Impacts of Large, Flexible Data Center Operations on the Future of ERCOT

By Joshua D. Rhodes, PhD
Thomas Deetjen, PhD
and Caitlin Smith

June 2021

Table of Contents

Executive Summary

This white paper seeks to quantify the impacts of large flexible data center loads located in West Texas on the overall carbon emissions of the Texas electricity grid. A capacity expansion model of the Texas electricity grid was utilized to examine how the grid would evolve out to 2030 under four scenarios; 1) a base case with no data centers, 2) 5 GW of **inflexible** data centers, 3) 5 GW of **flexible** data centers, and 4) 5 GW of **more flexible** data centers. The results of this analysis indicate that:

- Adding additional data center load to the grid incentivizes more wind and solar to be built than the base case of no data centers.

- If the additional data center load is **inflexible**, the model also builds more natural gas than the base case, but if the data centers are **flexible**, less natural gas is built.

- Operating the data centers in an **inflexible** manner results in more carbon emissions than the base case.

- Operating the data centers in a **flexible** manner can result in a *net-reduction* of carbon emissions from the base case.

- This analysis estimates that, to reduce carbon emissions, data centers will have to shed at least 13-15% of their load annually in an intelligent way.

- Beyond the possible carbon emissions reductions, this analysis also found that the additional flexibility of the data centers can increase the resiliency of the grid by reducing demand during high-stress times (low reserves) on the grid.

Introduction

With the growth of the digital services industry comes increased demand for data center operations, including even faster recent growth related to increased teleworking during the COVID-19 pandemic. While efficiency gains and hyperscale technologies have lowered some of the initial data center energy growth estimates, data centers still consume roughly 3% [1] of total delivered electricity globally, and close to 2%[1] in the US [2].

While data center operations drive many of the real-time needs of modern life, other operations, such as cryptocurrency mining have received significant criticism for their large amount of energy use and resulting carbon emissions [3]. All things equal, increasing energy use will increase carbon emissions, unless the additional energy consumed is met by, or offset with, carbon-free energy.

This analysis seeks to understand the carbon emissions tradeoffs of additional data center operations on an electric grid. In particular, we test how adding flexible load shed capabilities to these additional demands impacts the evolution of an existing electric grid and its associated total emissions and emissions intensity of its electricity. We present this analysis as a case study and utilize the Texas grid as the testbed.

The deregulated electricity grid of Texas is dynamic and evolves quickly. Texas consumes almost twice as much electricity as the next-highest state [4], and while overall electricity growth in the US has been relatively flat, Texas is expected to see considerable growth in the electricity sector. The Electric Reliability Council of Texas (ERCOT), the grid that serves roughly 90% of Texas, expects electricity consumption to increase over 25% from 2018 to 2033 [5].

Texas is also home to multiple data-driven companies and include massive campuses for Facebook, Apple, Amazon, Tesla, among others. Companies often find it advantageous to move operations to Texas given the relatively lower cost of living and the ease of procuring low-cost, wholesale, and renewable and clean energy through the deregulated market structure.

Methodology

The following is a brief description of the model and methodology description used for the analysis in this report. A more detailed description of each can be found in Appendix A.

The model

This analysis modeled the ERCOT grid by utilizing a customized version of the SWITCH 2.0 open-source capacity expansion model [6]. A capacity expansion model is an optimization program that makes decisions about the operation, retirement, and construction of power plants, transmission lines, and other electric grid assets. It accomplishes this on both short (grid operations) and long (system planning) timescales. On the short time scale, the model dispatches the power plant fleet so that electricity generation and electricity demand are

[1] As of 2014.

balanced for each hour of the simulation. On the long-time scale, the model builds new power plant and transmission capacity to 1) provide enough power plants so that electricity generation and demand can be balanced in future years, and 2) enable the composition of the power plant fleet to evolve in ways that minimize the total system cost.

To increase the level of detail in the final results, the power plant capacities from the SWITCH capacity expansion model are used to run an 8,760-hour unit-commitment and dispatch simulation for each of the simulated years using the PyPSA open-source dispatch model [7]. By increasing the time resolution to include every hour of the year, the PyPSA output provides greater detail about the power plant dispatch, production cost, transmission development, transmission congestion, renewables curtailment, emissions, and other data.

ERCOT-specific data

The baseline year for the grid optimization analysis was 2018. Baseline year data include both spatial load and renewable generation profiles from the same year, which is important because the same meteorological conditions that drive renewable generation also impact load. All data, including the existing power plant fleet, used in this analysis are based on public ERCOT reports. Future fuel price and technology costs are based on the National Renewable Energy Laboratory's NREL Annual Technology Baseline (ATB) and the US Energy Information Administration's (EIA) Annual Energy Outlook.

Transmission

Figure 1 shows the 16 zone ERCOT model and transmission network used in this analysis. These types of reduced-order transmission models are commonly used in these types of analyses to keep the problem tractable [8]. The transmission limits between each of the connected zones were calculated based on physical infrastructure, historical power flows, and Generic Transmission Constraints [9]. While not the focus of this analysis, the ability to build new transmission capacity is important when considering how other aspects of the grid might evolve.

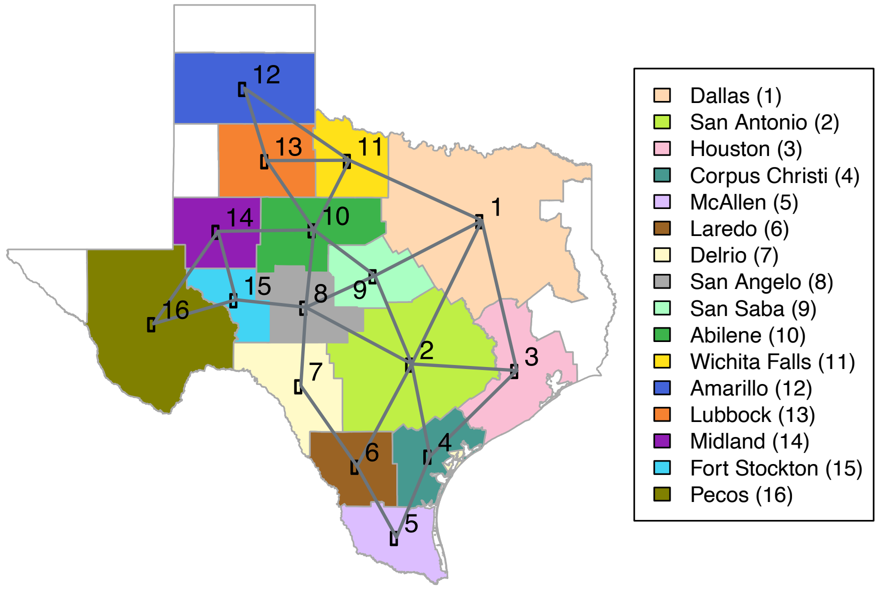

Figure 1: The 16-zone ERCOT model and transmission network used in this analysis.

Time horizon

Given the relative uncertainty about the costs of different technologies the further into the future one gets, this analysis analyzed the impacts out to 2030. This time horizon followed closely with the most recent ERCOT Long-Term System Assessment (LTSA) [5], from which load forecast data were utilized to model the evolution of the ERCOT grid in three time periods (2021-2024, 2025-2028, 2029-2032).

Data center scenarios

This analysis analyzed four (4) scenarios of data center growth in ERCOT;

1) no additional data center growth (Scenario 1, S1),
2) 5 GW of inflexible data center growth (Scenario 2, S2),
3) 5 GW of flexible data center growth (Scenario 3, S3), and
4) 5 GW of more flexible data center growth (Scenario 4, S4).

These additional loads (scenarios 2-4) were added to the demand growth assumptions already in place out to 2030. The data centers were assumed to be flexible based on wholesale market prices and would reduce operations when those prices exceeded a certain price. Table 1 shows a breakdown of the locations of each of the four date center locations, their size, number of load reduction tiers, the ERCOT real-time price at which the tier comes into effect, and the amount of load reduction at that tier.

Table 1: Table showing the data center locations (Texas cities), size, number of flexibility tiers, the price at which each tier comes into effect, and the size of demand reduction at each tier. The North Texas 1 data center is located in Figure 1's region 11, the Central Texas data center is located in Figure 1's region 10, the North Texas 2 data center is located in Figure 1's region 12, and

West Texas data center is located in Figure 1's region 8. These energy price tiers are based on future projections of cryptocurrency mining efficiencies and were provided to IdeaSmiths by Lancium, see Appendix B.

	Location	Date center size (MW)	Tier	Tier cutoff ($/MWh)	Tier demand reduction (MW)
Scenario 3 (S3)	North Texas 1	750	1	$35	363
			2	$95	363
	Central Texas	1750	1	$35	848
			2	$95	848
	North Texas 2	750	1	$35	363
			2	$95	363
	West Texas	1750	1	$35	848
			2	$95	848
Scenario 4 (S4)	North Texas 1	750	1	$24	242
			2	$41	242
			3	$82	242
	Central Texas	1750	1	$24	565
			2	$41	565
			3	$82	565
	North Texas 2	750	1	$24	242
			2	$41	242
			3	$82	242
	West Texas	1750	1	$24	565
			2	$41	565
			3	$82	565

The process

The model was fed the ERCOT-specific data, the transmission network, and the time horizon to allow the model to determine the power system architecture that will minimize system costs overtime. To accomplish this task, the model simulated the dispatch and retirement of existing power plants, as well as the construction of new generation, energy storage systems, and transmission capacity to meet future demand growth. The model will not build new infrastructure unless it reduces overall system costs. For example, new energy storage could create an opportunity to build newer, more affordable generation resources or allow existing resources to be dispatched in a way that reduces system costs enough to offset the additional capital investment requirements for the new infrastructure.

This analysis did not include any goals or targets for any particular type of technology, such as a Renewable Portfolio Standard, or a tax on any type of pollutant, such as CO_2. The analysis also did not include any subsidies, such as the Production Tax Credit.

Resiliency analysis

Because a grid model is inherently deterministic and "knows" the future, it designs systems that will not fail, under the circumstances that it is given. In reality, grid components (power plants, fuel supply chains, transmission lines, etc.) break (have outages) and actual experiences will differ from modeled ones. In general, grid modelers force models to not only match supply and demand, but also carry a certain level of reserves designed to cover any expected or unexpected outages. These levels of reserves are not sized to reduce the chances of scarcity and blackouts to zero, but to some acceptable level, generally such that you only have one lost load event every ten years [10].

After the grid modeling analysis was complete, we ran a Monte Carlo analysis using probability distributions that any power plant might be offline due to maintenance, planned outages, or forced outages based on NERC Generating Availability Data System (GADS) data [11]. Outages were calibrated such that 100 out of 1,000 annual simulations yielded levels of operating reserves would fall below 1,375 MW, thus resulting in firm load shed [12]. We then assess the impact of any additional load reduction (non-dispatched price-sensitive data center load) on the levels of reserves available to the system operator.

Results

The crux of this analysis focused on how differently the ERCOT grid would evolve given how flexible the future data centers in question were willing to be. The results of this analysis are presented below in three subsections: 1) generation capacity and energy evolution, 2) carbon emissions, and 3) resiliency.

Capacity and generation

As demand increases and prices change through time (see Appendix A), so does the optimal grid mix. Figure 2 compares the 2018 capacities of each type of generation with the 2030 capacities for each scenario.

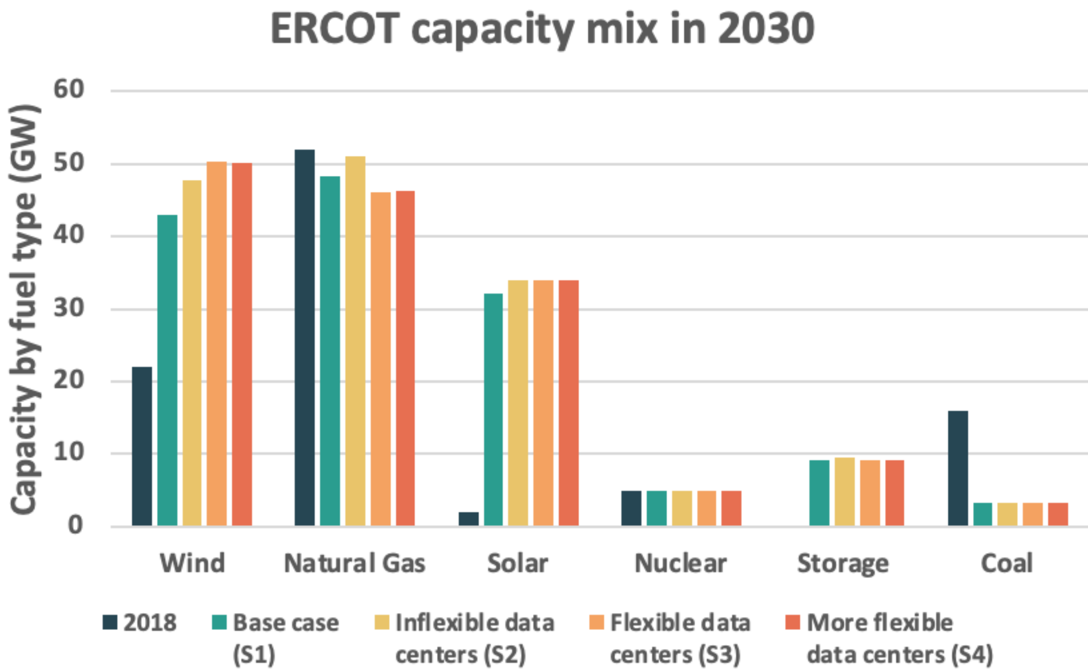

Figure 2: ERCOT generation capacity mix in 2018 and 2030 for all scenarios (GW).

Across all scenarios, wind, solar, and storage capacity increases while natural gas and coal capacity decreases. The additional loads introduced by the data centers further increase the amount of wind and solar built, but the inflexible data center scenario (S2) does result in more natural gas capacity than the base (no data center) case (S1). Both flexible data center cases (S3-S4) result in less natural gas capacity than the base case. Figure 3 shows the amount of energy generated by each type.

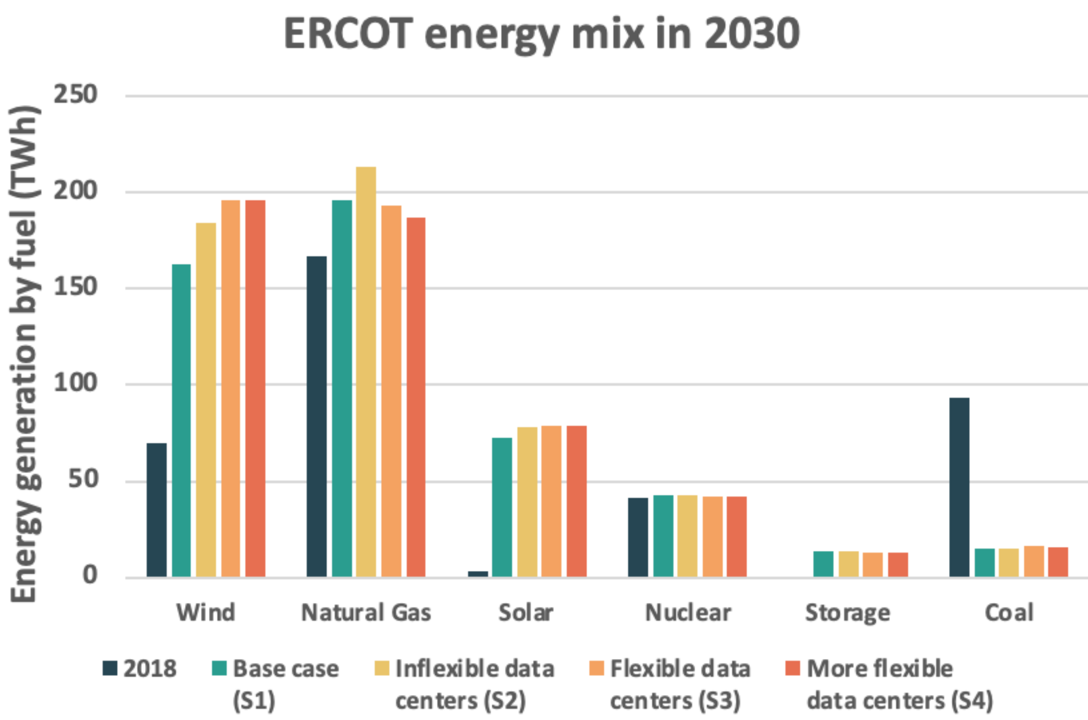

Figure 3: ERCOT energy mix in 2018 and 2030 for all scenarios (TWh).

As compared to 2018, all future scenarios get more overall energy from wind, natural gas, and solar and less energy from coal. As compared to the 2030 base case (no data centers, S1), each of the data center scenarios (S2-S4) see more energy generated from wind and solar. The inflexible data center scenario (S2) sees more energy from natural gas than the base case (S1) in 2030, but the flexible data center cases (S3-S4) see less.

Data center operations

The differences in scenario results are a direct result of the flexibility of the data centers themselves (see Data center scenarios). Each of the scenarios led to the data centers being dispatched (temporarily turned down) differently. The inflexible data center (S2) was modeled as having no flexibility and thus achieved an uptime of 100%[2].

The four data centers that in the flexible data center scenario (S3) are dispatched down (curtailed) slightly differently depending on their location in the ERCOT grid and range in uptimes of 85%-87%, meaning that, on average across the fleet, 14% of their capacity has been dispatched down. The actual dispatch has periods where the data centers are operating at full capacity and others where they are operating lower, depending on the real-time price of electricity. Figure 4 shows an example week in which the data center fleet is responding to ERCOT market prices.

[2] We did not consider outages of the data centers themselves and assumed across all scenarios that the only reason for downtime would be if they voluntarily turned down.

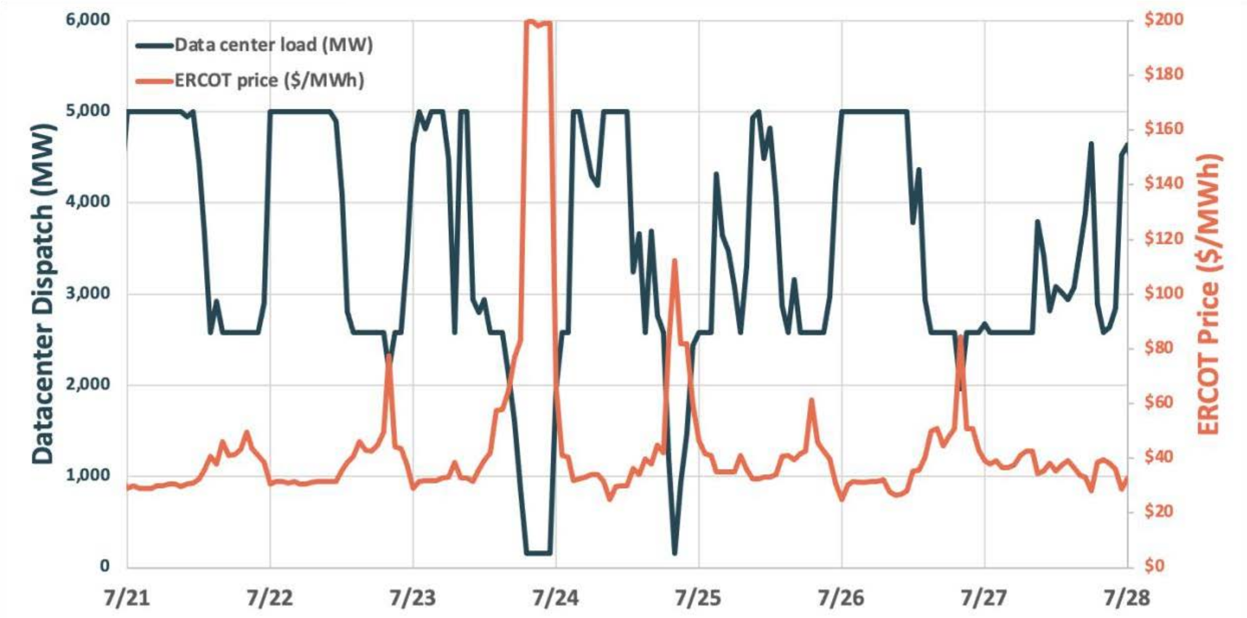

Figure 4: Example week of the data center load responding to real-time ERCOT prices for Scenario 3.

In general, the data centers operate at high or full capacity when prices are low (most summer mornings in Figure 4) and then reduce their operations as prices increase (most summer afternoons). When prices go very high (above $95/MWh in Scenario 3), the datacenters dispatch down to their minimum operating level, approximately 3% of their nameplate capacity.

If the four data centers are operated more flexibly (S4), they are also dispatched down (curtailed) more often. On average over the year, the data center fleet in Scenario 4 achieves a 70% uptime. This additional flexibility has a negligible impact on wind, solar, and natural gas capacity, but it does reduce the amount of energy generated by natural gas power plants.

Carbon emissions

How the data centers are operated impacts how much energy is generated from each type of fuel (see Section Capacity and generation) as well as the carbon emissions associated with the electricity grid. Figure 5 shows the differences in million tonnes of carbon emissions per year (in 2030), relative to the base case (S1) for Scenarios 2-4.

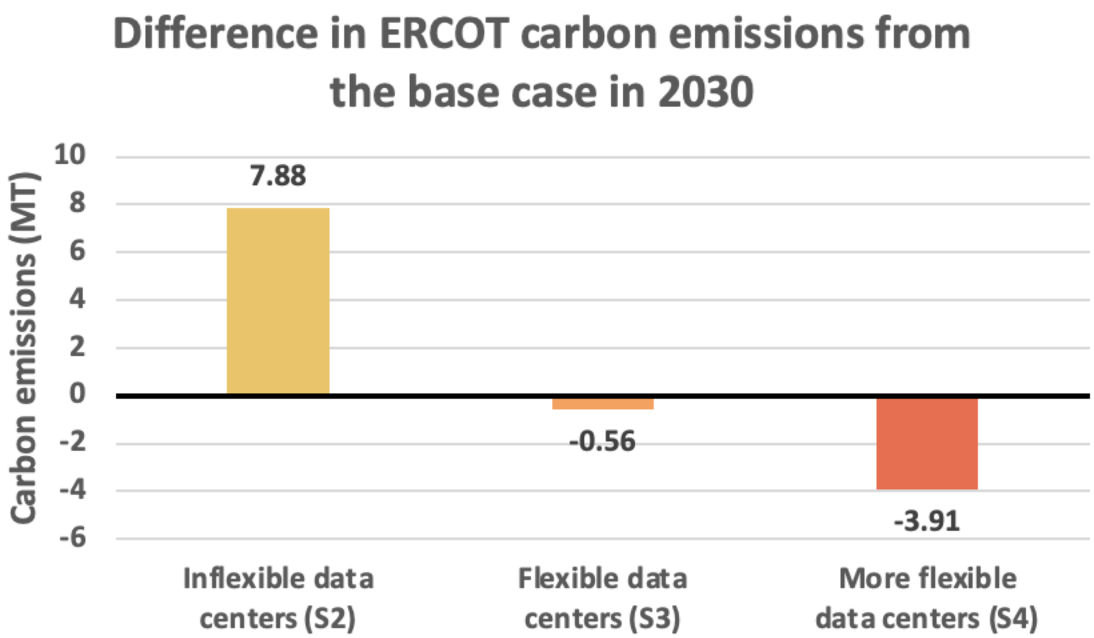

Figure 5: The difference in carbon emissions from the base case for each of the data center scenarios, in 2030. A positive value means that carbon emissions are higher than the base case (no data centers) and a negative number indicates that the carbon emissions for that scenarios are lower than the base case.

Figure 5 shows that, in 2030, we estimate that deploying the inflexible data centers (S2) would result in more carbon emissions than the no data center base case (S1). However, the flexibility of the data centers (cases S3 & S4) results in lower overall carbon emissions, even with the additional load added to the system to power the data centers.

In this analysis, the inflexible data center scenario (S2) adds approximately 43.8 million MWhs of energy consumption to the Texas Grid, while supporting the deployment of an additional 27 million MWhs of wind and solar energy and 17.1 million MWh of natural gas (than the base case, S1).

However, for an additional load to result in lower carbon emissions, the additional energy consumed by the load must be more than offset by zero carbon energy. The flexible data centers in scenario (S3) consume about 35.5 million MWh, but support the deployment of an additional 39.5 million MWhs of wind and solar energy. The data centers in the more flexible case (S4) consume about 30.6 million MWhs of energy while supporting the deployment of an additional 39.2 million MWhs of wind and solar. In both cases, the energy generated by natural gas power plants is lower than in the base case (S1).

The flexibility of the datacenters is key to these results and can be seen through the differences in the inflexible (S2) and flexible (S3, S4) data center cases. The flexibility of the demand reduces the need for firm generation. This allows the model to build more wind and solar energy providing capacity but less natural gas capacity to maintain the target reserve margin of 13.25%.

This analysis only considered a small number of flexibility scenarios. It is likely that a data center with less flexibility than S3 (i.e., higher bids) would result in net positive CO2 emissions. Thus, for a data center to have a net negative impact on CO2 emissions, it needs to operate its load shed strategically. Using these scenarios as a guide, it is likely that such a data center could expect to shed about 13-15% of their annual load.

Resiliency

Even with adequate reserve margins, higher levels of generation outages can impact the ability of any system to reliably match supply and demand. This analysis sought to assess the resiliency impact of high levels of demand flexibility by assessing how much additional flexible data center capacity would be available to mitigate low levels of reserves based on a Monte Carlo analysis of power plant failures.

ERCOT goes into an Energy Emergency Alert – Level 3 with firm load shed when operational reserves cannot be maintained above 1,375 MW. The grid strives to keep these situations to a minimum, with a target of it only happening once in ten years. Giving each power plant a probability of an outage, we tuned the levels of outages such that, for the base case (S1), only 100 out of 1,000 simulations resulted in reserves dipping below 1,375 MW [12]. We then took these same runs and (hour by hour) added back any additional flexible data center capacity that had not already been taken offline based on prices to see how many runs out of the original 1,000 still dipped below the 1,375 MW threshold.

Results for the flexible data center cases (S3 & S4) were similar and resulted in a lower number of simulations where reserves dropped to critical levels. For the flexible data center scenario (S3), the percentage of "years" reaching critical levels dropped (from 10% in the base case) to about 3.4% and for the more flexible data center case (S4) to 3.2%[3].

Additional Discussion

This section discusses additional aspects of controllable load, such as data centers, but was not considered in the above summarized analysis.

Additionality

This analysis assumed that the data center load considered was new. It is possible that, if existing data center load were moved from a currently carbon-intensive grid to the EROCT grid, the scenarios considered could result in an even greater reduction in global overall carbon emissions. However, we did not consider that as part of this analysis, so our estimates could be conservative depending on if the data center load considered is new or existing.

[3] We recognize that this comparison is not perfect as each scenario evolves differently with different capacities of different types of generation, but it does indicate that high levels of flexible demand contribute to greater levels of grid reliability.

Nodal pricing considerations

This analysis assumed that the data centers were acting based on the wholesale price of electricity at their location. In reality in ERCOT, load pays for electricity at the load zone whereas generation is compensated at the resource node. Loads such as the data centers considered in this analysis could qualify as Controllable Load Resources (CLRs), which have the ability of curtailing demand and providing energy and ancillary services. CLRs are able to function as generators, so it would be advantageous to have them responding to nodal instead of zonal prices. Also, if these controllable resources were able to receive nodal prices, it could encourage them to be sited in congested locations that have historically depressed prices. This optimal, price-based siting could reduce congestion and reduce the original need to build new transmission to alleviate it, further reducing overall system costs.

Conclusions

This analysis sought to assess the energetic and environmental impacts of deploying 5,000 MW of data centers in the ERCOT grid by 2030. The results indicate that operating the data centers in an inflexible manor spurs the deployment of more wind and solar than in the base case of no data centers, but results in a net increase in carbon emissions. However, operating the data centers in a flexibile manor during times of high grid prices could lead to the deployment of even more wind and solar and—if they are operated with enough flexibility—could result in lower overall carbon emissions. This analysis indicates that a data center in ERCOT would need to be willing to ramp down about 13-15% of its capacity per year, like Scenario 3 above, to achieve such a goal. A post grid modeling Monte Carlo analysis of grid operating reserves levels further indicates that the higher levels of flexible demand results in a lower probability of the grid reaching critical levels of reserves that would require firm load shed. While this work focused on the ERCOT grid, it is possible that other grids with a low cost of entry for new generation resources could see similar results.

Acknowledgements
This work was funded by Lancium[4].

About Us
IdeaSmiths LLC[5] was founded in 2013 to provide clients with access to professional analysis and development of energy systems and technologies. Our team focuses on energy system modeling and assessment of emerging innovations, and has provided support to investors, legal firms, and Fortune 500 companies trying to better understand opportunities in the energy marketplace.

[4] https://lancium.com/
[5] https://www.ideasmiths.net/

Bibliography

[1] R. Danilak, "Why Energy Is A Big And Rapidly Growing Problem For Data Centers," Forbes, 2017.

[2] J. Koomey, E. Masanet, N. Horner, I. Azevedo and W. Lintner, "United States Data Center Energy Usage Report," ERNEST ORLANDO LAWRENCE BERKELEY NATIONAL LABORATORY, Berkeley, California, 2016.

[3] R. Browne, "Bitcoin's wild ride renews worries about its massive carbon footprint," CNBC, 05 02 2021. [Online]. Available: https://www.cnbc.com/2021/02/05/bitcoin-btc-surge-renews-worries-about-its-massive-carbon-footprint.html. [Accessed 27 05 2021].

[4] US Energy Information Administration, "State Electricity Profiles," 02 11 2020. [Online]. Available: https://www.eia.gov/electricity/state/. [Accessed 16 12 2020].

[5] ERCOT, "2018 Long-term System Assessment for the ERCOT Region," 2018.

[6] J. Johnston, R. Henriquez-Auba, B. Maluenda and M. Fripp, "Switch 2.0: a modern platform for planning high-renewable power systems," *SoftwareX,* vol. 10, p. 100251, 2019.

[7] T. Brown, J. Horsch and D. Schlachtberger, "PyPSA: Python for Power System Analysis," *Journal of Open Research Software,* vol. 4, no. 1, 2018.

[8] T. A. Deetjen, H. Martin, J. D. Rhodes and M. E. Webber, "Modeling the optimal mix and location of wind and solar with transmission and carbon pricing considerations," *Renewable Energy,* vol. 120, pp. 35-50, 2018.

[9] Electric Reliability Council of Texas, "Use of Generic Transmission Constraints in ERCOT," Austin, TX, 2020.

[10] S. A. Newell, K. Spees, J. P. Pfeifenberger, I. Karkatsouli, N. Wintermantel and K. Carden, "Estimating the Economically Optimal Reserve Margin in ERCOT," The Brattle Group, Washington DC, 2014.

[11] North American Electric Reliability Corporation, "Generating Availability Data System (GADS)," 2021. [Online]. Available: https://www.nerc.com/pa/RAPA/gads/Pages/GeneratingAvailabilityDataSystem-(GADS).aspx. [Accessed 27 05 2021].

[12] ERCOT, "ERCOT's use of Energy Emergency Alerts," [Online]. Available: http://www.ercot.com/content/wcm/lists/164134/EEA_OnePager_FINAL.PDF. [Accessed 02 06 2021].

[13] ERCOT, "Resource Adequacy," [Online]. Available: http://www.ercot.com/gridinfo/resource. [Accessed September 2020].

[14] U.S. Environmental Protection Agency, "Emissions & Generation Resource Integrated Databse," 2018. [Online]. Available: https://www.epa.gov/egrid/emissions-generation-resource-integrated-database-egrid. [Accessed September 2020].

[15] National Renewable Energy Laboratory, "Annual Technology Baseline: Electricity," 2019. [Online]. Available: https://atb.nrel.gov/electricity/2019/. [Accessed September].

[16] J. B. Garrison, "A Grid-Level Unit Commitment Assessment of High Wind Penetration and Utilization of Compressed Air Energy Storage in ERCOT," *University of Texas Dissertation,* 2014.

[17] ERCOT, "Resource Adequacy," [Online]. Available: http://www.ercot.com/gridinfo/resource.

[18] National Renewable Energy Laboratory, "Global Radiation at Latitude Tilt - Annual Texas," [Online]. Available: https://openei.org/wiki/File:NREL-eere-pv-h-texas.pdf. [Accessed September 2020].

[19] S. Ong, C. Campbell, P. Denholm, R. Margolis and G. Heath, "Land-Use Requirements for Solar Power Plants in the United States," *National Renewable Energy Laboratory,* 2013.

[20] P. Denholm, M. Hand, M. Jackson and S. Ong, "Land-Use Requirements of Modern Wind Power Plants in the United States," *National Renewable Energy Laboratory,* 2009.

[21] National Renewable Energy Laboratory, "Texas 80-Meter Wind Resource Map," [Online]. Available: https://windexchange.energy.gov/maps-data/122.

[22] Texas Comptrolloer, "Summer 2020 Economic Forecast," Texas Comptroller, [Online]. Available: https://comptroller.texas.gov/transparency/reports/forecasts/2020-07/. [Accessed September 2020].

[23] Clean Energy Solutions Center, "Transmission Planning for a High Renewables Future," September 2017. [Online]. Available: https://cleanenergysolutions.org/training/texas-crez-success. [Accessed September 2020].

[24] ERCOT, "Quick Facts: April 2017, January 2018, January 2019, January 2020".

[25] American Society of Farm Managers and Rural Appraisers, "Texas Rural Land Value Trends," 2019.

[26] J. Rhodes, "The Economic Impact of Renewable Energy in Rural Texas," Austin, TX, 2020.

[27] W. Short, P. Sullivan, T. Mai, M. Mowers, C. Uriarte, N. Blair, D. Heimiller and A. Martinez, "Regional Energy Deployment System (ReEDS) NREL/TP-6A20-46534," National Renewable Energy Laboratory, 2011.

[28] H. I. F.-L. Data, "Electric Power Transmission Lines," [Online]. Available: https://hifld-geoplatform.opendata.arcgis.com/datasets/electric-power-transmission-lines. [Accessed November 2020].

[29] pandapower. [Online].

[30] U.S. Energy Information Administration, "Annual Energy Outlook, Reference Case, Table 3," 2020.

[31] U.S. Energy Information Administration, "Form 923," 2015-2019.

[32] ERCOT, "Hourly Load Data Archives," [Online]. Available: http://www.ercot.com/gridinfo/load/load_hist. [Accessed September 2020].

[33] ERCOT, "2020 ERCOT System Planning Long-Term hourly Peak Demand and Energy Forecast," 2019.

[34] National Renewable Energy Laboratory, "System Advisor Model," [Online]. Available: https://sam.nrel.gov/.

Appendix A

1. General Model Summary: Capacity Expansion Modeling in Switch

The analysis for this project is completed using the capacity expansion model called, "SWITCH" [6] and the unit-commitment and dispatch (UC&D) model called, "PyPSA" [7]. We use SWITCH to optimize the construction of power plants over multiple forecasted time periods that may span many years or decades. We use PyPSA to model the dispatch of those forecasted power plant fleets at a greater time-resolution and to determine optimal transmission capacities during each of those future time periods.

A capacity expansion model is an optimization program that makes decisions about the operation and construction of power plants, transmission lines, and other electric grid assets. It accomplishes this at two different time scales:

- Short Time Scale: the model dispatches the power plant fleet so that electricity generation and electricity demand are balanced for each hour of the simulation.
- Long Time Scale: the model builds new power plant capacity to 1) provide enough power plants so that electricity generation and demand can be balanced in future years, and 2) enable the composition of the power plant fleet to evolve in ways that minimize the total system cost.

The model solves for the Short and Long Time Scales simultaneously to meet the modeling objective. The objective for this model is to minimize the net present value of all investment and operation costs. Thus, the model will

- dispatch power plants in the Short Time Scale so that the least expensive power plants are turned on first, to balance the hourly generation and demand at the lowest possible cost, and
- build new power plants if the upfront investment cost of constructing those power plants will reduce the total net present value by reducing the cost of the Short Time Scale power plant operation during future time periods.

This objective is subject to a number of constraints and input variables. For example, power plant operational characteristics, fuel prices, power plant construction costs, renewable energy generation profiles, transmission capacity, and many other variables described in the following sections constrain the model's solution.

"Switch" is a unique grid planning model that is built using capacity expansion modeling theory. Switch is developed and maintained by Professor Matthias Fripp at the University of Hawaii, and has been in development since 2012. It is an open source model built on the Python programming language. For more details about the model, its validation, calibration, and equations, see [6].

A UC&D model is an optimization program that makes decisions about the operation of power plants. When compared to a capacity expansion model, a unit-comment and dispatch model operates at short time scales only, but at greater time resolution, simulating every hour of the year.

"PyPSA" is a unique UC&D model that is build using UC&D modeling theory. PyPSA is developed and maintained by T. Brown and a team of other developers. It has been in development since 2018. PyPSA is also able to determine optimal capacities for a given transmission grid. For more details about the model, its validation, and equations, see [7].

2. Time Series

Because a UC&D model operates at only Short Time Scales, it can solve for every hour of the year. In the PyPSA model, we solve the power plant commitment and dispatch for all 8,760 hours of the year.

Because a capacity expansion model operates at both Short and Long Time Scales, it must use simplified time series so that the model is tractable and can be solved. For example, a capacity expansion model that solves a 2020-2050 scenario will not solve for all 8,760 hours of all 30 analysis years. Instead it will use a few representative days for each year, and a few representative years for the whole 30-year time scope.

In the SWITCH model, we use 9 representative days and 4 representative years.

2.1. Representative Days

This model uses 9, 24-hour periods to represent the annual electricity market. Those 24-hour periods include:

- Annual Peak: we use the 24-hour profile of the day with the greatest instance of hourly system demand. The Annual Peak time series is scaled up to represent 3 of 365 days for each model year.
- Annual Net Peak: we use the 24-hour profiles of the two days with the greatest instance of hourly net system demand—i.e. demand minus renewables output. Each of the two Annual Net Peak time series is scaled up to represent 3 of 365 days for each model year.
- Seasonal and Monthly Averages: the electricity demand for each of these 6 profiles equals the average electricity demand of all of the days in that season/month. For example, hour 1 of the March/April profile is equal to the average demand of the first hour of the day for all 61 days in the March/April data. The model uses average profiles for the following seasons and months:
    - July – represents a high-demand summer profile. Scaled up to represent 30 of 365 days for each model year.

- June/August – represents the bulk of summer energy needs. Scaled up to represent 59 of 365 days for each model year.
- January – represents a high-demand winter profile. Scaled up to represent 30 of 365 days for each model year.
- February/November/December – represents the bulk of winter energy needs. Scaled up to represent 89 of 365 days for each model year.
- March/April – represents spring energy needs. Scaled up to represent 59 of 365 days for each model year.
- May/September/October – represents summer shoulder season energy needs. Scaled up to represent 89 of 365 days for each model year.

When compared to a complete, 8,760-hour demand profile, the 9 representative days outlined above have 5% greater annual energy consumption. Figure 6 below compares the 8,760 and 9-representative-day time series using a duration curve—where the demand for each hour of the year is sorted in decreasing order. The peak demand of the 5-representative-day curve is 98.5% of the peak of the 8,760 hour curve. When compared to the 8,760-hour series, the 9-representative-day series has higher demand for the lowest-demand hours of the year but is otherwise very similar.

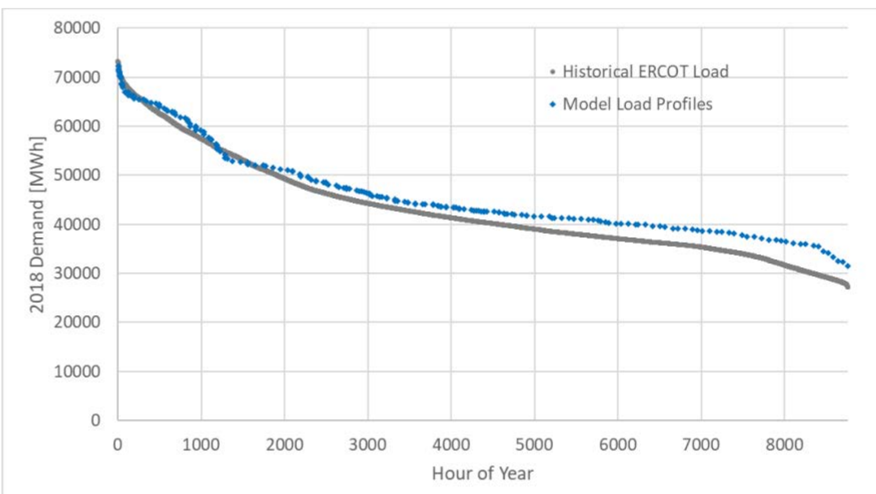

Figure 6: Duration curve of 8,760-hour time series (Historical ERCOT Load) and the 5-representative-day time series used in this model.

2.2. Representative Years

The model simulates these 9 representative days a total of four times each. Each of the four time periods represents a 5-year span: 2020-2025, 2025-2030, 2030-2035, and 2035-2040.

For each of these 5-year time periods, we average the input values across those years. For example, the natural gas price for the 2020-2025 time period equals the average of the 2020, 2021, 2022, 2023, and 2024 forecasted natural gas prices.

3. Generator Data

Our model represents each individual power plant in the ERCOT system. To parameterize each of these power plants, we compile data from a variety of sources as outlined below.

3.1. ERCOT Capacity Demand and Reserve Report, 2018 [13]

Twice a year, ERCOT releases a report that includes some data for all of the operating generators in the ERCOT market. We use this report to gather data on each existing generator's:

- capacity,
- construction year, and
- county.

3.2. Emissions & Generation Resource Integrated Database (eGRID), 2018 [14]

eGRID is maintained by the EPA and contains information about the existing U.S. power plant fleet. We use it to gather data on each ERCOT generator's:

- fuel type, and
- technology type.

3.3. Annual Technology Baseline (ATB), 2019 [15]

The ATB is published annually by NREL and contains a set of assumptions and futures to inform electric sector analyses in the U.S. The data provides operational and cost characteristics for different types of generators projected from 2018-2050. We use it to gather data for each generator's:

- scheduled outage rates,
- forced outage rates, and
- fixed operation and maintenance cost.

We also use the ATB to provide the following data for characterizing new generators:

- capital cost of construction,
- fixed operation and maintenance cost,
- heat rate, and
- roundtrip efficiency for battery charge/discharge cycles.

3.4. Garrison Dissertation, 2014 [16]

In addition to the sources above, which are used broadly for modeling the U.S. power sector across many different regions, we also refer to the dissertation of Dr. Jared Garrison, which contains data compiled specifically for modeling the ERCOT region. Those data include the following.

3.4.1. Heat Rates

Heat rates for existing generators are calculated by diving each generator's monthly fuel consumption by its monthly electricity generation. These data come from the US EIA 923 database. We average these monthly heat rates over multiple years to approximate each generator's full load heat rate.

3.4.2. Startup Costs

Startup costs for existing and new generators are based on data from the Power Plant Cycling Costs report. This report lists startup cost for cold, warm, and hot startups. For the ERCOT power plants, the startup costs for each generator type were selected based on whether that generator type tends to startup from warm or cold conditions.

3.4.3. Min Up and Down Time, Min Output, and Variable Operation & Maintenance Costs

These characteristics come from the assumptions that ERCOT uses for the capacity expansion model used to create the ERCOT Long Term System Assessment report. Based on conversations with different stakeholders, Garrison updated some of these original data for a few of the generator types.

3.5. Coal Retirements

Based on age, the majority of coal plants are expected to retire in Texas by 2035, we force coal retirements for any coal plants that have been operating for 43 years or longer, based on historical lifetimes and recent retirement rates. This requirement has the following impact on overall coal capacity:

- 2018: 13.1 GW
- 2020-2025: 11.5 GW
- 2025-2030: 5.5 GW
- 2030-2035: 3.3 GW
- 2035-2040: 3.0 GW

4. Wind and Solar

4.1. Profiles

We use hourly wind and solar generation profiles for hundreds of sites around ERCOT. These generation profiles were developed by AWS TruePower for ERCOT and are available for public download [17].

The hourly profiles are simulated using historical weather data. A generation profile is created for each existing wind and solar site in ERCOT along with many potential sites where wind and solar capacity have not yet been installed.

For developing future wind capacity, we let the model expand the capacity of simulated sites (modeled at a hub height of 90m) and existing sites with hub heights of 80m or greater. For existing sites with hub heights below 80m, we use their profiles to represent existing wind generation resources available for dispatch, but do not let the model expand their capacity. For counties without existing or simulated wind generation, we average the profiles of sites with similar wind resources in neighboring counties.

For developing future solar capacity, we let the model expand the capacity of the simulated sites. Texas solar resources [18] generally improve as one travels west. We observe this trend in the capacity factors of the simulated solar sites, but not consistently in the capacity factors of the existing solar sites. Thus, we use the profiles of existing sites to represent existing solar capacity resources available for dispatch, but do not let the model expand their capacity.

4.2. Site Limits for Wind and Solar Capacity

Since wind and solar plants require a significant amount of real estate, we limit the amount of wind and solar development that the model can build in each Texas county.

For solar, we assume single-axis tracking arrays built at a density of 30 MW/km^2 (77.7 MW/mi^2). [19]

For wind, we use the appendix data from [20] to divide the total Texas wind capacity by the total developed land area of that wind capacity to get a density of 7.14 MW/mi^2.

We then multiply these development densities by the square mileage of land in each county that is available for development[6]. The result is the maximum amount (MW) of wind and solar capacity that could be built in the developable land in each county.

The wind limit is, on average, 6.5 GW per county. But that capacity can only be realized if all of the county's available land area has suitable wind resources. However, in most counties, the wind resource quality varies across the county's geography. To account for this, we use data from [21] to estimate the amount of land in each county that has wind resources with wind speeds of 7.0-7.5, 7.5-8.0, and 8.0+ m/s. We use those estimates to cap the amount of capacity that each wind site may develop, depending on its capacity factor.

The solar limit is, on average, 70.4 GW per county. In practice, this solar limit never constrains the model. Thus, we assume that, because of its density, solar development has little impact on wind development—i.e., if a county builds many GW of solar capacity, this requires a relatively small amount of land and we assume that it does not meaningfully diminish the county's wind capacity limit.

4.3. Annual Limits for Wind and Solar Capacity Growth

Wind and solar development are also limited by materials supply chains, manufacturing capabilities, and construction capabilities. To capture this, we impose an annual limit on how much wind and solar can be built in the model.

For both wind and solar, we establish a baseline limit on GW/year that can be installed. Then, assuming that these limitations will increase with GDP, we scale the installation limits up according to the forecasted Texas GDP growth through 2050 [22].

For the baseline wind limit, we take data on annual wind development in Texas from 2009-2019 [23] [24]. We take the average of these numbers—1.45 GW/year—as the baseline for the wind development limit.

[6] Personal communication with the University of Texas at Austin Bureau of Economic Geology.

We assumed the same deployment rate for utility-scale solar.

4.4. Land Lease Rates for Wind and Solar

The fixed operating cost of each wind and solar site varies depending on which county it is built in. To accomplish this, we first compile a lease rates for rangeland, native pasture, and hunting leases in 33 Texas regions [25]. Then we normalize those lease rates, multiply them by wind and solar lease costs from [26], and assign them to the counties contained in each region. Note that wind land lease costs vary from 1,100 to 24,500 $/MW-year with an average of 8,960 and solar land lease costs vary from 630 to 14,400 $/MW-year with an average of 8,960.

We then use these land costs to adjust the fixed operation and maintenance costs from section 3.3 by:

- for wind sites: subtracting the average wind land lease cost from the wind FOM. Then adding back the county-specific wind land lease cost.
- for solar sites: because the ATB does not include solar land lease costs in its solar FOM, we simply add the county-specific solar land lease cost to the ATB FOM.

4.5. Tax Credits

We give solar an investment tax credit of 10% by reducing its overnight capital costs by 10%.

5. Transmission

As electricity travels from region to region it incurs losses and must not exceed the capacity of the transmission lines. The model can increase the capacity of the existing transmission lines by paying the capital cost to build new lines.

5.1. Losses

We assume losses of 1% per 100 miles of transmission. This aligns with the assumption used by the National Renewable Energy Laboratory's ReEDS model [27]—a capacity expansion model of the continental United States.

5.2. Regions and Capacities

The model comprises 16 regions with transmission capacity between many of the regions' borders. The regions and transmission locations were determined using geographic transmission data from the Department of Homeland Security [28].

Existing transmission capacities were determined by running the historical 2018 hourly load and generation in a power flow model [29].

- Hourly Load: see section 7.1 and 7.2.
- Hourly Thermal Generation: comes from aggregating CEMS data to the county level, and the aggregating those county-level generation profiles up to the transmission-region level
- Hourly Wind and Solar Generation: see section 4.1

- Nuclear Generation: we assume constant nuclear generation at 95% of total capacity to match the annual nuclear generation capacity factors.
- Existing transmission: we connect regions with transmission lines if they have existing transmission connections already. And we add multiple lines between regions when there are multiple 345-kV lines that connect those regions in the existing transmission grid. For example, we connect the 1Dallas—10WichitaFalls regions with (3) 345-kV lines based on their existing transmission connections, but connect the 15FortStockton—16Pecos regions with (1) 345kV line.

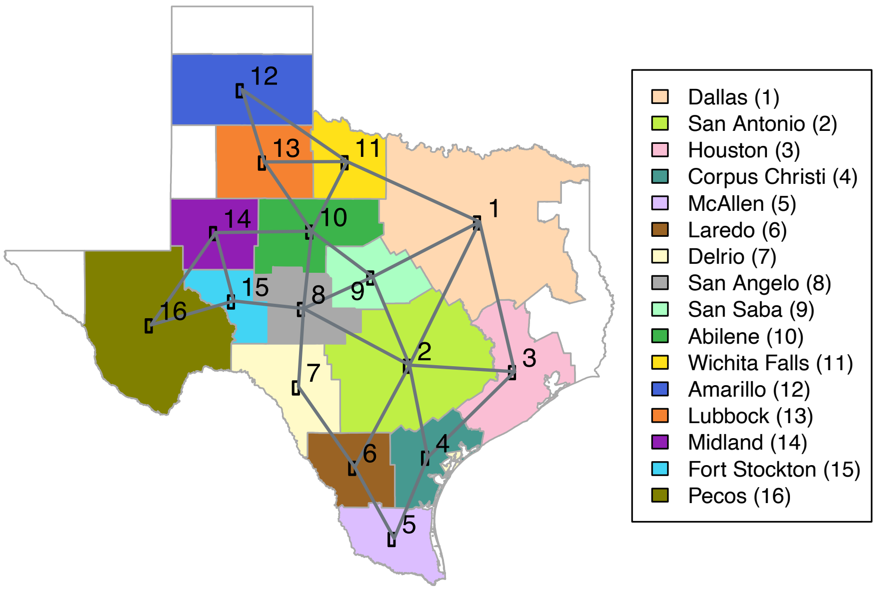

Figure 7: The 16-zone ERCOT model and transmission network used in this analysis.

5.3. Construction Cost

Transmission construction costs are based on data from the Competitive Renewable Energy Zones (CREZ) project—a large-scale transmission construction project carried out in ERCOT from 2008-2013. We use a transmission construction cost of 1500 $/MW-mile (932 $/MW-km) as described in [8].

6. Fuel Prices

Fuel price data come from the EIA's 2020 Annual Energy Outlook (AEO) [30]. This report contains future projections out to 2050 of energy consumption, emissions, and fuel prices. We use the

- forecasted AEO coal prices for our model's subbituminous coal prices,
- forecasted AEO coal prices plus 0.72 $/mmBtu for our model's lignite coal prices, and
- forecasted AEO natural gas prices for our model's natural gas prices.

The lignite prices are increased by 0.72 $/mmBtu so that the average of the forecasted 2020-2030 prices equal the average of the historical 2015-2020 Texas lignite prices [31].

7. Load

7.1. Load Data

We use 2018 hourly load data provided by ERCOT [32]. This load data is separated out for each ERCOT's 8 weather regions.

7.2. Scaling Load Data by Region

We scale this 8-region ERCOT data to our 16-region transmission model in two steps.

First, we distributed the ERCOT load down to the county level by assuming that county population is directly related to energy consumption. That is, if Region 1 has a demand of 12,000 MWh in a specific hour, and County 1—one of a number of counties in Region 1—has 15% of the population of Region 1, then we assume that County 1 also represents 15% of that hourly demand—or 1,800 MWh. The result is an hourly 2018 load profile for each Texas county.

Second, we aggregate these county-level load profiles up the regional level using the region boundaries in our model. The result is an hourly 2018 load profile for each of our 15 transmission regions.

7.3. Load Growth

We assume that load increases at a rate of 1.8% annually. This load growth rate was determined by calibrating the model's future loads against the energy forecasts in Figure 2 of the 2020 ERCOT System Planning Forecast [33].

We implement this load growth assumption by starting with our baseline hourly 2018 load profiles for each region and multiplying every hourly demand datum by 101.8% for each year after 2018.

7.4. Electric Vehicles

We include electric vehicle energy demand using the following steps.

First, we use a 24-hour profile from the LTSA that forecasts ERCOT electric vehicle charging behavior in 2033. We assume that electric vehicles will charge according to this 24-hour pattern for each day of the year.

Second, we scale the profile up and down for different model years. We assume that the charging pattern scales linearly, where the electric vehicle load in 2015 equals zero. Under this assumption, the electric vehicle load in 2015 is zero, in 2024 is 50% of the 2033 ERCOT profile, in 2042 is 150% of the 2033 ERCOT profile, etc.

Third, we distribute the total electric vehicle charging profile amongst the 16 transmission regions. We take the 2018 population for each of the transmission regions and divide by the total Texas population to calculate that region's load fraction. Then we multiply each region's load fraction by the total EV charging profile for each year to produce each region's hourly EV profile for each year.

Finally, we add the EV charging profile to each region's hourly load profile.

7.5. Distributed Solar

We simulate distributed solar generation for each region and subtract it from that region's hourly load. That is, the model does not treat distributed solar as power plant that can be dispatched, but as a distributed resource that reduces the amount of load that the model's power plants must provide.

First, we create hourly 2018 solar generation profiles for the largest city in each region using the NREL System Advisor Model (SAM) [34]. The SAM model uses historical weather and solar insolation data to calculate the hourly electricity generation of a photovoltaic panel depending on that panel's orientation, tilt, efficiency, and other parameters. We use the default SAM settings for the solar panel—180 degree azimuth, 20 degree tilt, 96% inverter efficiency, and 14.08% system losses. The result is a normalized, hourly 2018 solar generation profile for each of the 16 transmission regions.

Second, we scale these solar profiles up to match the forecasted capacities of distributed solar in each region. We calculate the forecasted solar capacities in two steps:

> Step 1: we forecast the total amount of distributed solar in all of ERCOT. We use 5 GW of distributed solar for 2033, based on Table I.1 of the ERCOT 2018 Long Term System Assessment (LTSA) [5]. Similar to the method used for electric vehicles, we assume that the distributed solar profile scales linearly, where the distributed solar in 2015 equals zero. Under this assumption, the distributed solar in 2015 is zero, in 2024 is 2.5GW (50% of the 2033 ERCOT profile), in 2042 is 7.5GW (150% of the 2033 ERCOT profile), etc.
>
> Step 2: we spread the distributed solar capacity amongst the 16 transmission regions. As for electric vehicles, we take the 2018 population for each of the transmission regions and divide by the total Texas population to calculate that region's fraction. Then we multiply each region's fraction by the total distributed solar capacity for each year to produce each region's distributed solar capacity for each year.

8. Financial

The Switch model uses an interest rate and discount rate for various financial calculations. We assume a discount rate equal to a weighted average cost of capital (WACC) of 7.17% and an interest rate of 6.01%. These align with the assumptions of the NREL ATB [15].

Appendix B

Bitcoin Assumptions		Today	Scenario 3 2030	Scenario 4 2030
Bitcoin Price	USD/BTC	37,500.0	150,000.0	150,000.0
Global Hashrate	EX/s	120.0	2,000.0	2,000.0
Reward per Block	BTC	6.25	1.56	1.56
Fees per Block	BTC	0.70	0.60	0.60
Bitcoins per Day	#	1,000.8	311.4	311.4
ExaHash per Bitcoin	#	10,359.71	554,913.29	554,913.29
Value per ExaHash	USD	$3.62	$0.27	$0.27

		Scenario 3 W/Th	Scenario 4 W/Th
Tier 1 ASIC - Newest Generation	EX/MWh	9.40	10.75
Tier 2 ASIC - Older Generation	EX/MWh	22.40	19.75
Tier 3 ASIC - Oldest Generation	EX/MWh		31.75