Grid-Aware Placement of Datacenters and Wind Farms*

Xiaoying Wang$^{\pm}$, Divya Kurthakoti Chandrashekhara$^{\delta}$, Md Haque‡, Íñigo Goiri†, Ricardo Bianchini†‡, Thu Nguyen‡

$^{\pm}$Qinghai University, China
wxy_cta@qhu.edu.cn

$^{\delta}$Siemens Corporate Research
divya_kc@ieee.org

†Microsoft Research
{inigog, ricardob}@microsoft.com

‡Rutgers University
{mdhaque, ricardob, tdnguyen}@cs.rutgers.edu

***Abstract*—Datacenters are being constructed at a rapid pace. Concurrently, offsetting renewable power plants are also being built to mitigate the environmental impact of the new datacenters' massive energy consumption. Research efforts that have studied how to best place new datacenters and renewable power plants have mostly neglected the impact on the electricity transmission grid. In this paper, we show that accounting for the impact on the transmission grid can be mutually beneficial to both datacenter owners and grid operators. Specifically, locating datacenters and renewable power plants at strategic places in the grid could help to minimize (i) overloading of transmission lines, (ii) grid voltage variations outside the acceptable range, and (iii) transmission system losses. We develop an optimization framework for placing a new datacenter and offsetting wind farm, and use it in a case study to show that considering transmission losses along with datacenter costs can lead to different placements and lower overall cost. Interestingly, co-locating the datacenter and wind farm does not always lead to lowest impact on the transmission grid and lowest overall cost. Thus, we conclude that the impact of location on the costs of *both* the datacenter owners and grid operators should be considered when placing new datacenters and offsetting renewable power plants.**

I. INTRODUCTION

Datacenters are being constructed at a rapid pace as computing is increasingly moving to the cloud (e.g., [1]). Some companies are also seeking to mitigate the environmental impact of the massive energy consumption of new datacenters by building offsetting renewable power plants (e.g., [2]–[4]). A large datacenter can require upward of 100MW of power, representing a significant load on the grid. Thus, attaching new large datacenters (and their offsetting renewable power plants) to the grid can lead to overloading of transmission lines, voltage variations outside the acceptable range, and increased transmission system losses. Since adding transmission capacity to the grid takes a long time (typically 7 to 10 years) and is extremely expensive [5], it is imperative to study the impact of the increasing penetration of such *renewable powered datacenters*[1] on the transmission grid.

Researchers have studied the placement of new datacenters since they are expensive to build and operate, and costs are partly location-dependent [6]–[10]. However, these studies have not considered the impact of new datacenters on the grid. Others have studied whether offsetting renewable power plants should be co-located with the datacenters [11], and have considered transmission loss when the datacenter and the renewable power plant are physically distributed. However, they assumed a fixed loss percentage.

In this paper, we first study (using simulation) the impact of placing a new datacenter at different locations within the real world transmission system of the New England Independent System Operator (ISO). This transmission system spans most of the North Eastern region of the United States and some parts of Canada. Our results show that different placements can lead to significant differences in the overloading of transmission lines, the number of voltage variations outside the acceptable range, and transmission system losses. Interestingly, co-locating the datacenter with a wind farm (we extended the base model of the New England ISO transmission network to include wind farms that have been developed since the model was created) does not always minimize the impact.

Motivated by the results of the above study, we then develop an optimization framework for placing a new datacenter and an offsetting wind farm.[2] Our framework is similar to previous work in that it considers various capital and operational costs, some of which are location-dependent, and is formulated to minimize overall cost. Unique to this work, however, is the added consideration of the cost of system loss in the transmission grid, as well as constraints for avoiding transmission line overloading and unacceptable voltage variations.

Finally, we use our optimization framework in a case study to demonstrate the potential benefits of our placement approach. Specifically, we study the placement of a new datacenter and an offsetting wind farm in the New England ISO system. Our results show that the strategic placement of a new renewable datacenter can indeed help to avoid the occurrences of transmission line overloading and unacceptable voltage variations. Further, strategic placement while considering transmission system losses can lead to 7.6% cost savings

---

*This work was partially supported by NSFC grant No. 61363019, NSF grant CSR-1117368, the Rutgers Green Computing Initiative, the BSC-CNS Severo Ochoa program, and the TIN2012-34557 project.

[1] We are calling a datacenter together with an offsetting renewable power plant a "renewable powered datacenter" for brevity, even though the two may *not* be physically co-located.

[2] Note that we are not proposing a new method for the placement of renewable power plants, which has been studied extensively in the industry. Rather, the point of our work is that costs for the placement/construction/operation of a new datacenter, renewable power plant, *and* their impact on the transmission grid should be studied together.

978-1-5090-0172-9/15/$31.00 © 2015 IEEE

compared to simply co-locating the datacenter and wind farm, and 3.3% compared to placement while ignoring transmission system losses. While these percentages are modest, they translate to savings of tens of millions of dollars per year.

Contributions. Our main contributions include: (i) demonstrating the potential impact of renewable powered datacenter placement on a grid transmission system, (ii) proposing an optimization framework for intelligent placement of renewable powered datacenters that accounts for the impact on the grid transmission system, and (iii) exploring the potential benefits of the placement framework using a realistic case study.

II. IMPACT OF DATACENTER PLACEMENT ON THE GRID

In this section, we first discuss different ways that the placement of renewable powered datacenters can impact the power grid. We then describe a simulation system that we have built to quantitatively explore this impact. Finally, we present some results to show that it is important to consider grid impacts when locating new renewable powered datacenters.

A. Potential Impacts

The rapidly increasing penetration of large renewable powered datacenters can potentially impact the power grid in several ways.

1) Overloading of transmission lines: Transmission lines (also referred to as "branches") are used to transport power from the large generators to the loads in an electric grid. The power carrying capacity of each line is limited to protect the line from overheating, mainly due to the line resistive losses, i.e., I^2R, where I is the current flowing through the line and R is the resistance of the line.

A transmission line typically has two ratings: a short term and a long term capacity rating. During certain wind and system load (including datacenter load) conditions, some of the transmission lines could get overloaded. If this happens during the normal operation of the grid, one of the following will be done: i) if an electronic power flow controller is available, it is used to control the power flow through the overloaded line; or ii) in extreme situations, the overloaded line is disconnected which may result in power supply interruption to the loads.

If major transmission lines are getting overloaded often during the year, new lines are planned and built. As already mentioned, this solution is very expensive and takes a long time. The need for such expensive grid retrofits may be minimized by planning the location of large new generators such as wind farms and large new loads such as datacenters.

2) Voltage variations: The voltage magnitude varies in the electric grid, and needs to be maintained within a narrow range (for example +/- 5% of nominal) to avoid damaging sensitive electronic loads. Unfortunately, there are times when changes in the power output of renewable power plants can cause the voltage to vary beyond the acceptable limits. Such over/under voltage problems can be mitigated by appropriately locating new renewable powered datacenters.

3) System losses: Historically, the electric grid was designed to have large central generating stations that are located far away from the load centers. The power from these central sources would be transmitted to the load centers over transmission lines. While designing such a grid, the generator location and the transmission line voltage level as well as the path would be optimized to minimize the line losses. With increasing penetration of renewable power, however, this scenario has changed: the generation sources are distributed and may be located near the load centers. Yet, we still use the existing transmission lines that were planned and built about 50 years ago or earlier. This may result in higher line losses and sub-optimal power transmission between generation sources and loads. Since we cannot re-design the entire electric grid to minimize line losses, we need to leverage the flexibility we have in locating new sources and loads, i.e., renewable powered datacenters in our specific case.

B. Simulation study

In order to quantitatively study the impact of renewable powered datacenter placement in a power grid, we consider the New England ISO transmission network. We choose this system for the following reasons:

Wind power expansion in New England: System studies carried out by the New England ISO show that there is a potential for integrating up to 12 Giga-Watts of wind power in this region. This enormous potential for wind power makes the region an interesting destination for new datacenters with offsetting wind farms.

Transmission network upgrades: A study carried out by the New England ISO shows that they could potentially integrate wind resources to meet up to 24% of the region's total annual electric energy needs in 2020 if the system includes transmission upgrades. The development of new wind farms is more economical if these transmission upgrades can be limited.

Positive impacts of wind power in New England: Introducing large amounts of low-marginal-cost wind generation tends to depress the spot price and reduce the price differential for bulk power between day and night. The studies above also show that there would be only a relatively small increase in the use of existing pumped-storage hydro power for large wind penetrations, mostly because the flexible natural-gas-fired generation fleet can provide most of the system balancing.

We will show in our study that wind power penetration can be increased within the New England ISO transmission system more cost effectively by strategically locating new datacenter loads. Intuitively, it may appear that co-locating datacenter loads and wind farms would minimize their impacts on the transmission grid. However, this is not necessarily true since *the impact on the transmission grid depends on how the placements affect the power flows in the entire transmission network.* For example, strategically placing a new load can *reduce* system losses in a transmission network.

Before we carry out the case studies, we describe the approximations we make and the models we use for each subsystem to be considered in our study.

New England transmission system model: Our study is based on a widely used model of the New England transmission network [12]. A single-line diagram of this test system is shown in Figure 1. As shown in the figure, the model lumps all the generators, loads and transmission lines in the New England ISO region to 10 generators, 19 loads and 46 lines and transformers. The 10 generator buses are numbered from 30-39 in Figure 1. Specifically, bus 39 represents the aggregation of a large number of generators interconnected to the rest of US/Canada. This model includes load data for periods of normal (nominal) and high loads.

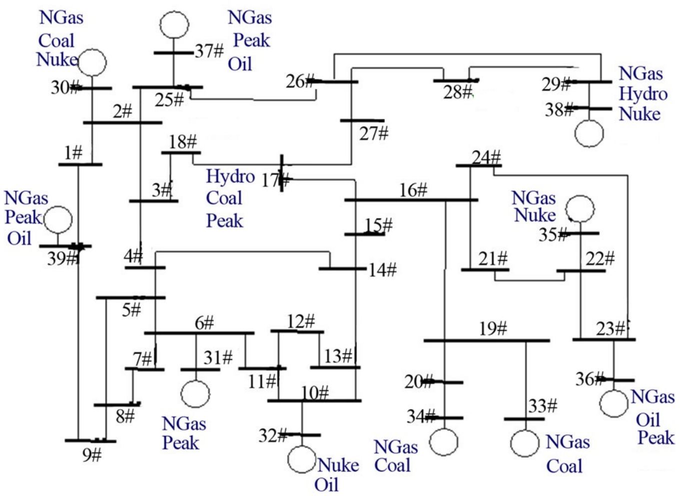

Fig. 1: The New England 39 bus test system

Background datacenters: Since the model of the New England system was created, large datacenters have been added to the region. To account for this added load, we collected information on datacenters in the area from [13]. Then, using a mapping of the buses in the model to geographic locations, we aggregate all the datacenters in the New England ISO region into six datacenters, each representing datacenters in a state, and placed each aggregated datacenter on a bus appropriate for that state. We estimate the size of each "aggregated" datacenter as $L_i = (n_i * 1.45 * 9.8GW)/1278$, where L_i is the aggregated load of the ith state, n_i is the number of datacenters in that state, 9.8GW is the upper bound of total electricity used by US datacenters in 2010 [14], 1.45 is a rough estimate of growth in datacenter electricity consumption from 2010 to 2014 (using information from [14]), and 1278 is the total number of datacenters in the US as reported in [13]. Table I shows the placement of these aggregated datacenters and their sizes. All datacenters are assumed to be operating at full capacity.

TABLE I: Background datacenters and their locations.

Aggregate DC	State	Number of DCs	Aggregate size (MW)	Location (Bus No.)
DC1	Connecticut	12	133.43	6
DC2	Maine	3	33.36	29
DC3	Vermont	4	44.48	25
DC4	Rhode Island	3	33.36	20
DC5	New Hampshire	4	44.48	16
DC6	Massachusetts	27	300.21	4

Background wind farms: We have similarly added wind farms to the model to represent recent penetration of wind energy. Specifically, lumped models of several wind farms within geographical regions have been connected to bus 18, 28, 36, 37 and 38. The locations and capacity settings of the five wind farms are presented in Table II. We assume that each farm can be represented by n identical wind turbines, where n is equal to the total farm rated capacity divided by the individual wind turbine rating. This approximation does not change any of our results because we are interested in studying the global impact of wind farm powered datacenters on the electric grid. Also, since most of the wind turbines in this region are GE 1.5MW machines, we use the wind speed versus power characteristics of this turbine model [15]. For this particular turbine model, the cut-in wind speed, i.e., the wind speed at which the turbine starts producing power, is 5m/s and the cut-off wind speed is 25m/s, beyond which the turbine will be shut down for safety reasons. The wind turbine produces rated output that grows linearly from near 0MW to 1.2MW for wind speeds from 6m/s to 10m/s, then more slowly to reach the top rated output of 1.5MW between wind speeds of 13-25m/s. We assume that wind farms are not provisioned with energy storage capacity since energy storage is expensive and not always cost effective.

TABLE II: Background wind farms and their locations.

Aggregate WF	State	Capacity (MW)	Location (Bus No.)
WF1	New Hampshire	100	18
WF2	Maine	90	28
WF3	Vermont	90	36
WF4	Maine	90	37
WF5	Massachusetts	90	38

C. Case study

We use the above "modified" New England model to study the impact of connecting a new datacenter to an electricity transmission system. Specifically, we assess the impact using the three metrics described earlier for three different case studies:

Case 1: The base modified New England system.

Case 2: The modified New England system with one additional 200MW datacenter co-located with one of the wind farms in Table II.

Case 3: The modified New England system with one additional 200MW datacenter connected to a bus away from all wind farms in Table II.

We compute the power flows through the transmission lines, the bus voltages, and the system losses by solving a set of power flow equations that model the power balance in a transmission system (i.e., net load + losses = total generation). Within an electric grid the power can be easily measured at the loads and at generators. Also, some generators have the capability to regulate the voltage at a bus at a constant preset reference value. The power flow equations are used to calculate the bus voltages (magnitude and angle), for a given network and a set of load and generation powers. The power flow equations for a generic n bus network with k branches are:

$$P_i = \Sigma_{j=1}^{n}(|Y_{ij}||V_i||V_j|cos(\theta_{ij} + \delta_j - \delta_i)) \qquad (1)$$

$$Q_i = -\Sigma_{j=1}^{n}(|Y_{ij}||V_i||V_j|sin(\theta_{ij} + \delta_j - \delta_i)) \qquad (2)$$

where P_i and Q_i are real and reactive powers at the i^{th} bus; $|V_i|\angle\delta_i$ are the voltage magnitude and angle at the the i^{th} bus; and $|Y_{ij}|\angle\theta_{ij}$ is the admittance of the branch between i^{th} and j^{th} bus. For a given pair of powers P_i and Q_i at the i^{th} load bus, the above power flow equations are used to solve for the voltage magnitude and angle at the i^{th} bus. Since the above equations are non-linear functions of voltage, they are solved iteratively using the Newton Raphson method. Once the bus voltages have been calculated, the line flows and system losses can be computed.

D. Results and Discussion

We now explore the impact of the above three cases under different load and wind conditions.

First, we simulate the three cases under peak system load (6,885MW in total) to explore differences in line overloading. The wind speed is set to LOW (4-6m/s), implying that the power generated by the wind farms is nearly zero (less than 2.3% of the rated capacity). For Case 2, the datacenter is connected to bus 18 (and so is co-located with wind farm WF1), and for Case 3, the datacenter is connected to bus 10 (away from all wind farms). Table III shows the number of line overloads in the system. These results highlight that different placements of the additional datacenter can affect line overloading. Although there are only a few cases of line overload in our experiments, annually the number of overloads and their durations will depend on the frequency and duration of occurrences of a particular wind speed and load condition. If it is too frequent or more persistent, then the overloads could be a serious problem and might require building new transmission lines. According to [5], the estimated cost of building new transmission lines of 345kV voltage level is about \$2.5M per mile, which is very expensive with total costs in the billions of dollars. Hence, it's important to choose the right place for datacenters in order to mitigate line overloading occurrences.

TABLE III: Overloaded transmission lines.

Case	No. of overloaded lines	List of overloaded lines
1	0	None
2	1	bus 4 - bus 5
3	0	None

Next, we investigate unacceptable voltage variations in the transmission system, as shown in Table IV. Here, the wind speed setting is MEDIUM (8-10m/s), which means the generated power of each wind turbine will be 33.3%-78.7% of its rated capacity. For Case 2 the datacenter is located at bus 38 (co-located with wind farm WF5), and for Case 3 the datacenter is located at bus 25 (away from all wind farms). The acceptable voltage range of a bus is set to [0.95p.u.,1.05p.u.]. We observe that under these conditions, there is already an unacceptable voltage deviation in the base modified New England system. Case 2 increases the number of unacceptable voltage deviations to four. In contrast, Case 3 has no unacceptable voltage deviations. These results show that unacceptable voltage variations can be mitigated by carefully choosing the place for the new datacenter.

TABLE IV: Unacceptable voltage variations.

Case	No. of buses with unacceptable voltage deviations	List of buses with unacceptable voltage deviations
1	1	bus 25
2	4	bus 25, bus 26, bus 28, bus 29
3	0	None

Finally, we compare the total system losses for the three cases under three different wind speeds: LOW, MEDIUM, and HIGH. The results are shown in Figure 2. Here, the load is set to normal (6,254MW in total). LOW and MEDIUM wind speeds are as before. HIGH wind speed corresponds to 11m/s, such that the output power of the wind turbine is very close to its rated capacity. For example, a 100MW wind farm would be generating at least 89.6MW of power under HIGH wind speed. The datacenter is located at bus 18 (co-located with wind farm WF1) for Case 2, and at bus 10 (away from all wind farms) for Case 3.

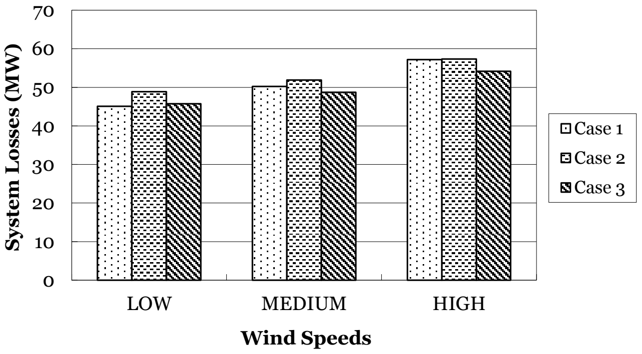

Fig. 2: Transmission system losses.

Results in Figure 2 show that the co-location case (Case 2) could lead to greater loss than Case 1 and Case 3. For example, at MEDIUM wind speed, the system loss for Case 3 is about 6% less than Case 2, which illustrates that system loss can be *reduced* by careful placement of the datacenter, and co-location with a wind farm is not necessarily the best choice. This may seem counter intuitive. However, adding a load can alter the power flows in the entire system, which can lead to significant changes in the transmission system losses.

Note that even though we have provided only one set of results for a specific set of conditions, we have simulated different wind speeds and load conditions. The results consistently show that the placement of a new datacenter can impact the number of transmission line overloads, the number of unacceptable voltage deviations, and system losses, and that co-location with a wind farm is often not the best placement strategy.

III. COST-BASED PLACEMENT

As shown in the last section, the placement of new datacenters and wind farms can have significant impact on a grid transmission system. As previously shown in [6], [10], the placement of datacenters can also significantly impact

TABLE V: Framework parameters. l is a location, and t is a time period.

Symbol	Meaning	Unit
dcCapacity	desired power capacity for computing in DC	kW
wfCapacity	desired power production capacity of wind farm	kW
$\mathit{pLand}(l)$	land price at l	\$/m^2
$\mathit{PUE}(l,t)$	PUE at l during t	
$\mathit{maxPUE}(l)$	maximum PUE at l	
dcArea	land needed per kW of DC compute capacity	m^2/kW
$\mathit{cLinePow}(l)$	cost to layout power line from l to the closest power plant	\$
$\mathit{cLineNet}(l)$	cost to layout optical fiber from l to closest network backbone	\$
$\mathit{pBuildDC}(c)$	per kW price of building a datacenter with c power capacity	\$/kW
serverPow	server peak power demand	kW/serv
switchPow	switch peak power demand	kW/switch
servsSwitch	number of servers per switch	servs/switch
pServer	price of a server	\$/serv
pSwitch	price of a network switch	\$/switch
pNBWServ	cost of external network bandwidth per server	\$/serv-month
$\mathit{pEnergy}(l)$	grid electricity price at l	\$/kWh
$\mathit{powNeed}(t)$	avg computing power demand of DC during t	kW
$\beta(l,t)$	avg generation efficiency of wind energy at l during t	%
wfArea	land needed per kW wind power	m^2/kW
pBuildWF	per kW price of building a wind power plant	\$/kW
$\mathit{revEnergy}(l)$	revenue for selling wind energy to grid at l	\$/kWh
$\mathit{transLoss}(t)$	avg system transmission loss in grid during t	kW
pTransLoss	the price for system transmission losses per kWh	\$/kWh

their costs. Thus, in this section, we develop an optimization framework for placing new renewable powered datacenters that unifies the impact of location on the costs of datacenters and losses in the transmission grid. For simplicity, we study the case of placing a single new datacenter and one offsetting wind farm. Our framework can be easily extended to place multiple datacenters and multiple wind farms; of course, solving the optimization problem can become much more challenging for such cases. Our optimization seeks to find locations for the datacenter and wind farm that lead to the lowest total cost for construction and operation. The overall framework is based on computing the cost throughout a year, using amortized capital costs, operational costs of the datacenter, operational revenue of the wind farm, and system losses in the transmission network.

A. Optimization Framework

Table V lists the set of parameters in our framework. Using these parameters, we define the optimization problem shown in Figure 3. The objective of this optimization problem is to minimize the total cost (*totalCost*) of building and operating a datacenter of a given size (*dcCapacity*) and a wind farm of a given size (*wfCapacity*) over a given time period (T). The datacenter and wind farm can each be placed at any location within a set of given locations. The total cost has three components, the cost of the datacenter ($dcCost$), the cost of the wind farm ($wfCost$), and the cost of losses in the transmission system ($transCost$).

Minimize *totalCost*, where

$$\begin{aligned}
\mathit{totalCost} &= \mathit{dcCost} + \mathit{wfCost} + \mathit{transCost} && (3)\\
\mathit{dcCost} &= \mathit{dcCAPEX} + \mathit{dcOPEX} && (4)\\
\mathit{dcCAPEX} &= \mathit{dcLandCost} + \mathit{dcBuildCost} + \mathit{dcITCost} && (5)\\
\mathit{dcOPEX} &= \mathit{dcNetCost} + \mathit{dcEnergyCost} && (6)\\
\mathit{dcLandCost} &= \mathit{pLand}(d) \cdot \mathit{dcArea} \cdot \mathit{dcCapacity} && (7)\\
\mathit{dcBuildCost} &= \mathit{dcTotalPow} \cdot \mathit{pBuildDC}(\mathit{dcTotalPow}) + {}\\
&\quad\ \mathit{cLinePow}(d) + \mathit{cLineNet}(d) && (8)\\
\mathit{dcTotalPow} &= \mathit{dcCapacity} \cdot \mathit{maxPUE}(d) && (9)\\
\mathit{dcITCost} &= \mathit{nServers} \cdot \mathit{pServer} + \mathit{nSwitches} \cdot \mathit{pSwitch} && (10)\\
\mathit{nServers} &= \mathit{dcCapacity} / (\mathit{serverPow} + \mathit{switchPow}/\mathit{servsSwitch}) && (11)\\
\mathit{nSwitches} &= \mathit{nServers} \;/\; \mathit{servsSwitch} && (12)\\
\mathit{dcNetCost} &= \mathit{nServers} \cdot \mathit{pNBWServ} && (13)\\
\mathit{dcEnergyCost} &= \sum_{t \in T} |t| \cdot \mathit{powNeed}(t) \cdot \mathit{PUE}(d,t) \cdot \mathit{pEnergy}(d) && (14)\\
\mathit{wfCost} &= \mathit{wfCAPEX} - \mathit{wfRev} && (15)\\
\mathit{wfCAPEX} &= \mathit{wfLandCost} + \mathit{wfBuildCost} && (16)\\
\mathit{wfLandCost} &= \mathit{pLand}(w) \cdot \mathit{wfArea} \cdot \mathit{wfCapacity} && (17)\\
\mathit{wfBuildCost} &= \mathit{pBuildWF} \cdot \mathit{wfCapacity} + \mathit{cLinePow}(w) && (18)\\
\mathit{wfRev} &= \mathit{revEnergy}(w) \cdot \sum_{t \in T} |t| \cdot \beta(w,t) \cdot \mathit{wfCapacity} && (19)\\
\mathit{transCost} &= \mathit{pTransLoss} \cdot \sum_{t \in T} |t| \cdot \mathit{transLoss}(t) && (20)
\end{aligned}$$

Fig. 3: Optimization problem. The datacenter is placed at location d and the wind farm is placed at location w. The objective is to minimize $totalCost$ for a given time period T (divided into epochs denoted by t) and a set of possible locations for d and w. $|t|$ denotes the length of epoch t.

1) Datacenter: The cost of the datacenter can be broken down into capital ($dcCAPEX$) and operational ($dcOPEX$) components. The capital costs are those investments made upfront and depreciated over the lifetime of the datacenter. These costs include the cost for buying land ($dcLandCost$), building the datacenter ($dcBuildCost$), and buying IT equipment ($dcITCost$). The cost of building the datacenter include datacenter construction cost as well as the costs of laying power and network lines to the datacenter. IT equipment includes servers and switches. Land price varies according to location ($pLand(d)$ for location d), whereas the other prices do not to a first approximation. Of course, the total cost of laying the power ($cLinePow(d)$) and network ($cLineNet(d)$) lines depends on location, as the distances to the closest power plant and network backbone are location-dependent. The datacenter construction cost is typically estimated as a function of the maximum power to be consumed by the datacenter. This maximum power is that required by the maximum number of servers and network switches when running at 100% utilization times the maximum expected PUE of the datacenter, where PUE represents the expected overheads of power losses within the datacenter and cooling. The PUE is higher when temperature and/or humidity are high, since cooling consumes more energy under those conditions. The maximum PUE varies with location.

The operational costs are those incurred during the operation of the datacenter over the time period T, and include costs for external network bandwidth use ($dcNetCost$) and the grid electricity ($dcEnergyCost$) required to run the datacenter. The electricity cost is computed based on the IT equipment's power demand over time ($powNeed(t)$), the PUE, and the electricity price. Both the electricity price and the PUE vary with location.

2) Wind farm: The cost of the wind farm is modeled as the capital cost ($wfCAPEX$) minus the revenue earned by selling the wind energy to the grid ($wfRev$). We assume that the operational cost of a wind farm is low, and so do not consider it here; of course, this cost can be easily added to the framework. The capital costs include the cost for buying land ($wfLandCost$) and building the wind farm ($wfBuildCost$), which in turn includes the construction cost and the cost of laying the power line from the wind farm to the closest power plant. The construction cost is assumed to be a linear function of the desired power generation capacity. Note that if the datacenter and wind farm are co-located, then the cost for laying power lines is incurred only once.

The revenue earned by the wind farm is computed over the time period T, where the energy generated within any time epoch t in T at location w depends on the efficiency of the wind turbines and the wind speed. The efficiency of today's wind turbine is close to 50%. We capture the efficiency and impact of wind speed in epoch t using the parameter $\beta(w,t)$, which gives the fraction of the wind farm's maximum capacity actually produced during t.

3) Transmission system: As discussed above, adding a datacenter and wind farm to an existing transmission system will alter the power flow of the network, thus affecting the loads on transmission lines, the voltage of buses, and system losses. We model these losses across the entire time period T, and assume that each unit of loss has a corresponding cost. For each potential placement of the datacenter and wind farm, we compute the loads on the transmission lines, the voltages of buses, and the system loss for each time epoch t using the approach described in Section II.

4) Constraints: The major constraints that must be observed when optimizing is that there must not be any voltage and transmission line capacity violations.

B. Solving the Optimization Problem

The optimization problem formulated in Figure 3 is non-linear, since the calculation of $transLoss(t)$ requires solving the non-linear power flow equations introduced in Section II-C. Currently, we perform a complete search over all possible placements of the datacenter and wind farm to find the best solution. This approach works well for the scale that we are studying (e.g., the New England system, which includes most of the US North East and parts of Canada). Studying larger systems may require a more scalable approach, but we leave this issue for future work.

C. Case Study

We now explore the use of our optimization framework in a case study. Specifically, we use the framework to study the placement of a 100MW datacenter and wind farm in the New England transmission network. Instead of specifying a desired capacity for the wind farm, we assume that we want sufficient wind energy to completely offset the energy consumption of the datacenter (making the size of the wind farm dependent on the locations of both the datacenter and the wind farm).

1) Instantiating the framework parameters: The production of wind power and datacenter cooling both depend on weather conditions. Thus, we obtained Typical Meteorological Year (TMY) information from the US Department of Energy[3] for 56 locations in the New England area as shown in Figure 4. A TMY is a 1-year dataset of hourly weather values selected to include a representative range of weather phenomena for a location, while still giving annual averages that are consistent with the long-term averages for the location. We use the TMY wind speeds and air pressures, conversion losses, and a model of the GE 1.5MW wind turbine [15] to compute $\beta(l,t)$ at each location l during time epoch t.

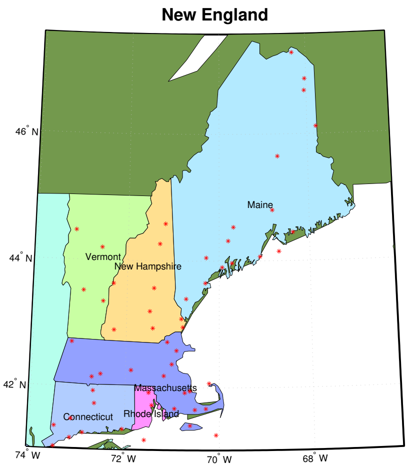

Fig. 4: Candidate locations in New England

We adopt the values and approaches for computing PUE, datacenter construction costs, wind farm construction costs, land costs, transmission lines and network connection costs, and grid energy costs from [10]. Table VI shows part of the location-dependent parameters for five sample locations used in our experiments, and Table VII shows the values of the location-independent parameters. Datacenter construction costs are financed and amortized across 12 years. Wind farm construction costs are financed over 12 years and amortized over 24 years. We assume that land cost is fully recoverable, so that the only incurred cost is that of financing, spread over 12 years. Finally, the cost for IT equipment is financed and amortized over 4 years. For financing, we use an annual interest rate of 3.25%.

TABLE VI: Location-dependent parameter values for five sample locations.

	Burlinton, NH	Springfield Hartnes, VT	Moose Land, ME	Nash Island, CO	Mount Washington, NH
$pLand$ (\$/m^2)	946.9	946.9	1034.1	946.9	946.9
$maxPUE$	1.645	1.645	1.681	1.619	1.676
$cLinePow$ (M\$)	64.5	77.2	19.0	16.6	107.1
$cLineNet$ (M\$)	13.5	16.7	42.4	16.9	21.3
$pEnergy$ (\$/kWh)	0.0941	0.0941	0.1281	0.1281	0.1257
Average β (%)	11.7	2.7	53.4	40.0	56.8

We use the nominal grid load, which totals 6,254MW across all the buses in the New England system. For each time epoch t, we compute the wind energy being generated by all the wind farms (existing ones and the new one being placed) using $\beta(l,t)$. We assume full loading of all datacenters (existing ones and the new one being placed). We then compute

---

[3] http://apps1.eere.energy.gov/buildings/energyplus/weatherdataabout.cfm

TABLE VII: Values of location-independent parameters.

Parameter	Value	Unit
$dcArea$	0.557	m^2/kW
pBuildDC	12000	\$/kW
$serverPow$	0.275	kW/serv
$switchPow$	0.48	kW/switch
$servsSwitch$	32	servs/switch
$pServer$	2000	\$/serv
$pSwitch$	20000	\$/switch
pNBWServ	1	\$/serv-month
$wfArea$	18.21	m^2/kW
pBuildWF	2100	\$/kW

the system loss for the time epoch for the placement of the new datacenter and wind farm at locations d and w, respectively, using the simulation approach described in Section II. We map the candidate locations in Figure 4 to buses using the approach previously discussed in Section II. For each possible pair of (d, w), we sum the transmission loss over the entire year. Finally, we set $pTransLoss$ to the maximum electricity price in the whole area.

2) Placement approaches: As already discussed, the primary novelty of our optimization framework is that it considers the cost of transmission system losses. In addition, it also simultaneously places a new datacenter and an offsetting wind farm. To assess the impact of these characteristics, we compare results for five different placement strategies as follows.

DC_WF_OPT: This strategy individually looks for the best locations to put the datacenter and the wind farm; i.e., it solves the optimization problem for the datacenter without considering the new wind farm, and then solves the optimization problem again for the wind farm without considering the new datacenter. This strategy also ignores transmission losses. Note that assuming a constant transmission loss would give the same results.

DC+G_WF+G: This strategy is the same as DC_WF_OPT except that grid transmission losses are considered when solving the optimization problem.

Min_Loss: This strategy finds locations for the datacenter and wind farm that minimize the cost of transmission system losses.

Co-location: This strategy assumes that the datacenter and wind farm should be co-located, and so finds a single location that minimize the overall cost, including the cost of the transmission system losses.

Jointly: This strategy considers the simultaneous placement of the datacenter and wind farm, and uses all of the costs and revenues in the optimization framework to find locations (the datacenter and wind farm may be co-located or placed at different locations) that minimize overall cost.

3) Results: Our results show that each placement strategy was able to find at least one placement of the new datacenter and wind farm that satisfies the constraints for avoiding the overloading of transmission lines and unacceptable voltage variations throughout the simulated year. This is an important finding since overloading of transmission lines and unacceptable voltage variations can have serious consequences as already discussed.

Figure 5 shows the resulting total cost when using the five different placement strategies. We observe that not considering transmission system losses (DC_WF_OPT), considering only transmission system losses (Min_Loss), and forcing the co-location of the datacenter and wind farm (Co-location) all lead to higher total cost. Specifically, DC_WF_OPT incurs the highest transmission cost, Min_Loss incurs high costs for building the wind farm, and Co-location incurs high energy cost. In comparison, considering both transmission system losses and datacenter construction/operation costs can lead to finding locations that best balance these factors, leading to lowest overall cost.

The locations found and the total cost achieved by the five strategies are listed in Table VIII. These results show that using the full optimization framework (Jointly) achieves savings of 7.6% compared to simply co-locating the datacenter and wind farm (Co-location) and 3.3% compared to not accounting for transmission system losses (DC_WF_OPT). Interestingly, the simultaneous placement of datacenter and wind farm (Jointly) gives the same results as the individual placement of datacenter and wind farm (DC+G_WF+G). Thus, our results do not provide evidence to support the importance of simultaneous placement, and this issue should be studied further. The optimization problem can be solved more efficiently if we can view the datacenter and wind farm placements as independent problems without increasing the total cost.

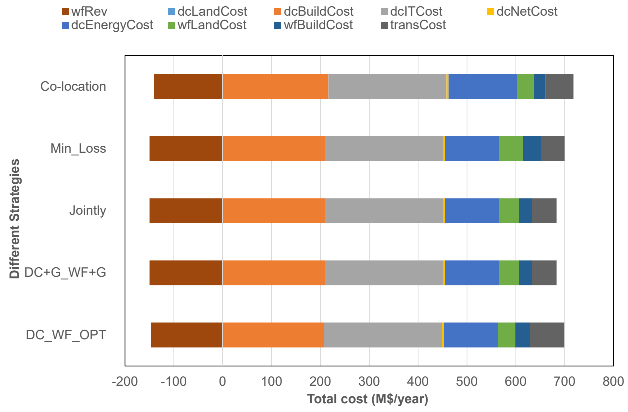

Fig. 5: Cost of building/operating a 100MW datacenter and an offsetting wind farm.

TABLE VIII: Locations of datacenter and wind farm, and the resulting total costs.

Strategy	Datacenter location	Wind farm location	Total cost (M\$/year)
DC_WF_OPT	Burlington,NH	Mount Washington, NH	552.2
DC+G_WF+G	Springfield Hartnes, VT	Moose Land, ME	533.7
Jointly	Springfield Hartnes, VT	Moose Land, ME	533.7
Min_Loss	Springfield Hartnes, VT	Nash Island, CO	550.0
Co-location	Mount Washington, NH	Mount Washington, NH	577.3

IV. RELATED WORK

As already mentioned, a number of previous efforts have studied the placement of new datacenters. Alger [7] explained how to choose an optimal location for a datacenter by considering hazards, accessibility, and scalability factors. Stansberr [16] ranked some cities by estimating the annual operation costs of a datacenter. Oley [8] considered looking for a proper location for a datacenter by investigating the power rates of different states. Goiri *et al.* [6] focused on intelligently finding the best places for building multiple datacenters to form a network for interactive Internet services. Berral *et al.* [10] considered selecting sites for datacenters and on-site power plants that support "follow-the-renewables" cloud services. Gao *et al.* [17] studied how to site datacenters near existing wind farms, and distributing load using a greedy online algorithm. None of these works have considered the impact of placing new datacenters on the transmission grid.

A previous work that has considered the interaction of datacenters and the grid is [18]. In this work, Liu *et al.* show that adding renewable power plants (solar in their study) can lead to voltage violations within a grid distribution system. They also show that datacenters can help avoid such voltage violations by dynamically adjusting their power demand based on signals from the grid. A number of research efforts have also studied how datacenters can participate in demand response programs and ancillary services to help ease the management of the grid [19], [20]. A key difference between these works and ours is the assumption of datacenters being able to dynamically adjust their power demands in response to grid signals. In addition, they did not actually consider the placement of datacenters, nor did they consider transmission system losses.

Mohsenian *et al.* in [21] proposed a request distribution policy among datacenters to ensure power load balancing. They tried to minimize the maximum power on any transmission line by distributing the computing requests to suitable datacenters. Their work assumes that a fairly large number of datacenters (e.g., 6) are connected to the same power distribution network. Further, they did not consider the impact of transmission system losses on the placement of datacenters.

V. CONCLUSIONS

In this paper, we proposed that new renewable powered datacenters should be placed intelligently while considering their impact on the electricity transmission system (along with other datacenter/wind farm capital and operation costs). Specifically, we studied the potential impact of new datacenters on the New England ISO transmission system. Using simulation, we showed that different placements of a new datacenter can change the overloading of transmission lines, the number of unacceptable grid voltage deviations, and transmission system losses. Even when transmission constraints are met, strategic placement can lead to lower overall costs for grid operators and datacenters owners. We thus developed an optimization framework for the placement of a new datacenter and offsetting wind farm that considers transmission constraints and system losses, and used it in a case study. Our results show that considering the cost of transmission system losses can indeed lead to different placements that achieve lower total cost.

REFERENCES

[1] ResearchAndMarkets, "Global Data Center Construction Market 2014-2018," 2014, http://www.researchandmarkets.com/reports/2757583/.

[2] Google, "Google Data Centers – Renewable Energy," 2014, http://www.google.com/about/datacenters/renewable/.

[3] Apple, "Apple and the Environment," 2013, http://www.apple.com/environment/renewable-energy/.

[4] R. Miller, "Data Centers Scale Up Their Solar Power," 2012, http://www.datacenterknowledge.com/archives/2012/05/14/data-centers-scale-up-their-solar-power/.

[5] PJM, "A Survey of Transmission Cost Allocation Issues, Methods and Practices," *Valley Forge*, 2010.

[6] I. Goiri, K. Le, J. Guitart, J. Torres, and R. Bianchini, "Intelligent Placement of Datacenters for Internet Services," in *Proceedings of the International Conference on Distributed Computing Systems (ICDCS)*, 2011.

[7] D. Alger, "Choosing an Optimal Location for Your Data Center," *InformIT*, Sep. 2006.

[8] B. Oley, "Power & Approaching the Best Place to Build a Data Center," *The Data Center Journal*, 2009.

[9] F. Larumbe and B. Sansò, "Optimal Location of Data Centers and Software Components in Cloud Computing Network Design," in *Proceedings of the IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing (CCGRid)*, 2012.

[10] J. L. Berral, I. Goiri, T. D. Nguyen, R. Gavalda, J. Torres, and R. Bianchini, "Building Green Cloud Services at Low Cost," in *Proceedings of the International Conference on Distributed Computing Systems (ICDCS)*, 2014.

[11] C. Ren, D. Wang, B. Urgaonkar, and A. Sivasubramaniam, "Carbon-Aware Energy Capacity Planning for Datacenters," in *Proceedings of the International Symposium on Modeling, Analysis & Simulation of Computer and Telecommunication Systems (MASCOTS)*, 2012.

[12] G. W. Bills, "On-line Stability Analysis Study," Edison Power Research Institute, Tech. Rep. RP90-1, 1970.

[13] Data Center Research, "Data Center Map," 2014, http://www.datacentermap.com/datacenters.html.

[14] J. Koomey, *Growth In Data Center Electricity Use 2005 To 2010*. Analytics Press, 2011.

[15] Y. Lei, A. Mullane, G. Lightbody, and R. Yacamini, "Modeling of the Wind Turbine with a Doubly Fed Induction Generator for Grid Integration Studies," *IEEE Transactions on Energy Conversion*, vol. 21, no. 1, 2006.

[16] M. Stansberr, "Data Center Locations Ranked by Operating Cost," *SearchDatacenter.com*, 2006.

[17] V. Gao, Z. Zeng, X. Liu, and P. R. Kumar, "The Answer is Blowing in the Wind: Analysis of Powering Internet Data Centers with Wind Energy," in *Proceedings of the INFOCOM Conference*, 2013.

[18] Z. Liu, I. Liu, S. Low, and A. Wierman, "Pricing Data Center Demand Response," in *Proceedings of the ACM SIGMETRICS Conference*, 2014.

[19] D. Aikema, R. Simmonds, and H. Zareipour, "Data Centres in the Ancillary Services Market," in *Proceedings of the International Green Computing Conference (IGCC)*, 2012.

[20] A. Wierman, Z. Liu, I. Liu, and H. Mohsenian-Rad, "Opportunities and Challenges for Data Center Demand Response," in *Proceedings of the International Green Computing Conference (IGCC)*, 2014.

[21] A.-H. Mohsenian-Rad and A. Leon-Garcia, "Coordination of Cloud Computing and Smart Power Grids," in *Proceedings of the IEEE International Conference on Smart Grid Communications (SmartGridComm)*, 2010.