Architecture for Web-Based Visualization of Large-Scale Energy Domains

Graham Johnson* NREL

Sam Molnar† NREL

Nicholas Brunhart-Lupo‡ NREL

Kenny Gruchalla§ NREL

Figure 1: Snapshot of the 100 megapixel high-resolution display with an interactive visualization in the browser. Two synthetic energy model topologies are shown: an electrical transmission system (blue lines) and a corresponding distribution system (orange points) in the San Francisco Bay area. These two models have over 12 million combined features.

Abstract

With the growing penetration of inverter-based distributed energy resources and increased loads through electrification, power systems analyses are becoming more important and more complex. Moreover, these analyses increasingly involve the combination of interconnected energy domains with data that are spatially and temporally increasing in scale by orders of magnitude, surpassing the capabilities of many existing analysis and decision-support systems. We present the architectural design, development, and application of a high-resolution web-based visualization environment capable of cross-domain analysis of tens of millions of energy assets, focusing on scalability and performance. Our system supports the exploration, navigation, and analysis of large data from diverse domains such as electrical transmission and distribution systems, mobility and electric vehicle charging networks, communications networks, cyber assets, and other supporting infrastructure. We evaluate this system across multiple use cases, describing the capabilities and limitations of a web-based approach for high-resolution energy system visualizations.

Index Terms: Visualization—Visual Analytics—Human Computer Interaction—Grid Modernization

1 Introduction

Historically, power systems transmitted electricity in a single direction from synchronous utility-owned generators to easily forecastable customer loads, requiring few assets to be measured or modeled to analyze and control a system. However, renewable energy is driving us to reconsider how we model, visualize, and ultimately support decision-making about the energy system, requiring analyses of higher-fidelity information. Inverter-based resources (IBRs), such as wind or solar farms, change the dynamics of the power system [23], requiring high temporal fidelity to capture the system behavior. Additionally, these generation resources are no longer exclusively owned and controlled by utilities; increasingly, they are Distributed Energy Resources (DERs) installed closer to end loads. Moreover, customer loads are changing as more elements traditionally powered by fossil fuels, such as transportation or heating, undergo electrification [38]. As a result, energy studies and operational tasks can require information on individual buildings or even individual plug loads within a system.

Modern energy systems and supporting infrastructure are more intertwined with other large-scale systems than ever before. Individual systems were previously modeled in isolation, but now are considered alongside adjacent domains so that the diverse groups of stakeholders can learn how actions in one domain affect the comprehensive system. For example, the placement of electric vehicle (EV) charging equipment is a decision that crosses multiple domains and can involve many stakeholders. In this setting, the location of an EV charging station affects electrical grid planners and operators, transportation authorities, community planners, and EV charging installers. We must consider travel patterns, plug-in times, EV charging profiles, and charging locations across a regional extent at a distribution-bus fidelity to understand how these installations will impact the power system [29].

The growing scale and complexity of multi-domain data for modern energy systems underscores the need for high-resolution visualization systems that can facilitate the necessary analysis and management for making decisions with this data. In this paper, we describe a web-based architecture that supports a 100-megapixel high-resolution large display wall for direct visualization of these multi-domain datasets. To effectively visualize datasets with mil-

*e-mail: graham.johnson@nrel.gov
†e-mail: sam.molnar@nrel.gov
‡e-mail: nicholas.brunhart-lupo@nrel.gov
§e-mail: kenny.gruchalla@nrel.gov

lions of features, we leverage data processing steps that include vector tiling pipelines, data transformations into efficient buffer formats, and layer definitions with high-performance WebGL2 rendering libraries. Vector tiling and efficient buffer formats need to work together to support the largest time-varying datasets. While vector tiles can enable scalable and interactive visualizations by statistically sampling large and spatially dense static datasets to generate appropriate levels of detail [24], merging time-series values to existing tiles is computationally expensive and limits interactivity. Conversely, encoding large and spatially dense features into efficient buffer formats can handle fast time-series updates but leads to significant over-plotting, which constrains interactivity when dataset feature sizes approach multiple millions.

2 Background

The energy sector has employed various visualizations for power systems, with the traditional one-line diagram serving as the standard representation of an electric network's static structure. This schematic uses symbols and lines to depict electrical components and their connections [35]. However, one-line diagrams are not scalable and generally lack geographic context, making it difficult to understand geographically-driven phenomena such as renewable energy resource distribution, electric vehicle travel patterns, weather-related events, and the interplay with non-power infrastructure. Researchers commonly use glyphs like ellipses and arrows to represent various system properties, including load and generation, in a geographical context [8, 12]. However, overplotting becomes a limitation at any significantly large scale [22]. Two-dimensional colored contour maps are a common aggregation technique for visualizing power systems data [26, 37]. However, colored contour maps can misrepresent data in dense, topologically complex distribution systems [9]. Lyons-Galante et al. [19] have proposed tessellation-based aggregation techniques that more faithfully represent the statistical distributions of dense electrical bus values, but these have yet to be empirically evaluated with power systems stakeholders.

Considering the scales of the emerging energy systems, analyses based on aggregation techniques alone are likely to be insufficient. Traditional electrical control rooms utilize large-scale displays with millions of pixels composed from multiple independent screens, which generally display multiple independent views [11]. Similar large-scale display designs have been found to be beneficial for collaboration [1, 34] and providing space to think [3]. Like many use cases in the visualization literature [5], our system creates a single, integrated visualization utilizing all the pixels as a single unified display. Large-scale high-resolution displays offer significant advantages, such as improved access to information within multi-scale, heterogeneous data sets [4]. For large, multi-scale data, zoom and pan interactions on traditional displays can be insufficient, whereas high-resolution displays are a tool to see the context and details simultaneously [14, 30]. These high-resolution displays may provide perceptual and cognitive benefits as empirical studies have shown increased user performance on spatial tasks, mental mapping, and memory on large-scale high-resolution displays compared to traditional desktop displays, even when controlling for visual angles between the display types [32].

2.1 Domain Models

Modern energy domain models span power systems (e.g., transmission and distribution electric systems with DER assets such as rooftop solar installations), transportation and mobility systems (e.g., vehicle trajectories and EV charging data), and cyber systems (e.g., communications, sensing, and supporting infrastructure) to name a few. A common starting place for many energy analyses in the context of grid modernization involves leveraging transmission and distribution electric system models [17]. For a reference of distribution model sizes and scales, the SMART-DS models [27] represent synthetic distribution electric systems and are shown in Table 1. For

Metric	SAF	GSO	SFO
Number of buildings	38,590	70,554	**2,265,594**
Medium voltage	31	144	**1,535**
Low voltage	38,558	70,407	**2,264,014**
Number of consumers	84,169	134,882	**4,299,805**
Medium voltage	31	495	**11,503**
Low voltage	84,138	134,387	**4,288,297**
Number of buses	88,886	181,631	**4,916,869**
Number of electrical nodes	168,005	375,334	**9,868,205**
Number of transmission substations	2	7	**148**
Number of sub-transmission substations	8	31	**632**
Number of distribution transformers	11,300	25,933	**559,151**
Length of power lines (km)	1,921	4,534	**116,837**
High voltage (sub-transmission)	27	167	**4,128**
Medium voltage	966	2,302	**64,460**
Low voltage (secondaries)	928	2,107	**48,249**
Number of feeders	28	98	**2,236**

Table 1: Details for three reference SMART-DS models [28].

example, the San Francisco Bay area model captures roughly 4.3 million consumers and contains close to 10 million electrical nodes and 116,837 kilometers of distribution lines. Traditionally, visualizing datasets of this size involves aggregation techniques such as vector tiling to statistically sample representative levels of detail based on client view-port location and zoom levels. Vector tiling methods can generate faithful representations of large datasets [31], and have been shown to outperform similar raster-based tiling methods [24]. However, vector tile approaches can require computationally expensive pre-processing steps as datasets grow and have interactive limitations when attempting to style assets based on dynamic time-series values. Conversely, the availability of high-resolution large display environments provides opportunities to visualize more of a domain model's assets simultaneously. However, there are multiple considerations when choosing visualization data formats and rendering parameters to reduce overplotting, enable dynamic time series, and allow for smooth interactions [36]. Additionally, geospatial indexing techniques can be advantageous in addressing the challenges between vector tiles and raw buffer approaches [6].

Figure 2: H3 hex aggregation of the San Francisco SMART-DS model feeders. Hexes are colored by electrical feeder and overlaid with the distribution buses.

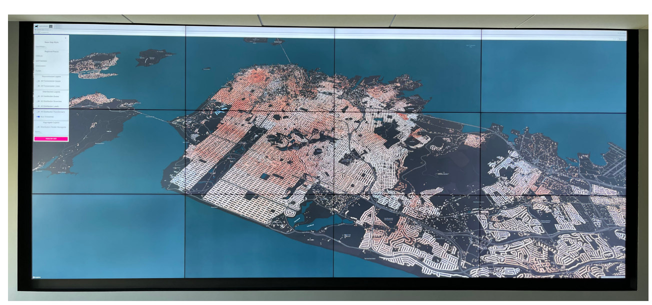

Figure 3: San Francisco SMART-DS model on the Insight Center 100 megapixel display wall. Assets are styled from time-series values with a diverging color-scale based on nominal per-unit voltage (white) and +/-5%, with red as higher and blue as lower voltages.

3 ARCHITECTURE

We are visualizing high-performance computing (HPC) energy system simulation data on a 100-megapixel, high-resolution display wall composed of twelve 4K displays with a total resolution of 15360 by 6480 pixels. A single Linux server uses NVIDIA Mosaic to drive the high-resolution large display with three Quadro RTX 8000 GPUs and 1TB of CPU RAM. The server connects to the NREL HPC data center [25] via a 100-gigabit fiber network, where the large-scale models are run and resulting data are stored. Once data has been processed through our pipeline, we use Google Chrome to render interactive web-based visualizations.

3.1 Data Processing

To support web-visualization of large-scale data, we begin our visualization pipeline with data processing steps of vector tiling, hierarchical geospatial aggregation, and buffer formatting.

3.1.1 Vector Tiling

When dealing with especially large or dense datasets, we generate vector tiles using Mapbox's *Tippecanoe* [20] command line interface. Written in C++, the pipeline can process large geospatial datasets into a custom *.mbtiles* file format or as a directory of corresponding protobuf (*.pbf*) files at specified zoom levels. There are many runtime parameters that need to be properly chosen–such as balancing sampling in datasets with high densities, managing representations over large geographical areas, and smoothing polygons with many features, to name a few. For example, the amount of data included in each tile at various zoom levels can significantly impact interactive performance and may even cause client browsers to crash due to memory overload. In our case, we process large energy domain models (e.g., the SMART-DS San Francisco model) through *Tippecanoe* and manage the tile size by extracting individual layers of asset types (e.g., electrical buses, lines, transformer, regulators). These layers can be rendered as references, upon which we can dynamically assign styles from data properties and simulation outputs. Tippecanoe's outputted directory of tiles can be hosted from a static file tile server that runs locally, on networked servers, or from a cloud resource.

3.1.2 Geospatial Aggregation

Often, there are properties of domain model assets that naturally group hierarchically. For example, in transmission and distribution electric system settings, a bottom-up construction would start with individual assets (e.g., electrical buses, transformers, lines, etc.), which group into distribution *feeders*. These *feeders* can be grouped to corresponding transmission *substations*, which are subsequently grouped into a utility's *service area*, onto *balancing authorities*, and eventually into *interconnections*. Rather than pre-computing convex or concave hulls, or Voronoi tessellations that can lead to elongated and irregularly shaped units, we process assets through an iterative geospatial indexing library, *H3* [6]. This process generates hexagonal partitions at desired levels of detail, while enabling fast hierarchical (parent-child hexagon) computations and on-the-fly styling. These features allow the same generated covering to be efficiently and dynamically styled. For example, we can dynamically style all hexagon bins that group into *feeders*, which *feeders* group together to a transmission *substations*, and generate boundaries for utility *service areas*. A sample of this type of covering is shown in Figure 2. Moreover, the styles are not limited to an asset's static properties, but can also be updated based on any function of their contents' time-series values (e.g., voltages and violation counts). While this geospatial aggregation allows for efficient client rendering, without needing to draw and style large amounts of individual assets, it also serves the purpose of anonymizing the underlying model and data when required. For example, when utilities share distribution feeder data, they often need to keep their feeder topologies and asset geospatial locations private. Passing this data through a *H3* partitioning process at an acceptable resolution allows researchers to dynamically compute aggregate values in ways appropriate for sharing information with stakeholders.

3.1.3 Buffer Formatting

When focusing on energy model simulations that prioritize visualizing frequent time-series updates, directly modifying asset styles in the vector tile format can be challenging and computationally costly–ultimately preventing levels of smooth interactivity from users. In settings where simulation data is regularly changing or needs to be animated, data processing steps that prioritize efficient memory formats become paramount for achieving smooth style updates. This memory formatting depends closely on the chosen rendering pipeline. Our architecture leverages the highly-performant *Deck.GL* library [36], wherein the specification for binary data formats is well-defined and can be immediately used for primitive layer definitions. This performance-optimization method makes heavy use of transforming raw data representations into JavaScript TypedArrays formats (e.g., Float32Array and UInt8Arrays etc.). These formats allow the library to minimize CPU overhead and pass data directly onto the GPU for rendering. When paired with web workers [2], this process can efficiently compute and no-copy transfer memory objects of updated asset positions, styles, and properties without impacting the main thread, which has allowed us to support interactive datasets with millions of assets.

3.2 Front-End

Our data pipeline feeds our interactive visualizations through the use of modern web frameworks, allowing researchers to design and develop applications that can scale across devices, including mobile, desktop, and high-resolution large displays. In this paper's context, we make use of the *Next.js* framework that wraps the *React.js* library for state management. We defined custom web-workers to handle connections and requests to external data APIs, load and parse local data into proper (binary) formats, and transfer updated buffers back to the main application thread to minimize memory footprints and unnecessary data copying. Additionally, we leverage *Deck.GL* specifications for formatting data to be used within layer definitions for vector tiles, GeoJSON, trips, and primitive points/lines. When combined with the data processing approaches above, our architecture supports interactive exploratory visualization for a wide array of energy domain models.

4 EVALUATION

We applied our architecture in a variety of different modeling contexts including power systems analyses, renewable resource assessments, situational awareness platforms, and cyber studies [13,15,21]. We outline two of these applications here. At a high-level, we design and implement our applications with domain users by considering a handful of properties: number and density of assets, level of interactivity, dynamic styling, time-series frequency, and desired frame rates. A current challenge arises when a large and dense dataset has frequent time-series updates that require dynamic styling. Using the vector tile approach, reference layers can be generated through *Tippecanoe* upon which time-series data, either requested from an external API or parsed in memory (e.g., parquet and csv), can be used to re-draw the defined *Deck.GL* layers. However, this approach is CPU intensive as *Deck.GL* must re-compute buffers for all tile assets in each animation frame and quickly limits interactivity. Conversely, time-series data can be baked into the tiles during their generation–outputting a new directory for each time-step. The rendering client can subsequently point to the corresponding directory at the current time-stamp. While pre-processing is computationally intensive and can require substantially more back-end tile storage, this approach moves the bottleneck from the CPU to network I/O, as each new time-step will request a fresh set of tiles and unable to leverage the browser's cache. In our experience, if a dataset has 5 million features, regardless of density, we can utilize the *Deck.GL* binary buffer formats to bypass the CPU-intensive bottleneck to achieve sub-second re-renders with the map remaining interactive. Ultimately, our applications leverage a combination of the data processing steps–vector tiles for rendering references of large and dense models, hierarchical aggregations for rendering group inclusions and statistics, and binary buffer formats for rapidly re-rendering models, selected assets in focus-areas, or assets that are flagged in simulation-defined events.

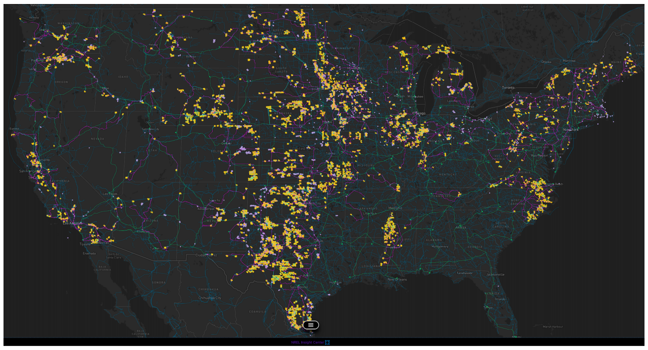

Figure 4: Example visualization of wind resource availability involving electrical transmission system components, existing wind turbine arrays, modeled future wind sites, and expanded capacity transmission lines across the continental United States [18].

4.1 Regional Distribution Systems Analysis

We used our architecture to analyze the adoption of DER and the impacts of EV charging on electric distribution systems in large metropolitan areas and the surrounding rural regions, with a detailed resolution down to the customer meter. The largest of these studies was conducted on the San Francisco Bay region [29]. The SFO model has over 12 million features and a density on the order of 1000 assets per half-mile radius–see Table 1 and Figure 3. This corresponded to a directory of 2198 feeder GeoJSON files at 8.7GB that were used as tiling input files. Due to feature densities and size, we processed each asset type (e.g., buses and lines) into its own set of vector tiles. We ran *Tippecanoe* on a Linux server with two Intel Xeon Gold 2.40GHz CPUs, two NVIDIA A40 GPUs, and 1 TB of RAM. *Tippecanoe*'s runtime and output directory size are highly dependent on the flags included in the processing. Our chosen method used the following flags for generating tiles for each asset type: max threads of 96, no line simplification, no feature limits, no tile size limit, no tile compression, and a zoom levels 9 - 14. Processing the input GeoJSONs for the bus assets took our system 5 minutes 46 seconds and produced a tile directory of 470MB. When using no optional flags for the same bus assets at the same zoom levels, *Tippecanoe* produced a directory of 198MB in 6 minutes 40 seconds. As a reference, processing all asset types together, with the same flags as above, *Tippecanoe* produced a 2.9GB tile directory in 6 minutes 50 seconds. To understand distribution impacts in DER scenarios, we ran power flow simulations, generating voltage and power time-series on each electrical bus. These time-series data were mapped to the reference tiles using *Deck.GL*, but are limited to infrequent updates to maintain interactivity. As a test, we encoded the 4.9 million bus assets into binary buffers using custom *web workers* and could maintain interactivity with a sub-second time-series update frequency; when the dataset is on the order of 500K assets, we achieved interactive re-renders in less than 250ms. Without *web workers*, the client is split between supporting UI events and processing the time-series data into binary buffers–ultimately blocking the user from interacting. Additionally, we leveraged *H3* in these studies to generate multiple resolutions of aggregate model representations–showing *feeders*, *substation mappings*, and aggregations of asset time-series at *H3* resolutions 4 - 12. *H3* aggregations enable on-the-fly styling of assets, as well as efficient representation of statistical trends (a sample is shown in Figure 2).

4.2 Renewable Energy Resource Assessments

The technical potential of renewable generation technologies across the continental United States [18] is based on several factors, including the availability and quality of renewable resources, technical system performance, topographic limitations, economics, and environmental and land-use constraints such as siting ordinances. The goal of this effort was to enable researchers to explore large amounts of scenario data without having to view layers independently. The data, while on the order of 1 million features, was considerably less dense than the SMART-DS example; allowing *Tippecanoe* to generate tiles, following the same process described in the previous section. The data was separated into eleven different input layers–such as transmission lines, substations, existing wind turbines, *ReEDS* model supply curves [10], modeled reinforcement lines, spur lines, and modeled future wind turbines. Each layer was processed into its own set of vector tiles using *Tippecanoe*, turning over 7 gigabytes of original features into a directory of vector tiles of 1.5 gigabytes in about 20 minutes. While it should be possible to employ the binary buffer format to render all features in all layers simultaneously, we opted for tile generation due to its simplicity, minimal data preparation, and re-use of layer definitions within *Deck.GL*. A sample of this application is shown in Figure 4.

5 Limitations and Future Work

Web technologies have steadily advanced through the evolution of WebGL, WebGL2, and WebGPU graphics APIs, and are capable of rendering millions of features in browsers. We have shown how we interactively render energy domain datasets with feature counts ranging from hundreds of thousands to 12 million elements, supporting both static data and time-series re-renders up to 250ms in a web browser running on a 100MP high-resolution display. However, we do have performance challenges and scalability limitations. There is a trade-off between vector tile approaches that need to support frequent time-series re-renders and binary buffer formats that are only supported on primitive layer types. While there has been community progress in developing a binary buffer format compatible with vector tiles, its evaluation remains open for future work. Additionally, we still need to invest in better aggregation algorithms [19] and targeted analytics. Our architecture helps provide an overview of the system but still requires application-specific development of visual analytics to answer user questions. At a lower level, Chrome-based browsers can only support 4GB file sizes, and maximum WebGL buffers can easily be exceeded if applications are being used in multiple tabs. While this can be solved by using separate browser windows with their own cache directories, large and dense datasets must be thoughtfully processed to avoid loss of WebGL contexts and related browser memory constraints. Our 100MP display web applications are ultimately limited by the current WebGL2 specification. We anticipate the continued adoption of WebGPU will support even larger datasets [33] and enable more general compute capabilities in the near future. Additionally, the standard mouse and keyboard interfaces lack the precision and fluidity to interact with large, dense pixel spaces efficiently, motivating a move beyond the traditional interfaces [16]. We are exploring integrating touch devices as controllers of applications on high-resolution large displays and, more generally, considering real-time interactive collaboration across distributed clients [7]. While web-based visualization platforms have current limitations, they provide a rich and capable ecosystem for interactively rendering a wide array of energy domain datasets. Web-based visualizations have an exciting future ahead–keeping pace to support analyses of growing datasets, offering more efficient GPU computing, faster render loops, deeper user interactions, and simulation-on-demand capabilities.

ACKNOWLEDGMENTS

This work was authored by the National Renewable Energy Laboratory, managed and operated by Alliance for Sustainable Energy, LLC for the U.S. Department of Energy (DOE) under Contract No. DE-AC36-08G028308. The research was performed using computational resources sponsored by the Department of Energy's (DOE) Office of Energy Efficiency and Renewable Energy (EERE) located at the National Renewable Energy Laboratory, and used resources at the Energy Systems Integration Facility, which is a DOE EERE User Facility. The views expressed do not necessarily represent the views of the DOE or the U.S. Government. The U.S. Government retains and the publisher, by accepting the article for publication, acknowledges that the U.S. Government retains a nonexclusive, paid-up, irrevocable, worldwide license to publish or reproduce the published form of this work, or allow others to do so, for U.S. Government purposes.

REFERENCES

[1] U. Afzal, A. Prouzeau, L. Lawrence, T. Dwyer, S. Bichinepally, A. Liebman, and S. Goodwin. Investigating cognitive load in energy network control rooms: Recommendations for future designs. *Frontiers in Psychology*, 13, 2022. doi: 10.3389/fpsyg.2022.812677

[2] T. Alatalo, T. Koskela, M. Pouke, P. Alavesa, and T. Ojala. Virtualoulu: collaborative, immersive and extensible 3d city model on the web. In *Proceedings of the 21st International Conference on Web3D Technology*, Web3D '16, p. 95–103. Association for Computing Machinery, New York, NY, USA, 2016. doi: 10.1145/2945292.2945305

[3] C. Andrews, A. Endert, and C. North. Space to think: large high-resolution displays for sensemaking. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '10, p. 55–64. Association for Computing Machinery, New York, NY, USA, 2010. doi: 10.1145/1753326.1753336

[4] C. Andrews, A. Endert, B. Yost, and C. North. Information visualization on large, high-resolution displays: Issues, challenges, and opportunities. *Information Visualization*, 10(4):341–355, 2011. doi: 10.1177/1473871611415997

[5] I. Belkacem, C. Tominski, N. Médoc, S. Knudsen, R. Dachselt, and M. Ghoniem. Interactive visualization on large high-resolution displays: A survey. *Computer Graphics Forum*, Apr. 2024. doi: 10.1111/cgf.15001

[6] I. Brodsky. H3 uber's hexagonal hierarchical spatial index., 2018. https://eng.uber.com/h3/.

[7] N. Brunhart-Lupo, B. Bush, K. Gruchalla, K. Potter, and S. Smith. Collaborative exploration of scientific datasets using immersive and statistical visualization. *Proceedings of the 2020 Improving Scientific Software Conference (No. NCAR/TN-567 +PROC)*, 2020. doi: 10.26024/p6mv-en77

[8] A. L. Figueroa-Acevedo, C. Tsai, K. Gruchalla, Z. Claes, S. Foley, J. Bakke, J. Okullo, and A. J. Prabhakar. Visualizing the impacts of renewable energy growth in the U.S. Midcontinent. *IEEE Open Access Journal of Power and Energy*, pp. 1–1, 2020. doi: 10.1109/OAJPE.2020.2967292

[9] K. Gruchalla, S. Molnar, and G. Johnson. Reevaluating contour visualizations for power systems data. *IEEE Transactions on Smart Grid*, 14(6):4877–4887, 2023. doi: 10.1109/TSG.2023.3252468

[10] J. Ho, J. Becker, M. Brown, P. Brown, I. Chernyakhovskiy, S. Cohen, W. Cole, S. Corcoran, K. Eurek, W. Frazier, P. Gagnon, N. Gates, D. Greer, P. Jadun, S. Khanal, S. Machen, M. Macmillan, T. Mai, M. Mowers, C. Murphy, A. Rose, A. Schleifer, B. Sergi, D. Steinberg, Y. Sun, and E. Zhou. Regional energy deployment system (ReEDS) model documentation (version 2020). 6 2021. doi: 10.2172/1788425

[11] E. Hénique, S. Lindegaard, and B. Hunt. Design of large and complex display systems. In T. Ivergard and B. Hunt, eds., *Handbook of Control Room Design and Ergonomics*, p. 50. CRC Press, 2nd ed., 2008. doi: 10.1201/9781420064346

[12] I. Idehen, B. Wang, K. Shetye, T. Overbye, and J. Weber. Visualization of large-scale electric grid oscillation modes. In *2018 North American Power Symposium (NAPS)*, pp. 1–6, 2018. doi: 10.1109/NAPS.2018.8600547

[13] M. Ingram, A. Florita, M. Martin, K. Gruchalla, X. Fang, M. Cai, G. Johnson, N. Guba, A. Hasandka, M. Mooney, M. Lunacek, N. Gilroy, T. Langan, S. Caruso, R. Cruickshank, B.-M. Hodge, and M. Marković. Situational awareness of grid anomalies (SAGA) for visual analytics—near-real-time cyber-physical resiliency through machine learning. Technical Report NREL/TP-5D00-85187, National Renewable Energy Laboratory, Golden, CO, 2023. doi: 10.2172/1957540

[14] A. Ion, Y.-L. B. Chang, M. Haller, M. Hancock, and S. D. Scott. Canyon: Providing location awareness of multiple moving objects in a detail view on large displays. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '13, p. 3149–3158. Association for Computing Machinery, New York, NY, USA, 2013. doi: 10.1145/2470654.2466431

[15] G. Johnson, K. Gruchalla, M. Ingram, N. Guba, R. Cruickshank, and S. Caruso. Vis-SAGA: Visual analytics for situational awareness of grid anomalies. In *2023 IEEE Symposium on Visualization for Cyber Security (VizSec)*, pp. 17–21. IEEE Computer Society, Los Alamitos, CA, USA, oct 2023. doi: 10.1109/VizSec60606.2023.00009

[16] B. Lee, P. Isenberg, N. H. Riche, and S. Carpendale. Beyond mouse and keyboard: Expanding design considerations for information visualization interactions. *IEEE Transactions on Visualization and Computer Graphics*, 18(12):2689–2698, 2012. doi: 10.1109/TVCG.2012.204

[17] H. Li, J. L. Wert, A. B. Birchfield, T. J. Overbye, T. G. S. Roman, C. M. Domingo, F. E. P. Marcos, P. D. Martinez, T. Elgindy, and B. Palmintier. Building highly detailed synthetic electric grid data sets for combined transmission and distribution systems. *IEEE Open Access Journal of Power and Energy*, 7:478–488, 2020. doi: 10.1109/OAJPE.2020.3029278

[18] A. Lopez, W. Cole, B. Sergi, A. Levine, J. Carey, C. Mangan, T. Mai, T. Williams, P. Pinchuk, and J. Gu. Impact of siting ordinances on land availability for wind and solar development. *Nature Energy*, August 2023. doi: 10.1038/s41560-023-01319-3

[19] I. Lyons-Galante, M. Karimzadeh, S. Molnar, G. Johnson, and K. Gruchalla. Alternatives to contour visualizations for power systems data. In *2023 Workshop on Energy Data Visualization (EnergyVis)*, pp. 11–15, 2023. doi: 10.1109/EnergyVis60781.2023.00009

[20] Mapbox. Tippecanoe. https://github.com/mapbox/tippecanoe, 2014. Version 1.36.0.

[21] S. McFly, J. Peterson, and T. Reynolds. Distributed energy resource cybersecurity framework and cyber range integration. 5 2022. doi: 10.2172/1868050

[22] C. Mikkelsen, J. Johansson, and M. Cooper. Visualization of power system data on situation overview displays. In *Information Visualisation (IV), 2012 16th International Conference on*, pp. 188–197, July 2012. doi: 10.1109/IV.2012.41

[23] S. Molnar, K. Gruchalla, and S. D. J. Tan. Visualization of the oscillatory dynamics of an island power system. In *2023 Workshop on Energy Data Visualization (EnergyVis)*, pp. 1–5, 2023. doi: 10.1109/EnergyVis60781.2023.00007

[24] R. Netek, J. Masopust, F. Pavlicek, and V. Pechanec. Performance testing on vector vs. raster map tiles—comparative study on load metrics. *ISPRS International Journal of Geo-Information*, 9(2), 2020. doi: 10.3390/ijgi9020101

[25] NREL. Advanced computing annual report 2023, 2024. https://www.osti.gov/biblio/2328111.

[26] T. Overbye and J. Weber. New methods for the visualization of electric power system information. In *Information Visualization, 2000. InfoVis 2000. IEEE Symposium on*, pp. 131–16c, 2000. doi: 10.1109/INFVIS.2000.885101

[27] B. Palmintier, T. Elgindy, C. Mateo, F. Postigo, T. Gómez, F. de Cuadra, and P. D. Martinez. Experiences developing large-scale synthetic U.S.-style distribution test systems. *Electric Power Systems Research*, 190:106665, 2021. doi: 10.1016/j.epsr.2020.106665

[28] B. Palmintier, C. Mateo Domingo, F. E. Postigo Marcos, T. Gomez San Roman, F. de Cuadra, N. Gensollen, T. Elgindy, and P. Duenas. SMART-DS synthetic electrical network data opendss models for SFO, GSO, and AUS, 12 2020. https://data.openei.org/submissions/2981.

[29] N. V. Panossian, H. Laarabi, K. Moffat, H. Chang, B. Palmintier, A. Meintz, T. E. Lipman, and R. A. Waraich. Architecture for co-simulation of transportation and distribution systems with electric

vehicle charging at scale in the san francisco bay area. *Energies*, 16(5), 2023. doi: 10.3390/en16052189

[30] F. Rajabiyazdi, J. Walny, C. Mah, J. Brosz, and S. Carpendale. Understanding researchers' use of a large, high-resolution display across disciplines. In *Proceedings of the 2015 International Conference on Interactive Tabletops & Surfaces*, ITS '15, p. 107–116. Association for Computing Machinery, New York, NY, USA, 2015. doi: 10.1145/2817721.2817735

[31] J. T. Sample and E. Ioup. *Tile-based geospatial information systems: principles and practices*. Springer, 2010. doi: 10.1007/978-1-4419-7631-4

[32] D. S. Tan, D. Gergle, P. Scupelli, and R. Pausch. Physically large displays improve performance on spatial tasks. *ACM Trans. Comput.-Hum. Interact.*, 13(1):71–99, mar 2006. doi: 10.1145/1143518.1143521

[33] W. Usher and V. Pascucci. Interactive visualization of terascale data in the browser: Fact or fiction? In *2020 IEEE 10th Symposium on Large Data Analysis and Visualization (LDAV)*, pp. 27–36, 2020. doi: 10.1109/LDAV51489.2020.00010

[34] K. Vogt, L. Bradel, C. Andrews, C. North, A. Endert, and D. Hutchings. Co-located collaborative sensemaking on a large high-resolution display with multiple input devices. In P. Campos, N. Graham, J. Jorge, N. Nunes, P. Palanque, and M. Winckler, eds., *Human-Computer Interaction – INTERACT 2011*, pp. 589–604. Springer Berlin Heidelberg, Berlin, Heidelberg, 2011. doi: 10.1007/978-3-642-23771-3_44

[35] A. von Meier. *Electric Power Systems: A Conceptual Introduction*. Wiley-IEEE Press, 1st ed., 2006.

[36] Y. Wang. Deck.gl: Large-scale web-based visual analytics made easy, 2019. doi: 10.48550/arXiv.1910.08865

[37] J. Weber and T. Overbye. Voltage contours for power system visualization. *IEEE Transactions on Power Systems*, 15(1):404–409, Feb 2000. doi: 10.1109/59.852151

[38] E. Zhou and T. Mai. Electrification futures study: Operational analysis of U.S. power systems with increased electrification and demand-side flexibility. Technical Report NREL/TP-6A20-79094, National Renewable Energy Laboratory, Golden, CO, 2021. doi: 10.2172/1785329