Issues Associated with the Development of a Wide-Area Analysis and Visualization Environment

John P. Stovall*, Brendan J. Kirby*
Engineering Science
and Technology Division
Oak Ridge National Laboratory
stovalljp@ornl.gov, kirbybj@ornl.gov

Thomas J. Overbye
Department of Electrical
& Computer Engineering
University of Illinois
at Urbana-Champaign
overbye@ece.uiuc.edu

James S. Thorp, Arun G. Phadke
Department of Electrical
& Computer Engineering
Virginia Polytechnic Institute
and State University
jsthorp@vt.edu, aphadke@vt.edu,

Abstract

This paper provides a discussion of issues associated with the development of an environment for wide-area power system analysis and visualization. In particular, the paper considers issues associated with the exchange of information between regional transmission operators, the use of a unified state estimation and contingency analysis solution, and the visualization of the results.

1. Introduction

Around the world the recent large-scale blackouts have dramatically demonstrated that even with modern energy management systems (EMSs) cascading blackouts are still not a relic from the past. Even in the early part of the 21st century a few untrimmed trees are still capable of putting tens of millions in the dark! In North America a list of recommendations for reducing this probability was included in the final report for the August 14, 2003 blackout [1], with a one year later progress report on the actions actually implemented given in [2]. These recommendations covered a wide range of topics, from institutional changes by government agencies to changes in protection system plans and practices. It is clear that a number of changes can be made in several different areas to decrease the risk of future blackouts. The purpose for this paper is to explore one of these areas – the role improved wide-area analysis and visualization can play in reducing blackout risk.

The need for this work arises because of how large interconnected power systems, such as the North American Eastern Interconnect, are actually operated. While such an interconnection is physically one big electric circuit, from an operational perspective it is subdivided into a number of autonomous control areas. Though operations in one control area influence reliability in the others there is only limited data exchange between them and operators in one control area have only limited awareness of conditions in other control areas.

Of course the development of regional transmission operators (RTOs), such as the expanded PJM Interconnection and Midwest ISO, have helped to mitigate this problem. But even these newer, larger organizations can have an inadequate view of the power system particularly when problems develop near their boundaries or when they span multiple RTOs. A germane example was the August 14th blackout in which problems in one control area were not significantly visible to the neighboring control areas and RTOs. Ultimately what will probably be required is the ability to observe the entire eastern and western interconnections in real time. The purpose of this paper is to discuss some of the issues associated with the development of a wide-area analysis and visualization environment (WAVE), where wide-area is defined to encompass more than a single RTO.

2. Towards a Wide-area Analysis and Visualization Environment

At first glance the development of a wide-area analysis and visualization environment might seem to be a straightforward extension of what is already done at the RTO level. That is, obtain data from multiple control areas, combine it into a unified model, analyze the model, and visualize the results. However, this approach requires the RTO to compile and keep up-to-date the unified power system node-breaker model. Since the node-breaker model is much more dynamic than the simpler bus-branch type model used for planning studies this is a significant task even at the regional level as currently done in North America. Furthermore, while some may argue it could be technically feasible to have a single entity overseeing the operation of an entire interconnected system, the formation of such an entity is unlikely to occur for either the Eastern or the Western North America interconnects for various reasons.

Therefore the basic premise of the wide-area analysis and visualization environment presented here is the absence of a single entity. Rather, the approach assumes a

* *Prepared Under DOE Contract DE-AC05-00OR22725 between U.S. Department of Energy and UT-Battelle, LLC.*

peer-to-peer relationship between the RTOs within an interconnect as shown in Figure 1.

The basic framework discussed here consists of monitored data, analysis tools, and a means to display or visualize the data and information. The framework is interdependent because visualization and display of information requires adjustments in the methods, processing and analysis of the data. An overview schematic of this framework is shown in Figure 2. The framework follows the process and analysis used in modern energy management control centers. Also, the framework draws upon the methodology and evaluation studies of the concept of physical and operational margins [3] and the concepts proposed for power system security assessment [4].

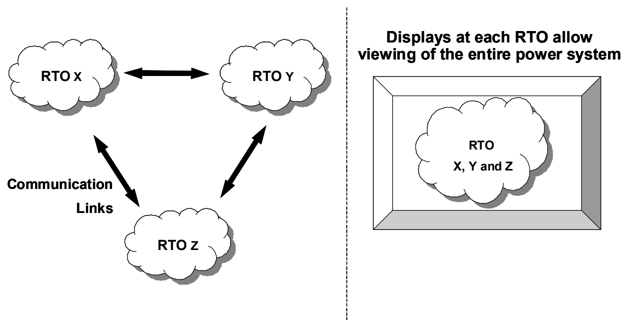

Figure 1: Peer-to-Peer RTO Information Exchange

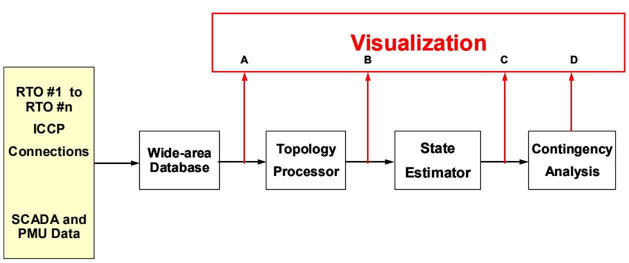

Figure 2: Schematic of the Wide-Area Analysis and Visualization Environment.

To implement such a wide-area framework one key issue to address is whether to exchange data using a breaker-node model, a bus-branch model, or perhaps some hybrid. The advantage of the breaker-node model is it represents the actual configuration of the electric grid. Such an approach is certainly required by individual RTOs for their own internal information. But the disadvantage of this approach is because the status of the system can be quite dynamic it can require significant maintenance to keep databases and displays up-to-date. In contrast, the bus-branch approach used in planning takes a higher level, somewhat less accurate view of the system. However, its advantages are it requires less maintenance and results in models with fewer overall buses. For wide-area data exchange a hybrid approach might be best in which data is exchanged usually a fairly detailed bus-branch model. Sufficient buses could be included in the model to represent all the reasonable bus configurations that might occur in an actual power grid. But the model would need not be sufficiently detailed to include the status of all individual objects such as disconnects. Each RTO would be responsible for the initial topology processing to determine the status of each branch, and to determine the voltage of each bus. Permanent ids (e.g. numbers) could be assigned to bus.

Another key issue is assigning a central point of data collection. Several approaches can be considered. First, each RTO could take on the task and build their version of the wide area view. This is essentially what is happening today with each entity adding selected data points to their EMS from their neighbors. If each followed a standard process, then presumably the same wide area view would be obtained. But such an approach would be an overwhelming task for the Eastern Interconnect in which even the largest RTOs cover a relatively small percentage of the entire interconnection. It would also involve significant duplication of effort. Second, one entity could be given the responsibility to be the integrator of the wide area view and then would transmit this view back to the other participating entities. This responsibility could be rotated among the group to share the burden over time. However, rotating the responsibility would require each organization to have sufficient in-house expertise to manage the task when it was their turn. A third approach would be to use a third party that would act as the integrator, returning the integrated data and analysis results to each of the system operations centers. Regardless of the location of the central data collection point in this framework data would be transferred from the RTO's to this point using the intercontrol center communications protocol (ICCP) standard [5]. The basic data would be the measured supervisory control and data acquisition (SCADA) data from each RTO's network. In addition, phasor measurement unit (PMU) data could be used to improve the analysis and determination of the power system state. The data would reside in a wide-area database for visualization and analysis. For the SCADA data, database schemes and software similar to those currently employed at individual RTO's could be used with the caveat that the size of the database would be larger and that names of objects would need to be unique across RTOs. Methods for including PMU data would need to be explored and may necessitate a separate database due to the significantly faster sampling rate of the data. Other data may be transferred to the wide-area database which to date has not been typically kept in an online historical database, such as the output of each RTO's state estimator.

3. Wide-area State Estimation and Contingency Analysis

Construction of a wide-area state estimator is essential to the overall analysis framework. The method in which the state estimator is constructed will determine its robustness. Two approaches are discussed below.

The first approach would be to build the wide area state estimator in a manner similar to the way it is done at the RTO [6]. At the RTO, the SCADA data from all members is brought to the RTO control center and then processed by the state estimator. A great deal of effort is expended in setting up and maintaining the data for the state estimator. For the state estimator to work the data must be of a good quality. The topology is dynamic as there are transformer and line changes due to both forced and maintenance outages, equipment additions with new construction and equipment retirements. The analog measurements vary in accuracy and are subject to errors resulting from equipment drift and failure. These factors contribute to the performance of the state estimator and require continual vigilance to ensure the performance. Because of the effort in maintaining and ensuring the data quality to the state estimator, there is a tendency to minimize the data used from neighboring systems in the state estimator. Great effort and time are needed to verify the data when it appears to be in error. Significant data failure in any part of the observed system can result in failure of the state estimator for the entire system. This problem was experienced during the August 14, 2003 blackout. As a result, the network model for the state estimator will typically reach several buses into the neighboring system and then rely upon a network equivalence of the neighboring system.

To adopt this procedure on a wide area basis would be a challenge. The effort would be duplicative of the effort at each of the RTO's and very labor intensive. Even if successful, there is a question weather a state estimator model constructed in this manner would reliably converge.

An alternate approach would be to build the wide area state estimator model by assembling the solved state estimator models from the individual transmission entities. Multi-area state estimation approaches have been examined [7]. The process would be similar to that of assembling a wide area power flow base case from individual power flow cases with the added challenge of completing the task in an automated fashion in real-time.

There are several issues, which need to be addressed for this method to work, and some of them are discussed below. Some of the issues can be resolved by establishing procedures for each transmission entity to follow while others will require new analysis methods.

The first issue is that of overlap and seams. The state estimator models from each transmission entity include their external network, which would need to be removed. This should be able to be done by a uniform area numbering system that would allow the external network to be easily removed from each state estimator model. The adjoining of the individual RTO network models would need to be automated. The procedure would need to ensure that buses and lines match up along the seams.

The second issue is time synchronization of the data. The overall solution of the wide-area state estimator can be improved by time synchronizing of the individual state estimators. By agreeing to run the individual state estimators at the same clock time, the time skew among the individual state estimators would be minimized. If available, PMU data could be used as an adjustment to improve time synchronization [8] [9]. However, doing such a synchronized state estimation could introduce additional time delay in the ultimate display of results.

Last, the question of missing data from an entire RTO would need to be addressed. When the state estimator at an individual RTO fails to converge, then that area would be missing from the wide-area state-estimator model. How to treat this missing area is not known. The wide-area state estimator may continue to converge with a missing area. If the wide-area state estimator does not converge then it may be possible to extract the missing area model from the last converging case that would allow the state estimator to converge. Of course, this would be an unusual solution with perhaps dubious results.

The wide-area contingency analysis would offer the opportunity to examine line, transformer and substation outages spread across multiple RTO's that are not presently being examined. Further study methods to perform such a large scale-contingency analysis may be needed.

4. Combining Adjoining State Estimates Phasor Measurement Techniques

Phasor measurements provide a unique opportunity to properly merge the state estimates of adjoining power systems. It is often the case that two neighboring power systems process their SCADA measurements and produce a state estimate which makes no reference to measurements from neighboring systems or their equivalents. However, in order to use the estimates for performing contingency studies or other simulations, it becomes necessary to have a good model of the neighboring systems. This could be in the form of equivalents, or could in fact be a state estimate obtained by its control center.

One approach used in some systems is to collect the SCADA data from both systems in one larger measurement set, and then estimate the state of the combined system by running an estimator for the combined system. This is wasteful of computational effort, and is often conducive to non-convergence because of data errors or system structure errors present in the enlarged data set. It is much more reasonable to take the results of the two state estimates and then devise an algorithm to provide a combined state estimate without running a large state estimator. The key to this approach is the judicious placement of a few phasor measurement units in the two systems, and using these measurements to place the two state estimates on a common reference.

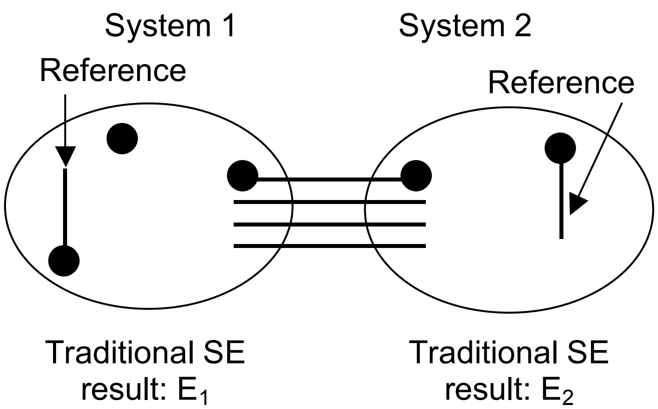

Figure 3: Concept of combining state estimates using phasor measurements.

The concept is illustrated in the Figure 3. The two bounded regions represent two adjoining power systems, with tie lines represented by the four connecting lines. Each system collects its own SCADA data, and produces a state estimate, E1 and E2 respectively. It is proposed that a certain number of PMUs be placed at buses marked by solid black dots. These phasor measurements are then brought to the control centers of the two systems. It is recognized that the two state estimates may be obtained at slightly different nominal times. The phasor measurements are of course obtained with synchronized clocks, so that these measurements are truly simultaneous. Not only that, they are obtained far more frequently (e.g. every 2 or 3 cycles of the nominal power system frequency) and a sequence of these measurements is available at the control centers. This is illustrated in Figure 4, which shows the time relationships. An algorithm will need to be developed to make use of the phasor data in order to put the two state estimate results on a common reference, thereby producing a single state estimate for the combined system. The application of the algorithm would not necessitate the additional solution of the combined state estimate.

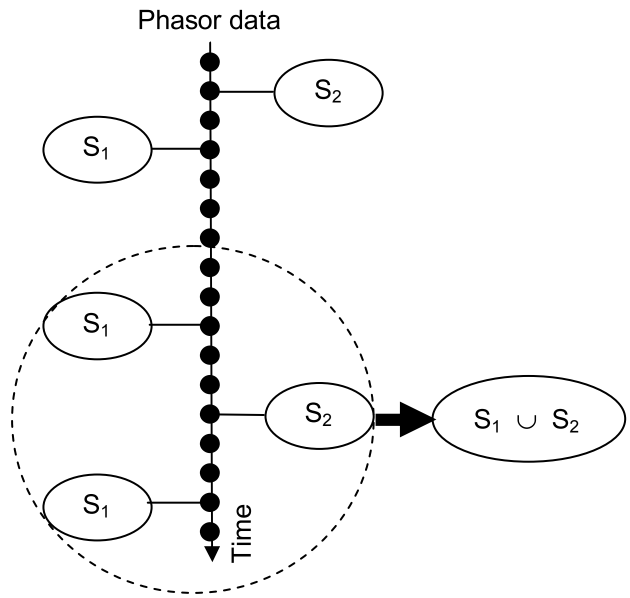

Figure 4: Synchronizing state estimates using PMU's.

5. Visualization of Results

The visualization of the wide-area power system would be based upon the different sources of data/information shown in Figure 2. To a large extent the visualization of the wide-area results would be similar to the display of other large power system datasets such as discussed in [10]. The most important issue that would need to be addressed is the ability to consistently map the data to be displayed to the associated display elements.

For planning models (i.e., those using a bus-branch model) data is mapped to the associated visualization displays using a fixed bus number approach. This approach assumes that the bus numbers remain essentially fixed over time, changing only when major system changes occur such as the construction of a new transmission line or generator. Hence the visualization displays would not need to be modified frequently due to system topology changes. For the wide-area visualizations considered here, in which one is mostly concerned with higher-level system monitoring, the bus-branch model could be used for both the real-time information and the state estimation/contingency analysis results. The remainder of this section discusses several techniques that could be quite useful for wide-area visualizations.

5.1 Line Pie Charts and Animated Flows

One visualization technique that has proven useful for quickly indicating the location of overloads/outages in a transmission network has been the use of pie charts in which the percent fill of each pie chart is equal to the percentage loading on the line [11]. Optionally, a numeric text could also be superimposed on the pie chart to indicate the exact percentage. For displays with relatively few lines, where each individual line's pie chart can be viewed with sufficient detail, such pie charts can quickly provide an overview of the system loading. If desired, different color shadings could also be used with the pie charts to highlight those devices loaded above some threshold percentage.

However, for the larger wide-area displays considered here there is insufficient space to show each individual pie chart with sufficient detail. Instead, a supplementary technique is to dynamically size the pie-charts based upon the line's percentage loading. In this approach the percentage fill in each pie-chart is still equal to the percentage loading on the line, but the size and color of the pie-chart can be dynamically sized when the loading rises above a specified threshold. By increasing the size and/or changing the color of the pie charts for only the small number of elements loaded above a critical threshold, the user's attention can be focused on those elements near or exceeding their limits even of displays with lots of other pie charts. The overloaded elements appear to "pop-out" on the display. This dynamic sizing/color changing can also be used to indicate open transmission lines.

Figure 5 shows an example of this technique for a case simulating the power flows in Northeast Ohio at the beginning of the sequence of events on August 14th, 2003. Overall the display shows slightly less than 400 buses (mostly 138 and 345 kV), slightly more than 400 transmission line/transformer pie charts, and the pie chart for a single flowgate. On a high resolution computer display the pie charts could be dynamically sized/colored to be 10 times their normal size and colored orange for a percentage loading above 85% of their emergency limit, and to be twelve times their normal size and colored red for a loading above 100%.

In particular, the figure shows the system state after the loss of the first 345 kV line (Chamberlin-Harding) at 15:05 EDT (Eastern Daylight Time) on August 14th. While the loss of this line did not cause any transmission line overloads, it did cause an overload on contingent flowgate #2265, which monitors the flow on the Star-Juniper 345 kV line for the contingent loss of the Hanna-Juniper 345 kV line. Note the large red pie chart immediately draws attention to this overload. The open transmission line's pie chart can also be dynamically sized and have its color changed to draw attention. In Figure 5, the open Chamberlin-Harding 345 kV line is indicated by the large black pie chart with the green "X". Of course, other color combinations and symbols could be used.

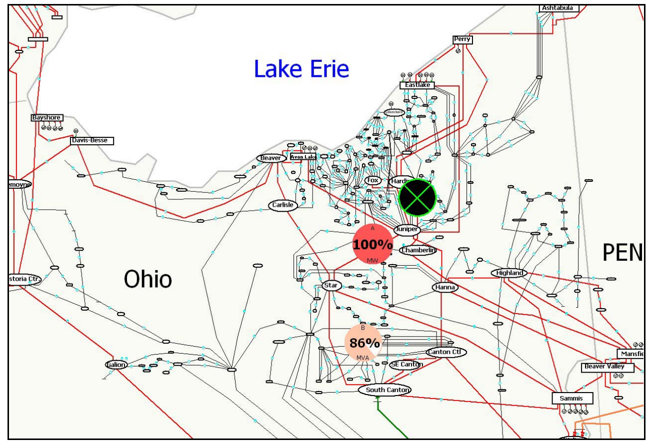

Figure 5: Northeast Ohio Transmission System Status after Loss of First 345 kV Line at 15:05 EDT

There are, however, several potential problems with the pie chart approach. First, care must be taken when applying resizing to devices that are designed to be regularly loaded at a high percentage level. Common examples of this are generator step-up transformers. Such transformers are designed to be regularly loaded at a high percentage of their ratings, but because of their radial connection, they are in no danger of overloading. A straightforward solution to this problem is to either not show pie charts on such devices or to specify that the pie charts should not be dynamically resized.

Second, the use of dynamically sized pie charts involves a tradeoff between making the pie charts large enough to draw attention, yet not too large so as to obscure other important one-line elements. For example on August 14th at 15:51 EDT after three 345 kV lines and a number of 138 kV lines have opened the display would contain a large number of oversized elements (see Figure 6). The large number of overloads/line outages would make it much more difficult for an operator or engineer to rapidly locate the most crucially overloaded devices. Of course, this is a characteristic of any approach that uses dynamically sized one-line elements, including the approach from [12] of dynamically increasing the line widths to indicate overloads.

One solution to this problem is to filter the display to highlight certain lines, and attenuate the display of other lines. For example a display could be created in which the lower voltage lines are blended into the background, helping to focus attention just on the more critical 345 kV lines. With such an approach, it would become clearer that the Sammis-Star 345 kV line is overloaded, and that four 345 kV devices are open. Yet, the lower voltage lines would still be partially visible, helping both to provide context for the higher voltage lines and to allow continued monitoring of these lines.

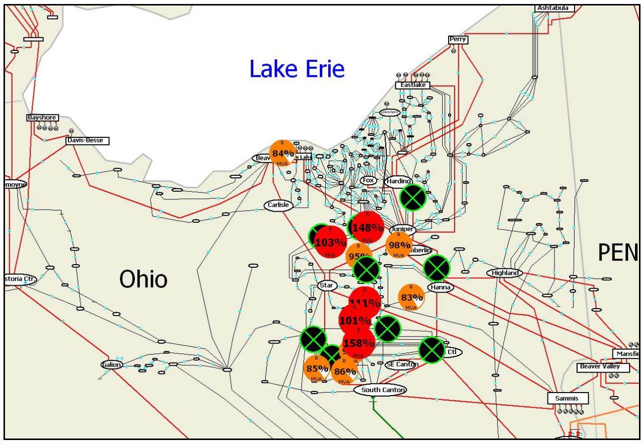

Figure 6: Transmission System at 15:51 EDT

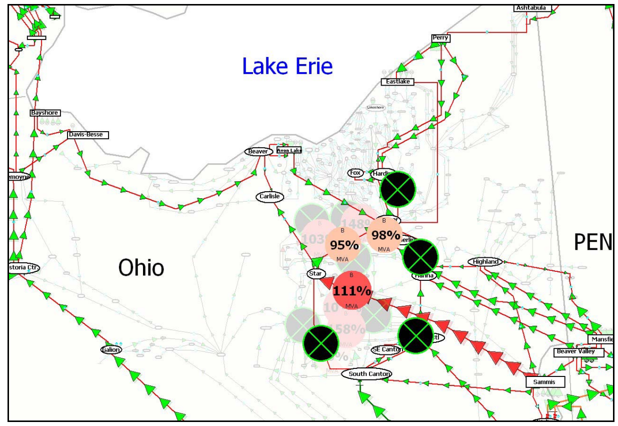

Figure 7: Transmission System at 15:51 EDT with Flow Arrows and Display Filtering

A second potential solution would be to proportionally reduce the amount of resizing as the display is zoomed. For example, a maximum zoom value for full resizing could be specified. Once the display is zoomed past this value the amount of pie chart resizing could be dynamically reduced it reached some threshold value, such as the size of a normal pie chart. The visual effect during increased zooming is the resized pie chart remains the same size on the screen, causing the overlap to eventually vanish. This technique could also be used with filtering.

Useful as the pie charts may be, they still to not provide any indication of the direction of flow. One technique for displaying line flows is to superimpose small arrows on the lines, with the arrow pointing in the direction of the MW flow, and the size of the arrow proportional to either the MW or MVA flow on the line [13]. The size and color of these arrows can also be used to provide a visual reinforcement with respect to severity of the problem. For example, in Figure 7 the arrows on the overloaded Sammis-Star line are colored red, while the net direction of the power flow into the Cleveland-Akron area could indicate a problem in that area.

To further emphasize the flow direction the flows themselves could also be animated. However, the application of such animated flows to larger systems has to be done with some care. It is certainly possible to create one-lines in which the presence of the animated flow arrows results in more clutter than insight. But if the view is filtered to just a particular voltage level, or zooming is utilized to focus on a particular portion of the grid, then the animated flow arrows can again be quite helpful. The use of animation could also be restricted to just those lines of particular interest, allowing one to take advantage of the fact that motion is also preattentively processed, provided the number of elements moving on the screen is small. Selective use of animation can also convey a sense of urgency, directing attention to important elements.

5.2 Contouring

In addition to visualizing flow-based quantities it is important to also convey the bus-base values such as voltage magnitudes and locational marginal prices (LMPs). One technique that has proven to be quite useful for wide-area visualization of such information has been the use of bus-based color contouring [14].

There are three main concerns with applying contouring to the display of power system bus-base quantities such as voltages and LMPs. First, the bus values are not spatially continuous – they only exist at distinct buses with no values in between. To resolve this concern virtual values must be created to span the entire two-dimensional contour region. The virtual value is a weighted average of nearby data points, with different averaging functions providing different results. Second, for voltage magnitude contours tap-changing transformers may introduce sudden changes in voltage. Contour plots normally imply a slow variation in value. However, this concern can be overcome by setting up a contour in which one a small group of voltages levels are contoured. Third, buses that are physically near one another on a one-line diagram may not be "near" one another electrically. This concern can also be overcome by only contouring particular voltage levels, and perhaps by judicious construction of the contouring diagram. If a strictly geographical layout is used, normally buses of the same nominal voltage that are near each other geographically are also near electrically.

Once these virtual values are calculated, a color-map is used to relate the numeric virtual value to a color for display on the screen. A wide variety of different color maps are possible, utilizing either a continuous or discrete scaling. One common mapping is to use red for lower voltage values and blue for higher values, while another common mapping is the exact opposite – standardization of contour color mappings would be beneficial particularly when displaying data across RTOs.

Perceptually, contouring works well because the human visual system is well designed for detecting patterns. For example, Figure 8 shows a voltage contour of about 5600 bus voltages in the Mid-Atlantic portion of North America at about 15:00 EDT on August 14th using a red-blue color mapping. Notice that at a glance the relatively low initial voltages in Ohio and Indiana are visible as darker regions, with the extremely low voltages in Southern Indiana due to line outages that occurred at about noon on August 14th. Another advantage of contouring for wide-area visualization is it is somewhat immune to the presence of unlinked data points – as long as enough nearby points are linked that lack of a few points probably won't have a significant impact.

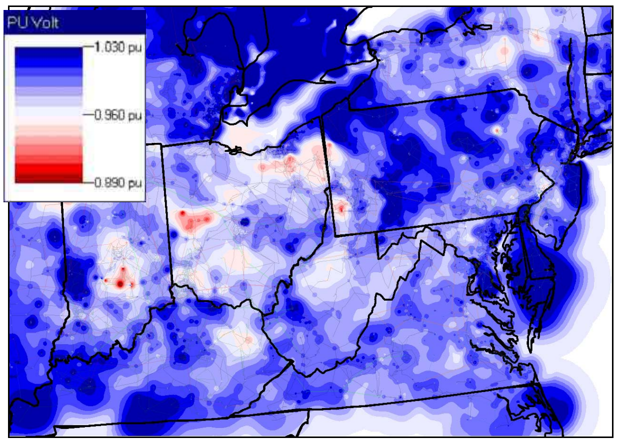

Figure 8: Pre-Blackout Ohio Region 115-230 kV Voltage Contour

Another advantage of contouring, which is impossible to show in a paper, is that the images can be readily animated to quickly show the progression of the system over time (similar to what is done with animated weather maps). During an emergency, this information could be used to rapidly convey the direction of the system state, and could also be used to rapidly bring late arriving people to the control room (such as engineers and management) up-to-speed on what is going on.

However, one does need to be careful about trying to convey too much information in a single display. For example, combining transmission system information (pie charts and flow arrows) with the contour as is done in Figure 9. While this may be tolerable if there are just a few problems, as more problems develop it would become increasingly cluttered.

Contouring can also be used to show other bus-based values, such as the voltage angle, the bus LMP or sensitivity information. Color contours could also be used in conjunction with the display of other related visualizations. For example, one might like to simultaneous visualize the loading of the transmission lines on a contour showing the bus LMPs. For such visualizations a more subdued (lighter) color contour would probably be appropriate to allow all information to be clearly seen.

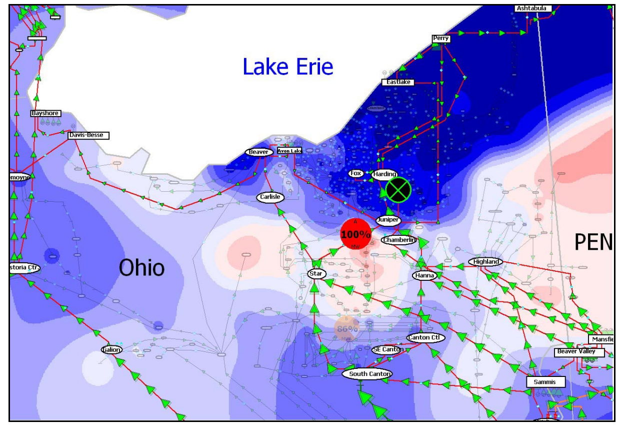

Figure 9: Northeast Ohio Voltage Contour at 15:05 EDT

5.3 Interactive 3D Visualization

The previous data visualization techniques can be quite useful when one is primarily concerned with visualization of one or perhaps two types of spatially oriented data such as bus voltages or transmission line flows. However, in power system analysis one is usually confronted with a large amount of multivariate data. Data of interest could include a potentially large list of independent and dependent variables, such as bus voltage magnitudes, transmission line loadings, generator real and reactive reserves, transformer tap and phase positions, scheduled and actual flows between areas, and interface loadings. This section presents results on the use of an interactive 3D visualization environment to assist in analyzing this vast amount of information [15].

In developing such an environment several key issues must be addressed. First, in visualizing power system data, there is usually no corresponding "physical" representation for the variables. For example, there is no physical representation for the reactive power output of a generator, or for the percentage loading of a transmission line. Rather, these values are typically shown as a numerical values on either a one-line diagram or in a tabular display. This contrasts with the use of interactive 3D for power system operator training, in which the 3D environment seeks to mimic, as closely as possible, an existing physical environment. It also differs from the use of interactive 3D for some types of scientific visualization, in which the purpose of the environment is to visualize physical phenomena, such as flows in a wind tunnel or molecular interactions. To address this issue, an environment based upon the common one-line representation serves as a starting point. The new environment differs from the one-line in that a one-line is a 2D representation, whereas the new one is 3D. How this third dimension can be exploited is covered in the following sections.

A second issue is the 3D environment must be highly interactive. In power systems, there is simply too much data to simultaneously display all the data that may be of interest. Rather, the user should be able to quickly and intuitively access the data of interest.

To introduce such an environment, Figure 10 shows a traditional 2D one-line for a small thirty bus system mapped into a 3D environment. To take advantage of the 3D environment, the generators are represented using 3D cylinders of potentially varying heights. The one-line has been mapped into 3D using a perspective projection. The one-line is now oriented in the xy-plane (horizontal plane), while the generators extend in the z (vertical) direction. In Figure 10, the height of each generator cylinder is proportional to its maximum real power capacity. The height of the lighter bottom portion of the cylinder is proportional to current output of the generation, while the upper portion indicates the available reserve capacity.

One of the advantages of 3D is its ability to show the relationships between variables. For example, in studying the voltage security of a system one is often interested in knowing both the location and magnitude of any low system voltages, and also the current reactive power output and the reactive reserves of the generators and capacitors. Such a situation is illustrated in Figure 11 where the height of each generator cylinder is now proportional to the maximum reactive capacity of the generator; the darker region on the lower portion of the cylinder is proportional to the current reactive power output, while the lighter top portion represents the reactive power reserves. The bus voltage values are indicated using a contour, with only voltage values below 0.98 shaded. Rapid navigation in such a visualization would be important to allow one to rapidly see the more distant parts of the display.

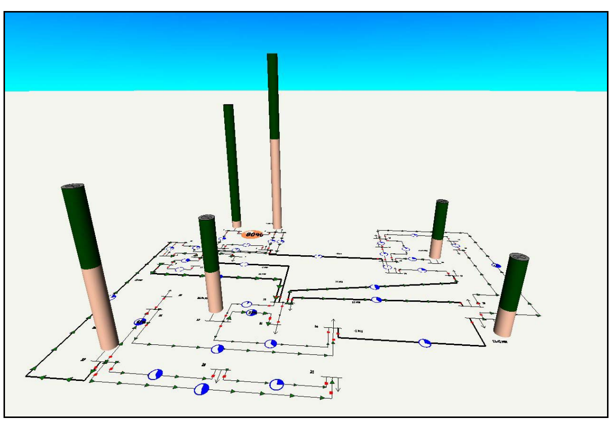

Figure 10: 3D View of 30 Bus System

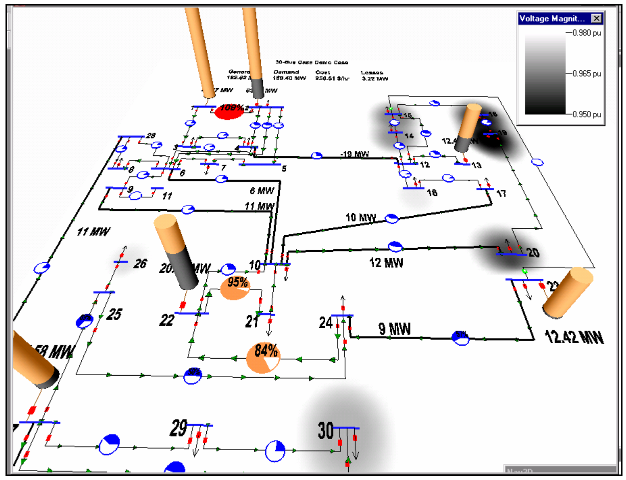

Figure 11: 30 Bus System Generator Reactive Reserves

Note in the figure it is apparent, almost at a glance, the location of the sources of reactive power generation and the reactive power reserves. The figure also does a good job of conveying qualitative information about the magnitude of these values. What it does not do is convey quantitative information. Thus while one learns from Figure 11 that the reactive power generation at bus 20 is about 50% of its maximum, one does not learn what the actual var output is, nor the maximum var limit. In some situations this could be a significant limitation. Thus, the use of 3D is advocated to supplement, rather than to completely replace, the existing one-line and tabular display formats.

However, in support of the 3D approach several observations are warranted. First, in many situations these qualitative relationships are sufficient. This is particularly true when one is studying a familiar system. For example, utility operators and engineers already know the reactive limits of their generators; what they need to know is how close those generators are to their limits in a qualitative sense. And one needs to know the value for only a few generators in order to get acceptable estimates for the remaining generators by comparing their relative sizes. Second, there is nothing that prevents displaying numerical values of system quantities in the 3D environment, just a bus numbers and some line flows are shown in Figure 11. A valid objection to this approach is that if the viewpoint is changed (i.e., rotated) these fonts could be shown at unappealing angles. However, this problem could be easily solved by having the 3D environment automatically rotate the fonts so they are always facing the viewer (provided the one-line is designed with sufficient room for the fonts to rotate). Alternatively, the fonts could be shown in the vertical plane, such as on the surface of the generator cylinders. An third approach, which is prefered and has been implemented, is to allow the user to get additional information about system objects by selecting the object with the mouse. This allows the display to be relatively uncluttered, yet still allows rapid access to a large amount of quantitative information. The hint functionality mentioned early for showing contour values could also be used.

A final issue associated with 3D visualization that needs to be addressed is performance. In order to give the user the feeling of interacting with the 3D environment it is crucial that the display refresh quickly and that there be little latency between the user issuing a command, such as desiring to change position, and the display being updated. How fast the display can be refreshed depends, of course, on a number of factors, including the speed of the computer's processor, the speed of the display card, whether the display card has hardware support for OpenGL, and software considerations such as the level of detail of the display and the lighting model employed. Given the trend toward supporting OpenGL directly in the display card hardware, newer microprocessors with new commands to directly support 3D graphics and faster processors in general, very good 3D performance is becoming widely available even for the display of relatively large systems. For benchmarks, relatively small displays like Figures 10 and 11 can be animated at about 60 frames per second on a 2.8 GHz PC. To demonstrate performance on a slightly more realistic system Figure 12 shows a one-line of the Midwest electric grid, again using cylinders to visualize the generator outputs and reserves. The refresh rate for this reasonably detailed display, which contains about 1100 buses, 375 generators and 1000 transmission lines, is around 4.5 frames per second.

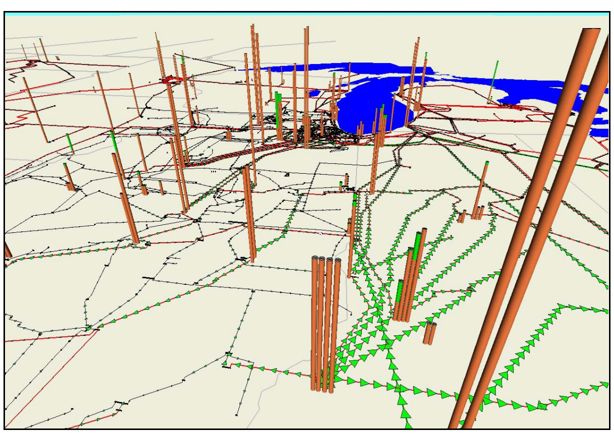

Figure 12: Large-Scale 3D Visualization

6. Conclusion

Restructuring in the electricity industry is resulting in a need for innovative new methods for analyzing and visualizing large amounts of system data spanning an entire interconnect. This paper has discussed several of the issues associated with the use of this data for wide-area analysis and visualization.

7. References

[1] "Final Report on the August 14, 2003 Blackout in the United States and Canada: Causes and Recommendations," U.S. Department of Energy, April 2004.

[2] "The August 14th 2003 Blackout One Year Later: Actions Taken in the United States and Canada to Reduce Blackout Risk," U.S. Department of Energy, August 2004.

[3] S. Lockwood, et. al., "Utility Experience Computing Physical and Operational Margins: Part II - Application to Power System Studies" *IEEE PES Power Systems Conference and Exposition, 2004*, Oct. 10-13, 2004, pp. 1446-1452.

[4] K. Morison, L. Wang, P. Kundur, "Power System Security Assessment, *IEEE Power and Energy Magazine*, Vol. 2, No. 5, pp. 30–39.

[5] J.T. Robinson, et. al., "Development of the Intercontrol Center Communications Protocol (ICCP)," *IEEE Conference Proceedings of Power Industry Computer Application Conference, 1995*, pp. 449 – 455.

[6] A. Monticelli, "Electric Power System State Estimation," *Proceedings of the IEEE*, Vol. 88, No. 2, pp. 262–282.

[7] R. Ebrahimian, R. Baldick, "State Estimation Distributed Processing," *IEEE Trans. on Power Systems*, Vol. 15, No. 4, Nov. 2000, pp. 1240–1246.

[8] Z. Liang Zhao, A. Abur, "Multi Area State Estimation using Synchronized Phasor Measurements," *IEEE Trans. on Power Systems*, Vol, 20, No. 2, May 2005, pp 611–617.

[9] R.P. Klump, R.E. Wilson, K.E. Martin, "Visualizing Real-Time Security Threats using Hybrid SCADA/PMU Measurement Displays," *Proceedings of the 38th Annual Hawaii International Conference on System Sciences, 2005,* pp.55c - 55c.

[10] T.J. Overbye, D.A. Wiegmann, "Reducing the Risk of Major Blackouts through Improved Power System Visualization," accepted for presentation at 15th Power System Computation Conference (PSCC), Liege Belgium, August 2005.

[11] T.J. Overbye, P.W. Sauer, C.M. Marzinzik, and G. Gross, "A User-Friendly Simulation Program for Teaching Power System Operations," *IEEE Trans. on Power Sys.*, pp. 1725-1733, November, 1995.

[12] P.M. Mahadev, R.D. Christie, "Minimizing User Interaction in Energy Management Systems: Task Adaptive Visualization," *IEEE Transactions on Power Systems*, vol. PWRS-11, No. 3, pp. 1607-1612, August 1996.

[13] T.J. Overbye, G. Gross, M.J. Laufenberg and P.W. Sauer, "Visualizing Power System Operations in the

Restructured Environment," *IEEE Computer Applications in Power*, pp. 53-58, January 1997.
[14]. D. Weber and T. J. Overbye, "Voltage contours for power system visualization," *IEEE Trans. on Power Systems*, vol. PWRS-15, pp. 404-409, February, 2000.
[15] T.J. Overbye, Y. Sun, R.P. Klump, J.D. Weber, "Interactive 3D visualization of power system information," *Electric Power Components & Systems*, pp. 1205-1215, December 2003.