A First Look at Node-Level Curtailment of Renewable Energy and its Implications

Diptyaroop Maji
University of Massachusetts, Amherst
Amherst, USA
dmaji@cs.umass.edu

David Irwin
University of Massachusetts, Amherst
Amherst, USA
deirwin@umass.edu

Prashant Shenoy
University of Massachusetts Amherst
Amherst, USA
shenoy@cs.umass.edu

Ramesh K. Sitaraman
UMass Amherst & Akamai Tech
Amherst, USA
ramesh@cs.umass.edu

Abstract

Electricity grids are trying to meet their demand by using more renewable energy as they move towards decarbonization. As the amount of renewables in the grid increases, there are periods when renewable supply exceeds the demand or when excess supply cannot be transmitted to a different location to satisfy the demand due to grid congestion. Consequently, renewable generators often need to be *curtailed* so that they operate below their maximum capacity. Such curtailment represents unutilized "green" energy that could have replaced energy produced from non-renewable "brown" sources. While prior works have studied curtailment at the grid level, curtailment is a local phenomenon that occurs at the level of a *generation node*, where each node is a power plant at a specific location feeding the grid. A grid may consist of hundreds of nodes, but curtailment may only occur in some nodes and at some times. Hence, understanding curtailment at the node level is important to evaluate its potential for decarbonization.

We study curtailment at the node level for the Texas grid operated by ERCOT, which consists of hundreds of nodes producing wind and solar energy. Using extensive node-level data for the year 2023, we show that curtailment is highly non-uniform and intermittent — 20% of the nodes account for 77% of the total curtailed energy, while 70% of the nodes are curtailed for less than 10% of the year. We find that although wind curtailment is more prevalent, a greater fraction of solar generation is curtailed than wind. We also develop a method to identify the cause of curtailment from the Locational Marginal Price (LMP), showing that 74.3% of the time, curtailment in Texas is due to grid congestion. Overall, our analysis of node-level curtailment implies that while curtailment can potentially be forecasted in only a small fraction of the nodes, a considerable amount of curtailed energy can be utilized by adding demand adjacent to these nodes.

CCS Concepts

• **Hardware** → **Renewable energy**; • **Social and professional topics** → *Sustainability*; • **General and reference** → *Estimation.*

E-ENERGY '25, Rotterdam, Netherlands

ACM ISBN 979-8-4007-1125-1/25/06
https://doi.org/10.1145/3679240.3734627

Keywords

renewable energy curtailment, grid congestion, locational marginal price, power grid

ACM Reference Format:
Diptyaroop Maji, David Irwin, Prashant Shenoy, and Ramesh K. Sitaraman. 2025. A First Look at Node-Level Curtailment of Renewable Energy and its Implications. In *The 16th ACM International Conference on Future and Sustainable Energy Systems (E-ENERGY '25), June 17–20, 2025, Rotterdam, Netherlands.* ACM, New York, NY, USA, 12 pages. https://doi.org/10.1145/3679240.3734627

1 Introduction

Electricity demand has been rapidly increasing over the last decade for several reasons ranging from the electrification of the residential, industrial, and transportation sectors [28, 29] to the growth of AI and data centers [22, 42]. At the same time, there is also a push towards decarbonizing the grid to reduce the carbon emissions associated with electricity generation, leading to a proliferation of renewable sources in many regions. The electricity grid needs a continuous supply to match the demand, and the supply in such regions is increasingly coming from intermittent renewable sources. However, renewable supply and demand are not always matched in time and space. Renewable energy, like solar and wind energy, is intermittent in nature and is available only when the sun shines or the wind blows. Moreover, solar and wind plants may be situated far from the demand location, and the grid may not have enough infrastructure to transmit such energy to the place of demand. To address this mismatch in supply and demand, grid operators often turn off renewable plants and ramp up non-renewable plants when required. Consequently, some renewable energy gets curtailed to maintain stable grid operations, where curtailment is defined as the difference between potential and actual renewable generation. As many grids are adding more renewables to decarbonize the electricity supply, renewable energy curtailment is increasing year after year in such grids [5, 7, 10].

Since a lot of clean energy is being wasted via curtailment, researchers have recently proposed optimization techniques to utilize curtailed renewable energy by modulating the electricity demand, thus reducing the carbon footprint of electricity consumption. Some works examine shifting flexible loads like EV charging or data center computing through time and space [8, 49, 51] to consume energy that would otherwise be curtailed. Such curtailment-aware optimizations need curtailment estimates and predictions at the node level rather than at the grid level, where a node is defined as one or more power plants at a specific location feeding the grid. This is because curtailment is inherently a local phenomenon. Due to transmission constraints, curtailment is often restricted to small regions in the grid. Hence, if such optimizations do not know when and where curtailment occurs at a node level, they may result in increased carbon emissions instead of reducing them.

For example, the Texas grid covers a large geographical area and is often congested. So, if there is curtailment in Southern Texas and a curtailment-aware optimization shifts demand to Northern Texas, the renewable supply from the South may not be able to reach the demand location, and some non-renewable nodes in the North may need to meet that extra demand, leading to more carbon emissions. Consequently, understanding and analyzing curtailment at a node level is essential for forecasting curtailment and leveraging curtailed energy to reduce carbon emissions.

Node-level curtailment estimates and predictions, although crucial, have not received much attention. Prior works have studied curtailment only at the grid level [3, 20, 23, 26, 39, 41]. In this paper, we take the first step to address this gap by providing an in-depth analysis of solar and wind curtailment at a nodal granularity.

Our Contributions. Our work uses a data-driven approach to analyze node-level curtailment in the Texas grid (ERCOT) in the United States. We analyze where and when curtailment happens at the different nodes across Texas. We also analyze how much renewable generation is curtailed and how long those curtailments typically last. Further, we develop a method to identify the cause of the occurrence of curtailment events from the nodal pricing data. Our specific contributions follow.

(1) We perform a detailed analysis of node-level solar and wind curtailment for the year 2023 using publicly available data from the Texas grid (ERCOT).
(2) We propose a method to identify the cause of curtailment events (oversupply versus grid congestion) from the Locational Marginal Price (LMP).
(3) We perform a price-based analysis to examine if curtailment can be detected and estimated from the LMP.
(4) Our datasets are derived by combining and curating multiple data sources. We release our datasets[1] to the research community to support further research in this area.

Our Key Observations. We make several observations while analyzing the node-level curtailment and pricing data. The key observations from our analyses are below.

(1) Curtailment in Texas is highly skewed geospatially. The majority of the curtailment (55%) occurs in the western part of the state that has the lowest population density. Overall, 20% of the nodes account for 77% of the total curtailment.
(2) Solar generation is curtailed much more than wind generation. During curtailment events, almost all solar generation is curtailed 25% of the time. In comparison, almost all wind generation is curtailed only 10% of the time.
(3) Curtailment is highly infrequent and intermittent at the node level — 70% of the nodes are curtailed for less than 10% of the year, and the median curtailment duration is 15 minutes.
(4) Curtailment is mostly weakly correlated even across geographically proximal nodes, and the correlation becomes even weaker as the distance between the nodes increases. Solar curtailment is more correlated than wind curtailment.
(5) Curtailment is caused due to grid congestion 74.3% of the time, making it the leading cause of curtailment in Texas. Curtailment is due to an oversupply of renewable energy the rest of the time.
(6) LMP can potentially be used to detect a curtailment event at the node level. However, deriving the amount of curtailed energy (in MWh) from LMP is not straightforward.

Implications of Our Work. Our future goal is to find ways in which the curtailed energy can be put to use to serve demand. To achieve this goal, one would need to *forecast* curtailment events and then *modulate the demand* during those events to consume the energy that would otherwise have been curtailed. The key implications of our analysis for these two tasks are below.

(1) Forecasting Curtailment. Most nodes are curtailed infrequently. Our analysis shows that 70% of the nodes are curtailed for less than 10% of the year. Forecasting curtailment in such nodes may be difficult since a node may have long intervals of time between two consecutive curtailment events. The remaining 30% of the nodes are curtailed more frequently, and future work could focus on forecasting curtailment in these nodes using existing ML techniques for carbon intensity forecasting [34]. Our analysis also shows that these frequently curtailed nodes account for 65% of the curtailed energy; hence, the events that account for a majority of the curtailed energy can potentially be forecasted.

Our work also sheds light on what features are likely useful in forecasting curtailment. We show that the nodal LMP values are a strong indicator for detecting curtailment events, making LMP a good feature for forecasting curtailment events. However, LMP is not a good indicator of the actual amount of curtailed energy.

Further, we found that curtailment events across nodes are only weakly correlated, even when the nodes are geographically proximal. This suggests that forecast models would need to be trained individually for each node based on its own historical data. Models trained for one node may not be effective in forecasting for another.

(2) Modulating Demand. Demand can be modulated at a location by moving or creating new workloads that consume energy during curtailment events. However, since the leading cause of curtailment is congestion, one may need to add demand adjacent to the nodes to bypass transmission constraints. For example, a cloud provider could deploy server clusters colocated with nodes and shift workloads in those servers during periods of curtailment. A considerable portion of the curtailed energy (77%) can be utilized by modulating demand only at a few (20%) locations (i.e., nodes).

2 Background

In this section, we provide background on the electricity grid, how the electricity market operates in many regions, and renewable energy curtailment.

[1] https://github.com/codecexp/nodal-curtailment-analysis

Electricity Grid. The electricity grid has three components — generation, transmission, and distribution [12]. In any region, electricity is usually generated from a mix of renewable and non-renewable sources. Once generated by the different power plants, electricity is transmitted via a network of transmission lines and finally distributed to end consumers. The grid consists of *generator nodes* (or *nodes*, for short), where each node is a collection of one or more power plants at a specific location that feed the grid. A node can generate electricity from a renewable source, such as wind or solar, or a non-renewable source, such as coal or natural gas. The nodes may offer electricity at different prices at different times of the day.

The electricity grid is managed by a grid operator who ensures that the electricity supply always matches the demand. The grid operator periodically solves the *Security Constrained Economic Dispatch (SCED)* problem [27] at the node level. SCED is an optimization to dispatch the existing set of electricity generators in a way such that the electricity demand is met with the lowest generation cost while adhering to grid and generator constraints like transmission constraints, ramp rates of different generators, generation capacity, intermittency of the renewables, etc. SCED outputs how much each generator should generate so that the total supply matches the total demand, along with the price to generate the next unit of electricity called the Locational Marginal Price (LMP).

Renewable Energy Curtailment. Renewable energy curtailment events are a by-product of SCED and occur when solar or wind nodes have the capability to generate electricity, but the grid operator instructs them to operate below capacity. A *curtailment event* has two aspects — *curtailment amount* (in MWh) and *curtailment duration* (in hours). Curtailments can be due to the following two reasons:

(1) Congestion. Sometimes, renewable energy supply cannot reach the demand location due to grid transmission constraints. In such cases, a nearby non-renewable node may feed the demand while the renewable node that is further away is curtailed so that total supply matches total demand. We refer to this as curtailment of renewable energy due to congestion.

(2) Oversupply. Even when there is no congestion, renewable energy may be curtailed when the total supply exceeds the demand. Such curtailment ensures that the total supply matches the demand and grid stability is maintained. We refer to this situation as curtailment due to oversupply. When there is an oversupply, renewable energy is curtailed for several reasons. First, renewable generation might be a large fraction of the total supply due to local factors, requiring that the renewable generation be ramped down. Second, the cost of ramping down non-renewable generators such as coal might be more expensive than turning off renewable generation. This can also lead to curtailing renewable energy, even when it could have replaced the non-renewable energy to meet the demand.

Note that during oversupply, nodes offering electricity at higher prices may get curtailed before nodes with lower prices.

Curtailment amount is estimated by computing the difference between the High Sustained Limit (HSL) and the Base Point (BP) [25, 31] of a node and multiplying that with the curtailment duration, where HSL is the maximum power production capability of a node at a particular time (in MW), and BP is the actual generation (in MW) by the node determined by SCED. Thus,

$$\text{Curtailment Amount } (in\ MWh) = (HSL - BP) \times \text{Duration} \quad (1)$$

Locational Marginal Price (LMP). LMP (in $/MWh) at a given node and time is the price of generating the next unit of electricity at that node and at that time [30]. LMPs are obtained after solving SCED, which considers the current demand, generation, and grid conditions. Depending on the above variables, LMPs across nodes and across times can vary significantly.

In recent years, LMPs in grids with an increasing renewable penetration are sometimes zero or negative, due to a combination of surplus renewable supply and various subsidies [11] offered for renewable generation. Some generators cannot turn off or on quickly, and hence may be willing to offer electricity at negative prices if the systemwide electricity supply exceeds the demand for short periods of time [9]. In those situations, solar and wind generators may be curtailed to decrease the systemwide supply. On the other hand, renewable generators like wind get *Production Tax Credits* per unit of electricity generated [17], and may sell electricity at negative prices to avoid curtailment as long as the net cost of generation remains positive.

3 Data Sources

Our goal is to analyze curtailment at the node level. However, while many grids publish nodal pricing data, it is still not common to publish electricity generation or curtailment data at the node level. In the US, ERCOT is one of the few ISOs that make nodal electricity data publicly available. So, we specifically focus on the Texas grid in this paper. This section lists the data sources we used. We also describe the data preprocessing and cleanup methodology that we used to curate the datasets from the raw data.

Nodal Electricity Generation. ERCOT provides node-level SCED reports with a 60-day lag [15]. The reports contain various generation-related data at a 15-minute granularity, including data regarding generation capability (HSL), actual generation (BP), type of generator, etc. We downloaded SCED reports for 2023 and built the nodal HSL and curtailment datasets for our analysis. Our datasets comprise 281 nodes across Texas — 112 solar and 169 wind nodes. We replaced any missing data with the data from the closest available time interval. We also used hourly ERCOT electricity data from the US Energy Information Administration (EIA) [13] to validate and clean our dataset. For example, we set any non-zero solar HSL in our SCED dataset to zero whenever there is no solar generation in the EIA dataset. We estimated curtailment amount using Eq. 1 and set any curtailment values less than 0.5 to 0 to minimize reporting errors. The HSL and curtailment data for any solar or wind node are time series vectors of length 35040.

Nodal LMP. ERCOT also provides nodal LMP data in 5-minute granularity in almost real-time [16]. We downloaded the 2023 LMP data for our analysis, filtered the LMPs for only solar and wind nodes, and used the values reported at each 15-minute interval since SCED data is every 15 minutes. We replaced any missing data with the closest available LMP values. Like the curtailment dataset, the LMP data set also has 35040 rows for each node.

Node Capacity and Location. We estimated the capacity of a node (in MW) by taking the maximum HSL for that node in 2023. Since HSL is the generation capability for a node, we posit that our

Hub	Solar nodes (%)	Solar curtailment amount (%)	Wind nodes (%)	Wind curtailment amount (%)
West	37.5	77.0	53.2	46.6
Pan-handle	1.8	5.1	11.8	18.2
South	21.4	6.9	23.7	32.3
North	30.3	9.0	10.7	2.7
Houston	8.9	1.8	0.6	0.02

Table 1: *Distribution of solar/wind nodes and yearly solar/wind curtailment amount across hubs. West has proportionally more curtailment than other hubs.*

estimate accurately reflects a node's capacity. Solar (resp. wind) node capacities go up to 457.1 MW (resp. 669.5 MW) in our dataset.

We obtained the locations (latitude and longitude) for 248 out of 281 power nodes from the US Solar Photovoltaic Database (USPVDB) [43], the US Wind Turbine Database(USWTDB) [44], the Energy Information Administration (EIA) power plant database [46], and ERCOTs Monthly Outlook for Resource Adequacy (MORA) report [18]. Due to insufficient information, we could not match node names to locations for the remaining nodes.

Population Density. We fetched population density data at the county level from the US Census Bureau website [45].

While the raw data is available publicly, collecting, compiling and aggregating the data is not straightforward and involves numerous manual steps, which adds to the challenge of finding good-quality node-level curtailment data. Hence, we release all the datasets used in our analysis to the research community to accelerate curtailment-related research. All the curated datasets are publicly available at **https://github.com/codecexp/nodal-curtailment-analysis**.

4 Nodal Curtailment Analysis

We begin our analysis using the curtailment data for 2023 and analyze where, when, how much, and how long curtailments occur in Texas. We enquire if nodal curtailments have any seasonal patterns and if they are correlated across the different renewable nodes.

Analyzing how curtailment is distributed geospatially would help understand the cause of curtailment — if there is more curtailment away from population centers, then the curtailment may be primarily due to grid transmission constraints. Also, since both curtailment and demand are mainly localized, optimizations trying to leverage curtailed energy by modulating demand should know where to add or shift demand at a nodal granularity. Adding demand at places where the supply cannot reach may increase carbon emissions instead of reducing them because a non-renewable source may need to meet that demand.

Such optimizations also need to know the amount of curtailed energy available at a specific location and how long it will be available. Adding more demand than available curtailed energy or scheduling demand when no renewable energy is curtailed may again be detrimental. Using system-wide curtailment estimates and predictions may not be sufficient since such estimates are usually an order of magnitude higher than nodal curtailment, and such predictions may hide the intermittent nature. For example, we find that the average system-wide curtailment in ERCOT in a 15-minute interval is 555.2 MWh — more than the capacity of all but one wind node

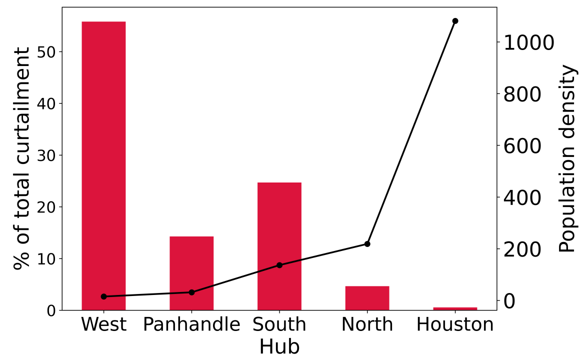

Figure 1: *Comparing total yearly renewable curtailment amount with population density. Curtailment is more at hubs with lower population densities.*

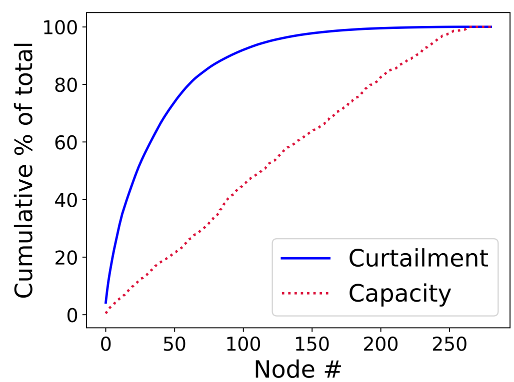

Figure 2: *CDF of nodal curtailment amount. 20% of the nodes account for 77% of the total curtailment amount, but represent only 23.1% of total solar and wind generation capacity.*

in our dataset. Also, some wind energy is always curtailed at the grid level in 2023, but nodal wind curtailment lasts less than 1 hour 90% of the time (§ 4.3). Hence, analyzing the amount and duration of curtailment at the node level is crucial to developing accurate curtailment-aware optimizations and curtailment forecasts that such optimizations can use.

Further, suppose optimizations need to add more demand than the available supply at one location. In that case, they need to know which other nodes are also experiencing curtailment simultaneously so that they can distribute the demand accordingly if feasible. For example, data center workloads like batch processing may need a considerable amount of electricity supply, and hence, can utilize such information to distribute the computation across multiple locations. Consequently, we try to answer the following questions:

Q1. How is curtailment distributed geospatially? How is it correlated to population centers? (§ 4.1)

Q2. How much energy is curtailed in a node on average? What fraction of potential generation is curtailed at any instant? (§ 4.2)

Q3. How long and how often do curtailments typically occur? (§ 4.3)

Q4. How do curtailment events vary across seasons? (§ 4.4)

Q5. When a node is curtailed, which other nodes are also curtailed? Is curtailment strongly correlated across nearby nodes? (§ 4.5)

4.1 Geospatial Distribution of Yearly Curtailment Across Hubs and Nodes

We first look at one level below the grid level: the hub level. A hub is a subregion in the grid comprising a collection of nodes. The Texas grid is divided into five hubs [19] — North, Panhandle, West, South, and Houston. We map all the nodes to their respective hubs using the MORA report [18] and add up their respective curtailments to get the yearly hub-wide curtailment. Once we analyze the geospatial distribution at the hub level, we proceed to the nodal level to see which nodes contributed more towards the total curtailment and where they are located.

Table 1 and Fig. 1 shows the geospatial distribution of the renewable nodes and solar/wind curtailment across the five hubs in 2023. Most solar and wind nodes are located in the West. Curtailment is also similarly skewed across the hubs. 55.8% of the total yearly curtailment in 2023 is in the West. When divided into solar and wind curtailment, this amounts to 77% (resp. 53.2%) of the annual solar (resp. wind) curtailment.

Fig. 1 also plots the relation between total yearly curtailment and the population density in the different hubs. We aggregated the population density of all counties in a hub using [19] to get the population density of a particular hub. In general, we see that curtailment is more prevalent in hubs with lower population density than in hubs with higher population density. That is, there is more curtailment away from population centers. West has the most curtailment but the lowest population density, whereas Houston has the highest population density but the least curtailment.

While high curtailment in the West may be partially due to an oversupply of renewable energy within the hub, it may also be because electricity cannot be transmitted to other hubs due to grid congestion. Similarly, low curtailment in Houston may be partially because most of the renewable supply is utilized by demand. On the other hand, renewable plants need large areas to operate, which may not be available in regions with high population density. Hence, low curtailment in Houston may also be due to low renewable supply in Houston. We need information about the local demand and grid topology to find the definitive reason, and such analysis is kept as future work.

Curtailment is also highly skewed at the node level. Fig. 2 shows the cumulative percentages of curtailment and capacity with respect to the total system-wide curtailment and solar/wind generation capacity. We find that 20% of the nodes are responsible for 77% of the total yearly curtailment in 2023, although they constitute only 23.1% of total solar and wind generation capacity in Texas. Thus, curtailment is not evenly spread out across the nodes. Instead, some nodes are curtailed much more than the rest. Among those 20%, 65% of the nodes are located in the West. The remaining are located in the South and the Panhandle.

Key Takeaways. Curtailment is highly non-uniform and mostly away from population centers. West has the lowest population density and accounts for 55% of the total curtailment in 2023, while Houston, having the highest population density, only accounts for 0.5% of the total curtailment. The top 20% of the nodes with the most yearly curtailment are in the West, South and the Panhandle and account for 77% of the total curtailed energy.

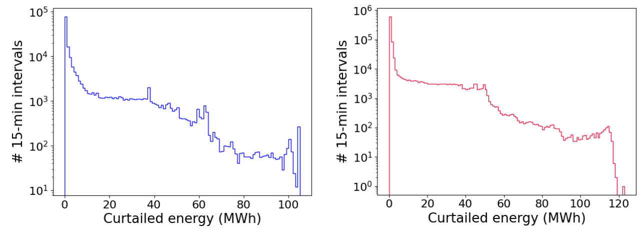

(a) Solar. **(b) Wind.**

Figure 3: ***Distribution of the amount of energy curtailed in 15-minute intervals. Solar (resp. wind) curtailment in a 15-minute interval can be as large as 105 MWh (resp. 122.5 MWh).***

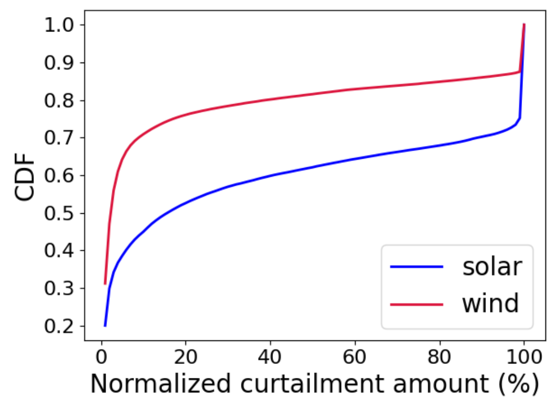

Figure 4: ***CDF of normalized curtailment amount over all 15-minute intervals with curtailment events. When curtailment occurs, a larger fraction of solar is curtailed than wind.***

4.2 Curtailment Amount Distribution

Next, we take our nodal curtailment dataset and consider the curtailment reported in any 15-minute interval for 2023, ignoring all intervals when there is no curtailment. We then analyze the distributions of solar and wind curtailment separately. During those intervals, we also normalize the curtailed amount with respect to the HSL to calculate how much of the potential generation is curtailed in that interval, where

$$Norm.\ Curtailed\ Amt.(\%) = \frac{Curtailment\ Amount}{HSL \times 0.25} \times 100 \qquad (2)$$

Fig. 3 shows the curtailment amount distribution across the nodes. Solar (resp. wind) curtailment in a 15-minute interval ranges from 0.25 MWh to 105 MWh (resp. 122.5 MWh). Most of those events are curtailments of small amounts — during periods of curtailment, solar (resp. wind) curtailment in a 15-minute interval is less than 1 MWh 43% (resp. 68%) of the time. Still, there are many intervals across nodes with a high amount of curtailed energy.

Fig. 4 shows the CDF of the normalized curtailment amount (in %) in a 15-minute interval during curtailment events. The median normalized curtailed amount during such events is 16% and 2%, respectively. We see that the normalized amount is more than 99% during 25% (resp. 10%) of the solar (resp. wind) curtailment events. Hence, although the total wind curtailment exceeds the total solar curtailment, typically, a larger fraction of solar generation

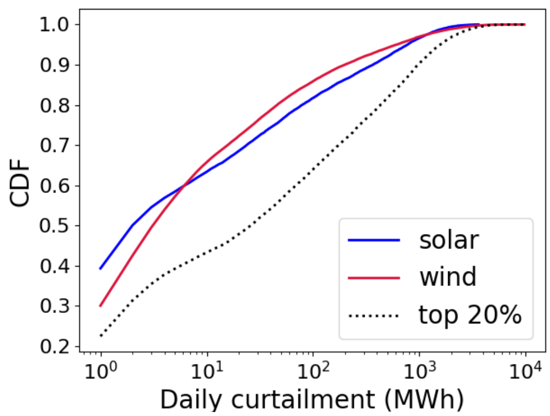

Figure 5: ***CDF of the daily solar and wind curtailment amount (only days with curtailment events are considered).***

is curtailed than wind. This may be because wind usually offers electricity at a lower price than solar [25, 47]. Hence, SCED would curtail solar before wind without any other constraints. However, more analysis is needed to find exactly why solar generation is curtailed more than wind.

Many workloads require more energy than the curtailed energy available in a 15-minute interval and run for longer periods but are delay-tolerant. Such workloads can often be executed successfully if there is enough curtailed energy available over a day. Hence, we also look at daily solar and wind curtailment distributions. Recall from § 4.1 that 20% of the nodes account for most of the curtailment. So, we also show the distribution of daily curtailment for the top 20% of nodes with the most amount of curtailment. Note that we only consider the days with at least one curtailment event since there are often days without curtailments. Fig. 5 shows the solar and wind curtailment distributions. Daily solar (resp. wind) curtailment ranges from 0.12 MWh to 3.6 GWh (resp. 9.7 GWh), with a median value of 0.47 MWh (resp 3.1 MWh). For the top 20% of the nodes, daily curtailment ranges from 0.12 MWh to 9.7 GWh, with a median value of 25.5 MWh. Thus, while the amount of daily curtailment across nodes is highly non-uniform, the top 20% of the nodes experience considerable curtailment over a day and can potentially execute delay-tolerant workloads with high energy requirements.

Key Takeaways. Curtailment amounts across nodes have a wide range. Although nodal curtailment amounts in a 15-minute interval are mostly small (< 1 MWh), they can be as much as 122.5 MWh. During curtailment events, a larger fraction of solar is curtailed than wind — more than 99% of the generation is curtailed during 25% (resp. 10%) of the solar (resp. wind) curtailment events.

4.3 Curtailment Duration and Frequency Distribution

We now analyze how long nodal curtailment lasts on average. To do so, we count the number of contiguous 15-minute intervals when curtailment events occur in 2023 and analyze the distributions. We do this separately for solar and wind. Additionally, we add up all such intervals over the whole year for individual nodes to calculate how frequently a node experiences curtailment over the year.

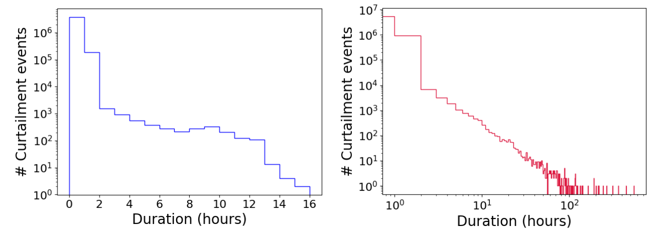

(a) Solar. **(b) Wind.**

Figure 6: ***Distribution of curtailment event durations. Solar (resp. wind) curtailment events can last up to 16 hours (resp. over a week).***

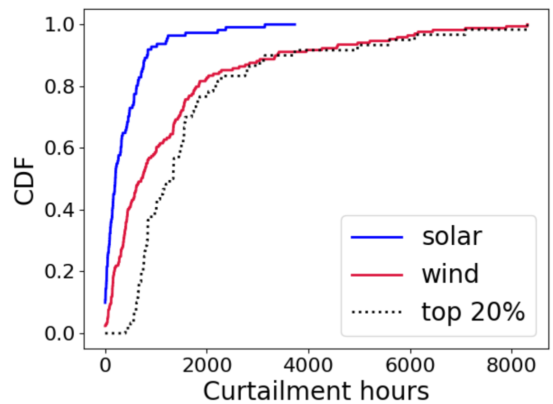

Figure 7: ***CDF of hours of curtailment across solar and wind nodes in Texas. Most of the nodes are curtailed infrequently.***

Fig. 6 shows the nodal solar and wind curtailment duration distribution. Curtailments at the node level can occur continuously from 15 minutes to 16 hours (resp. over a week) in solar (resp. wind) nodes. However, the median curtailment duration is 15 minutes, and solar (resp. wind) curtailments last less than 2 hours (resp. 1 hour) 90% of the time. Hence, curtailment at the node level is highly intermittent most of the time. However, we observe a significant number of curtailment events lasting several hours across the nodes, even though they constitute only a small percentage of the overall number of curtailment events. Thus, there is considerable potential to leverage curtailed energy for long-running workloads.

We observe that when all curtailment events over a day are aggregated, the duration of wind curtailment is typically more than that of solar. This is because while solar curtailment is limited by the hours of sunlight available, wind curtailment can occur during both day and night. Hence, there are more intervals available for potential curtailment. In extreme cases, wind curtailment can occur in a node during all hours of the day.

The frequency of curtailment is highly non-uniform across nodes. Fig. 7 shows the CDF of the number of hours of curtailment events across the nodes in 2023. 90% (resp. 57%) of the solar (resp. wind) nodes experienced some curtailment for less than 876 hours, which is equivalent to less than 10% of the year. Overall, 70% of the nodes

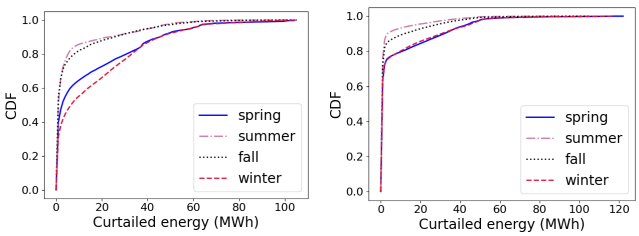

(a) Solar. **(b) Wind.**

Figure 8: *Solar and wind curtailment amount distribution (in MWh) across seasons in ERCOT. Curtailments are usually more in the Spring and Winter than in the Summer and Fall.*

are curtailed for less than 10% of the year. In this paper, we refer to the remaining 30% nodes as *frequently curtailed nodes*.

In general, nodes with high yearly curtailment are frequently curtailed. However, some nodes among the top 20% nodes with the most curtailment are not curtailed frequently (Fig. 7). On the other hand, we also observe nodes that are curtailed very frequently but have significantly low yearly curtailment. More analysis is needed to know why some nodes are curtailed more frequently than others.
Key Takeaways. Nodal curtailment is highly intermittent, with a majority of them lasting less than an hour and the median duration being 15 minutes. However, there are still a significant number of curtailment events that can last several hours, even though they constitute only a small percentage of the total number of events. Over the year, only 30% of the nodes are curtailed frequently, with wind nodes usually getting curtailed more frequently than solar.

4.4 Seasonal Patterns of Curtailment

Since weather conditions and electricity demand influence curtailment — both of which have distinct seasonal patterns — curtailment is likely to vary across seasons. Lin et al. [33] observed that system-wide wind curtailment in ERCOT is more during the Spring and Winter than during Summer or Fall. They concluded that higher electricity demand coupled with lower wind generation during Summer and Fall results in less wind curtailment.

In this section, we analyze if nodal curtailment also shows similar patterns. We divide our dataset into four seasons — Spring (March to May), Summer (June to August), Fall (September to November), and Winter (December to February). We see that the seasonal trends for the amount of nodal wind curtailment are similar to the findings reported by Lin et al. [33]. Interestingly, the amount of nodal solar curtailment also follows the same seasonal patterns even though solar production peaks during Summer and is less during Winter. Fig. 8 shows the CDF of node-level solar and wind curtailment amounts per 15-minute interval across seasons in ERCOT. Although the maximum curtailed energy in a 15-minute interval is similar across the seasons, median wind curtailment can be up to 1.2× higher in Spring and Winter than in Summer or Fall. Similarly, the median amount of curtailed solar energy in a 15-minute interval can be up to 2.8× (resp. 7.1×) higher in Spring (resp. Winter) than in Summer or Fall. We posit that more solar curtailment during Spring may be due to the same reason (rising generation combined with less demand). However, determining the reason behind higher solar curtailment during Winter needs more analysis and is kept as future work.

Additionally, we observe that wind curtailment durations do not show any significant difference across the seasons.
Key Takeaways. Solar and wind are curtailed more in Spring and Winter than in Summer or Fall. Median curtailment amount can be up to 7.1× more in Spring or Winter over other seasons. However, there is no significant seasonal difference in curtailment durations.

4.5 Curtailment Correlation Across Nodes

Finally, we examine the curtailment correlation across the different nodes in ERCOT. To do so, first, we partition all pairs of nodes based on their pairwise distances. Each partition has a radius of 50 miles; that is, we cluster all pairs located within 50 miles, then we cluster all pairs more than 50 miles apart but within 100 miles, etc. We continue this till we exhaust all possible pairs of nodes. Then, we use Pearson's correlation [48] on all pairs of nodes, where each node has a time series vector containing the curtailment for the whole year. This gives us the correlation between nodes at different distances. If Pearson's correlation coefficient (r) for a pair of nodes is > 0.6 (resp. < -0.6), we say that the pair has a high positive (resp. negative) correlation.

We analyze two types of correlation: (1) temporal correlation of curtailment events: in this case, each time series is a binary vector of 1s and 0s, with "1" denoting a curtailment event, and (2) temporal correlation of curtailment amount: in this case, the vectors also have the amount of curtailment.

Fig 9 shows the temporal correlation of curtailment events between pairs of nodes (ignoring the curtailment amount). Solar pairs are more correlated than wind pairs, and the correlation decreases as the distance between them increases. However, only 9.9% (resp. 0.5%) of solar (resp. wind) pairs are highly correlated within a 50-mile radius. Most pairs of nodes show a weak positive correlation ($0 < r \leq 0.3$). The high correlation in some nodes may be due to similar weather patterns and electricity demand. However, grid conditions surrounding the nodes can be considerably different even when two nodes are geographically proximal. Since curtailment is also affected by grid conditions like congestion, this may be the reason why most nodes show only weak correlations.

Fig. 10 shows the temporal correlation of curtailment amounts between pairs of nodes. More pairs of nodes have stronger correlations in this case — within a 50-mile radius, 29.2% solar and 6% wind nodes are strongly correlated. This implies that there are a lot of weakly correlated curtailments of low amounts, and curtailments of higher amounts are more correlated, which increases the Pearson coefficient in the second case. Most node pairs are still weakly correlated, and the correlations still weaken as the distance between the nodes increases.

When solar and wind nodes are paired together, 23.9% of the node pairs within a 50-mile radius have a weak negative temporal correlation ($-0.3 \leq r < 0$). The percentage of pairs with a weak negative correlation increases to 31.3% when the curtailment amount is also considered. This means there are periods when there is either solar curtailment or wind curtailment, but not both. Hence, there is potential to get a longer supply of curtailed energy by aggregating solar and wind curtailments.

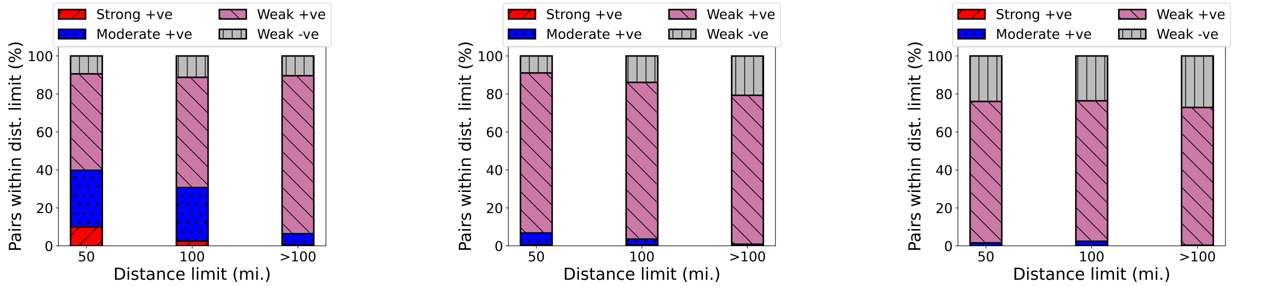

(a) Solar-Solar. (b) Wind-Wind. (c) Solar-Wind.

Figure 9: ***Temporal correlations of the occurrence of curtailment between pairs of nodes. Solar pairs have a higher correlation than wind pairs. Most pairs of nodes are weakly correlated, and correlation weakens with the distance between the nodes.***

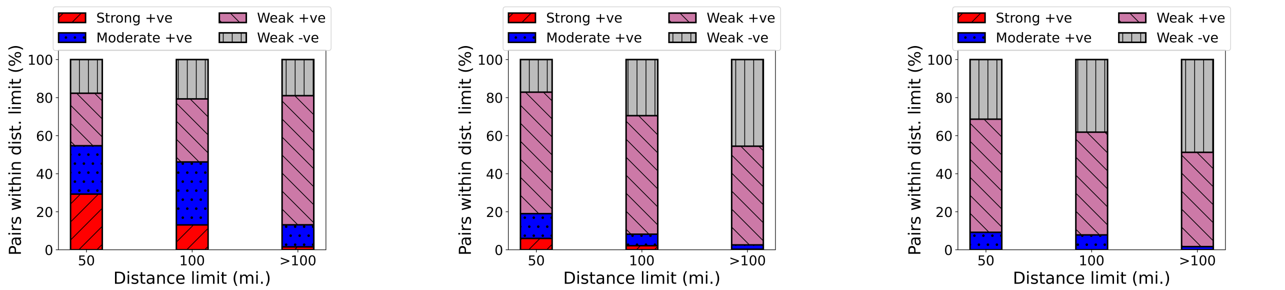

(a) Solar-Solar. (b) Wind-Wind. (c) Solar-Wind.

Figure 10: ***Temporal correlation of curtailment amount between the pairs of nodes. More pairs are strongly correlated, but most are still weakly correlated. This suggests many weakly correlated small curtailment events and some strongly correlated larger events.***

Key Takeaways. There are many weakly correlated small curtailment events and a few strongly correlated larger curtailment events. Consequently, most pairs of nodes are weakly correlated even when they are geographically proximal. In general, solar nodes are more correlated than wind nodes.

5 Price-Based Curtailment Cause Identification

Recall that curtailment can be broadly due to an oversupply of renewable energy or due to grid congestion when renewable energy cannot reach demand locations. Identifying the cause of curtailment is crucial to reduce curtailment or utilize curtailed energy effectively. For example, if the curtailment is due to congestion, adding demand to a place where the renewable supply cannot reach may result in a non-renewable node meeting that additional demand and increasing the system-wide carbon emissions. In this section, we develop a method to identify the cause of curtailment in a node at a specific time by looking at the nodal price (LMP).

LMP varies both spatially across nodes and temporally depending on the supply, demand, and grid conditions like transmission congestion. Thus, we posit that we may be able to derive the cause of curtailment from LMP signals. Specifically, we ask the question:

Can we identify the cause of curtailment in a node at a particular time by looking at the LMP signals?

To that end, we analyze the LMP at a node during curtailment. The LMP at any node is the sum of three components:

(1) Marginal cost of generating the next unit of electricity (MEC).
(2) Marginal cost of grid congestion (MCC).
(3) Marginal cost of electricity transmission losses (MLC).

$$\therefore LMP\ (\$/MWh) = MEC + MCC + MLC \tag{3}$$

MCC becomes zero when the grid has no congestion [38]. Hence, if MCC is zero when there is curtailment, then the curtailment is due to oversupply. Otherwise, curtailment is due to congestion. Since many grids provide information about MCC publicly [14, 24], the reason for curtailment in a node at a specific time can be determined by simply looking at MCC.

In our case, while ERCOT provides only the LMP values and not the value of each component, ERCOT ignores transmission losses in their LMP data [14, 24]. Thus, for ERCOT,

$$ERCOT\ LMP\ (\$/MWh) = MEC + MCC \tag{4}$$

ERCOT provides LMPs at both nodal and hub levels. The LMP of a hub is an average of the LMPs of all the nodes in that hub. When there are no transmission losses, LMP is the same across all nodes when there is no congestion [30]. Consequently, the LMP averaged over all the nodes in a hub will be equal to the nodal LMP. Thus, we can indirectly determine whether MCC is zero for a particular node by looking at the difference between the nodal LMP and the LMP of the hub where that particular node resides. If the LMP of a node is not equal to the LMP of a hub when there is curtailment, then MCC is non-zero, and that curtailment is due to congestion somewhere in the hub (and hence, in the grid). On the other hand, if the LMP of a node is equal to the LMP of all the other nodes in

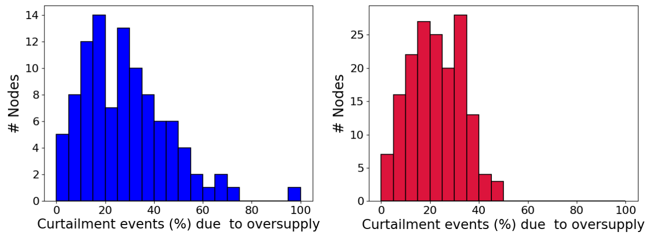

(a) Solar. **(b) Wind.**

Figure 11: ***Distribution of % of curtailment events occurring due to oversupply. Most solar and wind nodes are curtailed primarily due to congestion.***

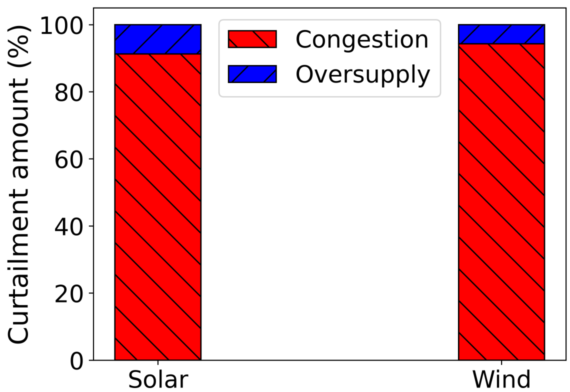

Figure 12: ***Comparison of amount of curtailment (%) due to oversupply and congestion. 93.4% of the total curtailed energy is curtailed due to congestion.***

the hub and hence the hub LMP when there is curtailment, MCC is zero, and that curtailment event is due to oversupply.

We evaluate the cause of curtailment across the nodes in Texas using our methodology. Figure 11 shows the histogram plots of the percentage of times solar or wind nodes are curtailed due to oversupply in 2023. During other times, curtailment is due to congestion. We see that in 79.4% (resp 100%) of the solar (resp. wind) nodes, less than 50% of the curtailment events are due to oversupply. Out of all the solar (resp. wind) curtailment events in 2023, 26.4% (resp. 25.6%) of the time, the nodal LMP is the same as the hub LMP. When added together, 25.7% of the time, the nodal LMP was the same as the hub LMP during curtailments. Thus, 74.3% of the curtailment events in ERCOT in 2023 are due to congestion.

Figure 12 shows the percentage of solar and wind curtailment that occurs due to oversupply or congestion. 91.2% (resp. 94.3 %) of the curtailed solar (resp. wind) energy is due to congestion, and the rest is curtailed due to oversupply. Together, 93.4% of the total curtailed energy in ERCOT is due to congestion.

Our method can also be used to pinpoint which transmission line in a grid is congested at a certain time. If the nodes at the start and end of a transmission line have different LMPs, then that transmission line may be congested (ignoring MLC). However, such analysis requires information about grid topology. Since ERCOT

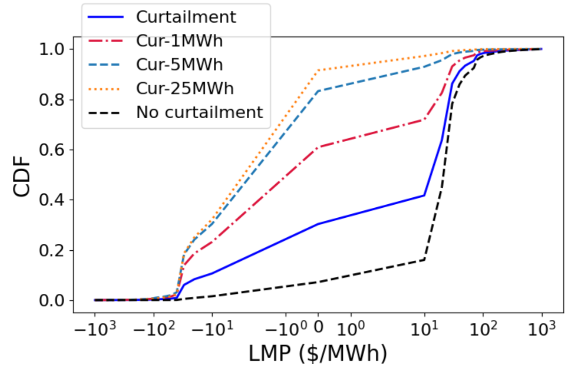

Figure 13: ***CDF showing LMP at different amounts of curtailment. LMP is typically less when there is curtailment and decreases with an increase in the amount of curtailment.***

does not publish any topology information, we are unable to include such an analysis in this paper.

Key Takeaways. During 74.3% of the curtailment events, MCC is non-zero. Based on our methodology, it means that those events are due to congestion. Curtailment events due to congestion account for 93.4% of the total curtailed energy. Thus, congestion is the leading cause of curtailment in Texas. Curtailment is due to oversupply the remaining 25.7% of the time (accounting for 6.6% of the total curtailed energy), as MCC is zero.

6 Price-Based Curtailment Detection

While curtailment data at the node level is scarce, many grids provide nodal LMP data. Thus, if we can somehow identify when and where curtailment occurs from LMP signals, researchers can use the easily accessible LMP data to overcome the data-related challenges associated with curtailment. For example, suppose there is a strong correlation between LMP and curtailment in a way that whenever the nodal LMP is below a certain threshold, it is likely that the node experiences curtailment. If that is true, then curtailment-aware optimizations can use LMP as an indicator to detect curtailment and modulate demand accordingly, even when curtailment data is unavailable. Hence, we ask the following question:

How is the curtailment in a node correlated to its LMP? Can we use LMP to detect or estimate curtailment?

To answer the questions, first, we investigate the LMP distribution when a node experiences curtailment versus when there is no curtailment. Fig. 13 shows the cumulative distribution of LMP at different amounts of curtailment. We see that the LMP distribution is typically less during curtailment than at other times. The median LMP when there is curtailment is 15.1 $/MWh. In contrast, the median LMP when there is no curtailment is 21.1 $/MWh. We also observe that LMPs decrease further as the amount of curtailment increases. If we consider curtailment of at least 5 MWh, the median LMP decreases to -1.5 $/MWh from 15.1 $/MWh.

Next, we examine the chance of curtailment given a certain price. Acun et al. [1] observed that in California and the South West, there are usually threshold prices below which curtailment is more likely to occur. If we can find a similar threshold price for ERCOT, that price may be a useful indicator of curtailment events.

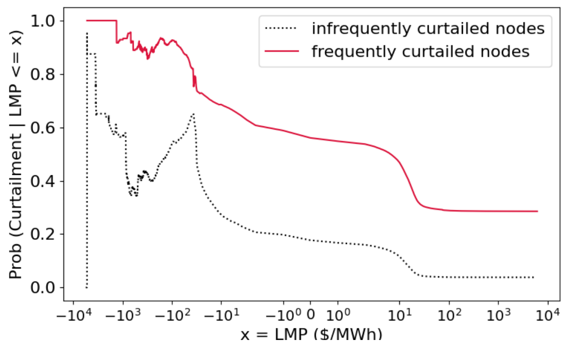

Figure 14: ***Probability of curtailment given the LMP is less than a certain amount (x). The probability of curtailment for frequently curtailed nodes increases as LMP decreases.***

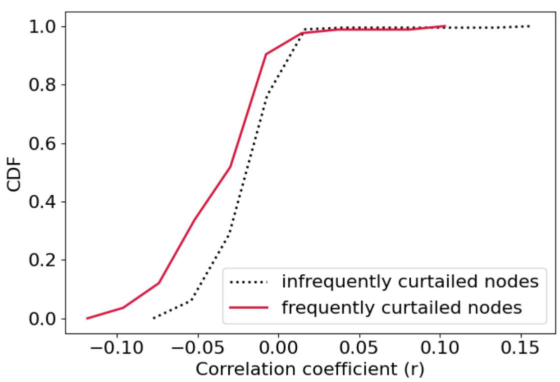

Figure 15: ***Distribution of correlation between LMP and curtailment across nodes. LMP has a weak correlation with the amount of curtailed energy.***

To do so, we compute the fraction of time an average node faces curtailment given the price is less than a specific value. Fig. 14 shows the results. In general, the probability of curtailment increases as the LMP decreases. However, for infrequently curtailed nodes (nodes curtailed less than 10% of the year), the probability of curtailment decreases when the price goes below -33$/MWh before increasing again at very low prices. On the other hand, the probability of curtailment increases for frequently curtailed nodes as the price decreases, and there is no significant dip. For those nodes, there is more than 70% probability of a curtailment event when the price is below -15 $/MWh. Thus, for frequently curtailed nodes, LMP may be a useful indicator for detecting curtailment events.

Finally, we analyze the temporal correlation between the amount of curtailment and LMP at a single node to determine if LMP can indicate curtailment amounts and trends in addition to detecting curtailment events. For each node, we treat the curtailment and LMP for the whole year as two vectors and calculate the year-long temporal correlation between them using Pearson's correlation [48]. Fig. 15 shows the CDF of the correlation coefficients across the nodes. All the nodes have a weak correlation between LMP and curtailment, regardless of how frequently they are curtailed.

Key Takeaways. LMP has the potential to indicate curtailment events in frequently curtailed nodes, as the probability of curtailment is more than 70% when the price is below -15$/MWh. However, estimating the curtailment amount from LMP is complicated since there is only a weak correlation, and needs more analysis.

7 Implications of Our Analyses

Our curtailment analyses shed light on the eventual goal of finding ways for using the renewable energy that would otherwise have been curtailed, thus decreasing the usage of "brown" energy. A natural way of achieving that goal is *forecasting curtailment events* and then *modulating the demand* during those events to consume the energy that would otherwise have been curtailed. We now discuss the implications of our analysis from these two perspectives.

(1) Forecasting Curtailment. Curtailment forecasts can be treated as either a classification or a regression problem. The classification problem (CP) forecasts whether there will be curtailment at a specific node and at a specific time, while the regression problem (RP) forecasts how much energy is likely to be curtailed. In general, RP is more difficult to solve than CP [23]. Specific implications follow.

(i) Our analysis in § 4.3 shows that most nodes are curtailed infrequently, with 70% of the nodes curtailed for less than 10% of the year. Many of these nodes have long periods of time — often spanning several days — without any curtailment events. Consequently, both CP and RP may be difficult in such nodes.

(ii) On the positive side, the remaining 30% of the nodes are curtailed more frequently. These nodes are likely amenable to ML-based forecasting techniques such as those used for carbon intensity or grid-level curtailment forecasting [23, 34, 41]. Although the fraction of frequently curtailed nodes is smaller, these nodes account for 65% of the curtailed energy. Hence, most of the curtailed energy can potentially be forecasted.

(iii) While factors like local weather conditions, renewable supply, and local demand may be helpful for forecasting curtailment, ML models may need additional features to increase their accuracy. Our work sheds light on what features are likely useful in forecasting curtailment. § 6 shows that nodal price is a strong indicator for detecting curtailment events. Hence, LMP can be a good feature for forecasting curtailment events (CP). However, since LMP only has a weak negative correlation with the amount of curtailed energy, it is not a good feature for forecasting curtailment amount (RP).

(iv) Finally, we found in § 4.5 that curtailment events across nodes are only weakly correlated, even when the nodes are geographically proximal. This suggests that forecast models would need to be trained individually for each node based on its own historical data. Models trained for one node may not be effectively transferred for forecasting curtailment at other nodes, thus limiting the efficacy of transfer learning for this problem.

(2) Modulating Demand. Specific implications are as follows.

(i) The geospatial distribution of curtailment (§ 4.1) implies that optimizers trying to reduce carbon emissions by using curtailed energy should primarily add or shift demand to the West, South, and the Panhandle rather than densely populated hubs like Houston.

(ii) Since 20% of the nodes account for 77% of the curtailment (§ 4.1), there is potential to significantly reduce curtailment and leverage curtailed energy by adding demand at only one-fifth of the nodes. However, the grid is congested 74.3% of the time when there is curtailment (§ 5). So, demand should be added carefully to avoid increasing the congestion. One way to avoid more congestion is by adding demand adjacent to the nodes — for example, by building data centers or deploying server clusters colocated with the nodes. When curtailment is due to oversupply and not congestion, demand can be added anywhere within a hub. In those situations, already existing data centers can scale up and add demand to use curtailed energy. Since renewable energy is typically cheaper, SCED would ensure that renewable nodes supply the additional demand.

(iii) While nodal curtailment is highly intermittent and infrequent, § 4.5 shows that when solar and wind nodes are paired together, 31.3% of those pairs have a weak negative correlation in terms of curtailment amount. Grouping such nodes may provide a much more stable supply of curtailed energy over longer periods of time. Hence, there is potential to execute long-running workloads using only curtailed energy by colocating server clusters in these nodes and then spatio-temporarily shifting such workloads.

(iv) Finally, many optimizations apply to workloads that can be shifted only temporally but not spatially. For example, researchers have looked at scheduling EV charging during curtailment events or low-carbon periods [8, 40]. However, EVs can only be charged at specific locations and cannot always be shifted to a location near a curtailed node. Our analysis extends to these types of optimizations and workloads, too. For example, our price-based curtailment cause identification methodology (§ 5) can enable such optimizations to schedule EV charging when curtailment is due to oversupply and there is no grid congestion. In this case, the renewable energy can flow from a curtailed source to a non-proximal node where there is demand.

Extensibility to Other Grids. Our analysis in this paper is restricted to ERCOT due to challenges in obtaining publicly available data for other regions. While our observations and implications are for the ERCOT grid, we hypothesize that some of them should apply to other grids as well. For example, similar to ERCOT, congestion is the primary cause of curtailment in the California grid (CAISO) [4]. However, verifying our observations and implications with data from other grids is future work.

8 Related Work

Curtailment Estimates and Studies. Curtailment is a well-known phenomenon in the electricity grids, and there are curtailment studies ranging over a decade. Bird et al. [2] report the curtailment practices in several US grids, examine the general causes of curtailment, and suggest practices that can potentially reduce curtailment. Chien et al. [6, 7, 33] characterize the growth of curtailment in different US grids and ways to leverage such energy. Nycander et al. [37] estimate curtailment in the Nordic grids by developing a power dispatch model. Numerous other works estimate curtailment both in the present as well as in the future in different geographical grids and countries [20, 39, 50]. All these works estimate curtailment at the grid level, whereas our work analyzes at a much finer node-level granularity. Frysztacki et al. [21] evaluate nodal curtailment by building a SCED model for a small-scale simulation of the German grid. In contrast, we use a data-driven approach to estimate and analyze curtailment. Additionally, we also analyze LMP to identify the cause of curtailment at different nodes.

Curtailment Forecasts. Recently, some works have forecasted solar and wind curtailment. Gorka et al. [23] detect and predict solar curtailment in California in real time. Hadian et al. [26], and Shams et al. [41] predict the amount of solar and wind curtailment in California, while Bunodiere et al. [3] predicts the same in Japan. Memmel et al. [35, 36] forecast curtailment in a simulated grid of Germany by detecting grid congestion. While these works forecast curtailment, our work is complementary to such works. We analyze the distribution and causes of curtailment by looking at node-level generation and pricing data. Our paper also describes how our data analysis insights can inform future prediction research and whether LMP can be used as an input to such forecasting methods.

Curtailment-Aware Computing. There is also growing attention to demand-side optimizations to utilize curtailed energy and reduce carbon emissions. Researchers have looked at shifting flexible loads like EV charging [8, 40] or data center computing [32, 49, 51] to places and periods of renewable curtailment. Our work is complementary to these works. All these optimizations need to know where, when, and why curtailment occurs at the node level, and our analysis tries to provide these answers.

9 Conclusions

As electricity grids move towards decarbonization, they are increasingly meeting their demand with renewable energy. As the amount of renewables in the grid increases, there are periods when renewable supply exceeds the demand or cannot reach the demand locations due to grid congestion. Consequently, a significant amount of renewable energy often needs to be curtailed. Since curtailment represents wasted renewable energy that could have potentially replaced "brown" energy, analyzing curtailment at the node level is crucial to understanding its potential for decarbonization.

In this paper, we study nodal solar and wind curtailment for the Texas grid. Using a data-driven approach, we show that nodal curtailment is highly non-uniform and intermittent — 20% of the nodes account for 77% of the total curtailed energy, while 70% of the nodes are curtailed for less than 10% of the year. We also develop a price-based method to identify the cause of curtailment and show that 74.3% of the curtailment events in 2023 are due to congestion. Overall, our analysis implies that while curtailment can potentially be forecasted in only a small fraction (30%) of the nodes, a considerable amount of curtailed energy (65%) can be utilized by adding demand adjacent to these nodes. Following this analysis, we plan to explore the feasibility of systems running exclusively on curtailed energy as future work.

Acknowledgments

This work is supported in part by NSF grants CNS-2105494, CNS-2325956, 23091241, 2213636, 2211302, and 2211888, DOE grant DEEE0010143, and a grant from VMware.

References

[1] Bilge Acun, Brent Morgan, Henry Richardson, Nat Steinsultz, and Carole-Jean Wu. 2024. Unlocking the Potential of Renewable Energy Through Curtailment Prediction. *arXiv preprint arXiv:2405.18526* (2024).

[2] Lori Bird, Jaquelin Cochran, and Xi Wang. 2014. *Wind and solar energy curtailment: Experience and practices in the United States*. Technical Report. National Renewable Energy Lab.(NREL), Golden, CO (United States).

[3] Alex Bunodiere and Han Soo Lee. 2020. Renewable energy curtailment: Prediction using a logic-based forecasting method and mitigation measures in Kyushu, Japan. *Energies* 13, 18 (2020), 4703.

[4] CAISO. 2024. Market Performance Report — December 2024. Retrieved May 1, 2025 from https://www.caiso.com/content/monthly-market-performance/dec-2024/renewable-resource.html

[5] CAISO. 2025. Managing the evolving grid. Retrieved Jan 13, 2025 from https://www.caiso.com/about/our-business/managing-the-evolving-grid

[6] Andrew A Chien. 2018. Characterizing opportunity power in the California independent system operator (CAISO) in years 2015-2017.

[7] Andrew A Chien and Liuzixuan Lin. 2024. As Grids Reach 100% Renewable at Peak, Growing Curtailment of 8 Gigawatts Looms as a Challenge to Decarbonization. *ACM SIGENERGY Energy Informatics Review* 4, 1 (2024), 3–10.

[8] James Dixon, Waqquas Bukhsh, Calum Edmunds, and Keith Bell. 2020. Scheduling electric vehicle charging to minimise carbon emissions and wind curtailment. *Renewable Energy* 161 (2020), 1072–1091.

[9] EIA. 2012. Negative wholesale electricity prices occur in RTOs. Retrieved May 1, 2025 from https://www.eia.gov/todayinenergy/detail.php?id=6730

[10] EIA. 2023. A Case Study of Transmission Limits on Renewables Growth in Texas. Retrieved Jan 13, 2025 from https://www.eia.gov/electricity/markets/quarterly/archive/2023/transmission_limits_07_2023.pdf

[11] EIA. 2023. Federal Financial Interventions and Subsidies in Energy in Fiscal Years 2016–2022. Retrieved Jan 25, 2025 from https://www.eia.gov/analysis/requests/subsidy/

[12] EIA. 2024. Electricity Explained. Retrieved Jan 14, 2025 from https://www.eia.gov/energyexplained/electricity/delivery-to-consumers.php

[13] EIA. 2024. Hourly Electric Grid Monitor. Retrieved Nov 20, 2024 from https://www.eia.gov/electricity/gridmonitor/dashboard/electric_overview/balancing_authority/ERCO

[14] Brent Eldridge, Richard P O'Neill, and Andrea R Castillo. 2016. *Marginal loss calculations for the DCOPF*. Technical Report. Sandia National Lab.(SNL-NM), Albuquerque, NM (United States).

[15] Electricity Reliability Council of Texas. 2024. 60-Day SCED Disclosure Reports. Retrieved Nov 20, 2024 from https://www.ercot.com/mp/data-products/data-product-details?id=NP3-965-ER

[16] Electricity Reliability Council of Texas. 2024. LMPs by Resource Nodes, Load Zones and Trading Hubs. Retrieved Nov 20, 2024 from https://www.ercot.com/mp/data-products/data-product-details?id=NP6-788-CD

[17] EPA. 2024. Renewable Electricity Production Tax Credit Information. Retrieved May 1, 2025 from https://www.epa.gov/lmop/renewable-electricity-production-tax-credit-information

[18] ERCOT. 2024. Monthly Outlook for Resource Adequacy (MORA), Reporting Month: November 2024. Retrieved Jan 8, 2025 from https://www.ercot.com/files/docs/2024/09/06/MORA_November2024.pdf

[19] ERCOT. 2025. Maps. Retrieved Jan 14, 2025 from https://www.ercot.com/news/mediakit/maps

[20] Bethany Frew, Wesley Cole, Paul Denholm, A Will Frazier, Nina Vincent, and Robert Margolis. 2019. Sunny with a chance of curtailment: operating the US grid with very high levels of solar photovoltaics. *IScience* 21 (2019), 436–447.

[21] Martha Frysztacki and Tom Brown. 2020. Modeling curtailment in germany: how spatial resolution impacts line congestion. In *2020 17th international conference on the European energy market (EEM)*. IEEE, 1–7.

[22] Goldman Sachs. 2025. AI is poised to drive 160% increase in data center power demand. Retrieved Jan 16, 2025 from https://www.goldmansachs.com/insights/articles/AI-poised-to-drive-160-increase-in-power-demand

[23] Joe Gorka and Line Roald. 2024. Classification Models for Forecasting and Real-Time Identification of Solar Curtailment in the California Grid. In *Proceedings of the 15th ACM International Conference on Future and Sustainable Energy Systems*. 158–169.

[24] Grid Status. 2024. Analyzing LMPs with a Nodal Price Map. Retrieved Jan 14, 2025 from https://blog.gridstatus.io/analyzing-lmps-with-a-nodal-price-map/

[25] GridStatus. 2024. Curtailment: When We Throw Away Clean Energy. Retrieved Nov 20, 2024 from https://blog.gridstatus.io/curtailment/

[26] Hengameh Hadian and Farnoosh Naderkhani. 2023. Deep Learning-Based Models for Wind and Solar Curtailment Forecasting. In *International Conference on Energy Harvesting, Storage, and Transfer*.

[27] Hong Chen (PJM). 2016. Security Constrained Economic Dispatch (SCED) Overview. Retrieved Jan 17, 2025 from https://www.aeso.ca/assets/Uploads/3.3-SCED-Overview-by-PJM.pdf

[28] IEA. 2021. Electricity Information: Overview. Retrieved Jan 18, 2025 from https://www.iea.org/reports/electricity-information-overview

[29] IEA. 2024. Electricity 2024 Executive summary. Retrieved Jan 18, 2025 from https://www.iea.org/reports/electricity-2024/executive-summary

[30] ISONE. 2025. FAQs: Locational Marginal Pricing. Retrieved Jan 12, 2025 from https://www.iso-ne.com/participate/support/faq/lmp

[31] Jan Dutton. 2024. ERCOT Curtailment and Renewable Power Generation. Retrieved Nov 20, 2024 from https://www.worldclimateservice.com/2024/10/18/ercot-curtailment/

[32] Kibaek Kim, Fan Yang, Victor M Zavala, and Andrew A Chien. 2016. Data centers as dispatchable loads to harness stranded power. *IEEE Transactions on Sustainable Energy* 8, 1 (2016), 208–218.

[33] Liuzixuan Lin and Andrew A Chien. 2012. Characterizing Stranded Power in the Electric Reliability Council of Texas (ERCOT) in Years 2012-2019: A Preliminary Report. *Energy* 2013 (2012), 2014.

[34] Diptyaroop Maji, Prashant Shenoy, and Ramesh K Sitaraman. 2022. CarbonCast: multi-day forecasting of grid carbon intensity. In *Proceedings of the 9th ACM International Conference on Systems for Energy-Efficient Buildings, Cities, and Transportation*. 198–207.

[35] Elena Memmel, Sunke Schlüters, Rasmus Völker, Frank Schuldt, Karsten Von Maydell, and Carsten Agert. 2021. Forecast of renewable curtailment in distribution grids considering uncertainties. *Ieee Access* 9 (2021), 60828–60840.

[36] Elena Memmel, Thomas Steens, Sunke Schlüters, Rasmus Völker, Frank Schuldt, and Karsten Von Maydell. 2023. Predicting renewable curtailment in distribution grids using neural networks. *IEEE Access* 11 (2023), 20319–20336.

[37] Elis Nycander, Lennart Söder, Jon Olauson, and Robert Eriksson. 2020. Curtailment analysis for the Nordic power system considering transmission capacity, inertia limits and generation flexibility. *Renewable Energy* 152 (2020), 942–960.

[38] NYISO. 2005. Congestion Cost Metrics. Retrieved Jan 14, 2025 from https://www.nyiso.com/documents/20142/1402684/agenda_13c_congestion_cost_description3.pdf/646044d2-77f0-18ab-134f-8e8c1df5f258

[39] Eric O'Shaughnessy, Jesse R Cruce, and Kaifeng Xu. 2020. Too much of a good thing? Global trends in the curtailment of solar PV. *Solar Energy* 208 (2020), 1068–1077.

[40] Sung-Won Park, Kyu-Sang Cho, Gregor Hoefter, and Sung-Yong Son. 2022. Electric vehicle charging management using location-based incentives for reducing renewable energy curtailment considering the distribution system. *Applied Energy* 305 (2022), 117680.

[41] Mohammad H Shams, Haider Niaz, Behzad Hashemi, J Jay Liu, Pierluigi Siano, and Amjad Anvari-Moghaddam. 2021. Artificial intelligence-based prediction and analysis of the oversupply of wind and solar energy in power systems. *Energy conversion and Management* 250 (2021), 114892.

[42] A. Shehabi, S.J. Smith, A. Hubbard, A. Newkirk, N. Lei, M.A.B. Siddik, B. Holecek, J. Koomey, E. Masanet, and D Sartor. 2024. *2024 United States Data Center Energy Usage Report*. Technical Report. Lawrence Berkeley National Laboratory, Berkeley, California. LBNL-2001637.

[43] United States Geological Survey. 2024. The U.S. Large-Scale Solar Photovoltaic Database. Retrieved Nov 20, 2024 from https://eerscmap.usgs.gov/uspvdb/data/

[44] United States Geological Survey. 2024. The U.S. Wind Turbine Database. Retrieved Nov 20, 2024 from https://eerscmap.usgs.gov/uswtdb/data/

[45] US Census Bureau. 2024. Quick Facts. Retrieved Jan 14, 2025 from https://www.census.gov/quickfacts/fact/table/TX/PST045219

[46] US Energy Information Administration. 2024. The US Energy Atlas. Retrieved Nov 20, 2024 from https://atlas.eia.gov/datasets/bf5c5110b1b944d299bb683cdbd02d2a_0/explore?location=36.153461%2C61.504000%2C4.21

[47] Wikipedia. 2025. Cost of electricity by source. Retrieved Jan 17, 2025 from https://en.wikipedia.org/wiki/Cost_of_electricity_by_source

[48] Wikipedia. 2025. Pearson correlation coefficient. Retrieved Jan 14, 2025 from https://en.wikipedia.org/wiki/Pearson_correlation_coefficient

[49] Fan Yang and Andrew A Chien. 2016. ZCCloud: Exploring wasted green power for high-performance computing. In *2016 IEEE International Parallel and Distributed Processing Symposium (IPDPS)*. IEEE, 1051–1060.

[50] Yoh Yasuda, Lori Bird, Enrico Maria Carlini, Peter Børre Eriksen, Ana Estanqueiro, Damian Flynn, Daniel Fraile, Emilio Gómez Lázaro, Sergio Martín-Martínez, Daisuke Hayashi, et al. 2022. CE (curtailment–Energy share) map: An objective and quantitative measure to evaluate wind and solar curtailment. *Renewable and Sustainable Energy Reviews* 160 (2022), 112212.

[51] Jiajia Zheng, Andrew A Chien, and Sangwon Suh. 2020. Mitigating curtailment and carbon emissions through load migration between data centers. *Joule* 4, 10 (2020), 2208–2222.