To model the growth of

Kibaek Kim is with the Mathematics and Computer Science Division, Argonne National Laboratory, Lemont, IL 60439, Email: kimk@anl.gov.

Fan Yang and Andrew A. Chien are with the Department of Computer Science, the University of Chicago.

Victor M. Zavala is with the Department of Chemical and Biological Engineering, the University of Wisconsin-Madison.

Andrew A. Chien and Victor M. Zavala are also affiliated with Mathematics and Computer Science, Argonne National Laboratory.

cloud computing, we consider the addition of twenty 200-MW data centers to the power grid in locations chosen for external purposes (e.g. optimized for cloud computing interests). Next, we consider the collocation of renewables with those data centers. Subsequently, we present a different model, introducing a new type of computing that is flexible and forms a dispatchable load. Such dispatchable loads can enable grid stability at higher RPS levels and also greater efficiency. We consider optimizing the location of these dispatchable loads and the resulting impact on power grid efficiency. Our findings include:

- Significant spillage and stranded power exists in current power grids, and this quantity increases as we scale to higher RPS levels.
- Collocating wind farms and data centers naively can be harmful to RPS goals, increasing stranded power and thermal generation.
- The use of dispatchable computing loads reduces stranded power and enables scaling to higher RPS levels by increasing price stability.
- Optimizing placement of dispatchable loads decreases overall system cost and achieves data center duty factors of 60-80%.

The rest of the paper is organized as follows. In Section II we discuss the issue of stranded power and introduce the concept of dispatchable computing loads. Section III outlines our optimization models, followed by experiments in Section IV. In Section VI, we summarize our results.

II. Stranded Power and Dispatchable Computing

In this section we describe the idea and definition of both stranded power and a dispatchable load.

A. Stranded Power

Power system operators must balance power flow across each bus in the power grid network. Generators offer their generation capability to the grid in real time (every 5 to 12 minutes), and the grid dispatches generation based on the demand and transmission. However, intrinsic variability of renewable generation (wind, solar, etc.) creates major challenges for power dispatch. Despite best efforts to match generation and demand, in the process of ensuring reliable power, there can be oversupply and transmission congestion that prevents generated power from reaching loads. Power grids call this excess power *spillage*, "curtailment," or "down dispatching."

Figure 1 shows the monthly wind generation and down-dispatched wind power (spillage) of the MISO system. Almost 7% of wind generation is curtailed because of transmission congestion. The total down-dispatched power of MISO for 2014 was about 2.2 terawatt-hours, corresponding to a 183 MW sustained rate. Comparable waste also exists in other independent system operators (ISO), including the Eastern Region Coordinating District of Texas (ERCOT), California ISO (CAISO) [21], and many European countries such as Denmark, Germany, Ireland, and Italy [22]. The amount of waste is projected to increase with higher RPS levels [19].

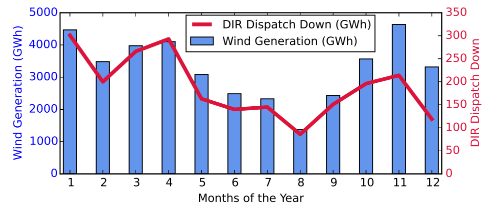

Fig. 1. Wind generation and down-dispatching (spillage) of MISO in 2014.

Modern energy markets dispatch generation by assigning locational marginal prices (LMP) that vary across generation sites, grid nodes, and time intervals. In situations of oversupply or transmission congestion, power prices can be low or even negative, causing power generators to dump power (spillage) or deliver it, and pay the grid to take it[1] Consequently, spillage can be significantly less than total uneconomic generation.

We define *stranded power* as all offered generation with no economic value, thus including both spillage and delivered power with zero or negative LMP. Figure 2 presents MISO's stranded power in 2014, breaking it down by month and type. It also compares the average stranded power from wind and non-wind generation. Overall stranded wind power, the sum of wind spillage and noneconomic wind dispatch (LMP$\leq$0), is about 7.7 TWh for 2014. Interestingly, non-wind sites, mostly thermal generators, have 10.1 TWh of stranded power. However, as fraction, because 90% of grid power is thermal, approximately 8-times less stranded power by percentage.

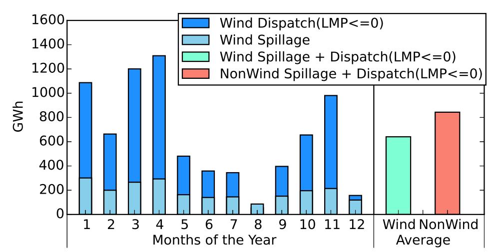

Fig. 2. MISO stranded power (wind) vs. stranded power (nonwind) in 2014.

B. Dispatchable Loads

We define *dispatchable loads* as adjustable demands that are dispatched in real-time by the power grid. Adjustable at every dispatch interval, dispatchable loads are an ideal form of demand response. Their flexibility makes them a good match to reduce congestion and stranded power, particularly due to generation variability. Key properties of dispatchable loads include:

- Dispatchable loads consumption can be increased to some limit under grid control.
- Dispatchable loads consumption can be decreased to zero under grid control.
- Grid control can be exercised at the dispatch interval (effectively instantly).

We consider one possible implementation of dispatchable loads: intermittent computing resources. Because computing

[1] Delivering power into the grid can be tied to "production tax credits", a financial incentive to sell power to the grid even at negative price.

operates on much shorter time scales, it has the potential to be agile, yet still productive. We call such intermittent computing resources zero-carbon cloud (ZCCloud) [23], [24], and have described and analyzed several possible models. For example, ZCCloud could be computing hardware deployed in shipping containers and directly connected to a wind farm or at a key transmission bottleneck. ZCClouds can transform stranded power into computing power with very short latency (in seconds) and can be easily turned on or shut down according to stranded power availability. Possible uses include data-center workloads such big data analysis, machine learning, or high-performance computing. The uptime and capabilities of intermittent computing resources deployed as dispatchable loads are determined by the quanity and temporal distribution of stranded power.

The idea of intermittent (or "volatile") computing resources is of growing interest. Cloud providers have begun to provide unreliable/revokable resources including Amazon's spot instances [25] and Google's preemptible VM Instance [26]. Several studies propose methods make such resources useful for high performance computing and more advanced cloud services [27]. We believe there is a broad application for intermittent computing resources.

Of course, many other realizations of dispatchable loads are possible, including energy storage. The two important differences between our dispatchable loads and energy storage are infinite capacity (our dispatchable loads can run forever without filling up) and externally subsidized economics (services generated by our dispatchable loads can defray their cost in part or in full).

III. Optimization Models

In this section we present two optimization models in order to assess the benefits of dispatchable loads. We also present various performance metrics for our analysis.

A. Economic Dispatch Model

To assess the economic benefits of dispatchable computing loads, we use the following economic dispatch (ED) model:

$$z_{ED} := \min \sum_{t\in\mathcal{T}} \left(\sum_{i\in\mathcal{G}} C_i p_{it} + \sum_{j\in\mathcal{D}} C_j^d d_{jt} + \sum_{i\in\mathcal{I}} C_i^m m_{it} + \sum_{i\in\mathcal{W}} C_i^w w_{it} + \sum_{i\in\mathcal{R}} C_i^r r_{it} \right) \tag{1a}$$

$$\text{s.t.} \quad \sum_{l\in\mathcal{L}_n^+} f_{lt} - \sum_{l\in\mathcal{L}_n^-} f_{lt} + \sum_{i\in\mathcal{G}_n} p_{it} + \sum_{i\in\mathcal{I}_n} (M_{it} - m_{it}) + \sum_{i\in\mathcal{W}_n} (W_{it} - w_{it}) + \sum_{i\in\mathcal{R}_n} (R_{it} - r_{it}) = \sum_{j\in\mathcal{D}_n} (D_{jt} - d_{jt}), \quad (\lambda_{nt}), \quad \forall n \in \mathcal{N}, t\in\mathcal{T}, \tag{1b}$$

$$f_{lt} = B_l(\theta_{nt} - \theta_{mt}), \quad \forall l = (m,n) \in \mathcal{L}, t\in\mathcal{T}, \tag{1c}$$

$$-RD_i \le p_{it} - p_{i,t-1} \le RU_i, \quad \forall i \in \mathcal{G}, t\in\mathcal{T}, \tag{1d}$$

$$-F_l^{max} \le f_{lt} \le F_l^{max}, \quad \forall l \in \mathcal{L}, t\in\mathcal{T}, \tag{1e}$$

$$\Theta_n^{min} \le \theta_{nt} \le \Theta_n^{max} \quad \forall n \in \mathcal{N}, t\in\mathcal{T}, \tag{1f}$$

$$0 \le p_{it} \le P_i^{max}, \quad \forall i \in \mathcal{G}, t\in\mathcal{T}, \tag{1g}$$

$$0 \le d_{jt} \le D_{jt}, \quad \forall j \in \mathcal{D}, t\in\mathcal{T}, \tag{1h}$$

$$0 \le m_{it} \le M_{jt}, \quad \forall i \in \mathcal{I}, t\in\mathcal{T}, \tag{1i}$$

$$0 \le w_{it} \le W_{jt}, \quad \forall i \in \mathcal{W}, t\in\mathcal{T}, \tag{1j}$$

$$0 \le r_{it} \le R_{jt}, \quad \forall i \in \mathcal{R}, t\in\mathcal{T}. \tag{1k}$$

Note that power is supplied from imports, (nonwind) renewables (e.g., bio-, hydro- and geo-) and wind generations as well as from conventional thermal generation units. Considering imports and nonwind renewables (we refer to these simply as renewables in the following discussion) is important because they account for a significant portion of the power generation in some systems. In the CAISO system, for instance, imports and renewables account for 27% and 25% of the total system generation, respectively.[2] Moreover, the analysis that we present later indicates that dispatchable loads can reduce spillage of imports and nonwind renewables. In the presented model we assume that import as well as renewable and wind power suppliers are not competitive agents in the market (they are high-priority suppliers). Consequently, their supplies are considered as negative demands for which we seek to minimize spillages at costs C_i^m, C_i^w, and C_i^r, respectively. We also allow for load shedding at certain nodes at cost C_j^d, which is set to the value of lost load (VOLL).

The objective function (1a) is to minimize the sum of the following items:

- Supply cost from conventional thermal generators
- Cost of the load shedding
- Cost of the import spillage
- Cost of the wind power spillage
- Cost of the nonwind renewable spillage

The objective function represents the *total dispatch cost*. Note that this objective can also be interpreted as maximizing social welfare, as defined in electricity market clearing models (e.g., [28]). In such formulations the objective function is

$$\max \sum_{t\in\mathcal{T}} \left(\sum_{j\in\mathcal{D}} [D_{jt} - d_{jt}] - \sum_{i\in\mathcal{G}} C_i p_{it} - \sum_{i\in\mathcal{I}} C_i^m m_{it} - \sum_{i\in\mathcal{W}} C_i^w w_{it} - \sum_{i\in\mathcal{R}} C_i^r r_{it} \right). \tag{2}$$

Here, d_{jt} is the unserved load (load shedding), and D_{jt} are (constant) requested loads. Consequently, $D_{jt} - d_{jt}$ is the served load that we seek to maximize. Because D_{jt} is a constant, this can be eliminated from the objective. Thus this problem seeks to minimize d_{jt}, as in (1).

Equation (1b) enforces the power flow balance of the network. Equation (1c) represents a lossless model of power flow equations that determines the power flow of line l by the

[2] http://content.caiso.com/green/renewrpt/20160201_DailyRenewablesWatch.txt

phase angle difference between two buses m and n. Equation (1d) represents the ramping constraints limiting the rate of change of generation levels. Constraints (1e) and (1f) represent the transmission line capacity and the feasible phase angle range, respectively. Constraint (1g) represents the generation capacity, and constraint (1h) is a bound for the unserved load. Constraints (1i)-(1k) bounds spillages of imports, wind, and renewable supply, respectively

B. Optimal Placement of Dispatchable Loads

We extend the proposed ED model to account for optimal placement (OP) of dispatchable loads at locations minimizing the expected total dispatch cost. The OP model is cast as a two-stage stochastic integer programming problem of the following form:

$$\begin{aligned} z_{OP} :=& \\ \min \sum_{s\in\mathcal{S}} \pi_s \sum_{t\in\mathcal{T}} & \left(\sum_{i\in\mathcal{G}} C_i p_{its} + \sum_{j\in\mathcal{D}} C_j^d d_{jts} \right. \\ & + \sum_{i\in\mathcal{I}} C_i^m m_{its} + \sum_{i\in\mathcal{W}} C_i^w w_{its} + \sum_{i\in\mathcal{R}} C_i^r r_{its} \\ & \left. + \sum_{n\in\mathcal{N}} C_n^u [U x_n - u_{nts}] \right) \quad (3a) \end{aligned}$$

$$\text{s.t.} \sum_{n\in\mathcal{N}} x_n \le K, \quad (3b)$$

$$0 \le u_{nts} \le U x_n, \quad \forall n\in\mathcal{N}, t\in\mathcal{T}, s\in\mathcal{S}, \quad (3c)$$

$$\begin{aligned} & \sum_{l\in\mathcal{L}_n^+} f_{lts} - \sum_{l\in\mathcal{L}_n^-} f_{lts} + \sum_{i\in\mathcal{G}_n} p_{its} \\ & + \sum_{i\in\mathcal{I}_n} (M_{it} - m_{its}) + \sum_{i\in\mathcal{W}_n} (W_{its} - w_{its}) \\ & + \sum_{i\in\mathcal{R}_n} (R_{it} - r_{its}) - \sum_{t\in\mathcal{T}} u_{nts} \\ & = \sum_{j\in\mathcal{D}_n} (D_{jt} - d_{jts}), \quad \forall n\in\mathcal{N}, t\in\mathcal{T}, s\in\mathcal{S}, \quad (3d) \end{aligned}$$

$$\begin{aligned} & f_{lts} = B_l(\theta_{nts} - \theta_{mts}), \\ & \qquad \forall l=(m,n)\in\mathcal{L}, t\in\mathcal{T}, s\in\mathcal{S}, \quad (3e) \end{aligned}$$

$$\begin{aligned} & -RD_i \le p_{its} - p_{i,t-1,s} \le RU_i, \\ & \qquad \forall i\in\mathcal{G}, t\in\mathcal{T}, s\in\mathcal{S}, \quad (3f) \end{aligned}$$

$$-F_l^{max} \le f_{lts} \le F_l^{max}, \quad \forall l\in\mathcal{L}, t\in\mathcal{T}, s\in\mathcal{S}, \quad (3g)$$

$$\Theta_n^{min} \le \theta_{nts} \le \Theta_n^{max} \quad \forall n\in\mathcal{N}, t\in\mathcal{T}, s\in\mathcal{S}, \quad (3h)$$

$$0 \le p_{its} \le P_i^{max}, \quad \forall i\in\mathcal{G}, t\in\mathcal{T}, s\in\mathcal{S}, \quad (3i)$$

$$0 \le d_{jts} \le D_{jt}, \quad \forall j\in\mathcal{D}, t\in\mathcal{T}, s\in\mathcal{S}, \quad (3j)$$

$$0 \le m_{its} \le M_{jt}, \quad \forall i\in\mathcal{I}, t\in\mathcal{T}, s\in\mathcal{S}, \quad (3k)$$

$$0 \le w_{its} \le W_{jt}, \quad \forall i\in\mathcal{W}, t\in\mathcal{T}, s\in\mathcal{S}, \quad (3l)$$

$$0 \le r_{its} \le R_{jt}, \quad \forall i\in\mathcal{R}, t\in\mathcal{T}, s\in\mathcal{S}, \quad (3m)$$

$$x_n \ge 0, \text{ integer} \quad \forall n\in\mathcal{N}. \quad (3n)$$

The objective function (3a) is to minimize the expected total dispatch costs where uncertainty arises from wind supply scenarios. The here-and-now decision is to determine the number and locations of dispatchable loads to be installed. The second-stage decisions involve flows, angles, supply, loads, and spillages for each scenario $s \in \mathcal{S}$. Equations (3b) and (3c) are budget and capacity constraints for dispatchable loads, respectively. The capacity of a dispatchable load is given by U. We note that the objective (3a) and constraints (3d) include dispatchable loads.

C. Performance Metrics

We define the metrics of interest for our analysis. The dispatch cost and the expected dispatch cost are given by z_{ED} and z_{OP}, respectively. The ED supply cost and the OP supply cost are defined respectively as

$$z_{ED} - \sum_{j\in\mathcal{D}} \sum_{t\in\mathcal{T}} C_j^d d_{jt} \quad (4a)$$

and

$$z_{OP} - \sum_{t\in\mathcal{T}} \sum_{s\in\mathcal{S}} \pi_s \left(\sum_{j\in\mathcal{D}} C_j^d d_{jts} - \sum_{n\in\mathcal{N}} C_n^u [U x_n - u_{nts}] \right). \quad (4b)$$

Here, the second term in both (4a) and (4b) is the load shedding cost. The total amount of load shedding and the expected amount of load shedding are defined respectively as

$$\sum_{j\in\mathcal{D}} \sum_{t\in\mathcal{T}} d_{jts} \quad (5a)$$

and

$$\sum_{j\in\mathcal{D}} \sum_{t\in\mathcal{T}} \sum_{s\in\mathcal{S}} \pi_s d_{jts}. \quad (5b)$$

The total amount of power supplied (produced) is defined as

$$P_{ED}^{Supply} := \sum_{t\in\mathcal{T}} \left(\sum_{i\in\mathcal{G}} p_{it} + \sum_{i\in\mathcal{I}} M_{it} + \sum_{i\in\mathcal{W}} W_{it} + \sum_{i\in\mathcal{R}} R_{it} \right). \quad (6)$$

The total amount of power spillage is defined as

$$P_{ED}^{Spillage} := \sum_{t\in\mathcal{T}} \left(\sum_{i\in\mathcal{I}} m_{it} + \sum_{i\in\mathcal{W}} w_{it} + \sum_{i\in\mathcal{R}} r_{it} \right). \quad (7)$$

Consequently, the total amount of power dispatched (absorbed into the system) is given by

$$P_{ED}^{Dispatch} := P_{ED}^{Supply} - P_{ED}^{Spillage}. \quad (8)$$

We differentiate the total power absorbed at positive LMPs and nonpositive LMPs. To do so, we define $\mathcal{N}_t^+ := \{n \in \mathcal{N} : \lambda_{nt} > 0\}$, where λ_{nt} is an optimal dual variable value of the ED model. The total power absorbed at positive LMPs and nonpositive LMPs for ED is given respectively by

$$\begin{aligned} P_{ED}^{LMP>0} := \sum_{t\in\mathcal{T}} \sum_{n\in\mathcal{N}_t^+} & \left[\sum_{i\in\mathcal{G}_n} p_{it} + \sum_{i\in\mathcal{I}_n} (M_{it} - m_{it}) \right. \\ & \left. + \sum_{i\in\mathcal{W}_n} (W_{it} - w_{it}) + \sum_{i\in\mathcal{R}_n} (R_{it} - r_{it}) \right] \quad (9) \end{aligned}$$

TABLE I
AVERAGE LOAD, IMPORTS, RENEWABLE SUPPLY, AND NET LOAD (MW)

	Load	Imports	Renewable	Net Load
SpringWD	26,868	7,478	6,681	12,708
SpringWE	23,980	7,608	6,998	9,373
SummerWD	31,089	7,678	6,672	16,737
SummerWE	28,184	7,400	7,124	13,659
FallWD	28,055	7,675	6,657	13,722
FallWE	25,186	7,108	7,065	11,012
WinterWD	26,352	7,663	6,634	12,054
WinterWE	23,708	6,800	5,581	11,399

and

$$P_{ED}^{LMP\leq 0} := P_{ED}^{Dispatch} - P_{ED}^{LMP>0}. \quad (10)$$

We define stranded power as

$$P_{ED}^{LMP\leq 0} + P_{ED}^{Spillage}. \quad (11)$$

Wind penetration levels (%) are defined as

$$\frac{\sum_{i\in\mathcal{W}}\sum_{t\in\mathcal{T}} W_{it}}{\sum_{j\in\mathcal{D}}\sum_{t\in\mathcal{T}} D_{jt}} \times 100. \quad (12)$$

The RPS is defined by the ratio of the renewable power absorbed to the total power loads,

$$\frac{\sum_{t\in\mathcal{T}}\left[\sum_{i\in\mathcal{R}}(R_{it}-r_{it})+\sum_{i\in\mathcal{W}}(W_{it}-w_{it})\right]}{\sum_{j\in\mathcal{D}}\sum_{t\in\mathcal{T}} D_{jt}} \times 100. \quad (13)$$

For the OP model the metrics are defined for each $s \in \mathcal{S}$ from which we compute expected values.

The duty factor (%) for the dispatchable computing loads is defined as the ratio of the total amount of dispatchable load served to the requested capacity:

$$\frac{\sum_{n\in\mathcal{N}}\sum_{t\in\mathcal{T}} u_{nts}}{|\mathcal{T}|\sum_{n\in\mathcal{N}} Ux_n}. \quad (14)$$

We note that duty factors can be used as a metric that represents profitability for the data center owners.

IV. COMPUTATIONAL RESULTS AND ANALYSES

We study a test system of CAISO interconnected with the Western Electricity Coordinating Council (WECC). The system consists of 225 buses, 375 transmission lines, 130 generation units, 40 loads, and 5 wind power generation units. We consider a 24-hour horizon with hourly intervals. We use network topology, import supply, renewables, wind production, and load data from [29]. Imports flow into the system through 5 boundary buses. Renewable power is generated from biogas, hydrothermal, and geothermal generators at 11 buses. For this system, the generation capacity that excludes imports, renewables, and wind power is 31.2 GW. We consider load profiles for 8 day types: spring, summer, fall, and winter; as well as weekday (WD) and weekend (WE). Table I reports average loads, imports, renewables, and net loads (load minus imports, renewables, and wind supply). Figure 3 shows the net loads for the different day types.

For each day type, we use the 1,000 specific wind power production scenarios taken from [29] where wind power

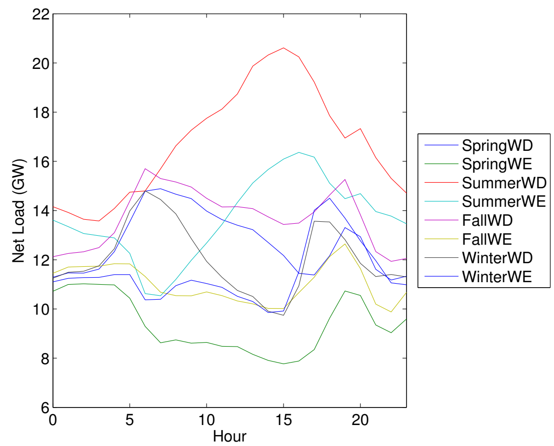

Fig. 3. Net load for each day type.

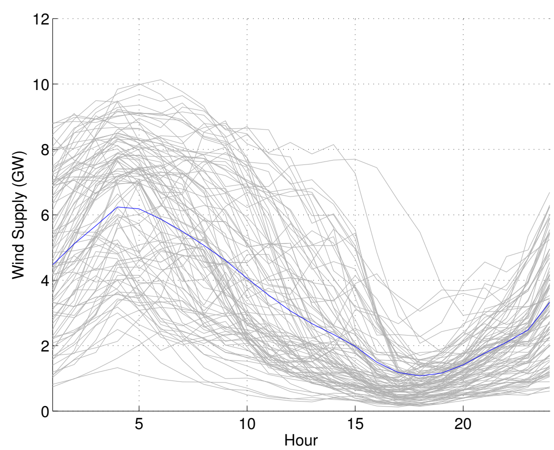

Fig. 4. Wind power supply (GW) at 15% level during summer day.

scenarios are 15% of load, representing the 2020 RPS target of California [29] as illustrated in Figure 4. We plot 100 scenarios for wind power (grey lines) to highlight the variability, and also plot the corresponding mean (blue line). While our simulations use 1000 scenarios, we plot only 100 scenarios for readability. The same wind scenarios for weekdays and weekends of the same season. We explore a range of additional wind penetration levels 5%, 15%, 30%, and 50% of load,

As is typical due to longer-term commitments and the goal a reducing carbon emissions, we assume that imports and renewables are higher priority (i.e., nondispatchable) but can be spilled, if necessary, at a cost of 1,000 and 2,000 $/MWh, respectively. We use a load shedding cost (VOLL) of 1,000 $/MWh and a wind spillage cost of 100 $/MWh. These values or higher are typical of ISO settings, and are chosen to impose a relative priority on different products.

We analyze the following cases:

- Case 1: Base, WECC configuration as described above.
- Case 2: Case 1 plus 20 additional 200 MW data centers that total 4 GW additional load (96 GWh per day). Each data center is a continuous 200 MW load and subject to VOLL penalties. Data-center locations were chosen arbitrarily to reflect choices driven by external business

considerations (e.g., networking, proximity to customers, and geographic diversity).

- Case 3: Case 2 plus collocated wind farms at each data center, sized match total load over 12 months. Due to typical wind capacity factor of 30%, the peak generation of these farms is typically 3x greater than the 200 MW data center load.
- Case 4: Case 1 plus 20 additional 200 MW data centers operated as dispatchable loads. The ISO determines the power consumption of each dispatchable load each hour at no penalty cost. The data centers are positioned optimally to minimize overall system dispatch cost across all wind and load scenarios by solving the OP model (3). Note that Case 4 is an extension of Case 1, not Case 2 or 3.

In Cases 1, 2, and 3, we solve the ED model (1), minimizing the total dispatch cost for all wind and load scenarios as we increase wind levels. In Case 4, the ED is subsumed within the OP model, and the solution minimizes the same metrics in addition to optimal placement of dispatchable loads.

The cases vary in total generation and load. To make the clearest comparison, we make comparisons reporting percentages relative to Case 1 (i.e., the base system). For example, the loads in Case 2 and 3 are higher due to addition of data centers, and Case 4 falls in the middle due to its variable dispatch. We use consistent wind penetration numbers, ignoring the additional wind generation in Case 3. The simple treatment of loads affects “real” wind penetration only $1 \sim 6\%$, far smaller than resulting spillage. Excluding data center wind power in Case 3, produces conservative estimates of spillage and stranded power in that case, painting it in the most favorable light.

Figure 5 presents a node-edge network representation of the test system with the 20 data-center locations of Case 2. We note that this network does not represent the actual geographical locations of the actual buses and the lines of the system. The network was generated by using the `gephi` package [30].

The network system under study is large, and we evaluate a large number of scenarios and system configurations. Thus, the analysis performed is computationally intensive. Cases 1, 2 and 3 solve the ED model (1) for 1,000 wind-power scenarios and for each day type and season. These represent a total of 32,008 linear programs (LPs) for each case. Each LP has 19,008 continuous variables and 17,544 linear constraints. Case 4 solves the OP model (3) including all the 1,000 wind scenarios; each OP instance is solved for the different day type and season and different wind power levels. We thus solve a total of 24 large-scale stochastic mixed-integer programs (for nonzero wind levels) and 8 deterministic mixed-integer programs (at zero wind level). Each OP model has 225 general integer variables, 24,408,000 continuous variables, and 22,944,001 constraints. Each deterministic mixed-integer program has 225 general integer variables, 24,408 continuous variables, and 22,945 constraints. The EDs are implemented in JuMP [31] and the stochastic ones in StochJuMP [32]. ED is solved with CPLEX 12.6.1, and OP is solved by using the parallel Benders decomposition implementation of the open-

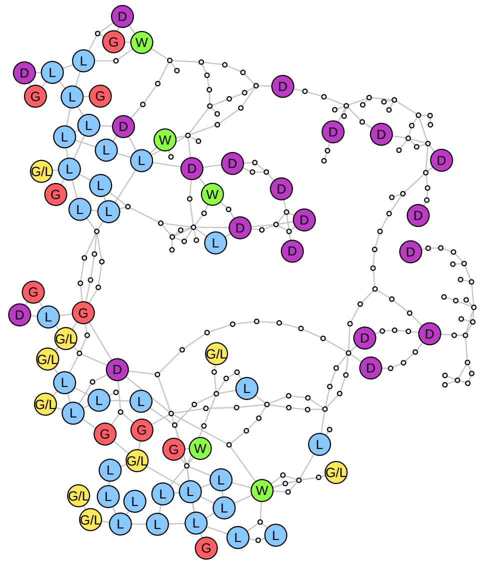

Fig. 5. Node-edge network representation of test system (Cases 2 and 3). The buses with thermal generation units, loads, both generation units and loads, wind generation units, and data centers are labeled as G, L, G/L, W and D, respectively.

source package DSP [33]. All computations were performed on `Blues`, a 310-node computing cluster at Argonne National Laboratory. Each computing node has two octo-core 2.6 GHz Xeon processors and 64 GB of RAM. Over 50,000 core-hours (six core years) of computing time were required for our analysis.

The cases studied reveal important trends that we summarize below. We then present numerical results to illustrate these trends.

- Case 1 reveals that there is significant spillage and stranded power in the base system due to imports and nonwind renewables. This finding is consistent with the observations of Section II-A. We also observe that, as expected, dispatch cost decreases initially as cheaper wind power displaces thermal generation, but eventually increases because of stranded power penalties. We also see that the variance of the cost increases dramatically as wind level is increased, indicating that the system becomes more vulnerable to uncertain wind-power variations.
- Case 2 reveals that positioning large data centers decreases system cost, even if locations are chosen arbitrarily and loads are inflexible. The reason is that the loads put stranded power to work, reducing penalties and moderately reducing system cost. We also observe that while cost is decreased, the variance of the cost is not improved (compared with Case 1).
- Case 3 reveals that collocating data centers at wind-farm locations gives little benefit to system cost. The slight benefit comes from wind power used to offset the data center loads, but stranded power is increased. Case 3 also reveals that collocation of data centers and wind farms does not reduce carbon emissions. In fact, increasing thermal generation (and thus emissions) because of increasing

stranded power. We find system cost variance decreased significantly compared with that of Cases 1 and 2. We attribute this decrease to better utilization of stranded power at data-center locations. This result thus highlights that stranded power can affect system vulnerability.

- Case 4 reveals the benefits of dispatchable loads, strategically positioned, can reduce power spillage from all sources (imports, nonwind renewables, and wind), not just wind (as in Case 3). Strategic positioning results in decreased system cost and decreased use of thermal generation (and thus emissions). We also find that cost variance is dramatically reduced compared to all other cases, indicating that the system can better manage wind-power fluctuations. Case 4 also reveals that duty factors of 60-80% be achieved for data centers at high wind penetration levels. This result occurs even if loads can be adjusted at no cost. Consequently, there is a natural economic incentive for data-center owners to provide flexibility. The incentives are provided by better stranded power utilization.

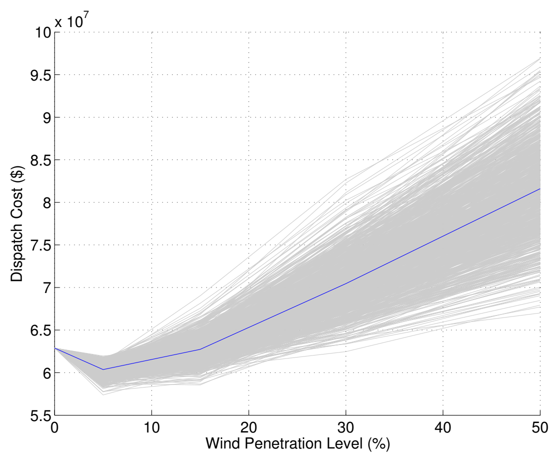

Fig. 6. Dispatch cost at different wind levels (Case 1).

A. Base System

We first analyze the impact of increasing wind levels in Case 1. Figure 6 shows that the average daily dispatch cost decreases below a 5% wind level, because of the use of cheap wind power. As wind levels increase, however, the dispatch cost is increased by 26% ($2.1 MUSD/day) relative to the system at 0% wind level. This increase is a combined effect of using more thermal generation to account for wind power variability and penalties induced by power spillage. Moreover, the system cost becomes more variable as we increase wind levels. In particular, the standard deviation is $5.2 MUSD/day at a 50% wind level and $0.7 MUSD/day at a 5% wind level. This variability indicates that the system cost becomes more vulnerable.

Figure 7a shows daily average power and spillage by generation source. As wind penetration is increased form 0% to 50%, the thermal generation decreases by 48% (from 349 to 182 GWh). To absorb the variability of the wind, the total

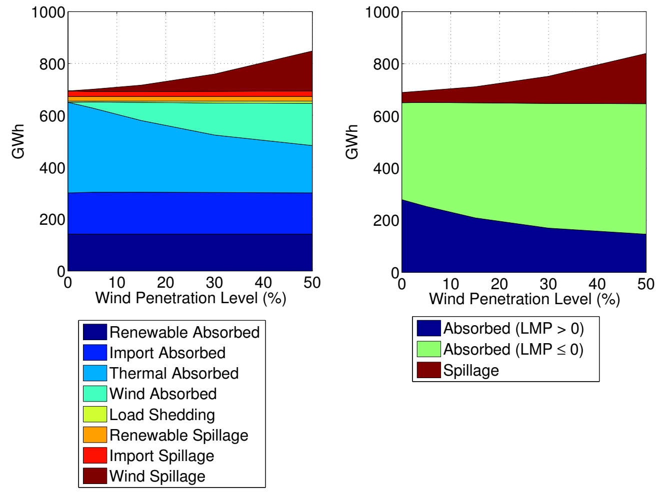

(a) Absorbed Power (b) Total Supply

Fig. 7. Absorbed power, spillage, and total supply for increasing wind level (Case 1).

supply increases by 21% (from 689 to 839 GWh), despite no increase in load. Consequently at 50% wind penetration level, 23% of generation is spilled (*not absorbed* into the system). This result reflects the difficulty in achieving high RPS because of wind variability as despite this overproduction, the daily load shedding is still 8 GWh.

Figure 7b shows daily average power by LMP value (price) and spillages. Recall that stranded power is the sum of spillage and power absorbed into the system at negative price (LMP$\leq$ 0), as defined in (11). We observe that as wind levels increase, both spillage and stranded power increase. Stranded power increases from 60% at a 0% wind level to 83% at a 50% wind level. While the total economic return for a generator may not be reflected in LMP alone, we note that the amount of power absorbed at positive price (profitable power) does not increase after a 15% wind level (see Figure 10b).

B. Adding Data Centers

Figure 8 compares dispatch cost for all cases. We note that the dispatch cost of the base system is decreased in all cases. The dispatch costs of the base system are decreased by Cases 2 and 3, respectively, by less than 5% at 0% wind level and by less than 13% at a 50% wind level. These results indicate that adding data centers has a beneficial effect and that this value increases as more stranded power is injected into the system. Case 4 reduces dispatch cost dramatically. A relative reduction of 98% is observed at 0% wind level and a relative reduction of 49% is observed at 50% wind level (compared with Case 1). As seen in Table II, the dramatic decrease in cost is due to minimization of spillage (penalized at large values). At a 0% wind level, in particular, dispatchable loads fully eliminate spillages. In Cases 2 and 3 we can see reductions in spillages, but these are much smaller than those observed in Case 4. In particular, we note that optimally placed dispatchable loads (Case 4) favor spillage reductions of nonwind renewable supply over wind supply. This result

TABLE II
THERMAL SUPPLY AND SPILLAGES AT DIFFERENT WIND LEVELS (GWH)

	Wind Level	RPS	Thermal Supply	Wind Spillage	Import Spillage	Renewable Spillage
Case 1	0%	22%	349	0	21	18
	5%	25%	325	8	19	18
	15%	32%	276	25	19	18
	30%	41%	221	66	20	18
	50%	47%	182	154	21	18
Case 2	0%	22%	439	0	14	18
	5%	25%	413	6	12	18
	15%	33%	361	19	12	18
	30%	43%	292	48	12	18
	50%	51%	238	122	13	18
Case 3	0%	22%	374	8	11	17
	5%	25%	355	17	11	17
	15%	32%	315	39	11	17
	30%	40%	272	86	11	17
	50%	45%	237	176	11	17
Case 4	0%	24%	358	0	0	0
	5%	28%	362	2	9	5
	15%	36%	309	14	11	7
	30%	45%	255	49	11	7
	50%	52%	220	130	11	8

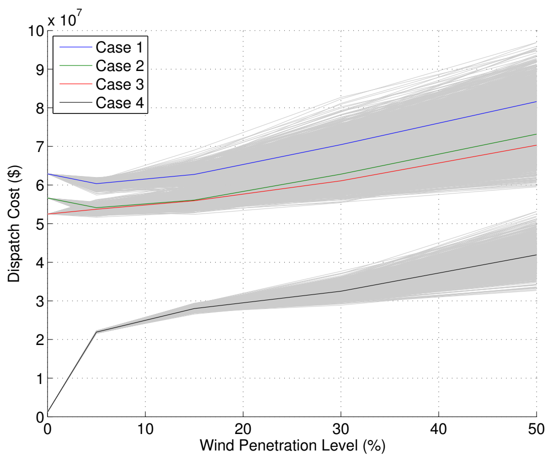

Fig. 8. Dispatch cost at different wind levels for Cases 1, 2, 3, and 4.

supports the conclusion that dispatchable loads, well placed, can provide much greater benefits for the power grid. The reason is that optimal placement allows them to eliminate spillages from a variety of generators in types and locations.

In Figure 8 we observe that significant value is obtained in Case 4 at all wind levels; with benefits as great at 40%. At low wind penetration, the dispatchable loads eliminate essentially all spillage, dramatically reducing associated penalties. As the wind penetration level increases, the gap with respect to the base system is reduced. The decreased benefit is due to the large amounts of spillage that are introduced at high wind levels and that cannot be fully eliminated even with optimally located dispatchable loads. This situation is observed in Table II where wind power spillage increases proportionally to the wind level. Fully eliminating spillage would require additional data centers.

A surprising result is that the system cost variance is dramatically reduced with dispatchable loads. Figure 9 illustrates this. In particular, the standard deviation for Case 4 is \$0.1 MUSD/day at a 5% wind level and \$3.6 MUSD/day at a 50% wind level. We recall that the standard deviations for Case 1 are \$0.7 MUSD/day at a 5% wind level and \$5.2 MUSD/day at a 50% wind level. For Case 2 the standard deviation is \$0.6 and \$5 MUSD/day at 5% wind level and 50% wind level, respectively. For Case 3 the standard deviation is \$0.6 and \$4.1 MUSD/day at 5% wind level and 50% wind level, respectively. For Case 4 we also note that variances are negligible for wind levels below 10% and remain small for wind levels below 20%. In contrast, the variances for Cases 1, 2, and 3 quickly increase with the wind level. The reduction in cost variance is the result of additional system flexibility.

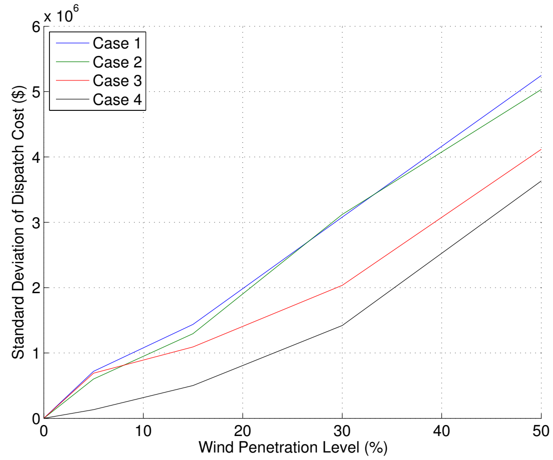

Fig. 9. Standard deviation of dispatch costs at different wind levels for Cases 1, 2, 3, and 4.

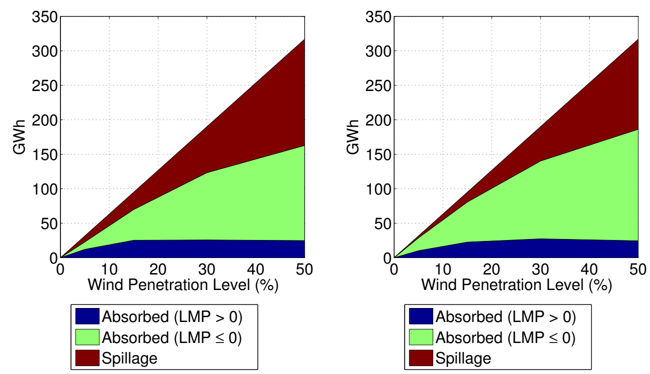

(a) Case 1 (b) Case 4

Fig. 10. Wind supply by LMP and spillage.

C. Impact on Wind Generators

Of particular interest are the generation, spillage, uneconomic and economic generation of wind power as the wind penetration level increases. Figure 10 shows the growing wind supply, what fraction is spillage, and what fraction is absorbed on both uneconomic and economic terms. For brevity we show results for only Case 1 and Case 4. Adding dispatchable loads obtains significant value and decreases wind spillage significantly, to zero at low wind penetration and by more than 15% at high penetration. However, there is a complementary

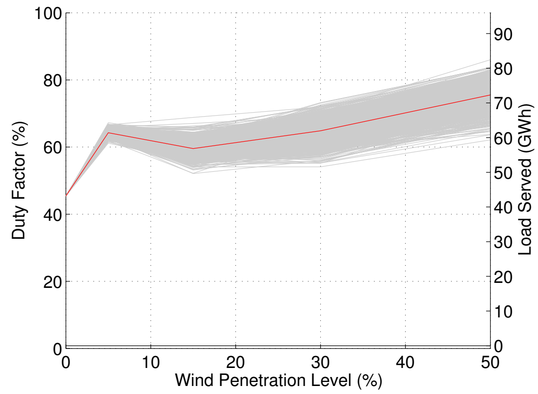

Fig. 11. Duty factors and total dispatched load for data centers in Case 4.

increase in the uneconomic power accepted by the grid, so the total stranded power remains large.

D. Data-Center Duty Factors

Figure 11 shows the daily dispatchable load served for all data centers and the corresponding duty factor. We recall that the total load capacity of the 20 dispatchable loads is 4 GW, which consumes 96 GWh a day. We can see that a duty factor of 45% is achieved at 0% wind level (duty factor is an average over scenarios and over time). This indicates that the dispatchable loads are used to decrease spillages of nonwind renewables and imports but the data center loads are far from fully served. The duty factor rapidly increases to 60% at a 5% wind level (indicating that stranded wind power adds value to the loads). The duty factor goes up to 75% at 50% wind penetration level and the trend is maintained (i.e., it does not settle). We have also found that the variance of the duty factors does not increase for wind levels higher than 15%. We note that unserved loads for flexible data centers are not penalized in Case 4 and that high duty factors can still be achieved. Consequently, we conclude that stranded power provides a natural economic incentive for flexible computing.

V. Discussion and Related Work

The computational analysis in Section IV illustrates that flexibility imposed by adding dispatchable loads can significantly reduce power spillage and dispatch cost by making better use of the stranded power. The flexibility cannot be achieved simply by consuming stranded power for nondispatchable data centers, which causes more thermal power generation and increasing carbon emissions. We acknowledge that the absolute value of adding dispatchable loads may vary depending on many factors including network topology (i.e., transmission line expansion), generation ramping capacities, and generation fuel mix. However, we note that changes in such factors are followed by significant capital expenses [34], whereas the dispatchable loads can be obtained by making the existing computing facilities flexibly operating.

A larger set of drivers is increasing volatility and scheduling challenges for the power grid. For example, San Diego's 100% renewable portfolio standard goal with a strategy based primarily on solar energy requires energy matching across temporal shifts of 8-16 hours; storage is likely a key element. Netmetering, widely practiced in residential solar systems, reduces demand during daylight, causing power grid demand to be anticorrelated with the availability of solar grid generation. Another example we explored, the addition of renewables on site with large loads such as industrial facilities, and army bases, also creates countervarying load with renewable grid generation. New sources of load such as electric vehicles will cause additional demand at times and locations whose predictability are not well understood.

Our results and those obtained from other systems such as batteries, buildings, and aluminum manufacturing facilities indicate that dispatchable loads provide more flexibility than do traditional models based on load shifting and demand response [35]–[37]. Such flexibility becomes relevant as we increase intermittency of renewable power, which patterns and peaks shift over time. Since small and large computing jobs can be scheduled in realtime, a dispatchable load of a data center can also be exploited by ISOs to provide frequency regulation at faster time scales [38]. A dispatchable load from a data center also provides a scalable alternative to battery systems. The reason is that dispatchable loads are linked to a primary service that makes profit (e.g., computing service) and therefore there is a natural economic driver driving investment in new installations. A battery system, on the other hand, is installed with the sole intention of being used as an ancillary service. This distinction is important because as electricity prices flatten out, less incentives will exist for battery installations while the need for data centers and computing capacity is expected to persist.

Our results and those of many others [21], [22], [39] highlight the challenges facing the power grid at a high level of wind penetration. For example, at 50% penetration, our studies show that 23% of the grid generation is wasted. Further, the study of stranded power suggests that the economics of the power grid are even more challenging at high RPS levels. At 50% wind penetration, we find that 83% of the power generation is uneconomic.[3]

Numerous efforts exploit renewable power for data centers (Microsoft, Facebook, and Amazon all buy renewable power through long-term purchase contracts; see [40]). Many researchers study the addition of renewable generators to data centers [13], [14], [41], matching our Case 3, but not Case 4, dispatchable loads, and thus these efforts will not deliver the grid benefits our results describe.

Our dispatchable data centers provide large-scale volatile computing resources, a significant departure from traditional data centers with nearly 100% availability and small-scale volatile resources such as in peer-to-peer [26], [42]–[44]. Effective exploitation of volatile resources requires new models of computing and resource management.

Numerous opportunities do exist for stranded power to either create economic benefit or reduce carbon emission. Examples include direct monetization through bitcoin mining,

[3]Based on the grid payments, setting externalities such as Production Tax Credits.

hydrogen generation, water desalination, and carbon recapture.

VI. Summary and Future Work

We examined the grid impact of adding data centers across a range of wind penetration levels. Our results show that increased wind penetration levels lead to high levels of spillage and uneconomic absorbed generation, which together we call stranded power. Significant at even moderate levels of wind penetration, these numbers grow even higher at high levels of wind penetration (83% at 50% penetration). In short, the quantity of stranded power in today's grid is large and grows rapidly with RPS levels.

Two of our scenarios added data centers in conventional ways, first alone and second with each data center paired with a wind power plant of equal average generation. However, our third scenario, adding data centers as a dispatchable load that the grid could turn on or off based on grid benefits, gave surprising results. Spillage was dramatically reduced, and average power cost also dramatically decreased; as much as 44%. Closer examination shows that such dispatchable loads enable better utilization of wind generation and significantly more efficient grid management. Our studies show that these dispatchable loads achieve duty factors as high as $70 \sim 80\%$, making them usable for a broad array of applications, including one we are considering, intermittent computing.

Directions for future work include exploring the impact of transmission changes and the types of computing services that might be feasible.

Appendix A
Notation

Sets:

$\mathcal{D}$	Set of demand loads
$\mathcal{D}_n$	Set of demand loads at bus n
$\mathcal{G}$	Set of all generators
$\mathcal{G}_n$	Set of all generators at bus n
$\mathcal{I}$	Set of import points
$\mathcal{I}_n$	Set of import points at bus n
$\mathcal{L}$	Set of transmission lines
$\mathcal{L}_n^+$	Set of transmission lines to bus n
$\mathcal{L}_n^-$	Set of transmission lines from bus n
$\mathcal{N}$	Set of buses
$\mathcal{R}$	Set of nonwind renewable generators
$\mathcal{R}_n$	Set of nonwind renewable generators at bus n
$\mathcal{S}$	Set of wind production scenarios
$\mathcal{T}$	Set of time periods
$\mathcal{W}$	Set of wind-farm locations

Parameters:

B_l	Susceptance of transmission line l
C_i	Generation cost of generator i $[\$/MWh]$
C_j^d	Load shedding penalty at load j $[\$/MWh]$
C_i^w	Spillage penalty at wind farm i $[\$/MWh]$
C_i^m	Spillage penalty at import point i $[\$/MWh]$
C_i^r	Spillage penalty at nonwind renewable i $[\$/MWh]$
C_n^s	Value of lost dispatchable load at bus n $[\$/MWh]$
D_{jt}	Demand load of consumer j at time t $[MWh]$
F_l^{max}	Maximum power flow of transmission line l $[MW]$
K	Maximum number of dispatchable loads
M_{it}	Power production of import i at time t $[MWh]$
P_i^{max}	Maximum power output of generator i $[MWh]$
R_{it}	Power production of nonwind renewable generator i at time t $[MWh]$
RU_i	Ramp-up limit of generator i $[MW]$
RD_i	Ramp-down limit of generator i $[MW]$
U	Dispatchable load capacity $[MW]$
W_{wt}	Wind power generation from generator w at time t $[MWh]$
W_{wts}	Wind power generation from generator w at time t for scenario s $[MWh]$
π_s	Probability of wind production scenario s
Θ_{nt}^{min}	Minimum phase angle at bus n at time t $[degree]$
Θ_{nt}^{max}	Maximum phase angle at bus n at time t $[degree]$

Decision variables:

d_{jt}	Load shedding at load j at time t $[MWh]$
d_{jts}	Load shedding at load j at time t for scenario s $[MWh]$
f_{lt}	Power flow of transmission line l at time t $[MWh]$
f_{lts}	Power flow of transmission line l at time t for scenario s $[MWh]$
m_{it}	Power spillage at import point i at time t $[MWh]$
m_{its}	Power spillage at import point i at time t for scenario s $[MWh]$
p_{it}	Power output of generator i at time t $[MWh]$
p_{its}	Power output of generator i at time t for scenario s $[MWh]$
r_{it}	Power spillage at nonwind renewable i at time t $[MWh]$
r_{its}	Power spillage at nonwind renewable i at time t for scenario s $[MWh]$
u_{nts}	Dispatchable load served at bus n at time t for scenario s $[MWh]$
w_{it}	Power spillage at wind farm i at time t $[MWh]$
w_{its}	Power spillage at wind farm i at time t for scenario s $[MWh]$
x_n	Number of dispatchable loads installed at bus n $[degree]$
θ_{nt}	Phase angle at bus n at time t $[degree]$
θ_{nts}	Phase angle at bus n at time t for scenario s $[degree]$

Acknowledgment

This material is based upon work supported by the U.S. Department of Energy, Office of Science, under contract number DE-AC02-06CH11357 and the National Science Foundation under Award CNS-1405959. We gratefully acknowledge the computing resources provided on Blues, a high-performance computing cluster operated by the Laboratory Computing Resource Center at Argonne National Laboratory. We thank Anthony Papavasiliou for generously sharing the Western Electricity Coordinating Council test system data.

References

[1] Intergovernmental Panel on Climate Change, "Climate change 2014: Synthesis report," 2014. [Online]. Available: http://www.ipcc.ch

[2] A. Gore, "An inconvenient truth," Documentary Film, 2006.

[3] U. N. F. C. on Climate Change. (1997) Kyoto protocol. [Online]. Available: http://unfccc.int/kyoto_protocol/items/2830.php

[4] ——. (2015) Paris climate change conference. [Online]. Available: http://unfccc.int/meetings/paris_nov_2015/meeting/8926.php

[5] H. S. Dunn. (2010) The Carbon Footprint of ICTs. [Online]. Available: https://www.giswatch.org/thematic-report/sustainability-climate-change/carbon-footprint-icts

[6] C. Pettey. (2007) Gartner estimates ICT industry accounts for 2 percent of global CO2 emissions. [Online]. Available: https://www.gartner.com/newsroom/id/503867

[7] S. Chinnadurai and K. Nandavarapu. (2015, November) Increasing carbon footprint of the ict sector. [Online]. Available: http://www.blueoceanmi.com/blueblog/increasing-carbon-footprint-of-the-ict-sector-2/

[8] Greenpeace. (2014, April) Clicking clean: How companies are creating a green internet. [Online]. Available: http://www.greenpeace.org/usa/wp-content/uploads/legacy/Global/usa/planet3/PDFs/clickingclean.pdf

[9] L. Barroso and U. Holzle, *The Data Center as Computer: An Introduction to the Design of Warehouse-Scale Machines*. Morgan-Claybpool, 2009.

[10] C. Patel and A. Shah, "A cost model for planning, development, and operation of a data center," Hewlet-Packard Corporation: HP Labs, Tech. Rep. HPL-2005-107R1, 2005.

[11] (2015) Google Makes Its Biggest Renewable Energy Purchase Yet. [Online]. Available: http://www.datacenterknowledge.com/archives/2015/12/03/google-buys-842-mw-of-renewables-for-google-data-centers/

[12] (2014) Apple to build a 3rd massive solar panel farm in North Carolina. [Online]. Available: https://gigaom.com/2014/07/08/apple-to-build-a-3rd-massive-solar-panel-farm-in-north-carolina/

[13] Í. Goiri *et al.*, "Parasol and GreenSwitch: Managing datacenters powered by renewable energy," in *ACM SIGARCH Computer Architecture News*. ACM, 2013, pp. 51–64.

[14] J. L. Berral *et al.*, "Building green cloud services at low cost," in *ICDCS'14*. IEEE, 2014, pp. 449–460.

[15] The Washington Post. (2014) Use the Web? Congrats, you're an environmentalist. [Online]. Available: https://www.washingtonpost.com/news/wonk/wp/2014/11/06/like-kite-surfing-the-internet-why-google-microsoft-and-yahoo-are-buying-up-wind-energy/

[16] X. Wang *et al.*, "Grid-aware placement of datacenters and wind farms," in *IGSC'15*. IEEE, 2015.

[17] Í. Goiri *et al.*, "Intelligent placement of datacenters for internet services," in *ICDCS'11*. IEEE, 2011, pp. 131–142.

[18] California Public Utilities Commission (CPUC). [Online]. Available: http://www.cpuc.ca.gov/PUC/energy/Renewables

[19] C. Megerian and J. Panzar, "Gov. Brown signs climate change bill to spur renewable energy, efficiency standards," Los Angeles Times Newspaper, September 2015.

[20] Renewable Energy13.4% of US Electricity Generation in 2014. [Online]. Available: http://cleantechnica.com/2015/03/10/renewable-energy-13-4-of-us-electricity-generation-in-2014-exclusive/

[21] "Investigating a higher renewables portfolio standard in california: Executive summary," Report from Energy and Economics, Inc., January 2014, http://ethree.com/.

[22] D. Lew, L. Bird, M. Milligan, B. Speer, X. Wang, E. M. Carlini, A. Estanqueiro, D. Flynn, E. Gomez-Lazaro, N. Menemenlis *et al.*, "Wind and solar curtailment," in *International Workshop on Large-Scale Integration of Wind Power Into Power Systems*, 2013.

[23] A. A. Chien *et al.*, "The zero-carbon cloud: High-value, dispatchable demand for renewable power generators," *The Electricity Journal*, pp. 110–118, 2015.

[24] F. Yang and A. A. Chien, "ZCCloud: Exploring wasted green power for high-performance computing," in *IPDPS 2016*. IEEE, May 2016.

[25] O. Agmon Ben-Yehuda, M. Ben-Yehuda, A. Schuster, and D. Tsafrir, "Deconstructing Amazon EC2 Spot Instance pricing," *ACM Trans. Econ. Comput.*, vol. 1, no. 3, pp. 16:1–16:20, Sep. 2013. [Online]. Available: http://doi.acm.org/10.1145/2509413.2509416

[26] "Google compute engine: Creating a preemptible VM instance," https://cloud.google.com/compute/docs/instances/preemptible.

[27] W. Guo, K. Chen, Y. Wu, and W. Zheng, "Bidding for highly available services with low price in spot instance market," in *Proceedings of the 24th International Symposium on High-Performance Parallel and Distributed Computing*, ser. HPDC '15. New York, NY, USA: ACM, 2015, pp. 191–202. [Online]. Available: http://doi.acm.org/10.1145/2749246.2749259

[28] V. M. Zavala, K. Kim, M. Anitescu, and J. Birge, "A stochastic electricity market clearing formulation with consistent pricing properties," *arXiv:1510.08335 [q-fin.EC]*, no. ANL/MCS-P5110-0314, 2015. [Online]. Available: http://arxiv.org/abs/1510.08335

[29] A. Papavasiliou and S. S. Oren, "Multiarea stochastic unit commitment for high wind penetration in a transmission constrained network," *Operations Research*, vol. 61, no. 3, pp. 578–592, 2013.

[30] M. Bastian, S. Heymann, M. Jacomy *et al.*, "Gephi: an open source software for exploring and manipulating networks." *ICWSM*, vol. 8, pp. 361–362, 2009.

[31] M. Lubin and I. Dunning, "Computing in operations research using Julia," *INFORMS Journal on Computing*, vol. 27, no. 2, pp. 238–248, 2015.

[32] J. Huchette, M. Lubin, and C. Petra, "Parallel algebraic modeling for stochastic optimization," in *Proceedings of the 1st First Workshop for High Performance Technical Computing in Dynamic Languages*. IEEE Press, 2014, pp. 29–35.

[33] K. Kim and V. M. Zavala, "Algorithmic innovations and software for the dual decomposition method applied to stochastic mixed-integer programs," *Optimization Online*, 2015.

[34] U.S. Energy Information Administration. (2013) Updated capital cost estimates for utility scale electricity generating plants. [Online]. Available: http://www.eia.gov/forecasts/capitalcost/

[35] S. Han, S. Han, and K. Sezaki, "Development of an optimal vehicle-to-grid aggregator for frequency regulation," *Smart Grid, IEEE Transactions on*, vol. 1, no. 1, pp. 65–72, 2010.

[36] P. Zhao, G. P. Henze, S. Plamp, and V. J. Cushing, "Evaluation of commercial building HVAC systems as frequency regulation providers," *Energy and Buildings*, vol. 67, pp. 225–235, 2013.

[37] D. Todd, M. Caufield, B. Helms, A. P. Generating, I. M. Starke, B. Kirby, and J. Kueck, "Providing reliability services through demand response: A preliminary evaluation of the demand response capabilities of Alcoa Inc," *ORNL/TM*, vol. 233, 2008.

[38] H. Hao, B. M. Sanandaji, K. Poolla, and T. L. Vincent, "Frequency regulation from flexible loads: Potential, economics, and implementation," in *American Control Conference (ACC), 2014*. IEEE, 2014, pp. 65–72.

[39] L. Bird, M. Milligan, and D. Lew, "Integrating variable renewable energy: Challenges and solutions," *National Renewable Energy Laboratory*, 2013.

[40] G. Cook *et al.*, "Clicking clean: how companies are creating the green internet," *Greenpeace Inc., Washington, DC*, 2014.

[41] M. E. Haque, I. Goiri, R. Bianchini, and T. D. Nguyen, "Greenpar: Scheduling parallel high performance applications in green datacenters," in *Proceedings of the 29th ACM on International Conference on Supercomputing*, ser. ICS '15. New York, NY, USA: ACM, 2015, pp. 217–227. [Online]. Available: http://doi.acm.org/10.1145/2751205.2751221

[42] A. Oram, Ed., *Peer-to-Peer: Harnessing the Power of Disruptive Technologies*. O'Reilly, 2001.

[43] "HTCondor." [Online]. Available: http://research.cs.wisc.edu/htcondor/

[44] "Spot Instance." [Online]. Available: http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/using-spot-instances.html