EuroVis Workshop on Visual Analytics (2013)
M. Pohl and H. Schumann (Editors)

Visual Analysis of Expert Systems for Smart Grid Monitoring

M. Steiger[1], T. May[1], J. Davey[1] and J. Kohlhammer[1]

[1]Fraunhofer IGD, Darmstadt, Germany

Abstract

This application paper introduces the combination of an expert system with a visualization system, specifically designed for SmartGrid control rooms. It supports operators in the efficient monitoring of electric grids with a focus on distributed, small-scale power generation. A rule-based expert system filters the stream of a large amount of incoming events, searching for potential problems. A multiple, coordinated view environment provides the required situational awareness for the user and presents the analysis results in a details-on-demand manner. Being a critical infrastructure, the electric grid is highly sensitive and any modification must be well justified. Our visualization system therefore also provides insight into the expert system enabling the user to validate and verify the expert system's analysis process. This provides the required decision support to assist the operator in keeping the grid in stable operating condition.

Categories and Subject Descriptors (according to ACM CCS): H.5.2 [Information Interfaces and Presentation]: User Interfaces—User-centered design I.2.1 [Computer Graphics]: Applications and Expert Systems—Industrial automation I.3.6 [Computer Graphics]: Methodology and Techniques—Interaction techniques J.2 [Physical Sciences and Engineering]: Electronics—

---

1. Introduction

Today, the majority of electricity is still produced at a few, large power plants. From there, the energy is then transferred over large distances using high-voltage lines. Low voltage distribution grids carry the electricity over the final mile to the customers. However, things are changing: over the past years, a noticeable trend towards renewable energies has been registered in most industrialized countries. Their continuous increase in market shares reduce the dependency on natural resources, but their integration in the electric grid that has been built up over the past century is not a trivial endeavor [IL08].

The majority of installed solar plants are rather small – consequently, the number of stakeholders is rather large and the site locations are distributed across the entire country. Taking into account that their input is directly dependent on the solar radiation and therefore highly fluctuating, new challenges arise not only on the engineering side but also for its management and monitoring. In a typical control room, a large amount of connected sensors deliver measurement readings for various kinds of data types in periodic time intervals. These events are then processed and refined by *Supervisory Control And Data Acquisition* (SCADA) systems. They are then presented to the user, often in tabular lists. In the electric engineering domain, an event is most commonly referred to as an alarm even if its content has purely informational character.

Over the years, the complexity of the network continuously increased leading to more and more event data. From these vast amounts of events, only very few require the operator's attention [HAB09], but identifying those can be quite laborious. The task of the control room operators is to monitor these events. They use their knowledge as domain experts to decide whether an event or a series of events indicates a problem in the grid or not. Where necessary, they must interfere and instruct counter-measures before the problem propagates through the grid causing cascading effects.

Different automatic tools have been proposed to assist the operators. Most of them are about alarm processing, i.e. the pre-processing, filtering and prioritization [KSH07] of events. However, they typically operate in a "black box" manner, prohibiting insight into the internal workings. Thus, the user is not able to verify the correctness and is bound to rely rather blindly on the assumption that the reduced data set contains all vital information.

In this paper we present a system that combines the

DOI: 10.2312/PE.EuroVAST.EuroVA13.043-047

strengths of an automated analysis with a visualization that enables the operator to track this process. We therefore define a knowledge base which contains a collection of so-called *facts*, that represent all known entities and properties of the domain model. With a set of *rules* that describe how these facts relate, the inference engine can then derive new facts from existing ones. The combination of knowledge base and inference engine is commonly known as *expert system*. We use such a system to be able to automatically process large amounts of sensor data in real time. The strict separation of business logic and visualization also allows us to adjust and modify the domain-specific content without touching the view component. Even switching to an entirely different domain is possible with this approach.

Our expert system follows the pipeline design and is split into several distinct analysis steps. Suspicious events are first filtered and grouped according to the similarity of different intrinsic properties such as location or severity. Then, diagnoses over these events are put up before user recommendations are generated. The system keeps track on items that are linked across the pipeline stages so that the visualization can display which facts and rules led to the generation of which diagnosis and recommendation. With this setup, the control room operator can verify the correctness of how these assumptions were made and decide which action should be taken.

The concept of Visual Analytics is about combining manual and automatic analysis in an iterative and alternating cycle which is an ideal choice for our analysis process flow. We think that it is crucial that no action is performed in fully automatic manner; the operator must have the final word on every decision.

2. Related Work

We locate our contribution between knowledge-assisted visualization as defined by Chen et al. [CEH*09] and, more related to the application domain, the monitoring and management of complex systems.

We start with the visualization perspective. In knowledge-assisted visualization, a knowledge base is added to Card's well-known visualization pipeline [CMS99] to control the visualization. Using this knowledge representation, [MHS07] et al. and Gilson et al. [GSGC08] select suitable mappings from the data and its schemata. Kohlhammer used knowledge representation for decision-centered visualization [KE05], focusing on situation awareness in emergency management. Going deeper into the internals of the reasoning process, Shi et al. [SQW11] display the internal reasoning network of a Rete-based expert system with an interactive animation. Possibly most related to our work, Garg et al. [GNRM08] and Xiao et al. [XGH06] combine data visualization with an interactive construction of its underlying models.

From the monitoring and management perspective, we split related work with respect to the time frame that is used. Typical analytical approaches focus on rather long-term knowledge building, while monitoring approaches often aim to support short-term situational awareness. A monitoring and diagnosis system for time-series data has been presented by McLachlan et al. [MMKN08]. We adopt their drill-down approach to diagnose the current state of nodes in the electric grid. For clinical data, [SGBBT06] propose an exploration system that features a knowledge base to derive a patient's status from status measurements. Similarly, Torralba-Rodriguez et al. [TRFBMBM10] work with a knowledge base of medical data and a custom-tailored evaluation language. Their rule-based system creates diagnoses for a patient's symptoms, but is also able to explain why a certain diagnosis has been put up.

Going into the domain of power system management, a wide range of aspects the control room operator has to deal with has been investigated in the community. For example, Marques et al. [MTF05] propose a control power system with rules written in fuzzy logic. Having acquired the domain knowledge from interviews with experts, the authors created a table with entries of the form *if x is A then y is B* with *A* and *B* being linguistic variables such as *HIGH* or *LOW*. Focusing on the transmission grid that transports electricity over large distances, Li et al. [LQS*10] explain lay out current shortcomings and sum up requirements for the future *Smart Grids*. The authors also advocate intelligent alarm management based on expert systems to be able to cope with the ever increasing amount of alarms.

Moving on to the visualization perspective, the display of this situational context has been investigated by Laufenberg [Lau05]. A web-based monitoring tool has been presented by Clark et al. [CPF09] that aims to integrate other data sources such as gas and water, but also communication structures to form a ”complete picture” for suppliers, but also for customers. With a focus on the flow and transfer capability of power lines in large-scale systems, the work of Overbye et al. [OMWC05] sheds light on different visualization techniques. In collaboration with Overbye, Klump et al. present a visualization systems that is specifically tailored to industry-demands [KSO02] and reports on practical experiences in the domain.

3. Expert Systems

Basically, an expert system contains a set of domain-specific facts and rules and an inference engine that derives new information from this knowledge base. The setup process that transforms the knowledge of domain experts into its adequate machine-readable representation is a complex and time-consuming task known as *knowledge engineering*. Software engineers can perform this operation either manually or with the help of domain-specific languages (DSLs) which are still machine languages, but closer to natural language and the user's domain.

In addition to the rule definition, the inference process must be specified. In its standard form, it comprises one or more premises – the condition – with a conclusion that is evaluated if all premises hold. These conclusions can insert, update or retract facts from the knowledge base. This modification leads to another evaluation of the rules' conditions and fires them where necessary.

4. Concept

In the following chapter we outline the concept of a monitoring system that combines the knowledge of an expert system with a visualization that conveys the gathered information from the sensor measurement network in a verifiable manner to the control room (See Figure 1). Based on that information, the operator can then control the actual model which influences future measurements.

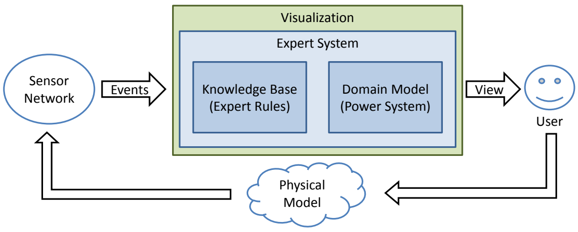

Figure 1: *Incoming events are delegated to the monitoring system that internally uses an expert system to process the data. With the help of the visualization, the user applies changes to the physical model which in turn influences the measurements.*

4.1. The analysis pipeline

The knowledge base provides the application context for the system. In our setup, it contains facts from the electric engineering domain such as transformer stations, overhead lines, voltage meters, measurement readings, etc. The collection of rules has been constructed with the knowledge of control room operators and defines how these elements relate with each other.

The expert system is designed as a pipeline with three distinct analysis steps. First, the large collection of alarms is filtered so that only suspicious items proceed to the next stage. All other events are discarded as early as possible which keeps the computational effort and memory requirements to a minimum. In the next step, the set of suspicious alarms is diagnosed and grouped by similarity leading to an even smaller set of diagnosis entries in the fact base. The information, which alarm contributed to which diagnosis is retained for the visualization. The third stage of the pipeline renders zero, one or more recommendations for each diagnosis. Again, the link between diagnosis and recommended action is recorded and displayed later on. See Figure 2 for an overview and an exemplary path through the expert system pipeline.

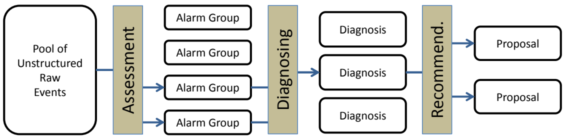

Figure 2: *Pipeline configuration of the expert system. The analysis path leads from the pool of alarms to two groups of similar, suspicious alarms. These alarms can be explained by one diagnosis which is linked to two alternative corrective actions.*

The expert knowledge in the system is used for the first, low-level analysis of the data and can be performed up to this point automatically – no user interaction is required. The generated set of recommendations can then be analyzed on a higher level by the user who can verify the correctness of the assumptions and conclusions that have been made. It is therefore essential to provide assistance for the human reasoning process, especially decision support to enable the user to decide which recommendation to follow.

4.2. Visualization & Interaction

The task of the visualization component of the system is to provide this decision support. The need for analysis in real-time and the provision of situational awareness – i.e. the context which is used to comprehend the meaning of status variables – impose additional requirements for the monitoring environment. The operator must be able to get an overview on the health status of the entire system quickly. We approach this challenge with a view environment that contains a set of coordinated views that display the data from different perspectives.

The first main contribution, the display of the overall status of the system is displayed in a large view (Figure 3, top). Initially, the view contains connected outlined rectangles that represent transformer stations at their respective geographic coordinates. The color indicates the overall state of the station as determined by the expert system. Green stands for *tested and o.k.*, gray for *undetermined* and red indicates problems. Zooming in on a station ”opens” it, giving more details on its distribution bars. Similarly, the font color of the voltage grid label indicates the status of all connected meters. One level deeper, individual meters are displayed. At the deepest level, the meters expand in size and display the latest measurement readings in a small linechart. Hovering over a defective meter with the cursor shows a tooltip that displays the latest alarm that relates to that meter.

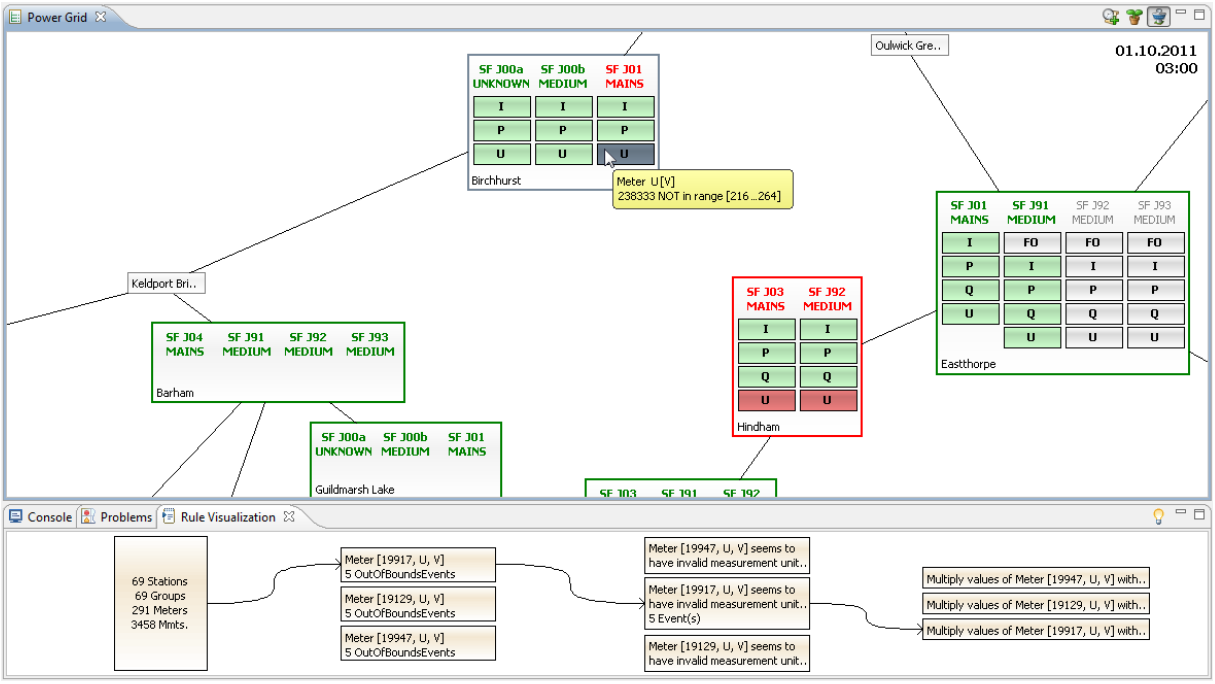

Meter U[V]
238333 NOT in range [216..264]

69 Stations
69 Groups
291 Meters
3458 Mmts.

Meter [19917, U, V]
5 OutOfBoundsEvents

Meter [19129, U, V]
5 OutOfBoundsEvents

Meter [19947, U, V]
5 OutOfBoundsEvents

Meter [19947, U, V] seems to have invalid measurement unit..

Meter [19917, U, V] seems to have invalid measurement unit..
5 Event(s)

Meter [19129, U, V] seems to have invalid measurement unit..

Multiply values of Meter [19947, U, V] with..

Multiply values of Meter [19129, U, V] with..

Multiply values of Meter [19917, U, V] with..

Figure 3: *Screenshot of the running system with its different views. In the overview (top), transformer stations are displayed together with all contained meters as rectangles, grouped by distribution bar (mains and medium voltage grids). The background color of the individual meters reflects the analysis result of the expert system, (green=ok, gray=unknown, red=problem). Tooltips give details on discovered problems. The status is propagated upwards in the hierarchy causing the frame of a station to become thick red if any contained meters report problems. Selected items are tinted with blueish color. In the analysis view(bottom) the user can get a detailed explanation on why a certain meter has reported problems and what the expert system recommends to cure this.*

Displaying the analysis process of the expert system is the second major task of the visualization. The analysis view displays the expert system pipeline. Starting on the left side with the pool of events, suspicious elements are transferred to the second stage. In Figure 3 (bottom), the alarms are grouped by the meter they are related to. Selecting a group opens it and displays all contained items, selecting a single alarm highlights the linked domain entity in the overview display. The third column displays all diagnoses and the last one shows the set of recommended actions.

The view draws connections from one stage of the pipeline to the next on demand only, because the total number of links can be quite large which makes them difficult to follow. Thus, the view allows to interactively track the path of specific alarms through the analysis pipeline. The two views are linked, i.e. selecting an element also selects it in the overview, thus providing a spatial context. Where necessary, the camera aligns the overview so that the selected meter becomes visible. Similarly, selecting a recommendation in the analysis view displays a preview of its effect in the overview. This can be, for example, the display of throttling values for energy providers.

5. Conclusion

In this paper we outlined a monitoring system that combines the real-time processing capabilities of an expert system with an operator-friendly multiple view environment. The expert system is designed to work as an analysis pipeline and its workflow is displayed in the visualization. This enables the operator to track the processing path of alarms and verify the correctness of the system. An overview visualization complements this approach by providing a temporal and spatial context to the discovered issues. An in-depth evaluation still needs to be performed, but first indicators show that the combination of expert system and visualization can provide valuable decision support for operators. The integration of renewable energy providers together with weather data has been implemented successfully, but an exhaustive description is beyond the scope of this short paper. This additional information aims to improve the situational awareness of the operators, but must be integrated in the system so that the user can easily understand them. The simulation of the effect that different action alternatives have on the system is an important aspect of our future work and will be further investigated.

References

[CEH*09] Chen M., Ebert D., Hagen H., Laramee R., van Liere R., Ma K.-L., Ribarsky W., Scheuermann G., Silver D.: Data, information, and knowledge in visualization. *Computer Graphics and Applications, IEEE 29*, 1 (jan.-feb. 2009), 12 –19. doi:10.1109/MCG.2009.6. 2

[CMS99] Card S. K., Mackinlay J. D., Shneiderman B. (Eds.): *Readings in information visualization: using vision to think*. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 1999. 2

[CPF09] Clark A., Pavlovski C., Fry J.: Transformation of energy systems: The control room of the future. In *Electrical Power Energy Conference (EPEC), 2009 IEEE* (oct. 2009), pp. 1 –6. doi:10.1109/EPEC.2009.5420865. 2

[GNRM08] Garg S., Nam J., Ramakrishnan I., Mueller K.: Model-driven visual analytics. In *Visual Analytics Science and Technology, 2008. VAST '08. IEEE Symposium on* (oct. 2008), pp. 19 –26. doi:10.1109/VAST.2008.4677352. 2

[GSGC08] Gilson O., Silva N., Grant P., Chen M.: From web data to visualization via ontology mapping. *Computer Graphics Forum 27*, 3 (2008), 959–966. doi:10.1111/j.1467-8659.2008.01230.x. 2

[HAB09] Hay S., Ault G., Bell K.: Control room scenarios on active distribution networks: Early results and next steps. In *Universities Power Engineering Conference (UPEC), 2009 Proceedings of the 44th International* (sept. 2009), pp. 1 –5. 1

[IL08] Infield D., Li F.: Integrating micro-generation into distribution systems; a review of recent research. In *Power and Energy Society General Meeting - Conversion and Delivery of Electrical Energy in the 21st Century, 2008 IEEE* (july 2008), pp. 1 –4. doi:10.1109/PES.2008.4596874. 1

[KE05] Kohlhammer J., Encarnação J. L.: *Knowledge Representation for Decision-Centered Visualization.* GCA-Verlag, 2005. 2

[KSH07] Kyriakides E., Stahlhut J., Heydt G.: A next generation alarm processing algorithm incorporating recommendations and decisions on wide area control. In *Power Engineering Society General Meeting, 2007. IEEE* (june 2007), pp. 1 –5. doi:10.1109/PES.2007.385537. 1

[KSO02] Klump R., Schooley D., Overbye T.: An advanced visualization platform for real-time power system operations. In *Proc. of the 14th Power Systems Computation Conference* (2002). 2

[Lau05] Laufenberg M.: Integration of large-scale visualization systems into a control center. In *Power Engineering Society General Meeting* (June 2005), vol. 3, IEEE, pp. 2712 – 2714. doi:10.1109/PES.2005.1489728. 2

[LQS*10] Li F., Qiao W., Sun H., Wan H., Wang J., Xia Y., Xu Z., Zhang P.: Smart transmission grid: Vision and framework. *Smart Grid, IEEE Transactions on 1*, 2 (sept. 2010), 168 –177. doi:10.1109/TSG.2010.2053726. 2

[MHS07] Mackinlay J., Hanrahan P., Stolte C.: Show me: Automatic presentation for visual analysis. *Visualization and Computer Graphics, IEEE Transactions on 13*, 6 (nov.-dec. 2007), 1137 –1144. doi:10.1109/TVCG.2007.70594. 2

[MMKN08] McLachlan P., Munzner T., Koutsofios E., North S.: Liverac: interactive visual exploration of system management time-series data. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (New York, NY, USA, 2008), CHI '08, ACM, pp. 1483–1492. doi:10.1145/1357054.1357286. 2

[MTF05] Marques A., Taranto G., Falcao D.: A knowledge-based system for supervision and control of regional voltage profile and security. *Power Systems, IEEE Transactions on 20*, 1 (feb. 2005), 400 – 407. doi:10.1109/TPWRS.2004.831672. 2

[OMWC05] Overbye T., Meliopoulos A., Wiegmann D., Cokkinides G.: *Visualization of power systems and components*. Tech. rep., Power Systems Engineering Research Center(PSERC), 2005. 2

[SGBBT06] Shahar Y., Goren-Bar D., Boaz D., Tahan G.: Distributed, intelligent, interactive visualization and exploration of time-oriented clinical data and their abstractions. *Artif. Intell. Med. 38*, 2 (Oct. 2006), 115–135. doi:10.1016/j.artmed.2005.03.001. 2

[SQW11] Shi J., Qiao Y., Wang H.: Visualizing inference process of a rule engine. In *Proceedings of the 2011 Visual Information Communication - International Symposium* (New York, NY, USA, 2011), VINCI '11, ACM, pp. 10:1–10:9. doi:10.1145/2016656.2016666. 2

[TRFBMBM10] Torralba-Rodríguez F. J., Fernández-Breis J. T., Martínez-Béjar R., Montagud V. B.: An incremental knowledge acquisition-based system for critical domains. *Expert Systems with Applications 37*, 4 (2010), 2838 – 2847. doi:10.1016/j.eswa.2009.09.007. 2

[XGH06] Xiao L., Gerth J., Hanrahan P.: Enhancing visual analysis of network traffic using a knowledge representation. In *Visual Analytics Science And Technology, 2006 IEEE Symposium On* (31 2006-nov. 2 2006), pp. 107 –114. doi:10.1109/VAST.2006.261436. 2