

SANDIA REPORT

SAND2011-0890

Unlimited Release

Printed February 2011

Final Report for “High Performance Computing for Advanced National Electric Power Grid Modeling and Integration of Solar Generation Resources”, LDRD Project No. 149016

David A. Schoenwald, Jason E. Stamp, Joshua S. Stein, Robert J. Hoekstra, Jeffrey S. Nelson, Karina Munoz, William C. McLendon, Thomas V. Russo, Laurence R. Phillips, Bryan T. Richardson, Andrew C. Riehm, Paul R. Wolfenbarger, Brian M. Adams, Matthew J. Reno, Clifford W. Hansen, and Ron A. Oldfield

Prepared by
Sandia National Laboratories
Albuquerque, New Mexico 87185 and Livermore, California 94550

Sandia National Laboratories is a multi-program laboratory managed and operated by Sandia Corporation, a wholly owned subsidiary of Lockheed Martin Corporation, for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-AC04-94AL85000.

Approved for public release; further dissemination unlimited.



Issued by Sandia National Laboratories, operated for the United States Department of Energy by Sandia Corporation.

NOTICE: This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government, nor any agency thereof, nor any of their employees, nor any of their contractors, subcontractors, or their employees, make any warranty, express or implied, or assume any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represent that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government, any agency thereof, or any of their contractors or subcontractors. The views and opinions expressed herein do not necessarily state or reflect those of the United States Government, any agency thereof, or any of their contractors.

Printed in the United States of America. This report has been reproduced directly from the best available copy.

Available to DOE and DOE contractors from

U.S. Department of Energy
Office of Scientific and Technical Information
P.O. Box 62
Oak Ridge, TN 37831

Telephone: (865) 576-8401
Facsimile: (865) 576-5728
E-Mail: reports@adonis.osti.gov
Online ordering: <http://www.osti.gov/bridge>

Available to the public from

U.S. Department of Commerce
National Technical Information Service
5285 Port Royal Rd.
Springfield, VA 22161

Telephone: (800) 553-6847
Facsimile: (703) 605-6900
E-Mail: orders@ntis.fedworld.gov
Online order: <http://www.ntis.gov/help/ordermethods.asp?loc=7-4-0#online>



Final Report for “High Performance Computing for Advanced National Electric Power Grid Modeling and Integration of Solar Generation Resources”, LDRD Project No. 149016

David A. Schoenwald, Jason E. Stamp, Joshua S. Stein, Robert J. Hoekstra, Jeffrey S. Nelson, Karina Munoz, William C. McLendon, Thomas V. Russo, Laurence R. Phillips, Bryan T. Richardson, Andrew C. Riehm, Paul R. Wolfenbarger, Brian M. Adams, Matthew J. Reno, Clifford W. Hansen, and Ron A. Oldfield

Sandia National Laboratories
P.O. Box 5800
Albuquerque, New Mexico 87185-MS1321

Abstract

Design and operation of the electric power grid (EPG) relies heavily on computational models. High-fidelity, full-order models are used to study transient phenomena on only a small part of the network. Reduced-order dynamic and power flow models are used when analysis involving thousands of nodes are required due to the computational demands when simulating large numbers of nodes. The level of complexity of the future EPG will dramatically increase due to large-scale deployment of variable renewable generation, active load and distributed generation resources, adaptive protection and control systems, and price-responsive demand. High-fidelity modeling of this future grid will require significant advances in coupled, multi-scale tools and their use on high performance computing (HPC) platforms. This LDRD report demonstrates SNL’s capability to apply HPC resources to these 3 tasks:

- High-fidelity, large-scale modeling of power system dynamics
- Statistical assessment of grid security via Monte-Carlo simulations of cyber attacks.
- Development of models to predict variability of solar resources at locations where little or no ground-based measurements are available.

CONTENTS

1. INTRODUCTION.....	9
2. SCALABLE HIGH FIDELITY ELECTRIC POWER GRID MODELING.....	10
2.1. Introduction.....	10
2.2. Power Grid Component Models using PSpice Circuit Elements.....	11
2.3. Design of Example Power Grid.....	15
2.4. Data Visualization.....	17
2.5. Summary.....	19
3. USING HIGH-PERFORMANCE COMPUTING TO CALCULATE RELIABILITY IMPACTS FROM CYBER ATTACK ON ELECTRIC POWER SYSTEMS	20
3.1. Introduction.....	20
3.2. Reliability Analysis for Cyber Attack.....	21
3.2.1. RICA Approach.....	21
3.2.2. Attacks against Protection.....	23
3.2.3. Attacks against SCADA.....	24
3.3. Rica Computation and Workflow.....	24
3.3.1. Load Flow Computation.....	24
3.3.2. Optimizer.....	25
3.3.3. Pre-Processor.....	26
3.3.4. Scheduler.....	26
3.3.5. Database.....	27
3.4. Rica Simulations and Results.....	30
4. PV OUTPUT VARIABILITY MODELING USING SATELLITE IMAGERY.....	31
4.1. Introduction.....	31
4.2. Modeling Approach.....	32
4.3. Sources of Ground Measurements of Irradiance.....	33
4.4. Sources of Satellite Imagery.....	34
4.5. Image Stabilization.....	37
4.5.1. Image Stabilization Algorithm.....	37
4.5.2. High Performance Computing for Image Stabilization.....	39
4.6. Cloud Detection.....	40
4.6.1. Thresholding.....	40
4.6.2. Movement Detection.....	41
4.6.3. Background Subtraction.....	41
4.7. Next Steps.....	45
5. CONCLUSIONS.....	47
6. REFERENCES	49
DISTRIBUTION.....	53

FIGURES

Figure 1. Equivalent circuit for d-axis of generator.....	12
Figure 2. Equivalent circuit for q-axis of generator.....	12
Figure 3. Block diagram for the turbine with governor.....	13
Figure 4. Block diagram for the excitation system.....	14
Figure 5. Block diagram relating different components of an example grid.....	15
Figure 6. Circuit diagrams for a Y-Y three phase transformer.....	16
Figure 7. Equivalent circuit of an induction motor used to represent dynamic load elements....	16
Figure 8. Screen shot of graphical user interface displaying power grid data for the State of New Mexico.....	18
Figure 9. Possible states for grid elements using the RICA approach.....	21
Figure 10. Executing RICA on the Red Mesa HPC.....	25
Figure 11. Network diagram of the 300-bus system.....	27
Figure 12. Histogram showing number of cases for served-load percentage intervals for the 300-bus system.....	28
Figure 13. Pie chart showing fraction of cases of served-load percentage for the 300-bus system.....	28
Figure 14. Line showing served-load percentage from simulation from beginning to end for the 300-bus system.....	29
Figure 15. Modeling approach for generating irradiance from two images 15 minutes apart.....	32
Figure 16. GOES 11 images of Las Vegas region for 6/04/2008 around 4PM (PST) with corresponding measured irradiance at two ground locations.....	31
Figure 17. Close-up of the Las Vegas Metropolitan Area showing the location of 1-minute irradiance sensors.....	34
Figure 18. Global horizontal irradiance measured at 1-minute intervals at Fort Apache station over a two day period.....	34
Figure 19. Map of Study Area. State boundaries in white, locations of irradiance stations in the Las Vegas area are shown.....	36
Figure 20. Example of GOES-11 Image with approximate state boundaries shown. The locations of two irradiance sensors are shown for reference.....	36
Figure 21. Illustration of image stabilization algorithm: (a) image after jitter correction; (b) difference between corrected and reference images; (c) difference between corrected and reference image.....	38
Figure 22. Histograms of difference in pixel values between reference image and uncorrected and corrected images.....	39
Figure 23. Example of Cloud Detection using Thresholding: raw image (left) and detected clouds (colored features, right).....	40
Figure 24. Example of Cloud Detection using Movement Detection: raw image (left) and detected clouds (white features, right).....	41
Figure 25. Example of Cloud Detection using Background Subtraction: raw image (left) and detected clouds (white features, right).....	42
Figure 26. Diurnal variation in pixel intensity during clear skies.....	43
Figure 27. Annual variation in average pixel intensity during clear skies.....	43
Figure 28. Comparison of statistics for pixel intensity between clear sky images and synthetic images: (a) annual average pixel intensity; (b) diurnal variation in pixel intensity.....	44

Figure 29. Images for 8/19/07 at 11:00AM (PST) centered around Fort Apache on a clear day from a) GOES-11 and b) Neural Network Simulation of the background image 45

Figure 30. Measured and Simulated Irradiance for Fort Apache at 1 minute resolution for May 27, 2008. Note: Inset images illustrate cloud patterns in vicinity of Fort Apache during periods of variability in irradiance..... 46

TABLES

Table 1. Device Forensics Intervals For Cyber Attack..... 29

Table 2. Restart Delays For Generators After Cyber Attack..... 29

Table 3. Restart/Reconnection Delays For Grid Elements After Cyber Attack 29

Table 4. Change In WECC Reliability Indices, With Cyber Attack (Forensics Intervals As Shown In Table I) 30

Table 5. Change In WECC Reliability Indices, With Cyber Attack (Forensics Intervals Halved Compared To Table IV) 30

Table 6. Irradiance Monitoring Stations used in this Study..... 33

Table 7. List of available data from GOES Satellites..... 35

Table 8. Image Scan Routine for GOES-11..... 35

NOMENCLATURE

AC	Alternating Current
APTB	Average Percentage of Tripped Breakers
C2P	Cyber to Physical
DC	Direct Current
DOE	Department of Energy
DOI	Duration of Interruption
EIR	Energy Index of Reliability
ENSI	Energy Not Served per Interruption
EPG	Electric Power Grid
GUI	Graphical User Interface
HL	Hierarchical Level
HPC	High Performance Computing
IEEE	Institute of Electrical & Electronics Engineers
IGFOV	Instantaneous Geometric Field of View
kV	kilo (1000) Volts
LCI	Load Curtailed per Interruption
LDRD	Laboratory Directed Research & Development
LOEE	Loss of Energy Expectancy
LOLE	Loss of Load Expectancy
MC	Monte Carlo
MTTA	Mean Time to Attack
MTTF	Mean Time to Failure
MTTR	Mean Time to Recover
NSTB	National SCADA Test Bed
RICA	Reliability Impacts from Cyber Attack
SCADA	Supervisory Control and Data Acquisition
SNL	Sandia National Laboratories
WECC	Western Electricity Coordinating Council

1. INTRODUCTION

Design and operation of the electric power grid (EPG) relies heavily on computational models. However, computational techniques in use today for modeling the EPG have not scaled beyond the workstation level. Current modeling entails accepting tradeoffs between model fidelity and electrical grid size. High-fidelity, full-order models are used to study transient phenomena and control interactions; computational limitations dictate that only a small part of the network be represented explicitly. On the other hand, reduced-order dynamic and power flow models are used when analysis involving thousands of nodes are required. The reduced order approach works, in part, because the system is normally operated with narrow voltage and frequency tolerances, where system behavior is approximately linear or well understood.

The level of complexity of the future EPG will dramatically increase due to large-scale deployment of variable renewable generation, nonlinear controls, active load and distributed generation resources, adaptive protection and control systems, price-responsive demand, and full integration of information networks. The additional complexity introduces a range of new forcing functions or system interactions, ranging from weather-driven generation variability to malicious cyber attacks, which can propagate through the system and degrade reliability over wide regions. Modeling tools may require the simulation of increasingly complex scenarios with higher-fidelity models and much larger numbers of nodes; modeling of this future grid will require significant advances in the use of HPC for higher fidelity, coupled, multi-scale tools.

This LDRD report demonstrates SNL's capability to apply HPC resources, including Sandia's Xyce software and the Red Mesa HPC platform, to three specific challenges related to power system modeling and simulation: :

- High-fidelity, large-scale modeling of power system dynamics
- Statistical assessment of grid security by evaluating impacts from cyber attacks on large networks, using Monte-Carlo techniques
- Development and validation of PV output variability models that would extend the state of the art for solar resource modeling.

Together, these three application areas represent fundamental challenges that need to be overcome to maintain system performance and reliability with increased complexity, high penetration of renewable resources, and full integration of information technology, which increases cyber-security concerns.

2. SCALABLE HIGH FIDELITY ELECTRIC POWER GRID MODELING

Traditional grid models for large-scale simulations assume linear and quasi-static behavior allowing very simple models of the systems. In this section, a scalable electric circuit simulation capability is presented that can capture a significantly higher degree of fidelity including transient dynamic behavior of the grid as well as allowing scaling to a regional and national level grid. A test case presented uses simple models, e.g. generators, transformers, transmission lines, and loads, but with the scalability feature it can be extended to include more advanced non-linear detailed models. The use of this scalable electric circuit simulator will provide the ability to conduct large-scale transient stability analysis as well as grid level planning as the grid evolves with greater degrees of penetration of renewables, power electronics, storage, distributed generation, and micro-grids.

2.1. Introduction

The scalability of simulation models for a wide range of power systems components has not been explored in significant detail. Dynamic models of the electric power grid (EPG) are divergent from the existing classes of electrical systems problems being solved in electric circuit simulators such as PSpice™. The dynamic analysis of large-scale power grids needs an advancement of high fidelity scalable tools capable of addressing the future architecture of the EPG. Currently, power grid models are either high level aggregated models (e.g. PSLF™, PowerWorld™) or low level high fidelity (e.g. Simulink™/SimPowerSystems™, PSpice™). The ability to analyze the impact of low level circuits (e.g. Photovoltaic arrays) on a large scale is missing.

By using a parallel electric circuit simulator, developed at Sandia National Laboratories, Xyce™, the ability to model individual electric power grid components and group them into successively larger circuits that can replicate a large scale grid has been achieved. This results in a unique analytical capability for the power grid simulation field. Xyce™ has the ability to model highly complex circuits with very large numbers of nodes. This ability is being leveraged to extend Xyce™ to the electric power grid by using electric circuit elements to model the various components of the power grid. The ability to analyze the impact of high levels of penetration of solar PV, wind, fuel cells, and storage can be analyzed with such a tool. This capability is needed to determine the impact of high levels of renewables, distributed generation, and storage in the future EPG.

The use of electric circuit elements in Xyce™ for the EPG has some distinct advantages:

- The ability to model EPG as a modular scale-up of electrical circuit components.
- The ability to handle a very large scale network via parallelizable solvers.
- The ability to interface with a graphical user interface to display simulation values on a grid map.

Though some of the EPG components (specifically generators) are not easily modeled as electrical circuits, the re-use of models developed in other platforms (e.g. Matlab™) is currently being addressed.

The basic EPG components modeled consist of generators (including the prime mover, governor, and exciter circuit), transformers (three phase and single phase) transmission lines (both AC and DC), and loads (static and dynamic). Constructing a power grid example begins with these basic components. To construct a realistic EPG, some assumptions were made on types of loads to be modeled as well as the size and types of neighborhood transformers, feeders, and substations to be represented.

There are three different load types being modeled: residential, commercial, and industrial. The only difference between the load types are the percentage of the load that is static vs. dynamic. Typically, residential loads are 80% static and 20% dynamic. Commercial loads are approximately 50% static and 50% dynamic. Industrial loads are generally 10% static and 90% dynamic. Static loads are represented as variable resistors. Dynamic loads are represented as induction motors. More sophisticated load models can be designed as well. Within each load type there are 3 residential sub-types: a medium home, a large home, and a medium apartment complex, 3 commercial sub-types: small, medium, and large, and 4 industrial sub-types: small, medium, large, and extra-large. For each of these sub-types, data for typical average power loads is used to determine how large the static and dynamic loads need to be in terms of power draw.

The build-up of the substation circuits progresses from loads to the distribution transformer level to the feeder circuit then to the substation level. For the higher levels, 9 distribution transformer types were defined, 6 feeder circuits were defined, and 7 substation types were defined. Finally, at the grid level, the substations are connected to generators via transmission lines using the circuit models corresponding to these components. A realistic EPG based on the state of New Mexico was constructed using the above procedure.

Data visualization was accomplished by developing a Google™ map based graphical user interface (GUI). The GUI can display actual EPG nodes and edges (generators, substations, transmission lines) overlaid on a geographical map while displaying voltage and power time series data for selected nodes.

2.2. Power Grid Component Models using PSpice Circuit Elements

The basic model to be used for the electric power generators is based on [1]-[2] and depicted in Figs. 1-4.

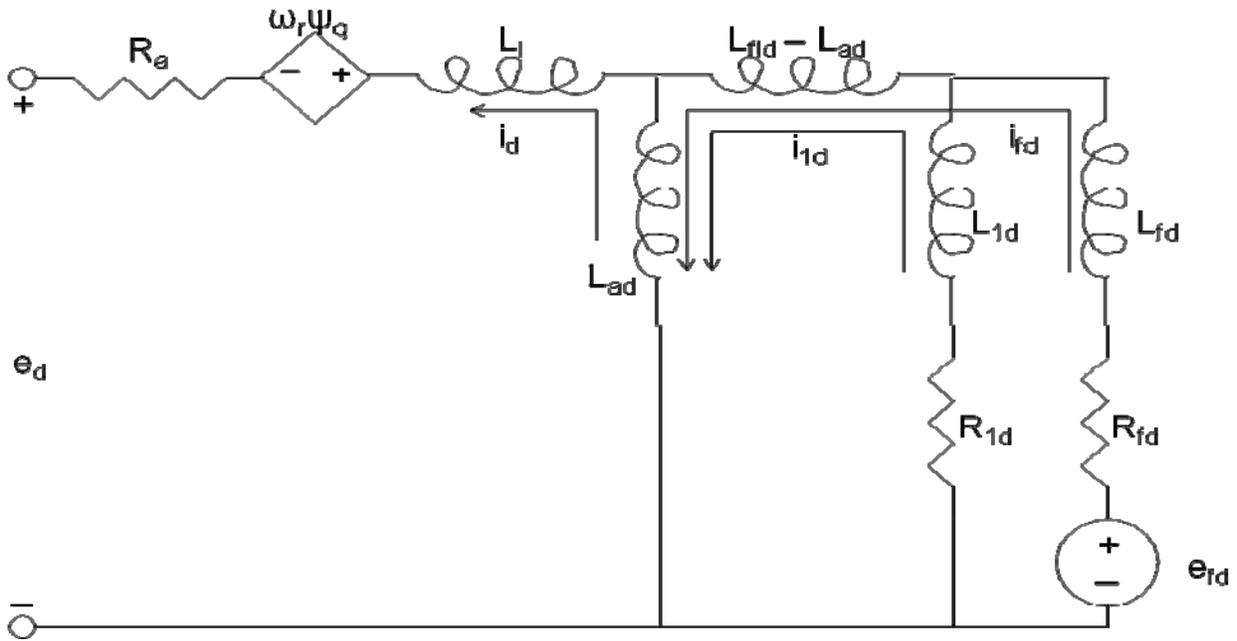


Figure 1. Equivalent circuit for d-axis of generator.

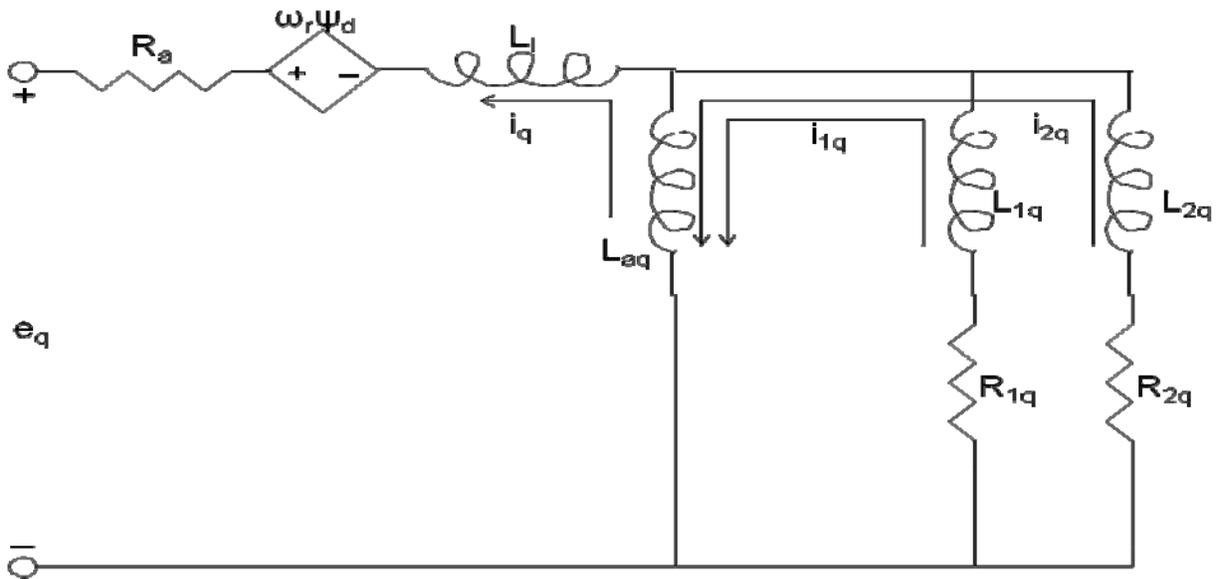


Figure 2. Equivalent circuit for q-axis of generator.

The inductances and resistances R_a , L_l , L_{ad} , L_{1d} , R_{1d} , L_{fd} , R_{fd} , $(L_{fd} - L_{ad})$, L_{aq} , L_{1q} , L_{2q} , R_{1q} , R_{2q} , will be known *a priori* and are constant.

Note that the currents i_d , i_{1d} , i_{fd} , i_q , i_{1q} , i_{2q} are defined in the circuit diagrams as loop currents. To determine the values of the dependent voltage sources in the above circuits, the following two equations are needed:

$$\Psi_d = - (L_{ad} + L_l) i_d + L_{ad} i_{fd} + L_{ad} i_{1d} \quad (1)$$

$$\Psi_q = - (L_{aq} + L_l) i_q + L_{aq} i_{1q} + L_{aq} i_{2q} \quad (2)$$

where ψ_d and ψ_q are the flux linkages of the d and q-axis circuits, respectively and the loop currents and inductances are as defined in Figs. 1-2.

Then, the generator swing equations (equations of motion) must be solved to determine ω_r , the angular swing velocity in rad/s to complete the expressions in the dependent voltage sources:

$$p\Delta\omega_r = (1/2H) (T_m - T_e - K_D\Delta\omega_r) \quad (3)$$

$$p\delta = \omega_0 \Delta\omega_r \quad (4)$$

where

$$\omega_0 = 2\pi 60 \quad (5)$$

$$\Delta\omega_r = \omega_r - \omega_0 \quad (6)$$

with p being the derivative operator d/dt.

Also needed is the generator output power. The active power is given by

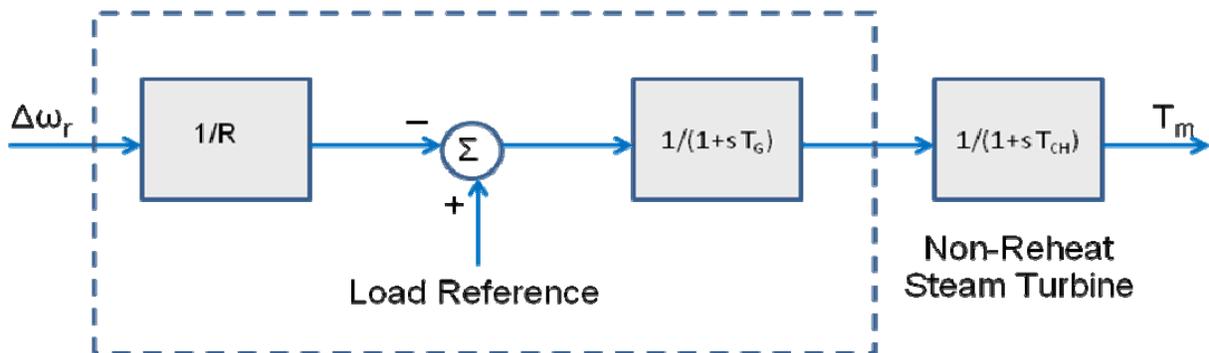
$$P_t = e_d i_d + e_q i_q \quad (7)$$

and the reactive power is given by

$$Q_t = e_q i_d - e_d i_q \quad (8)$$

T_e is the electromagnetic torque of the generator and is given by:

$$T_e = \psi_d i_q - \psi_q i_d \quad (9)$$



Proportional Speed Regulator
Governor with Load Reference

Figure 3. Block diagram for the turbine with governor.

T_m is the generator input power and is the output of the prime mover model (e.g. turbine and governor). It is being modeled as a non-reheat steam turbine with a proportional speed regulator governor and optional load reference setting as shown in Fig. 3 where

$$(T_{CH}T_G)p^2T_m + (T_{CH} + T_G)pT_m + T_m = -(1/R)\Delta\omega_r - (1/R)(LoadRef). \quad (10)$$

Parameters and initial conditions for the above equation include:

$$R=0.05, T_{CH} = 0.3 \text{ sec}, T_G = 0.2 \text{ sec}, LoadRef = 0.0, T_m(0) = 0.0, pT_m(0) = 0.0.$$

Finally, e_{fd} is the generator field voltage and is the output of the excitation system as shown in Fig. 4. We'll use an automatic voltage regulator (AVR) to determine e_{fd} . In many large electric power generators, a power system stabilizer (PSS), may be incorporated, but we omit that here. The following equations describe the excitation system:

$$pv_1 + (1/T_R)v_1 = (1/T_R)E_t \quad (11)$$

where $E_t = \sqrt{e_d^2 + e_q^2}$ and

$$E_{fd} = K_A (V_{ref} - v_1) \quad (12)$$

subject to saturation, e.g., if $E_{fd} \geq E_{FMAX}$ then $E_{fd} = E_{FMAX}$ and if $E_{fd} \leq E_{FMIN}$ then $E_{fd} = E_{FMIN}$. Then

$$e_{fd} = (R_{fd}/L_{ad})E_{fd}. \quad (13)$$

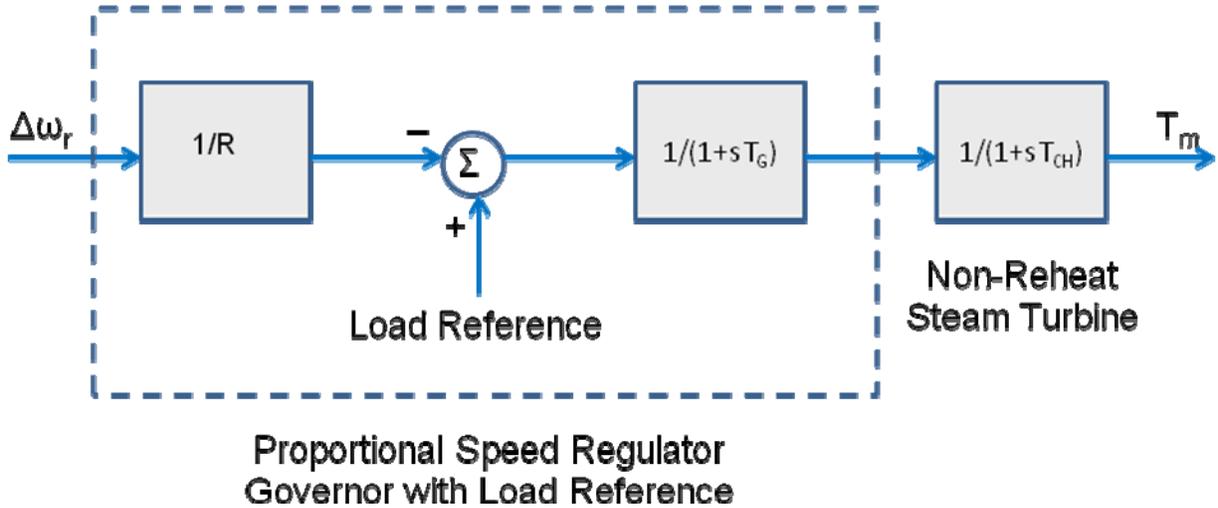


Figure 4. Block diagram for the excitation system.

For this exciter, the following parameters and initial conditions are used:

$$V_{ref} = (1/K_A)(L_{ad}/R_{fd}) e_{fd}(0) + v_1(0), T_R = 0.015 \text{ sec}, K_A = 200, E_{FMAX} = 7.0, E_{FMIN} = -6.4, e_d(0) = 0.631, e_q(0) = 0.776, e_{fd}(0) = 0.000939, v_1(0) = E_t(0) = \sqrt{e_d^2 + e_q^2} \text{ at } t=0.$$

The following are parameter values used in the testing of the generator equivalent circuit model, assuming a 60 Hz 3-phase round rotor (2 pole) synchronous generator rated at 555 MVA at 24 kV, with a power factor of 0.9. All of these values are in per unit, $H = 3.5$, $K_D = 0.3$, $R_a = 0.003$, $R_{1d} = 0.0284$, $R_{fd} = 0.0006$, $R_{1q} = 0.00619$, $R_{2q} = 0.02368$, $L_l = 0.15$, $L_{ad} = 1.66$, $L_{1d} = 0.1713$, $L_{fd} = 0.165$, $L_{aq} = 1.61$, $L_{1q} = 0.7252$, $L_{2q} = 0.125$, $L_{fld} = L_{ad} = 1.66 \implies$ the inductance in the d-axis subcircuit, $L_{fld} - L_{ad} = 0$ (which is a typical assumption).

2.3. Design of Example Power Grid

In order to build up an example power grid from basic electric circuit components, the following four elements will form the basic building blocks of our example grid: loads, transformers, transmission lines, and generators as shown in Fig. 5. Generators have already been defined from basic circuit elements in Sec. 1.2. Transmission lines are often defined as basic circuit elements in circuit modeling software such as PSpice™. More sophisticated transmission line models can be developed but are omitted here. The basic three-phase transformer is modeled using the Y-Y connection as shown in Fig. 6 [3]-[4]. The loads are built up from static and dynamic load elements. A static load is modeled here as a voltage controlled current source using PSpice™

$$\text{gload n1 n2 value} = \{1/(R_{MIN}/v(n1,n2) + v(n1,n2)/PLOAD)\} \quad (14)$$

which behaves like a resistor of value R_{MIN} at low voltages and a constant power load of $PLOAD$ at high voltages. A very high percentage of industrial loads behave like induction motors, thus we model a dynamic load as an induction motor depicted in Fig. 7 [5]-[6].

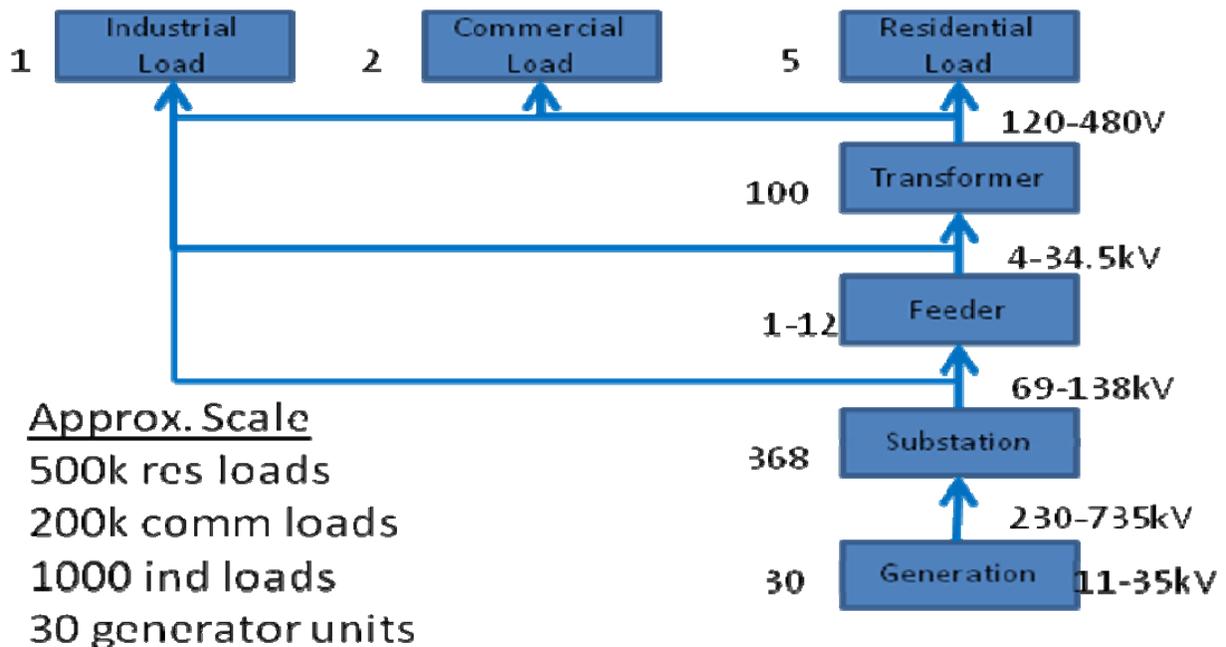


Figure 5. Block diagram relating different components of an example grid.

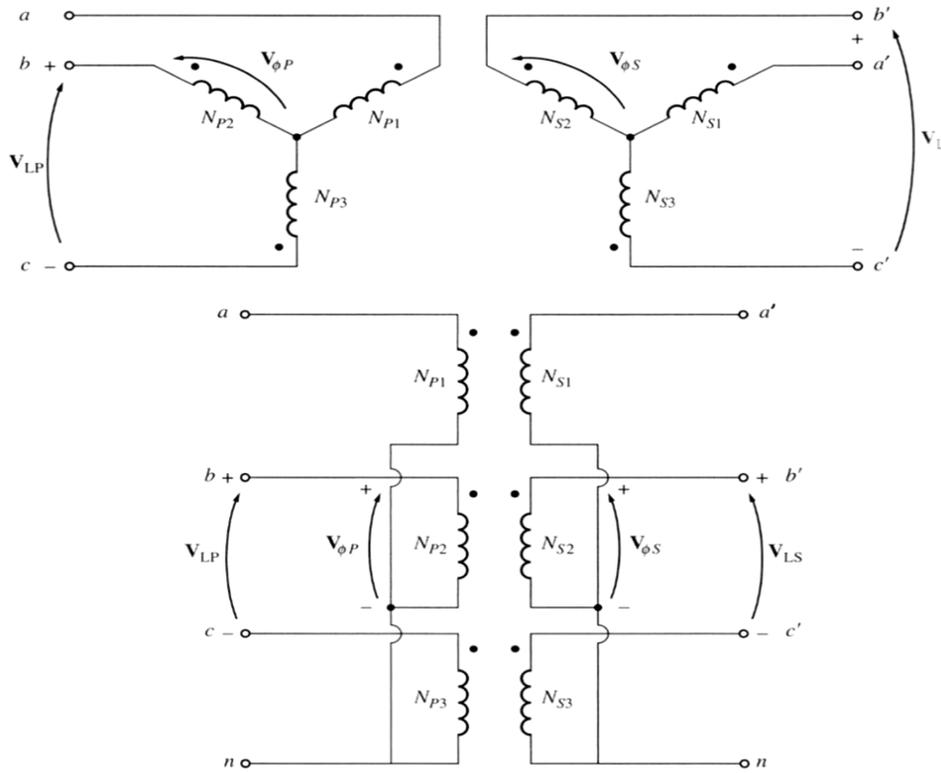


Figure 6. Circuit diagrams for a Y-Y three phase transformer.

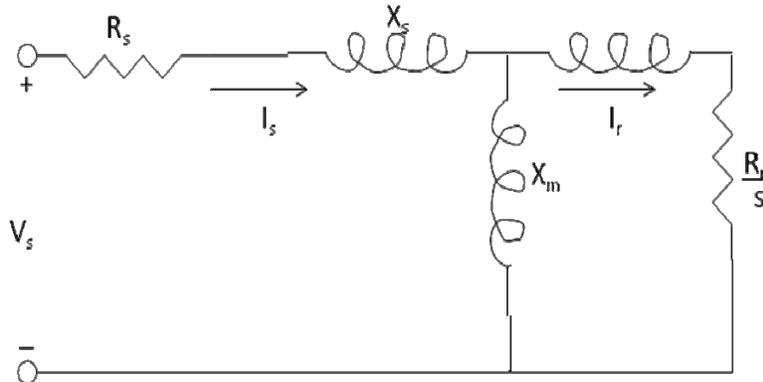


Figure 7. Equivalent circuit of an induction motor used to represent dynamic load elements.

From (14) and Fig. 7, three different load types are defined: Residential, Commercial, and Industrial. Subtypes are also defined within each type. For instance, residential loads are defined as 20% dynamic and 80% static with an average power factor of 0.95. Subtypes within the residential load include a medium-sized home (2300 sq. ft. with 1.75kW average power load), large-sized home (3000 sq. ft. with 2.25kW average power load), and a medium-sized apartment complex (75000 sq. ft. with 56.25kW average power load). Commercial loads are defined as 50% dynamic and 50% static with an average power factor of 0.9. Subtypes within the commercial load include a small load (1500 sq. ft. with 0.675kW average power load), a medium load (50000 sq. ft. with 22.5kW average power load), and a large load (200000 sq. ft. with 90kW average power load). Examples of these sub-types include a gas station (small), a grocery store (medium), and a big box store (large). Industrial loads are defined as 90% dynamic

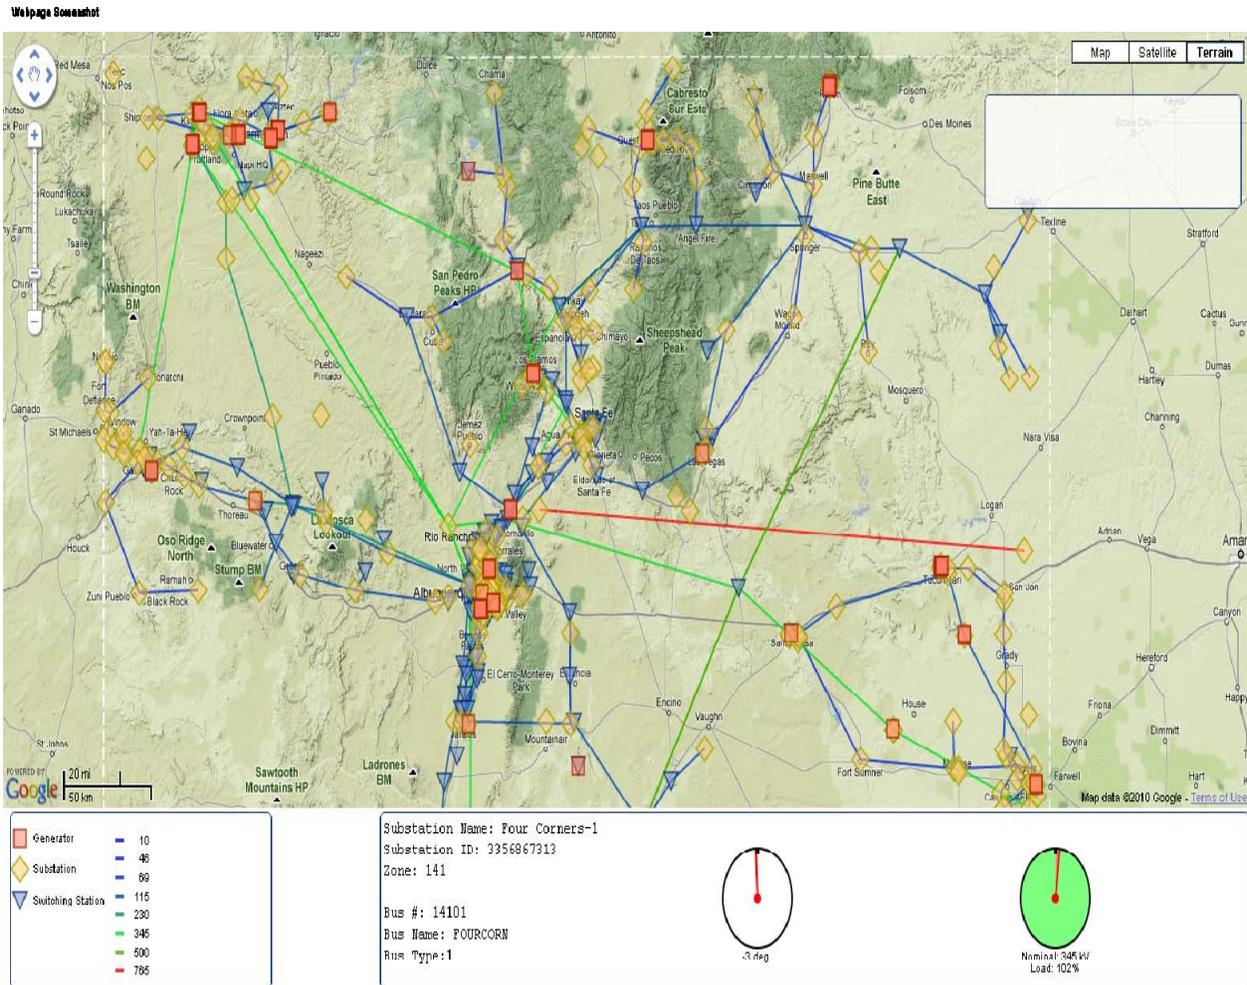
and 10% static with an average power factor of 0.85. Subtypes within the industrial load include a small load (200kW average power load), a medium load (500kW average power load), a large load (1MW average power load), and an extra-large load (3MW average power load). The advantage of using circuit elements for these loads is that the only difference between the load types is the % of the load that is static vs. dynamic.

The build-up of the substation circuits proceeds from load level to the distribution transformer level to the feeder level to the substation level as shown in Fig. 5. In order to build up realistic power grids, we define multiple types for each level similar to what was done with loads. That is, there are nine distribution transformer types, DTR1, DTR2, DTR3, DTC1, DTC2, DTC3, DTI1, DTI2, and DTI3, and six feeder types, Residential, Commercial, Industrial 1, Industrial 2, Industrial 3, and Mixed, and seven substation types, A, B, C, D, E, F, and G. Details of how each of these levels is assembled from the lower levels are addressed in the Appendix.

2.4. Data Visualization

A graphical user interface (GUI) was developed using a graph structure with the edges representing transmission lines and the nodes representing substations, switching stations, and generators. Fig. 8 depicts a screen shot of the GUI with node and edge data corresponding to power grid data from the State of New Mexico.

To differentiate node types, generators are denoted as red squares, switching stations are depicted as blue triangles, and substations are drawn as yellow diamonds. Each node has longitude and latitude coordinates allowing the nodes to be overlaid on a map application. In this case, a Google™ map based application is employed. The topology information is input separately from the simulation variables, which will vary with different runs. The simulation variables consist of 4 or 5 variables over time, with the duration being a couple of minutes sampled at sub-second intervals (potentially at much shorter sampling times). The simulated variables at each node include real and reactive power, voltage magnitude and phase angle, and frequency.



file:///C:/Users/04/dp/dp/gid.html

Figure 8. Screen shot of graphical user interface displaying power grid data for the State of New Mexico.

In the screen shot of Fig. 8, the lower left portion of the GUI contains a legend for the nodes and for the transmission lines. The transmission lines are color coded according to voltage category (e.g. 115 kV, 230 kV, 345 kV, etc.). In addition, if the screen cursor hovers over an edge of interest, the voltage rating for that transmission line will be displayed on the screen. The nodes are selectable, allowing one to display the time series plots in a control panel at the bottom of the screen. In Fig. 8 above, nominal values for voltage magnitude and phase angle are displayed in a console format representing a nominal base case simulation. The GUI is still in a state of further development. In the future, animation controls to view simulation values dynamically change over time will be added. In addition, some of the simulation runs may want to remove edges and nodes and evaluate the impact of a transmission line outage or a generator failure. Further, since the simulation captures detail down to the electric circuit level, a drill down capability will be added to the GUI allowing for time series values of individual loads and transformers.

2.5. Summary

A modeling technique that combines the high fidelity of electric circuit models with the scalability of grid level elements was developed as a tool for the use of transient as well as steady state simulation of electric power grids. Circuit simulation for electric power grid networks has some distinct advantages:

- Ability to model grid as a modular scale-up of electrical components.
- Ability to handle a very large scale network via parallelizable solvers using Xyce™ [7].
- Ability to interface with GUI to display simulation values on a grid map.

Some disadvantages of using circuit simulation include:

- Some of the components (e.g. generators) are not easily modeled as electrical circuits.
- Models developed in other platforms (e.g. differential/algebraic equations or Matlab™) can be re-used here with some difficulty -- this is currently being addressed.

The developed graphical user interface not only allows the representation of data over a geographic display but can also display time series values of selected nodes. Further development will include the ability to run simulations with the user specifying which nodes and edges are to be removed (e.g. generator failure, transmission line outages, and substation blackouts).

3. USING HIGH-PERFORMANCE COMPUTING TO CALCULATE RELIABILITY IMPACTS FROM CYBER ATTACK ON ELECTRIC POWER SYSTEMS

3.1. Introduction

The goal of the work described here is to quantify the impact of cyber attack on the US electric power infrastructure. The approach is termed Reliability Impacts from Cyber Attack (RICA). RICA estimates the effect of cyber attack on reliability by computing reliability index values for simulated power system operation in two environments, one with cyber attack and one without, the difference in reliability being the impact of cyber attack. The work has aimed at producing results for the Western Electricity Coordinating Council (WECC) region. Significant effort was expended to accommodate the size of the WECC simulation by utilizing the Red Mesa high-powered computing (HPC) platform at Sandia National Laboratories. The work reported here describes the results of running the RICA simulation for the standard IEEE reliability test systems (RTS-96) [31]. The Roadmap to Secure Control Systems in the Energy Sector [8] identifies a critical need to understand the possible impact of attacks on electric power systems in order to better prioritize mitigation investment to control risk. The approach taken in the Reliability Impacts from Cyber Attack (RICA) project provides a means to gauge the impact of cyber attacks independent of variations in weather, load, ordinary local outages, and time of day and year. The roadmap also points out that “asset owners are hard-pressed to justify control system security” because they are unable to “quantify and demonstrate the potential impacts of cyber attacks on energy sector control systems.” RICA analysis enables explicit quantification of cyber attack effects in terms of unserved load, enabling literal cost-benefit assessment of proposed cyber security measures.

The goal of the RICA work is to usefully quantify the effects of security measures on grid reliability. This entails the subsidiary goal of harnessing sufficient computational power to enable production of RICA results for large-scale power grids. The work reported herein describes significant progress towards achieving these goals. Current approaches to prioritizing security are weak because they rely on expert opinion and checklists to rank mitigation strategies and certify system security. These methods are flawed because they rely on untestable assertions, do not permit comparison among alternatives, and have been shown to overlook fundamental problems. The RICA approach provides testable, comparable, and comprehensive results concerning the efficacy of cyber security measures. RICA results are intended to inform strategic decisions and have value to industry planners, investors, and other researchers deliberating among investments to improve cyber security. To isolate the effects of cyber security, outage metrics are reported on an annual basis to “average out” the effects of the time of year and local outages.

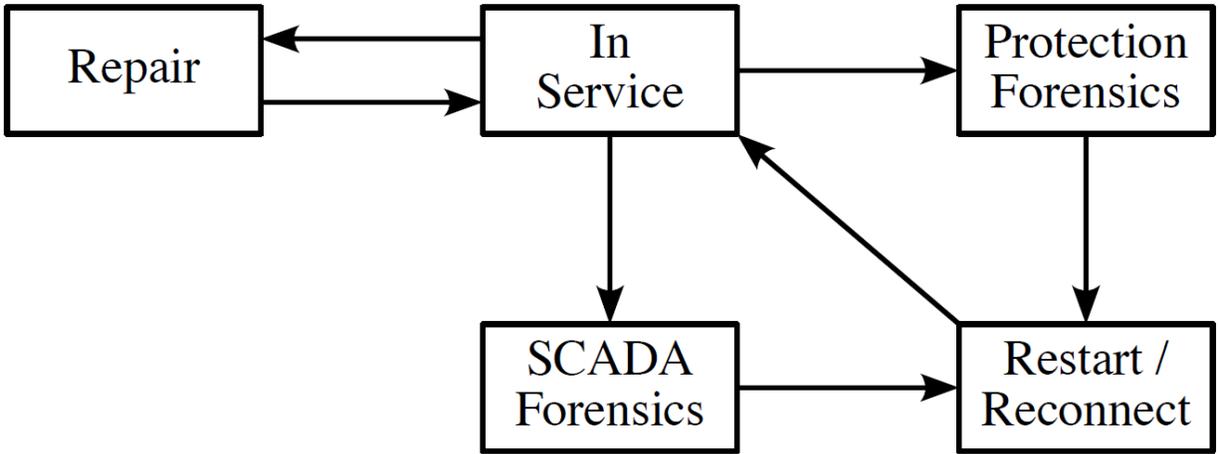


Figure 9. Possible states for grid elements using the RICA approach

Observed changes in RICA-produced reliability results are due to cyber security measures (and not, for example, the fact that it is currently winter).

Work described herein represents a significant capacity increase over our previous efforts, described in [9]: results therein were for a notional five-bus system, whereas the results reported here are for the IEEE RTS-96 72-bus system, and the utilized computing platform has the scaling capacity to produce results for the 30,000-bus Western Electric Coordinating Council (WECC) transmission grid. Several researchers ([10], [11], [12], [13], [14]) consider the ability to quantify system availability crucial to determining what kind of cyber hardening is needed to protect critical infrastructures. These researchers have examined reliability analysis as an approach to security analysis because it does exactly that: the prominent reliability measures are essentially system availability and outage characterization. The literature on this subject provides examples ([15], [16]) that specifically apply Monte-Carlo (MC) analysis to information security issues in a manner very similar to the current RICA approach. Although the analogy between reliability and security is not entirely apt, we note the repeated statement in the work cited above that the probabilistic framework of reliability analysis is particularly suited for critical infrastructure cyber security.

3.2. Reliability Analysis for Cyber Attack

3.2.1. RICA Approach

Given a grid—an electric power system comprising transmission lines, breakers and other switches, generators, and loads—reliability is defined as the efficacy of the grid in delivering power to the loads. In the RICA approach, reliability is represented by reliability indices, such as Loss of Energy Expectation (LOEE) and Frequency of Interruption (FOI), calculated using data from a simulation of the grid. RICA determines the reliability of a grid as follows: The grid of interest is modeled by integrating model of individual power system elements and the load-satisfying behavior of this grid model is observed over several thousand simulated years using probabilistically determined outages for each individual power system component and empirical demand patterns. Load flow, unserved load, and any additional outages (e.g., line tripping due to overload) are computed at each time step. The amount of unserved load (and other data needed

to compute the relevant reliability indices) is accumulated [17], [18] and, generally, averaged over time and reported per unit time. This approach is termed MC reliability analysis [19].

The high-level diagram for constituent grid components is shown in Figure 9. Each grid element (generator, line, or shunt) starts the simulation in service, and failure may cause it to go out of service. Failure may be random or purposely caused, e.g., by an attack on the system. After failure, an element returns to service after an interval during which repair, forensics, and so on ostensibly occur. This presumably results in the protection scheme being restored or a workaround developed, after which the grid element is restarted and reconnected to the grid (which may involve an additional delay, such as for a warm restart of a generator). If a Supervisory Control and Data Acquisition (SCADA) attack causes de-energization of the component, then it enters the SCADA forensics interval, before being reconnected. Several metrics based on unserved load and outage characteristics are computed. Overall, system reliability is measured using indices, including FOI (in occurrences per year), Loss of Load Expectancy (LOLE) (in hours per year), LOEE (in MW hr per year), Duration of Interruption (DOI) (in hours per interruption), Energy Not Served per Interruption (ENSI) (in MW hr per occurrence), Load Curtailed per Interruption (LCI) (in MW per occurrence), and Energy Index of Reliability (EIR) (the ratio of energy served to yearly demand).

During MC simulation, unserved load is calculated once per simulated unit of time (every hour, in RICA; generally, times discussed herein refer to simulation time, not actual “wall clock” time). For each such calculation, each system element is independently determined to be in or out of service as follows: At the beginning of the simulation¹, the time interval until the next failure of each piece of equipment is determined by scaling a sample from an exponential distribution by the relevant Mean Time to Failure (MTTF). This interval is added to the current simulation time to give the item’s “failure time”, i.e., the time at which it will fail. All active items whose failure time is less than or equal to the current time (i.e., whose failure time has passed) are marked “Failed” and do not contribute to generation or transmission. Failed equipment returns to service after an idle period (during which it is ostensibly being repaired) determined by a similar sample scaled by the equipment’s Mean Time to Recover (MTTR). Each type of equipment has its own MTTF and MTTR. This is conventional for reliability studies; note that there are no de-rated states used in this approach currently.) Once it has been determined which elements are functioning, load flow is calculated based on the remaining transmission lines and their capacities, the structure of the remaining network (i.e., what’s still connected to what), the capacities and setpoints of the remaining generators, and the loads to be served at that moment. Load magnitudes are based on empirical demand statistics. If the load flow indicates possible line overloads or other violations of limits, then an optimization routine is applied to shift generation in order to maximize served load. (This last step mimics the realistic restorative actions that would be taken by system operators when presented by a contingency situation; failure to model this would result in unreasonably low reliability calculations.) [15] calls the generation-transmission system, also referred to as the bulk power system, the “Hierarchical Level II (HL-II)”; HL-III includes distribution.

Distribution is not included in the RICA model because aggregation of load at the substation level provides sufficient resolution to develop an informative load picture at the regional and national levels. The approach as described so far provides a measure of system reliability that

accounts for random equipment failure and recovery on an item-by-item basis. We refer to such outages as natural to distinguish them from failures caused by cyber attack. To understand the impact of cyber attack, attacks and their effects must be modeled and added to the process of natural outages described above. Cyber attacks happen in addition to natural outages; both degrade grid performance, but they are represented independently in the model because they are expressed and mitigated in different ways. For each component the time until the next successful cyber attack is currently modeled using an exponentially distributed random variable that's independent of the natural outage variable and a selected mean time to attack (MTTA). The MTTR for a cyber-attacked piece of equipment is based on the time required for cyber forensics, control system restoration, and device restoration. The separate contributions of generation and transmission to whole-system reliability can be examined using RICA because both are explicitly represented. This means, for example, that RICA can be used to assess whether a cyber security budget would be better spent protecting generators or protecting transmission lines. Alternatively, the algorithm can show that sources would be best invested shortening the recovery time from cyber attacks, as opposed to potential modest improvements in prevention.

3.2.2. Attacks against Protection

To study protection attacks, we focus on attacks against three common types of protective relaying: Generator protection, Bus protection, and Line protection. Although generation protection attacks are generally unlikely, they are included since loss of generation is frequently modeled in reliability studies (although not, at least in the prior literature, from cyber attack). Based on a certain linked set of parameters from the Cyber-to-Physical (C2P) bridge, a successful attack is modeled as causing the generator to go offline as its breakers trip. The substation marks the boundary between “distribution” and “transmission”; power transport between generation and the substation is considered transmission, everything below the substation is distribution. Mean time to attack is similar to mean time to failure, except that MTTA is the average interval between successful cyber attacks and MTTF is the average interval between random outages. Most generators reside in generating stations, where protection significantly limits the possible attack paths. The interval between successful attacks is modeled using an exponentially-distributed random variable and some selected MTTA. The MTTR for the generator depends on the time required for cyber forensics and control system restoration, and the restart time for the generator, which depends on its classification as hydro, thermal, or nuclear (see Table II). A successful bus protection attack will presumably activate the differential scheme and completely isolate all connected devices. Bus protection is also considered a difficult scenario for an attacker given limited connectivity of its constituent relays. However, the large loss in connectivity will likely cause significant losses in grid performance. Line protection attacks assume advanced relays that allow cyber control of the protective breakers at both ends of a transmission line. A successful cyber attack opens the breakers and removes the line from service. The line is successfully attacked at random intervals denoted by an MTTA and is out of service for an MTTR, as above. This is again similar to the process for natural outages, but with the recovery interval dependent on cyber remediation and line restoration.

3.2.3. Attacks against SCADA

Adversarial penetration and use of a grid's SCADA system to send trip signals to system breakers was also modeled. Open breakers can isolate generators, open lines, and disconnect shunts and loads. The interval between successful attacks is determined by sampling from an exponentially-distributed random variable scaled by a specified MTTA, in a manner similar that used in the protection modeling. The MTTR, or downtime, is in this case the duration of the interval during which the breaker is open. We assume the duration would accommodate the cyber forensics activity any component restoration time. Another parameter for analysis is the Average Percentage of Tripped Breakers (APT_B) stemming from a successful attack, which is modeled using a Bernoulli random variable for each breaker with a selected mean value p_{SCADA}. The APT_B parameter enables quantification of the concept that not all breakers will be sent a trip signal by the attack, not all breakers sent a trip signal will receive it, and not all breakers receiving a trip signal will trip.

3.3. Rica Computation and Workflow

The elements of RICA computation, shown in Figure 10, are: Load Flow Computation, Optimizer, Pre-processor, Scheduler, Database, Data Harvester, and Visualization.

3.3.1. Load Flow Computation

The general relevance of the RICA analytical tool stems from the notion that the load-flow simulation determines fairly accurately the degree to which the modeled power grid satisfies a given set of loads. The load-flow calculation is an iterative algorithm based on Kirchoff's laws, and takes as input a grid, a maximum generation level at each power-producing node, and a requested load at each power-using node. The algorithm determines the amount of power carried by each line in the modeled grid, up to its capacity, as generated power is transported to the loads. Note that any network, whether real or simulated, may be unable to satisfy a given set of loads because of failures and/or lack of transport capacity (it may not be possible to supply a given set of loads no matter how much power is being generated if lines have failed, because every line has a maximum capacity). In general, a real power grid will have evolved specifically to supply the loads encountered in everyday operation, bringing effectively to zero the probability an expected load will not be satisfied under normal conditions. To accommodate new loads, either generation and transport capacity is added or the network structure is altered. And, of course, in a real power system, loads can simply disappear for any number of reasons, e.g., by being satisfied from a non-grid power source, or when the building housing the load burns down, or when the company using the energy goes out of business.

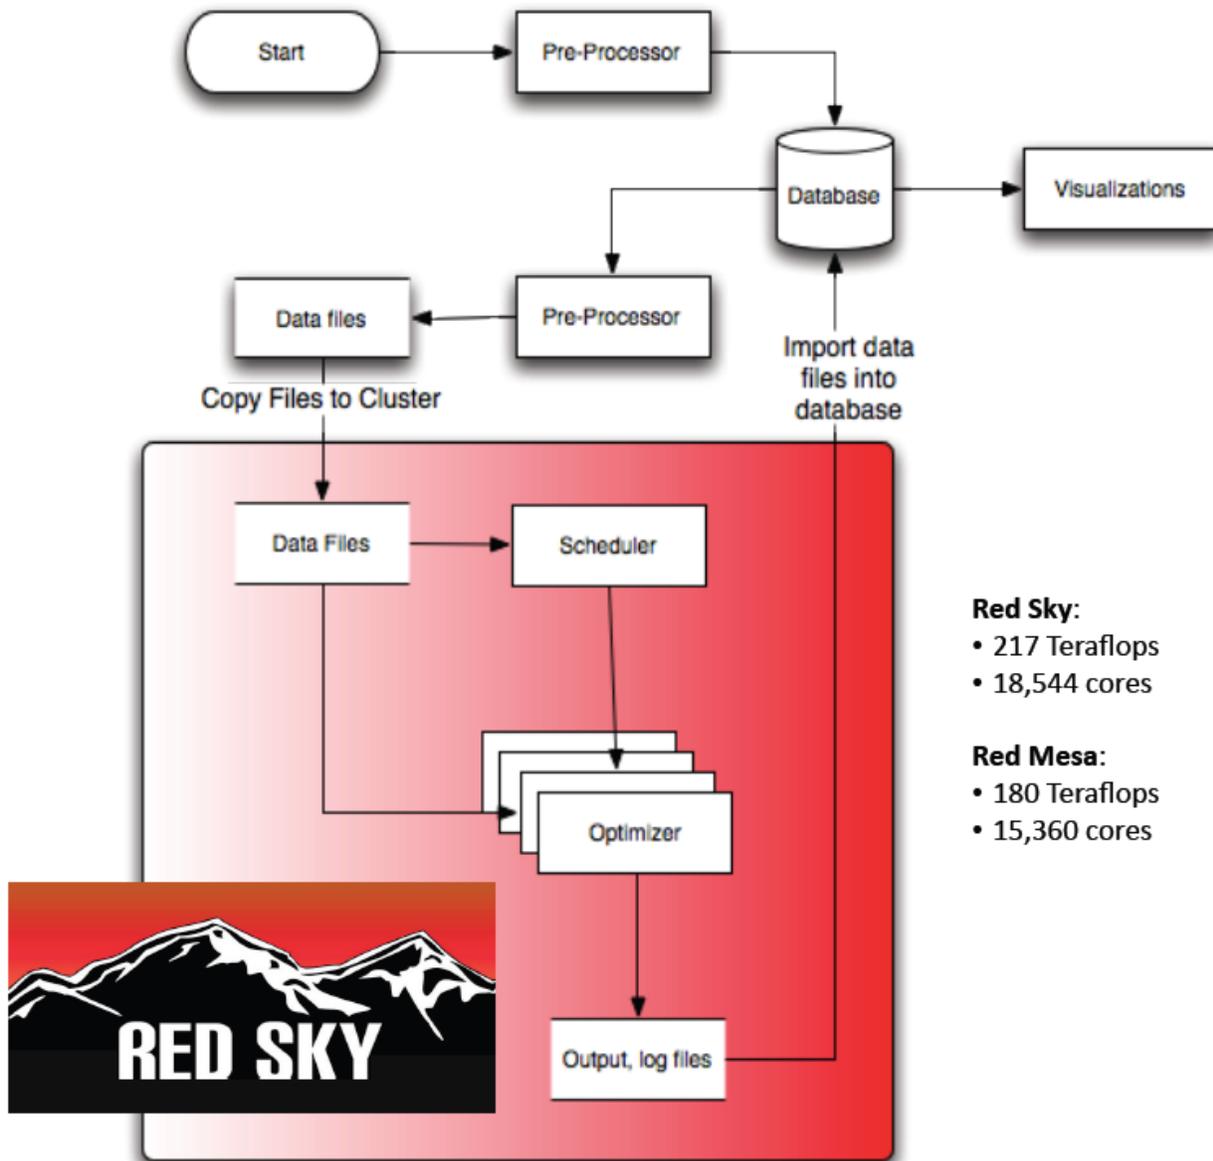


Figure 20. Executing RICA on the Red Mesa HPC

3.3.2. Optimizer

In a simulated power system, there is no foreknowledge that the network is capable of satisfying a given set of loads. In RICA, an optimization process acts to maximize the percentage of requested load that is satisfied. RICA optimization is based on the idea that generation can be adjusted within some capacity range to supply additional power to loads that are undersupplied because of power transport limitations. In RICA, if significant load is unsatisfied after load flow calculation, the optimizer adjusts generation to maximize the amount of satisfied load. This is understood to be a fairly realistic simulation of human power system operators, who attempt to do the same thing in real life when elements of a power system fail.

3.3.3. *Pre-Processor*

The pre-processor parses power system data provided in multiple formats, conditions the power system initial state to avoid a large difference in generation and consumption, calculates the outage and repair times for each power system component based on its MTTF and MTTR, stores the resulting power system and outages in the RICA database, and dumps the outage data to files for use by the optimizer. The pre-processor places object models in the RICA database for each power system component (bus, branch, generator, load, and shunt) along with object models that represent simulations and outages. As power system data input files are parsed, models are created and placed in the database for each element of the power system. These are then used to construct scenarios, which consist of outages for the different power system components. Once the power system models are created from the input files, the pre-processor creates a new simulation object, which includes a length of time for which the simulation is to conduct Monte Carlo Analysis. It then loops over all the generators and branches in the power system and, given the length of time for the simulation, calculates outage and repair times for each of the generators and branches using a uniform random variable based on MTTF and MTTR values. For each of the outage instances calculated, a record is saved to the database using the outage object model. Once the simulation is configured with outage and repair times, the outages are analyzed to determine ranges of hours where at least one component is out of service, an overall range file is created with the given ranges, and a config file containing which components are coming on or off line at each hour within a range is generated for all the ranges detected. The pre-processor makes it possible to regenerate the outage files described above (i.e., a file of ranges with a config file for each range) using data in the database and different MTTF and MTTR values.

3.3.4. *Scheduler*

The scheduler assigns processing tasks to Red Mesa processors. In RICA, each processing task is a time interval during which a particular set of outages, loads, and generation holds. The percent of load being satisfied, having been maximized by the optimizer, does not vary over the time interval since the network configuration (outages, nodes, and generation), and therefore the maximized load flow, do not change. The larger the power network, the less likely it is that there will be no changes in load, generation, or outage conditions, which implies a greater number of scheduled calculations. Thus, the larger the network the more separate tasks need to be scheduled. The scheduler acts to distribute computational burden efficiently to the available processors.

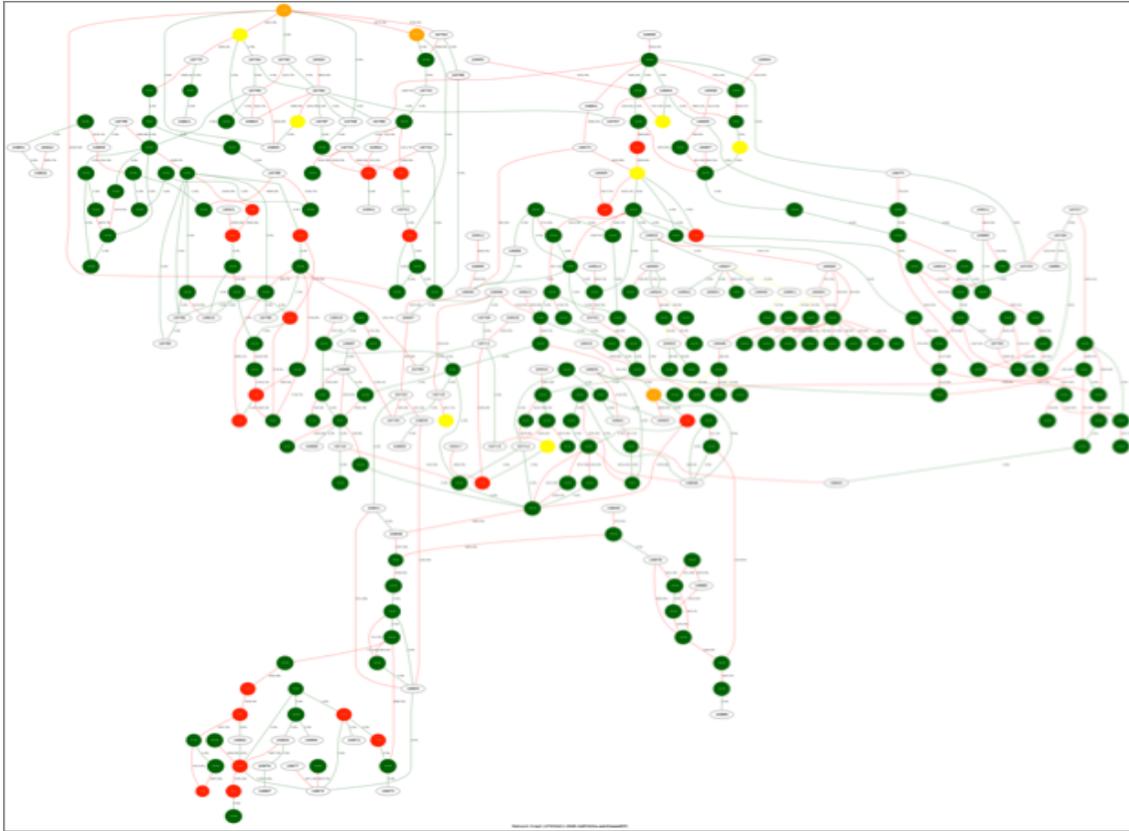


Figure 31. Network diagram of the 300-bus system

3.3.5. Database

The database stores all the input network/outage data and all output in the form of network load flow. The relevant output from a reliability perspective is the amount of requested load service that is actually provided by the network. RICA runs on the 180-teraflop Red Mesa6 HPC platform at Sandia National Laboratories (see Figure 11). RICA has successfully run a 300-node (see Figure 11), 1000-year hourly- hour simulation, which required a minimum of 8,760,000 one-hour load flow calculations, not counting the multiple load-flow computations needed for each optimization, and processed the output to obtain visualization images (see Figures 12, 13, and 14). RICA on Red Mesa also has calculated optimal load flow for the 30,000-bus model of the WECC power grid. Database population is automatic using data produced during RICA operation. Construction of database entries can be turned on and off to avoid data records during acceptance testing. Current database entries have been produced from the IEEE test suites. The largest RICA outputs demand more disk space (non-volatile memory) than is currently available on the database server. This could be alleviated by providing additional storage or output postprocessing. A web interface enables simulation results to be viewed and provides summary statistics and visualization to allow comparison of database entries. The web interface allows ready visualization of data in the database in near-real time following code execution. Figures 11, 12, 13, and 14 are taken directly from the website. The visual output is a starting point for further visualization work. Note that Red mesa is a subset of the Red Sky High-Powered Computation (HPC) platform, as indicated in Figure 10.

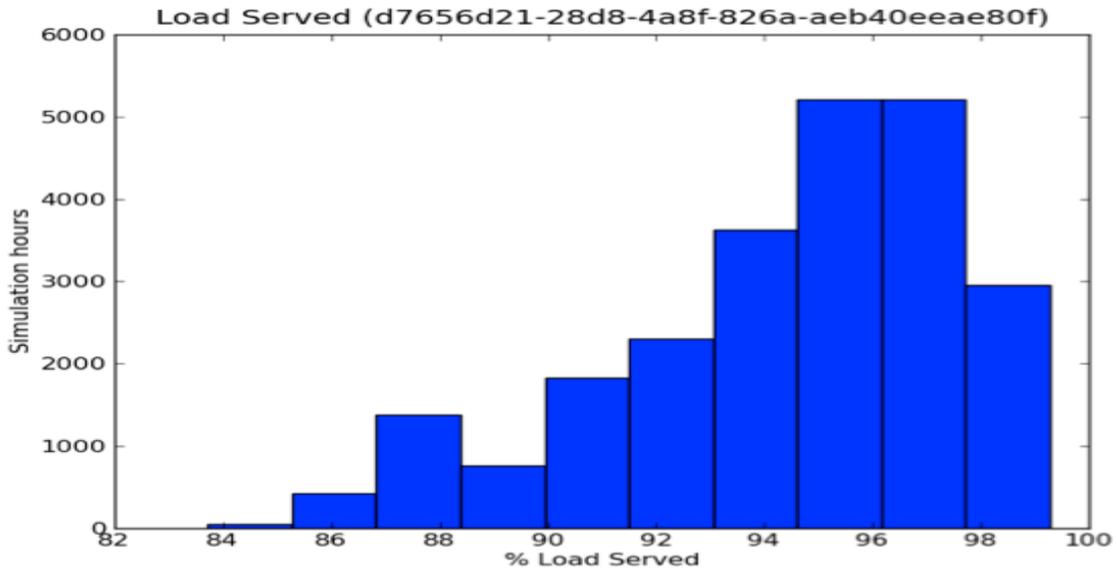


Figure 42. Histogram showing number of cases for served-load percentage intervals for the 300-bus system

Branch loadings (d7656d21-28d8-4a8f-826a-aeb40eeae80f)

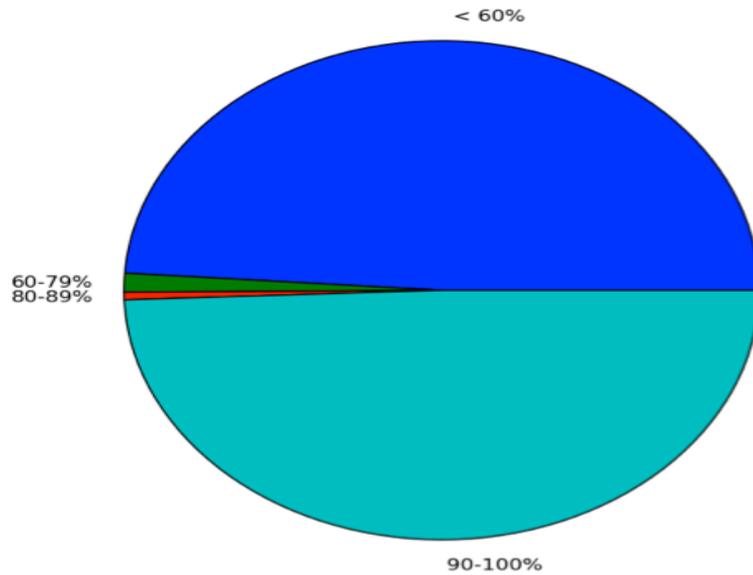


Figure 53. Pie chart showing fraction of cases of served-load percentage for the 300-bus system

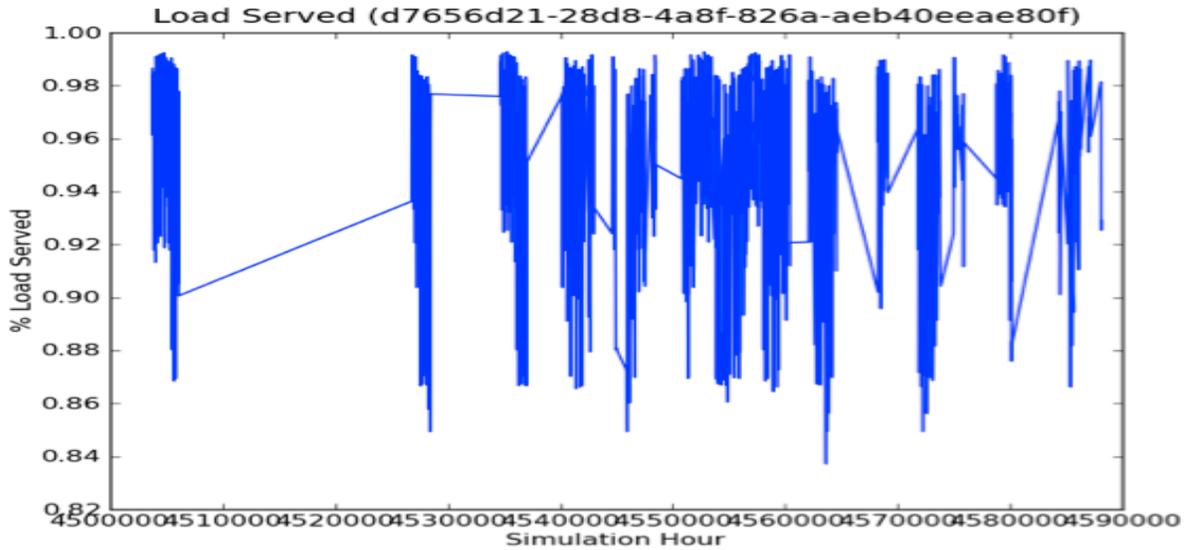


Figure 64. Line showing served-load percentage from simulation from beginning to end for the 300-bus system

Table 7. Device Forensics Intervals For Cyber Attack

Attack Vector	Symbol	Interval
Generation protection	T_{GEN}	12 hours
Transmission protection	T_{LINE}	4 hours
Bus protection	T_{BUS}	8 hours
SCADA	T_{SCADA}	8 hours

Table 2. Restart Delays For Generators After Cyber Attack

Unit Type	Size	Delay
Hydro	Any	+0 hours to restart
Thermal	$\leq 250\text{MW}$	+2 hours to restart
Thermal	$> 250\text{MW}$	+8 hours to restart
Nuclear	$> 250\text{MW}$	+36 hours to restart

Table 3. Restart/Reconnection Delays For Grid Elements After Cyber Attack

Component	Delay
Lines	+0.5 hours to reconnect
Shunts	+1 hours to reconnect
Loads	+0 hours to reconnect

Table 4. Change In WECC Reliability Indices, With Cyber Attack (Forensics Intervals As Shown In Table I)

Parameter	Base Case	Sim 1	Sim 2	Sim 3	Sim 4	Sim 5	Sim 6	Sim 7	Sim 8	Units
λ_{GEN}	0	0.1	0.2	0.1	0.1	0	0	0	0.1	occ/yr
λ_{LINE}	0	0.2	0.2	0.5	0.2	0	0	0	0.2	occ/yr
λ_{BUS}	0	0.1	0.1	0.1	0.2	0	0	0	0.1	occ/yr
λ_{SCADA}	0	0	0	0	0	0.1	0.2	0.1	0.1	occ/yr
ρ_{SCADA}	0	0	0	0	0	0.2	0.2	0.4	0.2	none
FOI										int/yr
LOLE										hr/yr
LOEE										MW·hr/yr
DOI										hr/int
ENSI										MW·hr/int
LCI										MW/int
EIR										none
Mflops										none
CPU time										seconds

Table 5. Change In WECC Reliability Indices, With Cyber Attack (Forensics Intervals Halved Compared To Table IV)

Parameter	Base Case	Sim 1	Sim 2	Sim 3	Sim 4	Sim 5	Sim 6	Sim 7	Sim 8	Units
λ_{GEN}	0	0.1	0.2	0.1	0.1	0	0	0	0.1	occ/yr
λ_{LINE}	0	0.2	0.2	0.5	0.2	0	0	0	0.2	occ/yr
λ_{BUS}	0	0.1	0.1	0.1	0.2	0	0	0	0.1	occ/yr
λ_{SCADA}	0	0	0	0	0	0.1	0.2	0.1	0.1	occ/yr
ρ_{SCADA}	0	0	0	0	0	0.2	0.2	0.4	0.2	none
FOI										int/yr
LOLE										hr/yr
LOEE										MW·hr/yr
DOI										hr/int
ENSI										MW·hr/int
LCI										MW/int
EIR										none
Mflops										none
CPU time										seconds

3.4. Rica Simulations and Results

The three-area RTS-96 system was simulated according to its given parameters (loading, topology, non-cyber outages rates) and augmented by cyber attacks against generator, line, and bus protection as well as SCADA. The additional cyber parameters are shown in Tables 1-5.

4. PV OUTPUT VARIABILITY MODELING USING SATELLITE IMAGERY

4.1. Introduction

One barrier to adding large amounts of photovoltaic (PV) solar energy to the electrical grid is the current uncertainty about the level of variability in power generation that large PV systems may introduce. Variability at unacceptable levels might threaten the stability of the electrical grid, decrease reliability, and possibly damage expensive equipment (e.g., transformers).

Utilities are required to match load with generation or risk being fined or causing damage to their grid. Because load is variable and not entirely predictable, utilities have methods to maintain balance, including running fast ramping generator units under a control scheme called automatic generation control (AGC), contracting for interchanges with neighboring balancing authorities, among others. It is unknown how these control methods will be affected by the presence of large PV plants on the grid.

Methods for controlling the variability introduced by large PV plants could be studied using models to simulate output of PV plants. However, current models are constrained by the lack of irradiance data with appropriate temporal and spatial resolution. Data from single irradiance sensors show that irradiance at a point can change from full sunshine to about 20% of full daylight within a few seconds. However, it is unclear how rapidly the output from a large PV plant will change. Studies have shown that the output of a PV plant changes more slowly than does the output from a single PV module (e.g., Mills et al. 2009). It is commonly accepted, however, that irradiance changes on the order of one minute may be significant when simulating a PV plant. Irradiance measurements at one minute or shorter intervals are generally available only at a few locations within a utility service area (if at all). Satellite data is currently used to estimate irradiance in places with no sensors, but the spatial and temporal resolution of this data is generally quite coarse. For example, the SUNY database supplies irradiance estimates for a 10 km grid at 1- hour intervals (Perez et al., 2002) and 3Tier advertises hourly irradiance data at ~3 km resolution (www.3tier.com). These datasets currently provide no information about irradiance within the hour or at a significantly finer spatial scale. There is no general consensus on how to estimate irradiance with high temporal and spatial resolution at locations where there are no sensors.

This section describes work building towards a method to estimate high frequency (i.e., one minute) irradiance values from satellite imagery for places at which ground measurements are not available. In concept, this method would estimate irradiance from satellite imagery that is available approximately every 15 minutes. Cloud patterns in each image would be used to determine the variability of irradiance within each 15 minute period between images. Here, we document analysis of satellite imagery, including jitter correction and cloud recognition algorithms, in preparation for training a neural network to simulate irradiance between images. Ultimately this work might provide the ability to predict one-minute PV plant output at these locations.

We focused the study on a small region in southern Nevada, where ground measurements of irradiance are available at a one-minute time resolution. Satellite data from the Geostationary Operational Environmental Satellite 11 (GOES-11) was downloaded from the National Oceanic and Atmospheric Administration (NOAA) web site for the study area for most of calendar year 2008. The GOES 11 was launched in June 2006 and provides imagery of the western North America and the Pacific Ocean.

Because the satellite imagery pixel size is approximately one square kilometer, satellite imagery does not offer sufficient resolution to estimate the exact time that a specific location within a square kilometer is occluded by cloud shadows. Consequently it is considered impossible to predict to the nearest minute when a cloud shadow will reach a certain location. However, we believe it may be possible to predict the number and severity of ramps within the hour based on the type of clouds and their spatial patterns as seen in the satellite imagery.

4.2. Modeling Approach

Figure 85 illustrates the modeling approach. The GOES 11 imagery is prepared for analysis by identifying cloudy and clear areas by means of a neural network, as discussed in the remainder of this section. The model would next use a second neural network to generate irradiance at a location for every minute between successive pairs of images. The second neural network would be calibrated with sample irradiance measurements from other ground locations. Although we were able to make preliminary tests of this second approach, these efforts will be further developed before irradiance results are reported in a later report.

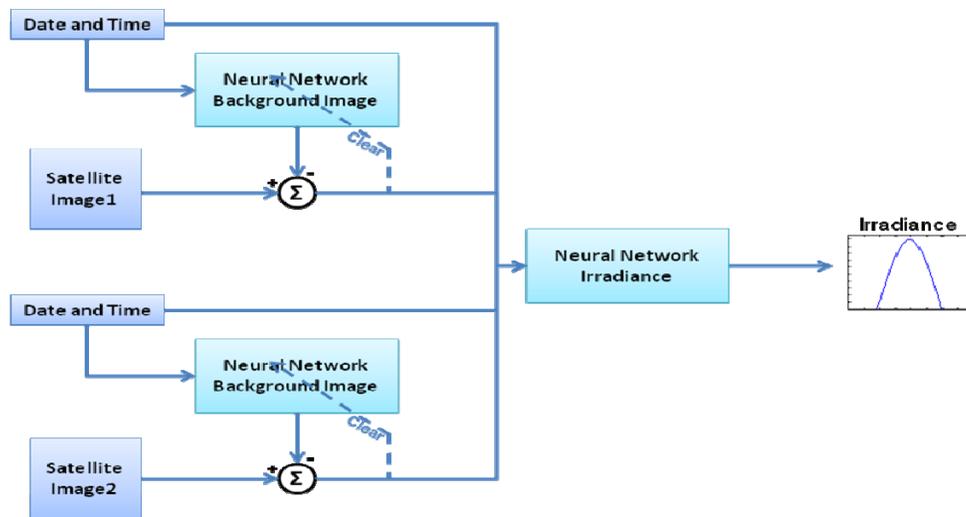


Figure 85. Modeling approach for generating irradiance from two images 15 minutes apart.

The basis for this modeling approach is illustrated in 16, which shows three successive images of the same region around Las Vegas, NV, , and the corresponding irradiance profile from two ground measurement sites. In the first image, clouds cover both sites, and the measured irradiance at both sites is low. In the second image, Fort Apache (blue) is no longer covered and

its measured irradiance has increased. In the final image, both Fort Apache and UNLV (red) are not covered by clouds in the satellite image.

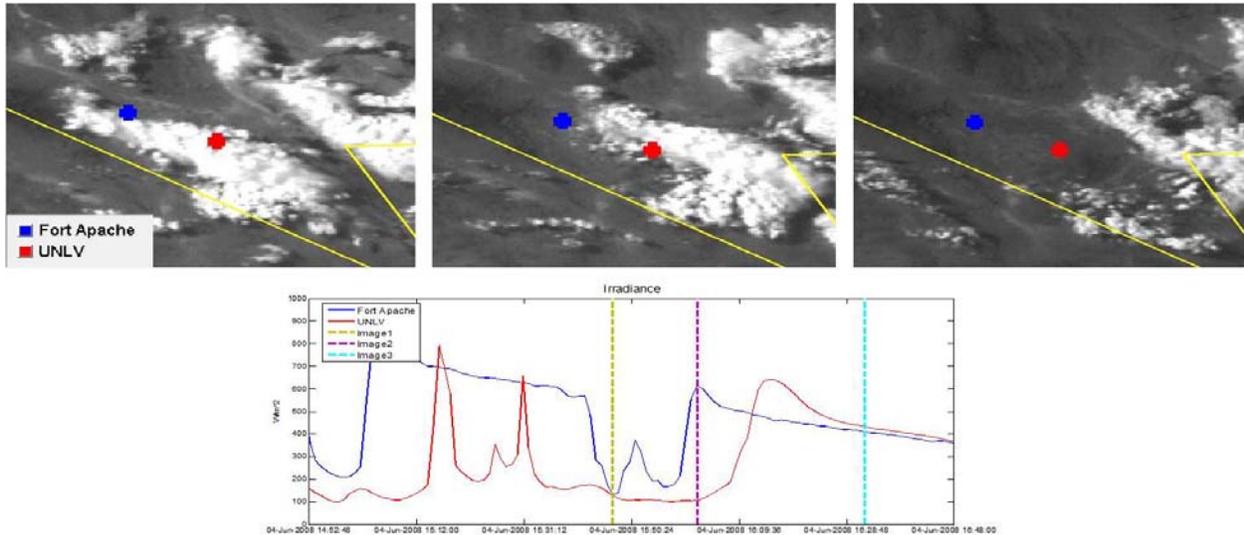


Figure 16. GOES 11 images of Las Vegas region for 6/04/2008 around 4PM (PST) with corresponding measured irradiance at two ground locations.

4.3. Sources of Ground Measurements of Irradiance

Within the Las Vegas, NV, region there are eight ground stations where measurements of global horizontal irradiance are available at one minute intervals (Figure 19). The stations are located within the Las Vegas metropolitan area. Six of the stations are collocated with PV generation plants operated by the Las Vegas Valley Water District (LVVWD) and two stations (UNLV and Clark Station) are operated by the National Renewable Energy Laboratory (NREL). Table 6 provides information on these monitoring sites.

One minute data available at these sites include measurements of global horizontal (GHI), direct normal (DNI), and diffuse horizontal irradiance (DHI). In addition, standard meteorological data such as temperature, relative humidity, wind speed, etc., are also available.

Table 6. Irradiance Monitoring Stations used in this Study

Site	Latitude	Longitude	Operator	Start Date
LVSP	36.17° N	115.19° W	LVVWD	7/26/2007
Luce	36.2° N	115.26° W	LVVWD	5/2/2007
Spring Mountain	36.12° N	115.29° W	LVVWD	11/30/2006
Grand Canyon	36.22° N	115.31° W	LVVWD	9/30/2006
Fort Apache	36.22° N	115.3° W	LVVWD	8/23/2006
Ronzone	36.19° N	115.23° W	LVVWD	4/27/2006
UNLV	36.06° N	115.08° W	NREL	3/18/2006
Clark Station	36.09° N	115.05° W	NREL	3/27/2006

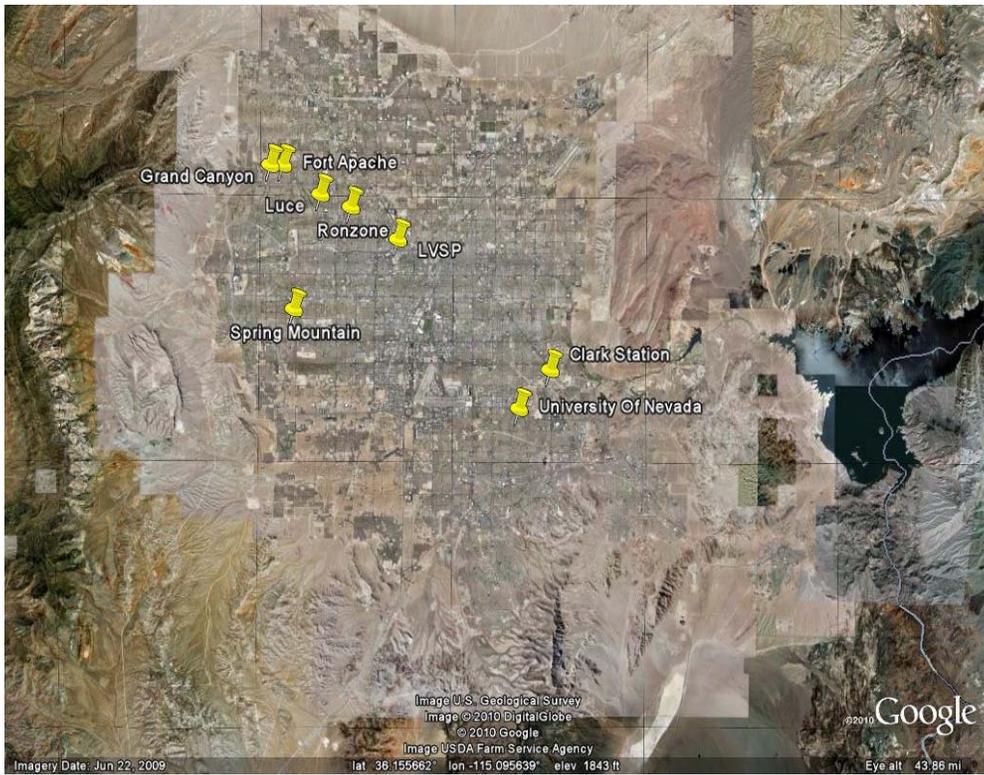


Figure 19. Close-up of the Las Vegas Metropolitan Area showing the location of 1-minute irradiance sensors.

Figure 18 shows GHI data from the Fort Apache station for a two day period (April 15-16, 2007).

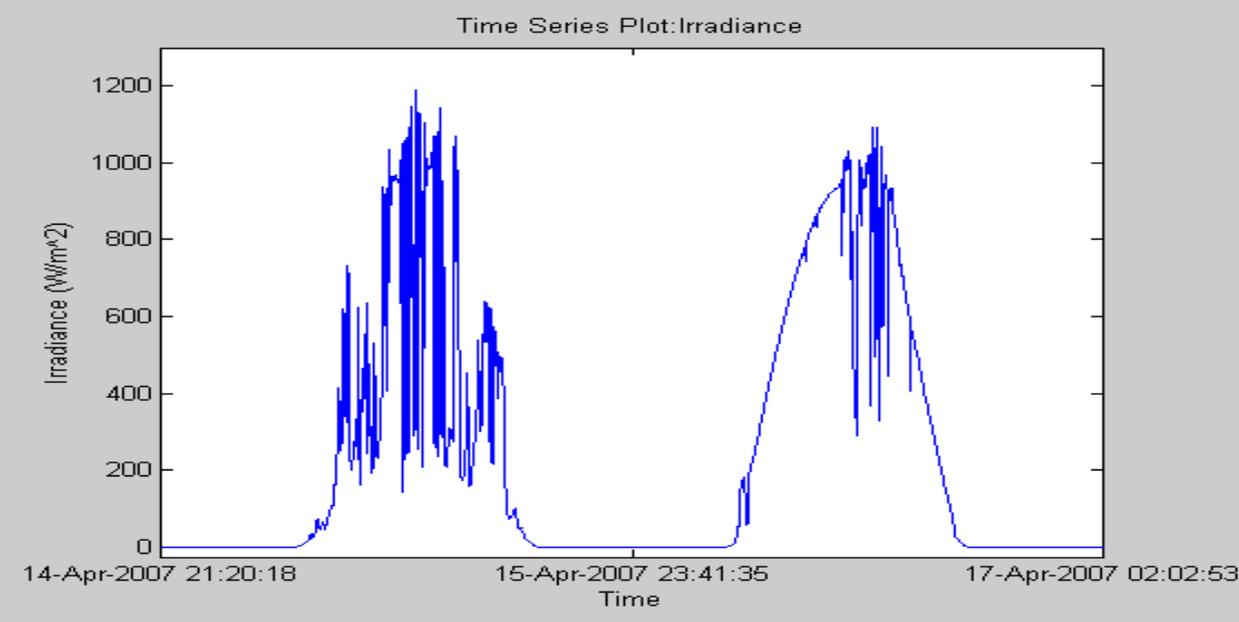


Figure 110. Global horizontal irradiance measured at 1-minute intervals at Fort Apache station over a two day period.

4.4. Sources of Satellite Imagery

Satellite data from the Geostationary Operational Environmental Satellite 11 (GOES-11) was downloaded from the National Oceanic and Atmospheric Administration (NOAA) web site for the study area for the calendar year 2008. The GOES 11 was launched in June 2006 and provides imagery of the Pacific United States (PACUS), northern hemisphere, southern hemisphere, and the full visible earth. The satellite has a geostationary orbit at 35,790 km above the equator at 135 deg W longitude. The satellite has a 5-channel imager, each connected to a different radiation detector (Table 7). Imagery is produced by scanning different parts of the Earth at different times. The scan schedule is shown in Table 8. The various image frames take different amounts of time to perform a scan (Duration).

The Study Area in Southern Nevada (Figure 111) is included in the Full Earth, Northern Hemisphere, and PACUS frames and therefore images of the Study Area are available approximately every 15 minutes with a gap every three hours when the Full Earth frame is scanned. Figure 12 provides an example image of the Study Area showing the location of two ground irradiance monitoring stations. Note that the image is distorted due to the angle of the satellite, which causes the resolution in the N-S direction to be different than in the E-W direction.

Table 7. List of available data from GOES Satellites

Channel Number	Wavelength Range (μm)	Detector Type	Nominal Square IGFOV at nadir
1	0.55 to 0.75	Silicon	1 km
2 (GOES 8/9/10) 2 (GOES 11/12)	3.80 to 4.00 3.80 to 4.00	InSb	4km 4km
3 (GOES 8/9/10/11) 3 (GOES 12/13/14/15)	6.50 to 7.00 5.77 to 7.33	HgCdTe	8 km 4 km
4	10.20 to 11.20	HgCdTe	4 km
5 (GOES 8/9/10/11) 5 (GOES 12/13/14/15)	11.50 to 12.50 12.96 to 13.72	HgCdTe HgCdTe	4 km 4 km

Table 8. Image Scan Routine for GOES-11

Frame Name	Boundaries	Duration (mm:ss)	Scan Times (UTC)
Full Earth	Earth Edge	26:10	0000, 0300, 0600, etc
Northern Hemisphere	0-66N/90W-170E	9:00	xx00, xx30
Southern Hemisphere	0-45S/115W-170E	7:00	xx22, xx52
PACUS	12-60N/90-175W	5:00	xx15, xx45

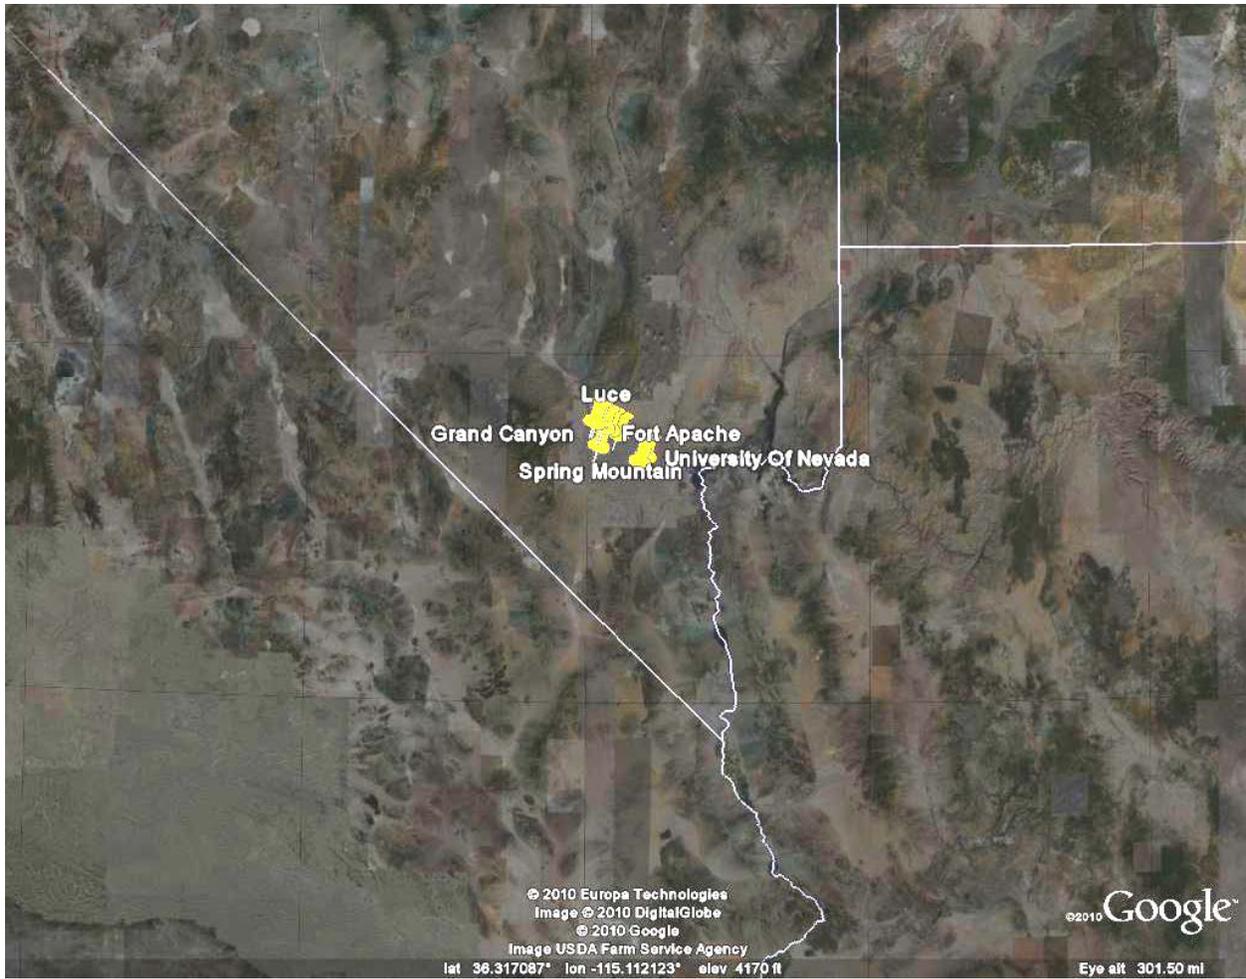


Figure 111. Map of Study Area. State boundaries in white, locations of irradiance stations in the Las Vegas area are shown.

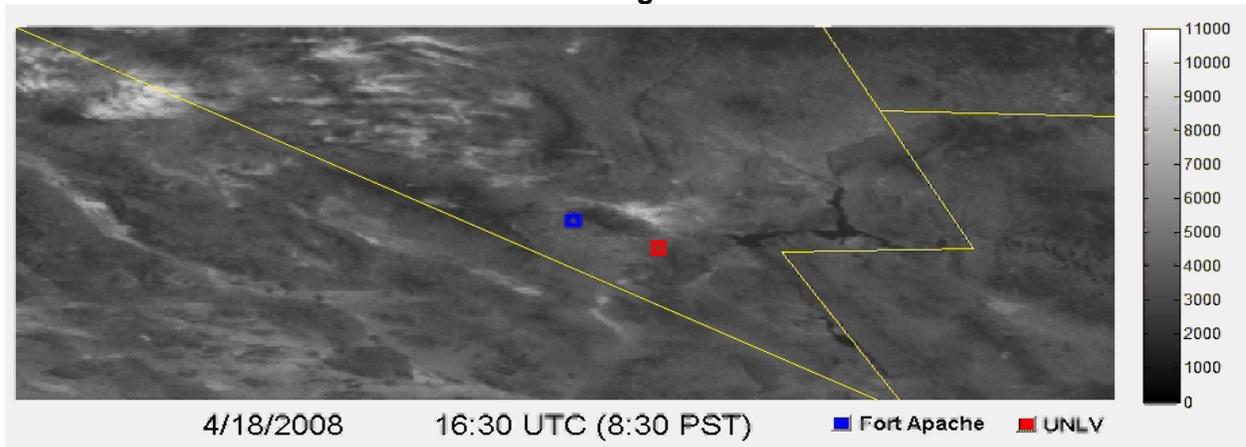


Figure 12. Example of GOES-11 Image with approximate state boundaries shown. The locations of two irradiance sensors are shown for reference.

4.5. Image Stabilization

NOAA geographically locates each pixel and calculates its closest latitude and longitude. This georeferencing algorithm is not perfect due to slight vibrations and variations in camera and mirror parameters. The result is that a visibly identifiable geographic location will not have a consistent latitude and longitude through time as calculated by NOAA. If uncorrected, the variability in a location's latitude and longitude would result in movement of the apparent position of objects in the geographically segmented image. This image jitter is typically less than ± 3 pixels (3km) in the x and/or y directions. Location variability includes both major movement of the image (>1 pixel) as well as sub-pixel image jitter. The images are processed through an image stabilization routine to correct for the jitter.

4.5.1. Image Stabilization Algorithm

Image stabilization is done by shifting an image to obtain the greatest correlation with a reference image constructed for each time. Construction of the reference images begins by identifying all clear sky images in the data set by applying the movement detection algorithm (Section 3.6.2) to exclude images with clouds. The brightest image as determined by average pixel intensity is selected from the identified clear sky images as the base reference image. Reference images at other times are constructed by scaling the image brightness and contrast by linear interpolation between clear sky images; scaling factors are determined by the neural network described in Section 3.6.3. Images are stabilized to the concurrent reference image by first expanding each image's resolution by four times by linearly interpolating in both x and y dimensions. Next, the expanded resolution image is shifted by ± 16 pixels in the x and/or y directions, and each shifted image is compared with the expanded resolution reference image. The shift exhibiting the best correspondence with the expanded resolution reference image is selected. Correspondence is determined by calculating correlation coefficients between two images. Finally, the shifted image is transformed to the original resolution through bilinear averaging and antialiasing.

The MATLAB (Image Processing Toolbox) function *normxcorr2* is used to compute the normalized cross-correlation of the two images using the following general procedure:

1. Calculate cross-correlation in the spatial or the frequency domain, depending on size of images.
2. Calculate local sums by pre-computing running sums.
3. Use local sums to normalize the cross-correlation to get correlation coefficients.

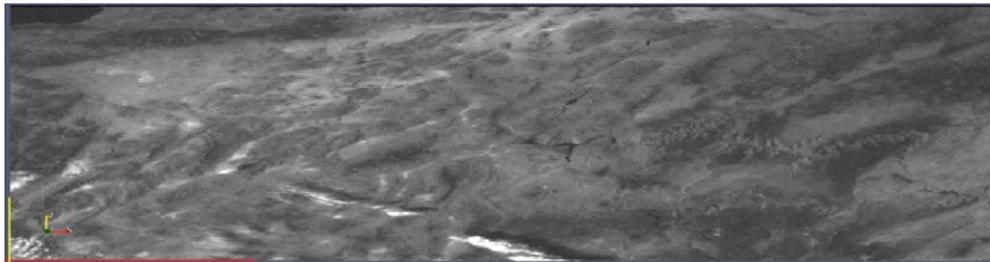
The implementation closely follows the following formula taken from the MATLAB documentation:

$$\gamma(u, v) = \frac{\sum_{x,y} [f(x, y) - \bar{f}_{u,v}] [t(x-u, y-v) - \bar{t}]}{\left\{ \sum_{x,y} [f(x, y) - \bar{f}_{u,v}]^2 \sum_{x,y} [t(x-u, y-v) - \bar{t}]^2 \right\}^{0.5}}$$

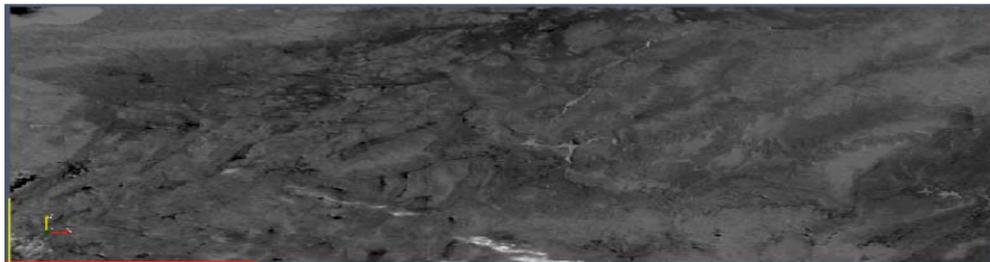
where

- f is the image.
- \bar{t} is the mean of the template
- $\bar{f}_{u,v}$ is the mean of $f(x, y)$ in the region under the template.

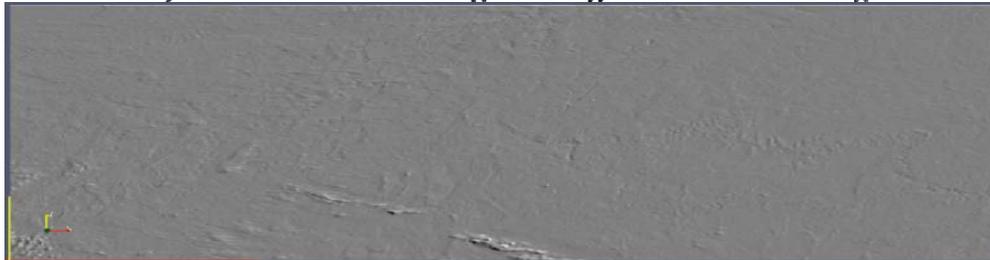
Figure 13 illustrates the results of the image correction process. The effect of jitter correction can be illustrated by the distribution of differences in pixel intensities between two images (Figure 14). Before correction, pixel intensity is biased compared to a reference image and exhibits wide variance; jitter correction results in unbiased pixel intensity with greatly reduced variance.



a) Aligned image



b) Difference between original image and reference image



c) Difference between aligned image and reference image

Figure 13. Illustration of image stabilization algorithm: (a) image after jitter correction; (b) difference between corrected and reference images; (c) difference between corrected and reference image.

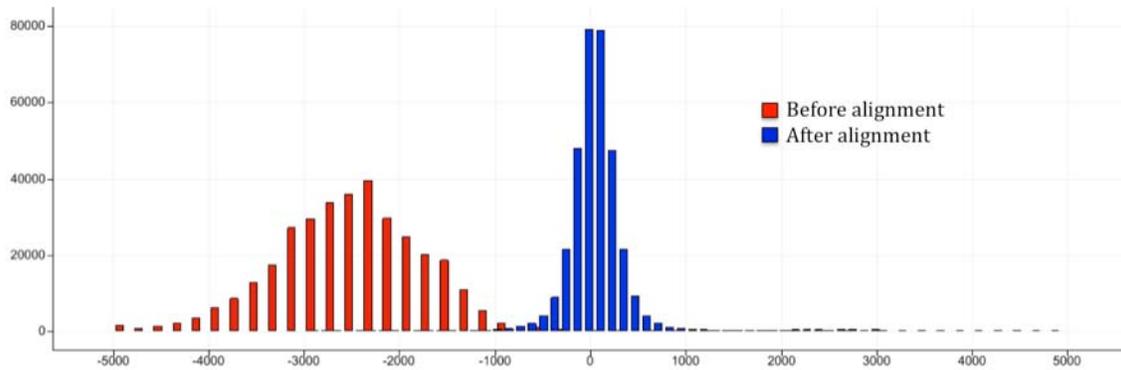


Figure 14. Histograms of difference in pixel values between reference image and uncorrected and corrected images.

4.5.2. High Performance Computing for Image Stabilization

Image stabilization is a great candidate for high-performance parallel computing because of the numerical requirements of the interpolation and cross-correlation, as well as the inherent parallelism in the algorithm. Since the stabilization algorithm operates on each image independently, it is trivial to write a parallel code to process a large set of images.

Our first implementation of a parallel stabilization code uses one process as a “task-manager” to keep track of the images that need to be processed. The remaining “worker” processes execute the stabilization algorithm for the set of images requested by the task-manager. This approach is more appealing than an evenly-distributed approach because some images take much less time to process. The algorithm works as follows:

Task Manager:

1. Read list of files to process.
2. Read reference image.
3. Broadcast reference image to all workers.
4. While image-list is not empty:
 - a) Wait for a request from a worker;
 - b) Pop an image file off the list;
 - c) Send image file name to worker.

Worker:

1. While there are still images to process:
 - a) Request an image file name from task-manager;
 - b) Load image from file system;
 - c) Perform stabilization algorithm on image.

Our implementation relies heavily on the Insight Segmentation and Registration Toolkit (ITK)¹, an open-source suite of libraries for image processing. ITK uses templated classes for registration, allowing rapid exploration of a variety of different distance metrics, interpolation methods, and optimization schemes. While we chose standard approaches for this prototyping effort, to adapt for cloud effects and other application-specific issues, we will likely need to implement custom metrics and optimization schemes to provide the data quality required for accurate modeling.

4.6. Cloud Detection

Identification of visible clouds in the images is not a straightforward process. Challenges include (1) variation in average image intensity with time of day and time of year due to the variable angle of the sun on the land surface and (2) variability of the brightness of ground features, such as dry lake beds, which can appear very similar to clouds.

4.6.1. Thresholding

One simple method of identifying clouds, called Thresholding, is based on simply finding pixels with intensity values above a certain value. This can be accomplished with a fixed threshold intensity value or with a moving threshold that depends, for example, on features within the image. A fixed threshold does not accurately represent the changes in brightness of the entire image throughout the day or seasonal variations. A dynamically adjusted threshold value must either be based on a full understanding of all variations affecting the images, or be set by the image itself by finding groups of brightest pixels. The brightest pixel method works well during any time or date, but only if there are clouds in the image. Moreover, the brightest pixel method works better under some weather conditions than others. For example, cumulus clouds are easily distinguished from the background because they are the brightest features in the image, whereas broad stratus clouds are more difficult to distinguish from background. Figure 15 shows an example of cloud detection using a threshold determined by the brightest pixel technique.

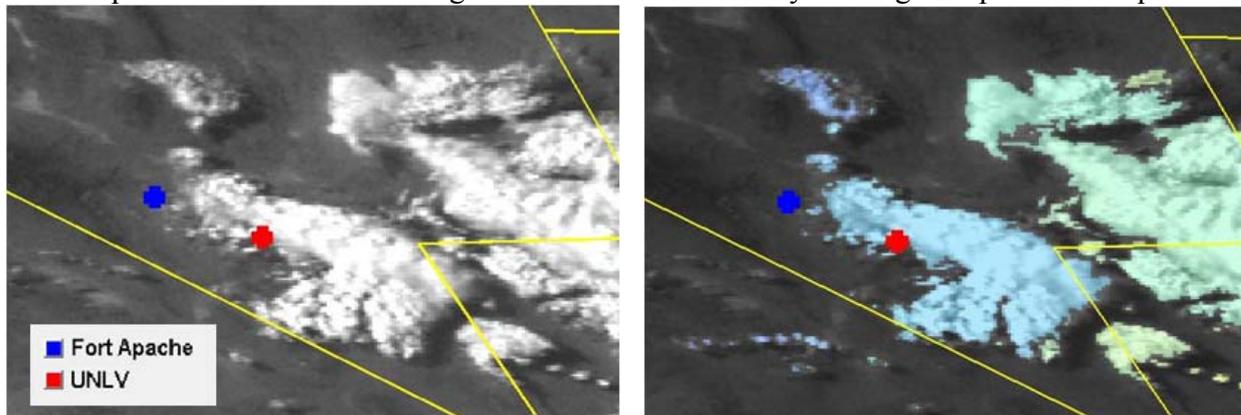


Figure 15. Example of Cloud Detection using Thresholding: raw image (left) and detected clouds (colored features, right).

¹ <http://www.itk.org>

4.6.2. Movement Detection

Another method of detecting clouds is referred to as Movement Detection. In this technique the intensity is compared between pairs of images at the pixel level. Where the difference in intensity exceeds a certain threshold, it is assumed that movement has occurred at that pixel. With accurate jitter correction, the only features in the image that can move are clouds, therefore pixels with movement are assumed to be clouds. One problem with this method is that it can only identify leading and trailing edges of clouds. This is because the middle of the cloud may cover the same pixels in both images and therefore the difference between intensity at these pixels may not exceed the threshold. Another problem with this method is that it depends on the accuracy of the image stabilization method, since errors in stabilization can lead to apparent movement of ground features and ensuing misidentification of these features as clouds. Figure 16 shows an example of cloud detection using the movement detection technique, and illustrates the main problem with this method. Parts of the clouds recognized in the image are not identified as clouds by the movement detection technique. Note how the large cloud in the center of the image appears thinner in the movement detection image. Several shadows on the ground are also identified as clouds since they move between images.

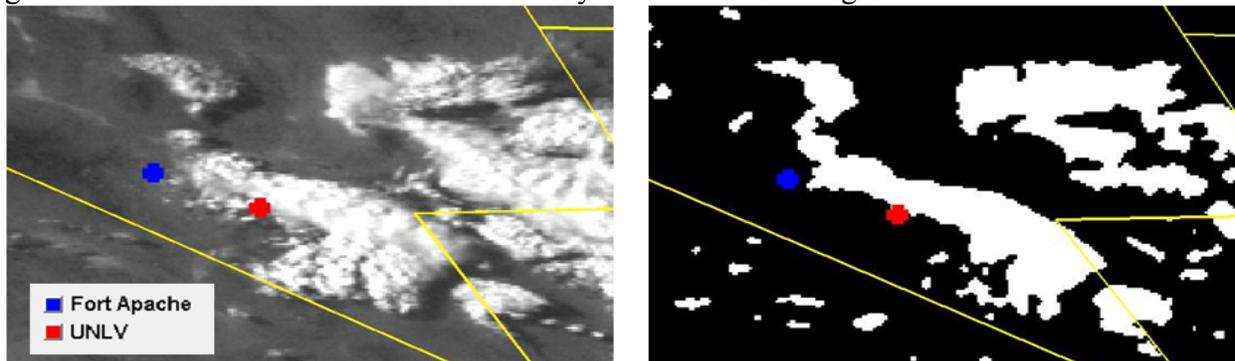


Figure 16. Example of Cloud Detection using Movement Detection: raw image (left) and detected clouds (white features, right).

Because of the challenges and the associated problems with thresholding and movement detection, we focused our attention on a third method called Background Subtraction. This is discussed in the next section.

4.6.3. Background Subtraction

The method of background subtraction is, as its name suggests, simply the process of estimating what an image of the ground would look like and subtracting this image from the actual image. Areas with clouds should then show up as areas where the intensity difference is above a certain threshold. Figure 17 shows an example of cloud detection by the method of background subtraction. Note how the background has essentially disappeared (i.e., is colored black in the right panel of Figure 17) and all that remains in the subtracted image are the clouds. This method allows for better detection of clouds represented by pixels with lower intensity. Even small changes in intensity can be detected between the image and the expected background. For example, with background subtraction a pixel could be detected if it is just slightly brighter than normal, even if it is still darker than another geographical feature in the image. As a result, image analysis depends only on the clouds in the image, and not any of the background content.

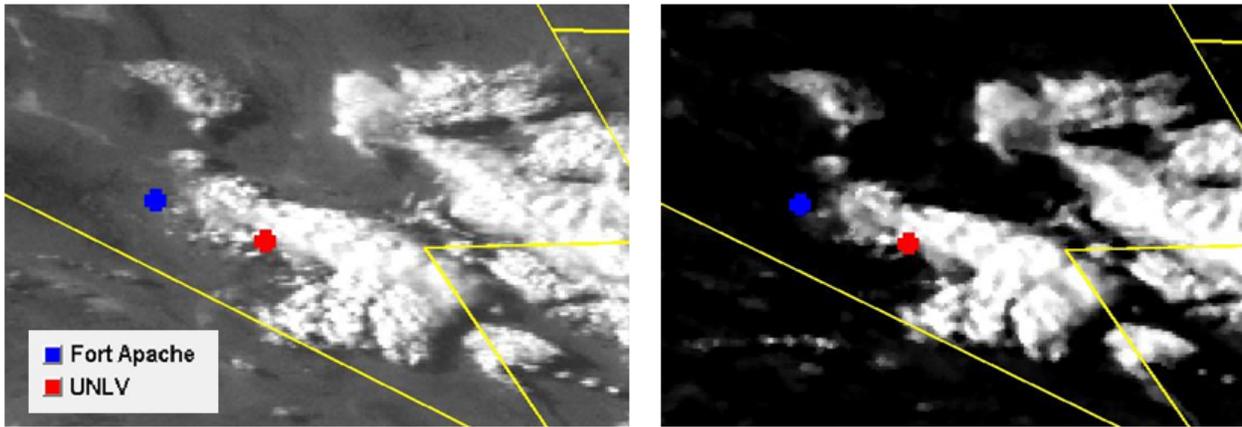


Figure 17. Example of Cloud Detection using Background Subtraction: raw image (left) and detected clouds (white features, right)

The method of Background Subtraction requires determination of a background image without the presence of clouds. Pixel intensity in an image of the background would vary due to diurnal and seasonal changes in solar illumination. Suitable background images are not available for all times of interest. We developed and trained a neural network to produce reference background images.

We first used Movement Detection to select a subset from available images that contain no clouds. For each of the images with clear skies, we computed statistics (mean, minimum and maximum) of the pixel intensity within each image. We noted that these statistics vary in a smooth manner during daylight hours and in a more complex but non-random manner annual. Figure 26 illustrates diurnal variation in minimum, average and maximum pixel intensity for one day of the year (left panel) as well as variation in average pixel intensity for different days of the year (right panel). Figure 27 shows the average pixel intensity over the course of a year for images that are classified as clear sky days. On each day, average pixel intensity begins low during morning hours, rises to a peak and then declines. Due to the display of an entire year in Figure 27 and the particular plotting program in use, average pixel intensity for individual clear sky days appears as a spike, and straight lines connect average pixel intensity values from the end of a clear sky day to the beginning of the next clear sky day. The variation of maximum average pixel intensity over the course of a year is visually indicated by the sequence of peaks.

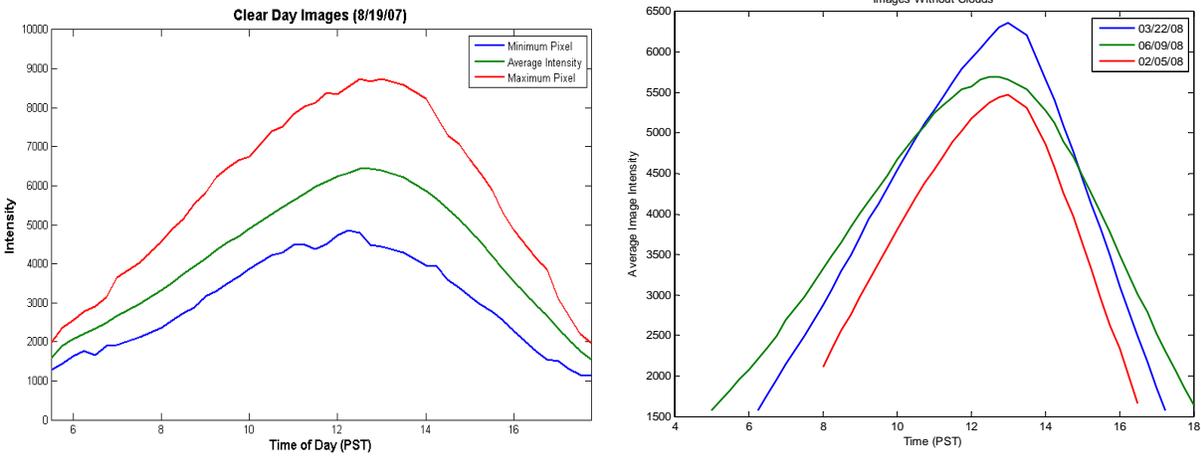


Figure 26. Diurnal variation in pixel intensity during clear skies.

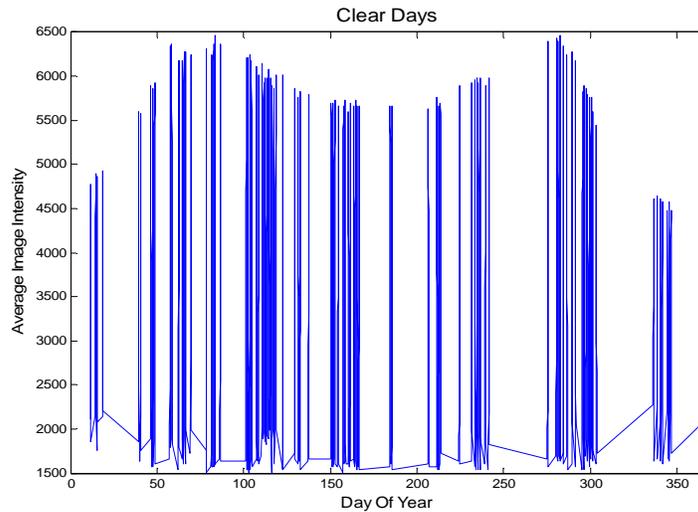


Figure 27. Annual variation in average pixel intensity during clear skies.

Because the background image varies by season and time of day, and clear sky images are not available at all times, we use a neural network trained on the clear images throughout the year to generate images for all other times during the year. Pixel intensities do not necessarily vary algebraically between clear sky images because of changes in earth's albedo, the occurrence of snowfall, and atmosphere properties. The trained neural network produces minimum, average and maximum pixel values for times when clear sky images are not available. One image is manually selected as the baseline image. A scaling algorithm is applied to adjust the pixels in the baseline image so that the synthetic image has the target minimum, average and maximum pixel intensities.

Figure 28 compares pixel intensities for synthetic images to those from images during clear sky days. The comparison shows that the neural network and scaling produce images for which the average pixel intensity follows the annual pattern. Moreover, for one clear sky day, the neural

network was used to produce synthetic images which were compared to actual images; the comparison shows that the statistics for images produced by the neural network and scaling are reasonably close to the statistics for the actual clear day images. Figure 29 compares one complete image from the GOES-11 satellite during clear sky conditions to the synthetic image generated by the neural network and scaling. Figure 29 shows that the synthetic image retains the general structure and characteristics evident in the GOES-11 image.

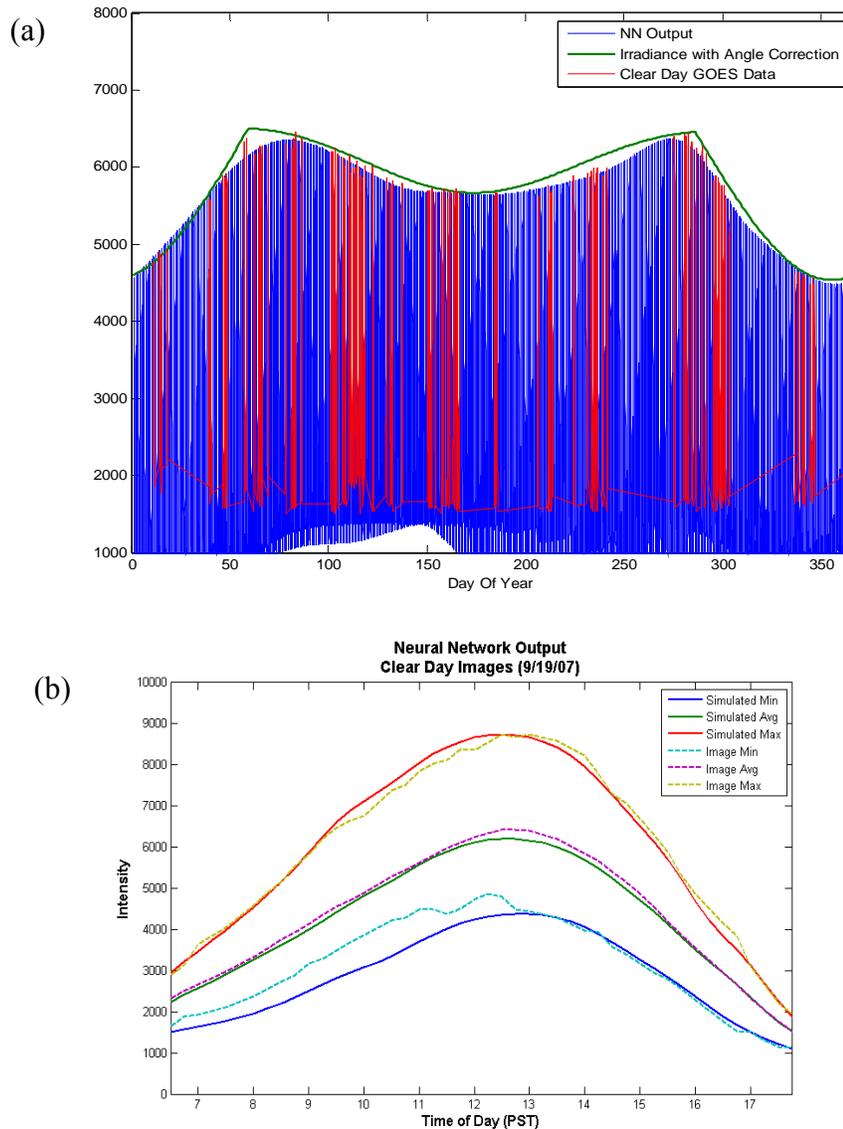


Figure 28. Comparison of statistics for pixel intensity between clear sky images and synthetic images: (a) annual average pixel intensity; (b) diurnal variation in pixel intensity.

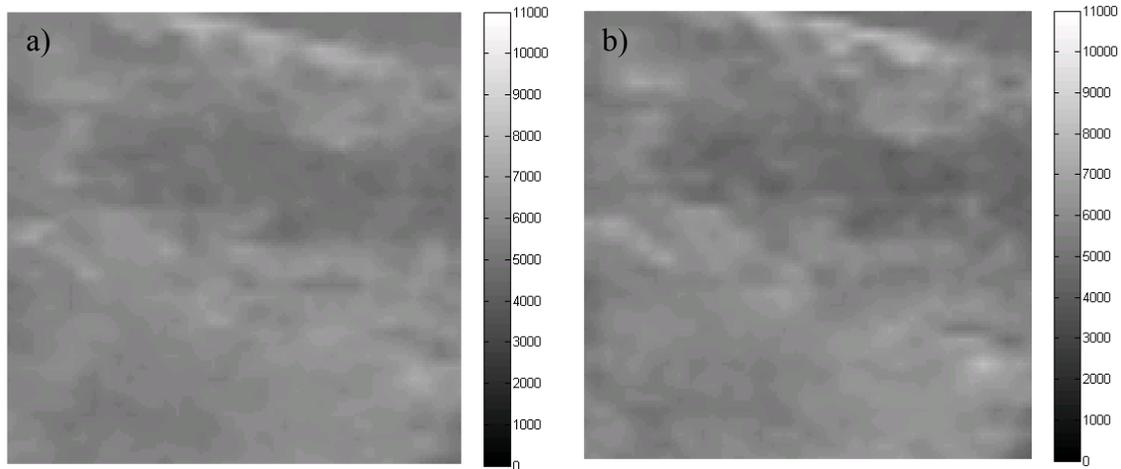


Figure 29. Images for 8/19/07 at 11:00AM (PST) centered around Fort Apache on a clear day from a) GOES-11 and b) Neural Network Simulation of the background image

4.7. Next Steps

In the previous sub-sections of this section, we describe progress toward a method for simulating one-minute irradiance from satellite imagery. We present a method for image analysis to identify and locate clouds in satellite imagery. The next step is to use the identified cloud patterns as input to a neural network that simulates ground irradiance. We began this effort by training a neural network to match cloud patterns observed in imagery with one-minute ground-based measurements of irradiance. Figure 30 shows a preliminary result of this effort. Note that the temporal patterns of irradiance from the ground sensor show that in the beginning of the day the variability has a different character (rapid shallow changes) than at the end of the day (longer duration and larger magnitude changes). It is also clear that this preliminary attempt at training a neural network to produce one-minute irradiance patterns worked quite well later in the day but was unable to match variability measured in the late morning. However, the results shown in Figure 30 represent a day with the closest match between measured and simulated irradiance; for many other days, our simulation produced irradiance predictions that were not nearly as close to ground measurements.

To further develop this approach we believe it would be more successful to train the neural network on the clearness index instead of irradiance. Clearness index is a normalized measure of irradiance that is not affected by variations to the time of day or time of year. By training the neural network on clearness index rather than irradiance we believe that we can arrive at more accurate irradiance predictions. In addition, we are exploring ways to calculate summary statistics to describe different spatial patterns of the identified cloud fields. For example, the two inset images in Figure 30 show the cloud conditions (imagery) during the beginning and end of the day shown. The difference in the cloud patterns is clear from the images and it may be that it would be more successful to correlate one-minute irradiance patterns to categorical variables relating to the type of cloud pattern rather than to raw images. We aim to investigate these ideas further in future work.

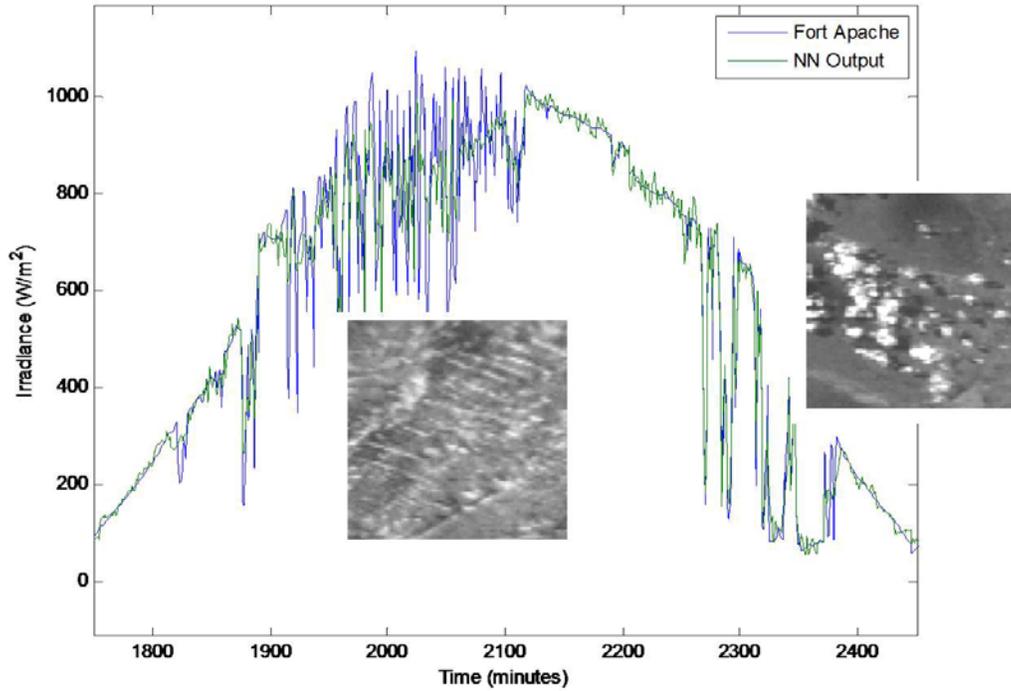


Figure 30. Measured and Simulated Irradiance for Fort Apache at 1 minute resolution for May 27, 2008. Note: Inset images illustrate cloud patterns in vicinity of Fort Apache during periods of variability in irradiance.

5. CONCLUSIONS

The scalable grid modeling task combines the high fidelity of electric circuit models with the scalability of grid level elements was developed as a tool for the use of transient as well as steady state simulation of electric power grids. Circuit simulation for electric power grid networks has some distinct advantages such as the ability to model the grid as a modular scale-up of electrical components, and the ability to handle a very large scale network via parallelizable solvers using Xyce™ [7], and the ability to interface with the GUI to display simulation values on a grid map. The developed graphical user interface not only allows the representation of data over a geographic display but can also display time series values of selected nodes. Further development will include the ability to run simulations with the user specifying which nodes and edges are to be removed (e.g. generator failure, transmission line outages, and substation blackouts).

The Reliability Impacts from Cyber Attack (RICA) task utilizes a high-powered computing platform to determine reliability results for complex electric power grids affected by various classes of cyber attack. The computed reliability degradation from the base case is due solely to the level of cyber attack and the network architecture, and it is independent of any particular outage, load, or calendar values. The RICA system has the capacity to generate results for the 30,000-bus Western Electricity Coordinating Council (WECC) region transmission grid. Results are shown for the IEEE RTS-96 72-bus system.

The PV output variability modeling task describes the preliminary steps of a method to estimate high frequency (i.e., one minute) irradiance values from satellite imagery in places where no ground measurements are available. The method proposes to use cloud patterns in each image to characterize the variability of irradiance within each 15 minute period between images. In this report, analysis methods are described to estimate cloud positions within images as a precursor to estimating high frequency irradiance.

6. REFERENCES

1. P. Kundur, *Power System Stability and Control*, New York: McGraw Hill, 1994.
2. T. R. Kuphaldt, *Lessons In Electric Circuits, Volume II – AC*, Sixth Edition, July 2007.
3. M. D. Ilić and J. Zaborsky, *Dynamics and Control of Large Electric Power Systems*, Hoboken, NJ: Wiley-Interscience, 2000.
4. L. G. Meares, and C. E. Hymowitz, “SPICE Models for Power Electronics”.
5. F. Khorrami, P. Krishnamurthy, and H. Melkote, *Modeling and Adaptive Nonlinear Control of Electric Motors*, Berlin: Springer, 2003.
6. A. Ellis, D. Kosterev, and A. Meklin, “Dynamic Load Models: Where Are We?” Proceeding of the IEEE/PES Transmission and Distribution Conference and Exhibition, pp. 1320–1324, May 2006.
7. E. R. Keiter, T. Mei, T. V. Russo, E. L. Rankin, R. P. Pawlowski, R. L. Schiek, K. R. Santarelli, T. S. Coffey, H. K. Thornquist, and D. A. Fixel, “XyceTM: Parallel Electronic Simulator, Reference Guide, Version 5.1,” Prepared by Sandia National Laboratories Albuquerque, New Mexico 87185 and Livermore, California 94550, November 2009.
8. Control Systems Roadmap Steering Group, Roadmap to Secure Control Systems in the Energy Sector, sponsored by the U.S. Department of Energy’s Office of Electricity Delivery and Energy Reliability in collaboration with the U.S. Department of Homeland Security’s Science and Technology Directorate (January 2006), <http://www.controlsroadmap.net/pdfs/roadmap.pdf>.
9. Jason Stamp, Annie McIntyre, and Bryan Richardson, “Reliability impacts from cyber attack on electric power systems”; Proceedings of the IEEE Power Systems Conference, (March 2009).
10. B. Littlewood, S. Brocklehurst, N. Fenton, P. Mellor, S. Page, D. Wright, J. Dobson, J. McDermid, and D. Gollmann, “Towards Operational Measures of Computer Security,” *Journal of Computer Security*, Vol. 2, No. 3, pp. 211-229 (1993).
11. C. Taylor, A. Krings, and J. Alves-Foss, “Risk Analysis and Probabilistic Survivability Assessment (RAPSA): An Assessment Approach for Power Substation Hardening,” in Proceedings of the 1st Workshop on Scientific Aspects of Cyber Terrorism, (November 2002).
12. S. Singh, M. Cukierz, and W. H. Sanders, “Probabilistic Validation of an Intrusion-Tolerant Replication System,” in Proceedings of the International Conference on Dependable Systems and Networks, pp. 615-624, (June 2003).
13. J. McDermott, A. Kim, and J. Froscher, “Merging Paradigms of Survivability and Security: Stochastic Faults and Designed Faults,” in Proceedings of the Workshop on New Security Paradigms, pp. 19-25, ACM:New York (2003).
14. N. F. Schneidewind, “Reliability – Security Model,” in Proceedings of the 11th International IEEE Conference on Engineering of Complex Computer Systems (ICECCS’06), pp. 269-278 (August 2006).
15. Ron Allan and Roy Billinton, “Probabilistic Assessment of Power Systems,” *Proceedings of the IEEE*, Vol. 88, No. 2 (February 2000).
16. J. R. Conrad, “Analyzing the Risks of Information Security Investments with Monte-Carlo Simulations,” in Proceedings of the 4th Workshop on the Economics of Information Security, Kennedy School of Government, Harvard University (June 2005).

17. Marcus Schilling, Armando Leite De Silva, Roy Billinton, and M.A. El-Kady, "Bibliography on Power System Probabilistic Analysis (1962- 1988)," *IEEE Transactions on Power Systems*, Vol. 5, No. 1, pp. 41-49 (February 1990).
18. Ron Allan, Roy Billinton, Art Breipohl, and Cliff Grigg, "Bibliography on the Application of Probability Methods in Power System Reliability Evaluation 1967-1991," *IEEE Transactions on Power Systems*, Vol. 9, No. 1, pp. 41-49 (February 1994).
19. Ron Allan and Roy Billinton, et al., "Reliability Assessment of Composite Generation and Transmission Systems," IEEE Power Engineering Society Tutorial, 90EH0311-1-PWR (1989).
20. Jason Stamp, John Dillinger, William Young, and Jennifer DePoy, Common Vulnerabilities in Critical Infrastructure Control Systems, Sandia National Laboratories report SAND2003-1772C: Albuquerque, New Mexico (2003); presented at SANS SANSFIRE 2003 and National Information Assurance Leadership Conference V (NIAL V), Washington, DC (July 14-22, 2003).
21. Michael Berg and Jason Stamp, A Reference Model for Control and Automation Systems in Electric Power, Sandia National Laboratories report SAND2003-1000C: Albuquerque, New Mexico (2005); published in the Proceedings of the 2005 Power Systems Conference, Clemson University (March 2005).
22. Nancy Spring, "Industry and Government Partnering for Cyber Security," *Electric Light & Power* (August 2006), http://uaelp.pennnet.com/articles/article_display.cfm?article_id=261363.
23. Annie McIntyre, Jason Stamp, and Ben Cook, I3P Risk Characterization Report, I3P Research Publication, Dartmouth College, New Hampshire (2007).
24. David Duggan, Categorizing Threat: Building and Using a Generic Threat Matrix, Sandia National Laboratories report SAND2007-5791, Albuquerque, New Mexico (September 2007).
25. Roy Billinton and Wenyuan Li, *Reliability Assessment of Electric Power Systems Using Monte Carlo Methods*, Plenum Press, New York (1994).
26. J.R. Ubeda and Ron Allan, "Sequential Simulation Applied to Composite System Reliability Evaluation," *IEE Proceedings C*, Vol. 139, No. 2, pp. 81-86 (March 1992).
27. Roy Billinton and J. Tatla, "Composite Generation and Transmission System Adequacy Evaluation Including Protection System Failure Modes," *IEEE Transactions on Power Apparatus and Systems*, Vol. PAS- 102, No. 6, pp. 1823-1830 (June 1983).
28. Kang Lin and Keith Holbert, "PRA for Vulnerability Assessment of Power System Infrastructure Security," Proceedings of the 37th Annual North American Power Symposium, pp. 43-51 (October 2005).
29. Roy Billinton, S. Kumar, N. Chowdhury, K. Chu, K. Debnath, L. Goel, E. Khan, P. Kos, G. Nourbakhsh, J. Oteng-Adjei, "A Reliability Test System for Educational Purposes — Basic Data," *IEEE Transactions on Power Systems*, Vol. 4, No. 3, pp. 1238-1244 (August 1989).
30. Billinton, R.; Hua Yang; "Incorporating maintenance outage effects in substation and switching station reliability studies," 2005. Canadian Conference on Electrical and Computer Engineering, 1-4 May 2005 Saskatoon, Page(s):599 - 602.
31. Reliability Test System Task Force of the IEEE Power Engineering Society, "The IEEE Reliability Test System – 1996," *IEEE Transactions on Power Systems*, Vol. 14, No. 3 (August 1999).

32. Reliability Test System (RTS) – 1996 files, “Power Systems Test Case Archive,” Department of Electrical Engineering at the University of Washington, https://www.ee.washington.edu/research/pstca/rtspg_tcart.htm (retrieved March 2008).
33. Mills, A., M. Ahlstrom, et al. (2009). Understanding Variability and Uncertainty of Photovoltaics for Integration with the Electric Power System, Berkeley National Laboratory.
34. Perez, R., P. Ineichen, et al. (2002). "A new operational model for satellite-derived irradiances: description and validation." *Solar Energy* 73(5): 307-317.

DISTRIBUTION

1	MS0316	Thomas V. Russo	01445
1	MS0321	Robert W. Leland	01400
1	MS0321	John L. Mitchiner	01460
1	MS0370	John D. Siirola	01465
1	MS0671	Laurence R. Phillips	05628
1	MS0671	Bryan T. Richardson	05628
1	MS0671	Andrew C. Riehm	05628
1	MS0721	Marjorie L. Tatro	06100
1	MS0932	Paul R. Wolfenbarger	09513
1	MS1033	Stanley Atcity	06113
1	MS1033	Abraham Ellis	06112
1	MS1033	Charles J. Hanley	06112
1	MS1033	Clifford W. Hansen	06112
1	MS1033	Matthew J. Reno	06112
1	MS1033	Joshua S. Stein	06112
1	MS1082	Anthony L. Lentine	01727
1	MS1104	Rush D. Robinett III	06110
1	MS1108	Ross Guttromson	06113
1	MS1108	Karina Munoz	06111
1	MS1108	Jason E. Stamp	06111
1	MS1108	Juan J. Torres	06111
1	MS1124	David G. Wilson	06122
1	MS1152	Steven F. Glover	01654
1	MS1315	Jeffrey S. Nelson	01130
1	MS1316	Mark Daniel Rintoul	01465
1	MS1318	Brian M. Adams	01441
1	MS1318	Bruce A. Hendrickson	01440
1	MS1318	Robert J. Hoekstra	01426
1	MS1318	Jean-Paul Watson	01465
1	MS1319	William C. McLendon	01461
1	MS1319	Ron A. Oldfield	01423
1	MS1321	David A. Schoenwald	01444
1	MS1321	Randall M. Summers	01444
1	MS1322	Sudip S. Dosanjh	01420
1	MS0359	D. Chavez, LDRD Office	01911
1	MS0899	Technical Library	09536 (electronic copy)



Sandia National Laboratories