Power Application Possibilities with Mission Critical Cloud Computing

This paper is a pre-fomatting version of the following:

1. David Bakken, Thoshitha Gamage, Mani Venkatasubramanian, Anurag Srivastava, Yilu Lui, Anjan Bose, Santiago Grijalva. "Power Application Possibilities for Cloud Computing", in *Smart Grids: Clouds, Communications, Open Source, and Automation*. David Bakken and Kris Iniewski, editors CRC Press, 2014.

Table of Contents

Overview

As we have seen in [13], cloud computing is not only coming to the grid, but mission critical implementations such as GridCloud can provide mission critical properties. This chapter explores new applications that such technology enables. While this chapter is only scratching the surface in what is likely to be routine in a decade, we hope that it provides a tantalizing glimpse of what is possible.

What, then, is mission critical cloud computing? To recap, in a nutshell it:

- Keeps the same fast throughput as generic commercial cloud platforms
- Does not deliberately trade off this throughput to allow “inconsistencies”, for example a replica that does a state update on a copy of the state but this update is “forgotten”
- Is much more predictable (and fast) in terms of rampup time, CPU performance per node, and number of nodes

Therefore, the questions for power application developers are how they can use:

- Hundreds of processors in steady state
- Thousands or tens of thousands of processors when a contingency is reached or is being approached. Note: often there are many minutes of advanced warning on this, sometimes an hour or more.
- Data from all participants in a grid that is enabled quickly when a crisis is approached (though, for market reasons, not necessarily during steady state).

With this in mind, we now present groundbreaking applications that can exploit such mission critical cloud platforms.

Robust Adaptive Topology Control

Pranavamoorthy Balasubramanian and Kory W. Hedman, Arizona State University

The electric power transmission system is one of the most complex systems available today. Traditionally, bulk power transmission systems (lines and transformers) are treated as static assets, even though these resources are controllable. However, it is known that transmission topology control has been used in the past and is still being used for corrective based applications, e.g., PJM uses corrective topology control as a special protection scheme [1]. These switching actions are primarily done on an ad-hoc basis, which are determined by the system operators based on past historical data rather than in an automated way based on decision support tools. Past research has demonstrated the ability of topology control to help improve voltage profiles, increase transfer capacity, improve system reliability, and provide cost benefits to the system [2-8]. Even though transmission topology control can provide these benefits, harnessing such flexibility from the transmission network in existing operational procedures is limited due to the computational challenges to optimize the transmission topology.

More recently, sensitivity based methods have been proposed as a mechanism to reduce the computational complexity [9-12]. The robust adaptive topology control method develops a sensitivity based heuristic, which reduces the computational time of the topology control problem. An expression is derived indicating the impact of changing the state of a transmission line on the objective. This expression is used to generate a line ranking system with the potential candidate lines for switching based on a DCOPF, which builds on the work of [12]. This approach selects a single feasible switching action per iteration, which provides an improvement to the system. The advantage of this method is that it solves linear programs iteratively to come up with a beneficial line switching solution, which is computationally simple as compared to other methods employing mixed integer programming. All the possible switching solutions are lined up in the ranked list with the most likely beneficial switching action placed at the top of the list. As the list is formed based on a sensitivity study, the switching action is not guaranteed to provide improvement to the system. Hence, the switching actions need to be checked for AC feasibility and whether they truly provide an improvement in the objective before it is to be implemented. This is done by selecting the first action from the ranked list and the switching is simulated to find the improvement in the system. If the switching is not beneficial, the next action in the rank list is checked for improvement. This process is continued until a beneficial switching action is found. While such a procedure is a heuristic, prior work has shown substantial economic savings [9] as well as a strong performance in comparison to global optimization techniques [12].

The processing time to come up with a beneficial switching action could be significantly reduced if this process is parallelized so th at all the proposed switching solutions could be checked at once. This opens up enormous opportunities for the application of cloud computing for transmission switching applications, which would drastically reduce the computational time and improve the solution quality as the best solution from the ranking list could be identified very quickly. With prior research demonstrating cost savings of close to 4% for a $500 Billion industry [3], there is a great opportunity for advanced decision support tools to fill this technological need both in terms of algorithm sophistication as well as advanced computing capabilities, like cloud computing.

References

[1] PJM, Manual 3: Transmission Operations, Revision: 40, 2012. [Online]. Available: http://www.pjm.com/~/media/documents/manuals/m03.ashx.

[2] W. Shao and V. Vittal, “Corrective switching algorithm for relieving overloads and voltage violations,” *IEEE Transactions on Power Systems,* vol. 20, no. 4, pp. 1877-1885, Nov. 2005.

[3] K. W. Hedman, M. C. Ferris, R. P. O’Neill, E. B. Fisher, and S. S. Oren, “Co-optimization of generation unit commitment and transmission switching with N-1 reliability,” *IEEE Transactions on Power Systems*, vol. 25, no. 2, pp. 1052-1063, May 2010.

[4] K. W. Hedman, R. P. O'Neill, E. B. Fisher, and S. S. Oren, “Optimal transmission switching with contingency analysis,” *IEEE Transactions on Power Systems,* vol. 24, no. 3, pp. 1577-1586, Aug. 2009.

[5] A. Korad and K. W. Hedman, “Robust corrective topology control for system reliability,” *IEE Transactions on Power Systems,* under review.

[6] K. W. Hedman, R. P. O'Neill, E.B. Fisher, and S. S. Oren, “Smart flexible just-in-time transmission and flowgate bidding,” *IEEE Transactions on Power Systems*, vol. 26, no. 1, pp. 93-102, Feb. 2011.

[7] E. B. Fisher, R. P. O'Neill, and M. C. Ferris, “optimal transmission switching,” *IEEE Transactions on Power Systems*, vol. 23, no. 3, pp. 1346-1355, Aug. 2008.

[8] K. W. Hedman, S. S. Oren, and R. P. O'Neill, “A review of transmission switching and network topology optimization,” *IEEE Power and Energy Society General Meeting,* July 2011.

[9] P. A. Ruiz, J. M. Foster, A. Rudkevich, and M. C. Caramanis, “On fast transmission topology control heuristics,” *IEEE Power and Energy Society General Meeting,* July 2011.

[10] J. M. Foster, P. A. Ruiz, A. Rudkevich, and M. C. Caramanis, “Economic and corrective applications of tractable transmission topology control,” *49th Annual Allerton Conference on Communication, Control, and Computing*, pp. 1302-1309, Sept. 2011.

[11] P. A. Ruiz, J. M. Foster, A. Rudkevich, and M. C. Caramanis, “Tractable transmission topology control using sensitivity analysis,” *IEEE Transactions on Power Systems*, vol. 27, no. 3, pp. 1550-1559, Aug. 2012.

[12] J. D. Fuller, R. Ramasra, and A. Cha, “Fast heuristics for transmission line switching,” *IEEE Transactions on Power Systems*, vol. 27, no. 3, pp. 1377-1386, Aug. 2012.

[13] Thoshitha Gamage, David Anderson, David Bakken, Ken Birman, Anjan Bose, Carl Hauser, Ketan Maheshwari, and Robbert van Renesse. “Mission Critical Cloud Computing for Critical Infrastructures”, in *Smart Grids: Clouds, Communications, Open Source, and Automation*. David Bakken and Kris Iniewski, editors CRC Press, 2014.

Adaptive Real-Time Transient Stability Controls

Vaithianathan “Mani” Venkatasubramanian, Washington State University

Power system is expected to undergo major changes in the next decade. These are from rapid growth in system loads (electric cars) and from increased dependence on renewable intermittent generation. To face up to these challenges, power utilities are making major upgrades to wide-area monitoring and control technologies with impetus from major Federal investments in the past few years.

Power system operation is designed to withstand onset of small and large-scale disturbances. However, when the system is subject to occurrence of several large disturbances in a short span of time, the system may become vulnerable to blackouts. Some of the recent blackouts such as the 2012 San Diego blackout and the 2003 Northeastern blackout point to the need for adaptive real-time transient stability control designs which are specifically designed on an adaptive premise of taking control decisions during the evolution of the event.

In the present day power system, wide-area transient stability controls such as Remedial Action Schemes (RAS) or Special Protection Schemes (SPS) are hard-coded control algorithms which are triggered by a central controller for the occurrence of specific contingencies based on preset switching logic. When the system is subject to any “unknown” set of contingencies that is not part of the RAS controller logic, the system operation typically switches to a “safe mode” where the inter-area power transfers are limited to low conservative settings. The tie-line transfers remain at the safe low values until the reliability coordinator completes a new set of transient stability simulation studies which results in significant economic losses by operating at non-optimal power transfer levels.

Cloud computing emerges as an ideal platform for handling transient stability mitigation issues both for the present day power system and for future control designs. In the present day operation, whenever the system operation is found to be one of the “unknown” operating conditions, the reliability coordinator can dial-in vast amount of cloud based processing power to carry out massive number of new transient stability simulations needed for determining the safe transfer limits.

In the future, we need to rethink the design of the transient stability controls such as RAS or SPS schemes. The massive computational capability offered by cloud computing opens up truly novel futuristic control schemes for mitigating transient stability events such as proposed in [Greg1]. In the present day power system, simulation studies are done off-line for a “guestimated” list of potential contingencies, and RAS schemes are implemented for a subset of problematic *N-2* or higher order contingencies whenever needed. Such RAS schemes then only work for a limited number of potential scenarios. Moreover, the respective control actions in these RAS schemes are also designed to be conservative being based on off-line studies.

[Greg1] proposes to select and implement transient stability controls based on simulations of the system in real-time during the evolution of the events themselves. Wide-area monitoring from abundance of PMUs in the future will pave the way for real-time monitoring of the large power system state and system topology. Combining this real-time state information with real-time simulations will allow us to evaluate what control actions are optimally suited for the system at the present time and the decisions are fully adaptive to whatever the system conditions are. Since the controller continues to monitor the system in a closed-loop fashion, the proposed control schemes are also robust with respect to simulation errors and communication/actuation failures. The formulation is not restrictive to any subset of contingencies and can handle low-probability events consisting of multiple outages such as the ones that have served as precursors to large blackouts of the past.

In this proposed formulation, denoted as Adaptive Real-time Transient Stability (ARTS) controls, massive processing power is needed to carry out “what if” simulations of many potential control candidates in parallel before deciding on whether any control action is needed and which specific action(s) will be implemented. The system monitoring and simulations of what-if scenarios will continue throughout during the event until the system has been stabilized. Once the controller recognizes the system to have returned to normal state, the controller returns to dormant system monitoring mode, and cloud resources can be released. Details of the control algorithms can be seen in [Greg1].

[Greg1] G. Zweigle and V.Venkatasubramanian, “Wide-area Optimal Control of Electric Power Systems with Application to Transient Stability for Higher Order Contingencies”, IEEE Trans. Power systems, Digital identifier 10.1109/TPWRS.2012.2233764.

Prosumer-Based Power Grid

Thoshitha Gamage, Washington State University

Santiago Grijalva, Georgia Institute of Technology

Introduction

The electric power grid, in a bid to improve its sustainability, is aggressively exploring ways to integrate distributed renewable energy generation and storage devices at many levels. The most obvious integration is at the generation level, where renewable generation sources such as large wind turbine and solar panel farms will supplement and eventually (hopefully) supplant traditional non-renewable power generation sources. Another natural integration is at the distribution level, where relatively smaller scale renewable energy generation by utilities and other power distribution entities offer cheaper and greener energy options to customers. While not at the same bulk scale as in generation or distribution level, an emerging trend in recent years is that the end consumers who are typically below the distribution level (e.g. households, microgrids, energy buildings, etc.) generating their own power using renewable sources and becoming self-sustainable and energy-independent from the grid.

A fascinating aspect of this changing energy landscape is the drastic changes in the roles of the players involved. For example, end consumers, in addition to their typical energy consumption role, are economically motivated to sell excess energy and provide energy storage services to the grid. Modern utilities also go beyond their traditional energy distribution role to buy energy from end consumers when available. Similar role augmentations can be observed at all levels of the modern electric power grid [1]. As a consequence, what traditionally has been a one-way energy transfer – from bulk generation through transmission and distribution to end user consumption – is transforming into a two-way energy exchange.

Prosumer-based Control Architecture

The key entities in this evolving electric power grid are *prosumers* who are economically motivated power system participants that can consume, produce, store, and transport energy [1] [2]. Examples for prosumers include a building equipped with an electric vehicle (EV) parking lot that provide storage services to a utility, an hybrid (generates own power but not completely off-grid) renewable energy microgrid who sells excess power at peak generation to a neighboring microgrid or utility, or two households with renewable resources exchanging power on-demand. In fact, almost all current power grid players can be conceptualized as prosumers [2]. A key characteristic of prosumers is that they assume different roles depending on the situation and underlying conditions; a prosumer who is a consumer at one instance can become a produce the next instance. This is similar in analogy to a P2P node vs. a typical client-server model in computer science. Prosumers interact with one another through the services they offer. For example, a utility prosumer who aggregates heterogeneous home user prosumers provide consumption and storage services to a distributed independent service operator (ISO) prosumer. Moreover, at any given instance, different prosumers may operate under different satisfaction/objective function (comfort, cost, efficiency, security, sustainability, reliability, etc.) that will directly influence their control decisions. Thus, the control architecture that revolves around prosumers is radically different from current centralized hierarchical control architecture. Instead of waiting for control decisions to trickle down the hierarchy, a prosumer-based control architecture promotes autonomy by taking proactive and distributed control steps that reflect local, internal, and external state of the network of prosumers.

Computational Challenges

An important but often overlooked trait of the future electric power grid is its multi-dimensional, multi-scale nature. The existing electricity industry is confronted by a massive invasion of energy-aware household ubiquitous devices that aim to empower consumers to make intelligent and mindfulness energy choices. The industry itself is invested in improving the grid's visibility and real-time sensing capabilities by deploying phasor measurement units (PMUs), smart meters, and other intelligent and highly accurate electronic devices. The inadvertent effect of this influx of new devices is that they produce data that have temporal, spatial, and scenario significance that are useful in planning, simulation, operation, and control of the future power grid [3]. If the next evolution of the power grid is a prosumer-oriented two-way energy exchange, the true benefit of such mass quantities of multi-dimensional data profoundly depends on the ability to extract and render usable, useful, and mission-critical information at the right time and deliver them to the right destinations. Inarguably, the computational complexity of this daunting task is beyond what today's general purpose computers can provide. Additional major driving forces such as federal and industry initiatives to incorporate renewable energy resources, various energy efficiency programs, and unconventional and novel system behaviors, coupled with consumer perceptions about the type of services they expect from the future power grids adds multiple scales granularities to this computational challenge.

Cloud Computing in the Future Electric Grid

There are numerous computational challenges and limitations associated with existing computational models and analytical tools [3]. A primary objective of today's power engineers and researchers alike is to design, develop, innovate new models and analytical tools that can support and fully exploit massive multi-dimensional and multi-scale data. Below are few such examples that have cloud computing implications in the potential solution domain.

Economic Dispatch of Stochastic Energy Resources

The inherent unpredictability and the variable nature of stochastic energy resources such as wind energy is making it difficult for utilities with large renewable penetrations to fully exercise real-time economic dispatch with a finer granularity [4]. The research community has already identified the need to replace current economic dispatch software with short-term stochastic scheduling software with seconds, minutes, and few hours granularity [5] [6]. Such software have highly elastic computational needs that scale up and down based on the conditions, the level of penetration, and availability of information, thus offer an invitation to exploit the cloud's rapid elasticity feature [7] to provision software, platform, and computational infrastructure with different configurations on-demand, over in-house computational stacks that offer little flexibility, scalability, and availability.

Cloud-based Apps

Most internal events experienced by prosumers are not visible to external entities. However, there are circumstances where knowledge of such events become significant beyond the immediate control domain. For example, conventional load forecasting within a 2% error bound by ISOs will need to be tightened even further with the deployment of distributed energy resources and utility-scale storage devices [8]. The enablers for such capability are widely available and acceptable new modelling and simulation needs – single- vs. three-phase modelling, non-standardization of operational and planning models, node/breaker vs. bus/branch modelling, data privacy, operational liability – which can be offered as Software-as-a-Service (SaaS) solutions.

Scenario Analysis and Transmission Planning

Fundamentally, the power grid has resorted to using an N-1 contingency analysis of bulk transmission system security. The reality however is that, to achieve a higher level of system reliability, the power grid must be capable of evaluating a large number of plausible contingencies beyond N-1 conditions. This is certainly a computationally very intensive process. A heuristically feasible options would be to systematically prioritize searches and reduce scenario spaces, and to make dynamically updated multi-dimensional information to a wider audience of power grid entities. These are still capabilities beyond general purpose computers and would require a sufficiently large high-performance computing facility. It is highly unrealistic to assume that every power grid entity would be excited about the proposition of individually housing such highly expensive facility, not to mention the highly skilled workforce required to maintain it.

A more likely solution is to utilize cloud computing infrastructures just on the basis of need and usage, while greatly reducing the capital investment needed and abstracting away the underlying maintenance cost and workforce requirements. Cloud computing can commission massive parallel computations to explore the huge permutation space on-demand and also to share the findings much faster to all entities of interest. Even the data aggregation, mining, and analysis tasks can be offloaded on to cloud-based computational platforms that offers better scalability features under changing conditions.

Model Integration

Power system at various spatial scales are modelled differently. Information exchange and seamless modelling integration requires significant processing capabilities. Often, the data from more detailed models has to be aggregated and transferred "by hand" to other applications. Cloud-based modelling solutions on the other hand can greatly improve this process by using in-the-cloud dynamic and highly automated aggregation and delivery capabilities. These capabilities can be encapsulated as a library of models with a consistent description of phenomena and various scales can be generalized and offed as Platform-as-a-Service (PaaS) options with well-defined APIs to system developers who can custom-build models that suite individual spatial scaling requirements of different power system entities. Moreover, entities such as utilities or microgrids that operate individual control systems with homogenous objective functions (e.g. cost minimization) can be aggregated for better control stability. Thus, the natural hierarchy of spatially separated power grid entities can be unified under common protocols while still abiding to any regulatory or compliance standards.

Exploiting and Abstracting Self-Similarity

Self-similarity is the notion of entities at various levels of the system hierarchy exhibiting similar characteristics [9] [10]. This is observable in the power grid where entities at lower level of the system are governed by physics similar to that of higher level entities, but at a lower scale. What this means for a prosumer-based power gird is that multi-scale prosumers at different levels can benefit from models that are based on the same reference model, which enables expansion and collapse of lower level information as well as seamless integration and aggregation at higher levels. Since similar physics takes place at various hierarchical levels of the power system, data at different aggregation levels can be theoretically collapsed or expanded dynamically depending on the required granularity. For example, a combined security constrained optimal power flow (SCOPF) application may start with transmission level congestion management, but certain buses (distribution systems) many be dynamically expanded depending on the solution state to address distribution system internal constraints or conditions. This again raises an important computational issue –

the ability to dynamically expand highly available computational capabilities on-demand – which is yet another cloud computing possibility.

Bibliography

[1] S. Grijalva, M. Costley and N. Ainsworth, "Prosumer-based Control Architecture for the Future Electricity Grid," in *2011 IEEE International Conference on Control Applications (CCA)*, 2011.

[2] S. Grijalva and M. Tariq, "Prosumer-based Smart Grid Architecture Enables a Flat, Sustainable Electricity Industry," in *2011 IEEE PES Innovative Smart Grid Technologies (ISGT)*, Anaheim, CA, 2011.

[3] S. Grijalva, "Grijalva, Santiago. Research Needs in Multi-Dimensional, Multi-Scale Modeling and Algorithms for Next Generation Electricity Grids," in *DOE conference on computing challenges for the next-generation power grid*, ithaca, NY, 2011.

[4] M. H. Albadi and E. F. El-Saadany, "Overview of Wind Power Intermittency Impacts on Power Systems," *Electric Power Systems Research*, vol. 80, no. 6, pp. 627-632, 2010.

[5] P. A. Ruiz, C. R. Philbrick and P. W. Sauer, "Wind Power Day-Ahead Uncertainty Management through Stochastic Unit Commitment Policies," in *IEEE/PES Power Systems Conference and Exposition (PSCE'09)*, 2009.

[6] W. B. Powell, A. George, H. S. Simão, L. A. W. and J. Stewart, "SMART: A Stochastic Multiscale Model for the Analysis of Energy Resources, Technology, and Policy," *INFORMS Journal on Computing*, vol. 24, no. 4, pp. 665-682, 2012.

[7] P. Mell and T. Grance, "The NIST Definition of Cloud Computing: Recommendations of the National Institute of Standards and Technology," The National Institute of Standards and Technology, Gaithersburg, MD, 2011.

[8] S. Fan, L. Chen and W.-j. Lee, "Short-Term Load Forecasting Using Comprehensive Combination Based on Multimeteorological Information," *IEEE Transactions on Industry Applications*, vol. 45, no. 4, pp. 1460-1466, 2009.

[9] J. Maver, "Self-Similarity and Points of Interest," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 32, no. 7, pp. 1211-1226, 2010.

[10] C. Vandekerckhove, "Macroscopic simulation of multiscale systems within the equation-free framework (PhD thesis)," Katholieke Universiteit Leuven , Leuven, Belgium, 2008.

[11] M. Tariq, S. Grijalva and M. Wolf, "Towards a Distributed, Service-Oriented Control Infrastructure for Smart Grid," in *2011 IEEE/ACM International Conference on Cyber-Physical Systems (ICCPS)*, Chicago, IL, 2011.

Wide-Area Frequency Monitoring

Thoshitha Gamage, Washington State University

Yilu Liu, University of Tennessee

Introduction

THE power system frequency monitoring network (FNET), proposed in 2001 and established in 2004, is a pioneering wide-area monitoring system (WAMS) in the United States. Composed of a network of over 200 frequency disturbance recorder (FDR) devices – a member of the phasor measurement unit (PMU) family – that are spread throughout the country, the FNET serves the entire North American power grid through advanced situational awareness techniques, such as real-time event alerts, accurate event location estimation, animated event visualization, and post event analysis [1] [2] [3].

An FDR is a single-phase PMU that is capable of measuring voltage phase angle, amplitude, and frequency from a single-phase voltage source, which can provide useful information for power system even recognition and status estimation. The frequency measurement algorithms in FDRs are highly accurate with virtually zero error in 52-70 Hz frequency range, and real hardware accuracy of ±0.0005 Hz, which outperform even some of the commercial PMUs.

The power grid operates in a very narrow frequency range such that frequency excursions outside this range give clues to the problems. Frequency is a universal parameter across the entire interconnected power system that provides insight into generation electro-mechanical transients, generation-demand dynamics as well as system operations such as load-shedding, breaker reclosing, and capacitor bank switching. This characteristic allows frequency monitoring to be as informative at the distribution level as it is at the transmission level.

FNET Architecture

The overall FNET architecture, illustrated in Figure Fig 1, comprises of a widely deployed network of FDRs that report phasor measurements either to a local processing unit or a remote data center through Ethernet. The outer most layer of the data center is a data concentrator for the incoming field FRD data. The data concentrator performs certain data preprocessing tasks such as creating GPS time-aligned records from the received data before forwarding them to FNET applications and storage agents in the interior data center levels.

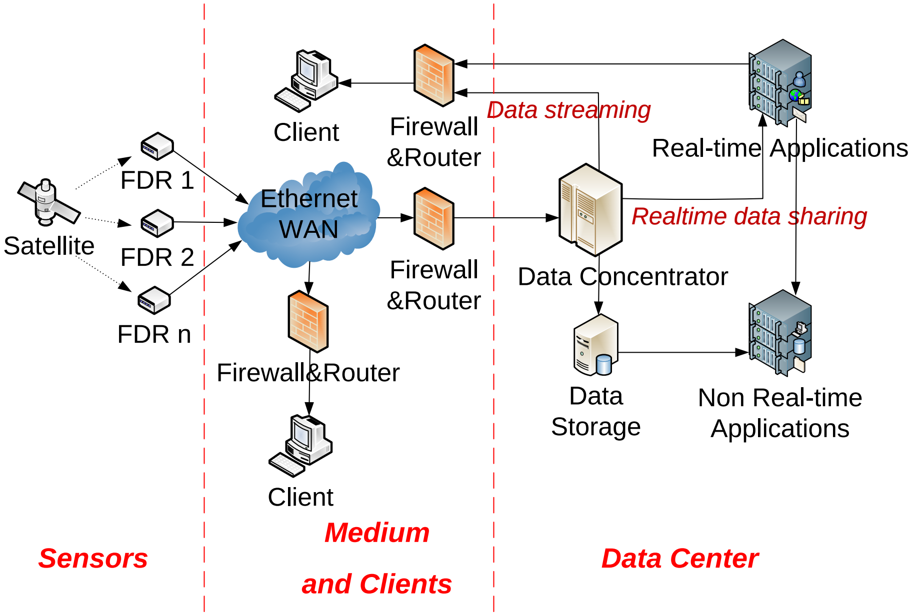

Fig 1: The FNET Architecture is composed of widely deployed FDRs that transmit phasor measurements to either a local client or a remote data center for processing and long term storage. Data centers host a multi-layer agent hierarchy that comprises of a top layer data concentrator, and a data storage agent, and real-time and non-real time application agents in subsequent layers.

FNET Applications

There are two types of FNET application classes – real-time applications, and non-real-time applications [4]. Real-time applications have stringent timing requirements where responses needs to be produced within seconds or sub-seconds after receiving the data. Examples for real-time FNET applications include frequency monitoring, event trigger, inter-area oscillation trigger, line-trip detection and identification, and event visualization. Non-real-time applications on the other hand have much more relaxed timing requirements. Event location, Inter-area oscillation modal analysis, data historians and associated visualizations are some examples for non-real-time applications.

Cloud Computing and FNET

Fundamental to the FNET is the ability to process and analyze measurements from FDRs in an efficient and effective manner. There is an inherent communications requirement even before data arrives at a data center for processing. Timely delivery of data with an adequately high availability is critical for an accurate depiction of the system state.

Rapidly Elastic Data Concentrators

The Data concentrator layer is pivotal to the underlying FNET applications since it acts as the main interface for the incoming phasor measurements to the data center. High availability and high consistency are two key data concentrator properties that ultimately dictate the accuracy of the system state representation for the FNET applications. One way to ensure high availability is by making the data concentrators fault-tolerant through replication. When data concentrators are sufficiently replicated, it also becomes important to make sure that the replica are consistent with each other. Inconsistent replicas adversely affect the underlying applications by forwarding out-of-sync or stale measurement records.

On the measurement side, the granularity of the sensor measurements can rapidly scale up and down based on specific grid conditions. Thus, the data concentrators must be capable of both absorbing large quantities of measurement data arriving at high rates when the grid is stressed as well as operate under normal conditions. This requires an underlying computational infrastructure that can rapidly scale on-demand, which is an inherent feature of the cloud [5].

Computational Requirement Flexibility

The multi-layer computational data center hierarchy has different computational needs at different layers. The outermost data concentrator units are mostly pre-processing units (aligning data as GPS-time synced records) that require more primary memory than secondary memory. Each pre-processed record is streamed immediately to real-time applications and are only buffered when the number of records reach the memory cache limits. Large memory caches results in faster processing times. The data collection and record padding are not interrupted by the saving procedure because the memory cache is designed as a stack; only the oldest portion of the data is saved.

On the contrary, FNET applications have higher CPU and processing requirements than that of data concentrators. Furthermore, different applications can have different configuration requirements (number of CPUs, memory, storage, network, etc.) based on the class of the application (real-time or non-real-time) and the type of spatial scope of the data that is being handled. Clearly, given the large capital investment, skilled workforce requirements, and the maintenance cost, supporting such diverse computational requirements using in-house solutions is not cost effective to any power grid entities. A much more feasible solution would be to take advantage of cloud computing facilities that can provide any type of platform or infrastructure requirement as a service, which greatly improves the overall value proposition.

Cloud-Based Applications

The most intriguing aspect of using cloud computing in the FNET is supporting FNET applications as software-as-a-service (SaaS) [5]. The FNET provides a rich pool of applications for power grid entities at all levels both spatially and temporally. Pre-built applications can be hosted as SaaS for interesting grid entities to use at any scale, based on their specific and individual organizational requirements. Furthermore, the cloud can provide a platform-as-a-service (PaaS) model for application developers to develop new applications under a unified reference model. The data delivery to these applications is already handled by the data concentration level and can be offered as a well-defined API. This not only greatly simplifies the instrumentation aspect of application development, but also helps to standardize the application development lifecycle and data processing pipeline.

Bibliography

[1] B. Qiu, L. Chen, V. Centeno, X. Dong and Y. Liu, "Internet based frequency monitoring network (FNET)," in *IEEE Power Engineering Society Winter Meeting*, 2001.

[2] Y. Liu, "A US-wide power systems frequency monitoring network.," in *IEEE PES Power Systems Conference and Exposition (PSCE'06)*, 2006.

[3] "FNET Web Display," Power Information Technology Lab, University of Tennessee, 2013. [Online]. Available: http://fnetpublic.utk.edu/sample_events.html.

[4] Y. Zhang, P. Markham, T. Xia, L. Chen, Y. Ye, Z. Wu, Z. Yuan, L. Wang, J. Bank, J. Burgett, R. Conners and Y. Liu, " Wide-Area Frequency Monitoring Network (FNET) Architecture and Applications," *IEEE Transactions on Smart Grid*, vol. 1, no. 2, pp. 159-167, 2010.

[5] P. Mell and T. Grance, "The NIST Definition of Cloud Computing: Recommendations of the National Institute of Standards and Technology," The National Institute of Standards and Technology, Gaithersburg, MD, 2011.

Oscillation Mitigation Strategies

Vaithianathan “Mani” Venkaatsubramanian, Washington State University

Real-time oscillation monitoring tools such as Oscillation Monitoring System (OMS) developed at Washington State University based on synchrophasors are being implemented in many control centers all over the world. Excellent progress has been made in the recent past on centralized [OMS, FDD] and distributed algorithms [Alex] for wide-area oscillation detection and analysis. Typically the monitoring tools issue alarms to the operators whenever the damping ratio of any of the local or inter-area modes goes below preset thresholds say 2%. Unbounded growth of negatively damped oscillations can lead to tripping of critical system components and may even lead to blackouts such as the August 10, 1996 western American blackout. Presence of poorly damped or sustained oscillations over a period of time can cause damage to rotor shafts and also impact on consumer power quality with potentially significant economic consequences. Therefore it is important to mitigate poorly damped oscillations as quickly as possible. However, the choice of what control actions should be used to mitigate different types of oscillatory modes is a nontrivial task. As the real-time oscillation monitoring tools get adopted by the industry, there is an urgent need to develop mitigation strategies for how operators can respond to oscillation alarms in the form of specific operator actions whenever such oscillation alarms occur.

Cloud computation is an ideal platform for developing such mitigation control strategies. The problem can be handled in two ways:

a) Base-line studies: Massive analysis of historical synchrophasor data can be carried out using oscillation monitoring engines such as OMS and the results can be recorded together with corresponding SCADA (EMS) data. Correlation studies can be conducted to identify patterns of how different modes are sensitive to different tie-flow flows and generation patterns. Cloud computing enables analysis of different groups of synchrophasor data in parallel for targeting different interarea modes and possibly all the observable local modes of relevant generators.
b) Online simulations: Whenever the oscillation engine issues an alarm, cloud resources can be called upon to carry out a parallel simulation of large number of feasible candidate operator actions to decide which actions are the most effective in mitigating the specific mode related alarm. For instance, a critical tie-line interface may be identified and the amount of tie-line flow to be reduced or increased can be determined based on the simulation studies for improving the damping of a problematic inter-area oscillatory mode.

[OMS] G. Liu, J. Quintero, and V. Venkatasubramanian, “Oscillation Monitoring System based on wide area synchrophasors in power systems,” *Proc. IREP symposium 2007*. Bulk Power System Dynamics and Control - VII, August 19-24, 2007.
[FDD] V. Venkatasubramanian and G.Liu, Systems and methods for electromechanical oscillation monitoring, US Patent 8,000,914, 2011.
[Alex] J. Ning, X. Pan, and V. Venkatasubramanian, “Oscillation modal analysis from ambient synchrophasor measurements using distributed frequency domain optimization”, IEEE Trans. Power Systems, May 2013.

Automatic Network Partitioning for Steady State Analysis

Hao Zhu, University of Illinois at Urbana-Champaign

Cloud computing is emerging as a new paradigm for contemporary computational tasks with explosive demand in computing resources [1]. It offers thousands of computers instant and affordable access to the powerful remote computing platforms. The cloud computing resources can match up to the highest instantaneous demand peaks, while provides the flexibility on the pay-by-use basis when the demand peaks are over. Thanks to the attractive computation it delivers, cloud computing is advocated as a great opportunity for diverse power system computational tasks, ranging from operations and control to economics and planning. [2]

One key factor to enable cloud computing for the power grid would be an efficient network partitioning scheme. Typically, the interconnected bulk electric system is divided into several zones based on specific applications. Perhaps the most immediate example is the Regional Transmission Organization (RTO) territories in the North America for market administration and transmission management purposes [3]. Partitioning the grid into zones with minimal inter-coupling serves as the basis for efficiently paralleling the computational tasks, since it leads to more succinct information exchange among the computers. However, with the growing size of the power grid, it would be an exciting future direction to develop an automatic tool to partition any electric network depending on the level of granularity as well as the application domain.

Benefited from the efficient parallelism, cloud computing can be immediately applied towards power system steady-state analysis, which is routinely performed in all control centers. In this realm, cloud computing with confidentiality is suggested in [4] for the optimal power flow problem – a crucial power system operational task. Nonetheless, it would be more interesting to perform the security analysis when the grid is approaching or reaching stressed conditions, leveraging on the increase of the number of processors from the existing order of dozens to thousands. In fact, the potential of parallel high performance computing has been preliminarily investigated in [5] for massive contingency analysis on the Western Electricity Coordinating Council power grid.

Nonetheless, it is not possible to just directly use cloud offerings to the power grid without some further considerations. For example, the GridCloud architecture project [6], developed by the Cornell and Washington State University researchers, aims to keep the throughput offered by the commercial cloud offerings while getting rid of a deliberate tradeoff of consistencies. This feature may not seem necessary for applications such as web browsing. However, this is of great importance when it comes to power system computational tasks which involve storing replica of e.g., some state variables, at different computers.

Another issue with the cloud offerings lies in that it can take many minutes to ramp up. This motivates the power grid control centers to seek close cooperation with some cloud provider. As another option, separate cloud installations for power grids may be more viable in this sense. For example, RTE in France has its own datacenter, and multiple ISOs in the eastern US grid are reportedly giving serious thought to develop their own GridCloud pilots in 2014.

[1] K. Birman. *Guide to Reliable Distributed Systems: Building High-Assurance Applications and Cloud-Hosted Services*. Springer, 2012.

[2] G. Dán, R. B. Bobba, G. Gross, and R. H. Campbell, “Cloud computing for the power grid: From service composition to assured clouds,” *Proc. 5th Usenix Workshop on Hot Topics in Cloud Computing (HotCloud)*, June 2013.

[3] North American Electric Reliability Corporation (NERC) [Online]. Available: http://www.nerc.com/.

[4] A. R. Borden, D. K. Molzahn, P. Ramanathan, and B. C. Lesieutre, “Confidentiality-preserving optimal power flow for cloud computing,” *Proc. Allerton Conf. Comm., Control, and Computing,* Oct. 2012.

[5] Z. Huang, Y. Chen, and J. Nieplocha, “Massive contingency analysis with high performance computing,” *Proc. IEEE Power & Energy Society General Meeting*, July 2009.

[6] Thoshitha Gamage, David Anderson, David Bakken, Ken Birman, Anjan Bose, Carl Hauser, Ketan Maheshwari, and Robbert van Renesse. “Mission Critical Cloud Computing for Critical Infrastructures”, in Smart Grids: Clouds, Communications, Open Source, and Automation. David Bakken and Kris Iniewski, editors CRC Press, 2014.