

Modeling Smart Grid Applications with Co-Simulation

Tim Godfrey, Sara Mullen,
Roger C Dugan, Craig Rodine
Electric Power Research Institute
Knoxville, TN and Palo Alto, CA, USA

David W. Griffith, Nada Golmie
National Institute of Standards and Technology
Gaithersburg, MD, USA

Abstract— Our analysis of a complex Smart Grid control scheme uses simulation to model both the communication network and the power system. The control scheme uses a wireless communication network to activate distributed storage units in a segment of the electrical grid to compensate for temporary loss of power from a solar photovoltaic (PV) array. Our analytical model of the communication network provides a means to examine the effect of communication failures as a function of the radio frequency (RF) transmission power level. We use these results in an open source event-driven simulator to determine the impact on the electrical power system.

I. INTRODUCTION

Engineers working on Smart Grid initiatives are modeling communications networks to gauge their ability to support next-generation power system applications. Progress has been made using analytic models and bulk traffic estimates based on Smart Grid use cases. However, event-driven simulation is required to model dynamic system behavior, and to probe the limits of network performance at “utility scale” and under challenging conditions.

It would also be desirable to link communication system modeling with simulation of emerging Smart Grid applications on actually deployed power systems. Such a co-simulation environment would allow engineers to assess the feasibility of using a given network technology to support communication-based Smart Grid control schemes on an existing segment of the electrical grid; and conversely, to determine the range of control schemes that differing communications technologies can support.

We constructed a co-simulation platform by linking the Open Distribution System Simulator (OpenDSS) with the ns-2 Network Simulator. We simulated a plausible deployment of distributed energy resources (DERs - a large photovoltaic solar source and 84 small-scale storage batteries) on a model of an actual distribution circuit (feeder). In this scenario, the underlying communications system is based on IEEE 802.11.

We present the baseline scenario without any communication, and show the response to power fluctuations. We then introduce the wireless communications network to provide real-time sensing and control. First, an analytical model examines the aggregate behavior of the wireless network and demonstrates the available capacity to support the application. Then an event model is used to focus on a single power interruption event, and perform a detailed analysis of device-to-device communication and the resulting mitigation of voltage variation in the time domain.

II. OVERVIEW OF THE SMART GRID SCENARIO

Renewable energy is an important element in proposed implementations of the Smart Grid. However, one challenging characteristic of many renewable energy technologies is their often unpredictable variation in power output.

One issue that is attracting attention recently is the so-called “cloud transient” or “solar ramping” phenomenon - when a cloud passes in front of PV panels, reducing power output to the point that inverter output ramps down quickly. This power variation has not been a significant problem with small rooftop PV units dispersed over a large area, where the geographic distribution of the units makes it unlikely that all the units will ramp up or down simultaneously. However, many electric utilities are receiving proposals to install large (2 MW and up) PV systems in a single location [1]. These locations are often remote from the nearest substation on weaker circuits. Voltage profile analysis sometimes indicates a significant voltage excursion if the PV output power were to drop suddenly.

The solar ramping impact on voltage regulation is being addressed on a number of fronts. One approach is to modify inverter behavior to increase reactive power output as the active power output drops. Another approach is to dispatch storage to replace the active power output of solar PV system. This is the approach that is simulated in this paper. The problem will be to compensate for output variation in a large PV generator with distributed storage units.




Fig. 1. Circuit diagram

The hypothetical example simulated here assumes a 2.5 MW solar photovoltaic system has been installed

approximately 3 km from the substation. This location is at the end of a feeder branch on which a standard utility 32-step voltage regulator has been installed to help maintain good voltage regulation to compensate for the slower variations in the PV generator output (see Fig. 1). This feeder also contains 2.1 MW of distributed storage implemented with 84 units of 25 kW capacity each. These units are mainly intended for substation peak shaving and residential reliability purposes, but the question we address here is: Can the distributed storage be exploited to compensate for cloud transients that cause the large PV unit to ramp down?

There are at least two issues that may work against this scheme:

1. The distributed storage units cannot be dispatched quickly enough, and
2. The storage units may not be sited in satisfactory locations to be effective.

The former issue is addressed as the principal subject of this paper, the co-simulation of communications systems and electric power distribution systems. The latter issue is naturally captured in the power system model, which has the storage units explicitly modeled in their actual locations.




Fig. 2. Assumed solar ramp characteristic

The characteristic assumed for the cloud transient is shown in Fig. 2. The power output is assumed to drop at 10 % per second and remains at zero for approximately 60 seconds, after which it returns to normal output at a slower rate. The sudden loss of generated power causes the voltage to drop almost 2 %, which causes the three single-phase regulators to begin tapping up to compensate (Fig. 3). Since this might happen many times a day on partly cloudy days, the mechanical regulator tap changer would operate much more than is typically expected, shortening its life. Utility customers in the vicinity would experience numerous voltage excursions of approximately 3 % as the regulator attempts to compensate for the fluctuations. While this is not extreme, it is likely to be visible, resulting in voltage quality complaints from consumers on the feeder.

The regulators in this case have a time delay of 15 s before they begin acting. This is a short delay for a typical utility regulator; they are more typically set for 30 s or 45 s. For the purposes of this simulation, it is assumed the regulators have been set to respond faster than usual to better compensate for

the PV output power variation. They still lag significantly behind the voltage changes induced by the cloud transient. As the PV output power ramps back up, there is an overshoot in the voltage until a regulator operates to compensate.




Fig. 3. Impact of solar ramp on feeder voltage (simulated with EPRI OpenDSS program)

III. AN ANALYTICAL MODEL OF THE COMMUNICATION LINKS

In the model that we developed to predict the communication network's performance, there is a single transmitter that broadcasts a message to a population of N passive receiver stations. The transmitter sends a copy of the message to each receiver in turn, and moves on to the next station when the message is successfully received or the maximum number of re-transmission attempts has been met.

Because the stations are passive and there are no other transmitters, there are never any collisions on the channel. Instead, frames are lost due to effects at the physical layer. In particular, a frame is considered lost if the signal to interference-plus-noise ratio (SINR) is smaller than some minimum acceptable value, which we denote as z_0 . SINR is the ratio of the message's signal power at the receiver, P_r , to the sum of the noise and interference powers which are respectively P_N and P_I . The power levels are often expressed in units of decibel milliwatts (dBm). The probability that a given transmission attempt fails is therefore

$$P_{\text{fail}} = \Pr\{P_r - (P_N + P_I) < z_0\}$$

where all powers are expressed in dBm, and the ratio threshold $z_0 = 2.5$ dB.

For this scenario, we do not have any active interference sources, so $P_I = 0$. The signal is subject to path loss, fading, and shadowing. For a given transmit power level P_t , expressed in dBm, the received signal strength in dBm is

$$P_r = P_t - L_{\text{path}} + X_{\text{shadow}} + X_{\text{fade}}$$

In the above equation, L_{path} is the path loss, which depends on the path loss exponent, γ , and X_{shadow} and X_{fade} are additive factors that are due to shadowing and fading effects [2]. The additive factor X_{fade} is a gamma distributed random variable with shape and scale parameters m_{fade} and $1/m_{\text{fade}}$, respectively, where m_{fade} is the Nakagami fading parameter. When $m_{\text{fade}} = 1$, we have Rayleigh fading. If $m_{\text{fade}} \rightarrow \infty$ there is no fading.

For a given set of model parameters, we can compute

performance metrics for the message transmission event once we have the message failure probability for each of the receivers in the network. We use a simplified model of the IEEE 802.11 backoff mechanism that was first studied in detail by Bianchi [3]. Because the transmitter is the only active sender in the network, the channel is never active while it is in backoff. Thus the mean time for the transmitter to wait in the i th stage, whose maximum number of backoff slots is W_i , is $(W_i - 1)s/2$, where s is the duration of a backoff slot when the channel is idle. Typically, $s = 20 \mu s$. There are $\alpha + 1$ backoff stages when finite retransmission are allowed [4]; the number of backoff slots in each stage is

$$W_i = \begin{cases} 2^i W_0, & i \leq n \\ 2^n W_0, & n < i \leq \alpha \end{cases}$$

For our system, we choose $W_0 = 16$, $n = 6$, and $\alpha = 7$. The expected delay given that the transmitter passed through i backoff stages to send the message is

$$E\{D | i\} = \mu_s + i\mu_F + E\{B | i\}$$

where μ_s is the mean time to complete sending a message successfully, μ_F is the mean time to complete sending a message that fails, and $E\{B | i\}$ is the mean time that the transmitter spends waiting in the 0th through i th backoff stages, given that no other transmitter is using the channel:

$$E\{B | i\} = \begin{cases} \left((2^{i+1} - 1)W_0 - (i+1) \right) \frac{s}{2}, & 0 \leq i \leq n \\ \left((2^n (2 + i - n) - 1)W_0 - (i+1) \right) \frac{s}{2}, & n < i \leq \alpha \end{cases}$$

We assume that the fading is sufficiently fast that we can treat consecutive transmissions as independent events. Summing up the conditional expected delays, we find that the average delay at the MAC layer for a given receiver that will not receive a message with probability P_{fail} is

$$\begin{aligned} E\{D\} &= \sum_{i=0}^{\alpha} (1 - P_{\text{fail}}^i) P_{\text{fail}}^i E\{D | i\} + P_{\text{fail}}^{\alpha+1} [(\alpha+1)\mu_F + E\{B | \alpha\}] \\ &= \left[\frac{1 - \left[1 + \left(\frac{P_{\text{fail}}^n}{1 - P_{\text{fail}}} + (1 - 2P_{\text{fail}}) \frac{P_{\text{fail}}^{\alpha}}{1 - P_{\text{fail}}} \right) 2^n \right] P_{\text{fail}}}{(1 - P_{\text{fail}})(1 - 2P_{\text{fail}})} W_0 - \frac{1 - P_{\text{fail}}^{\alpha+1}}{1 - P_{\text{fail}}} \right] \frac{s}{2} \\ &\quad + \frac{1 - P_{\text{fail}}^{\alpha+1}}{1 - P_{\text{fail}}} P_{\text{fail}} \mu_F + (1 - P_{\text{fail}}^{\alpha+1}) \mu_s. \end{aligned}$$

The time to transmit the message to the full ensemble of receivers is the sum of the expected delays for each of the receivers. It also follows that transmission attempts to different receivers are independent trials. Thus the expected number of stations that do not receive the message in $\alpha + 1$ attempts is given by

$$E\{N_{\text{failures}}\} = \sum_{n=1}^N P_{\text{fail},n}^{\alpha+1}$$

where $P_{\text{fail},n}$ is the message failure probability at the n th receiver.

In Fig. 4, we plot the expected number of messages that do not reach the receivers versus the transmit power level, for path loss exponent values of 2.5, 3.0, and 4.0. A path loss exponent of 2.5 corresponds to a relatively open environment that is free of physical obstructions; larger path loss exponents correspond to more obstructions in the physical environment. The plot shows that when the path loss exponent is relatively small, we have good performance for transmit power levels greater than 40 mW, with noticeable performance degradation occurring at transmit power levels of around 20 mW and below. Increasing the path loss exponent to 3.0 results in unacceptable performance unless the transmit power is at least 900 mW. A path loss exponent of 4.0 results in completely unacceptable performance even at relatively high power levels.




Fig. 4. Mean number of lost messages versus transmit power

IV. EVENT MODEL OF THE SCENARIO

We now examine the solar ramp event in more detail using event modeling. A co-simulation approach is used. The OpenDSS simulator is used for the electrical system, and ns-2 is used for the wireless communications network.

A. OpenDSS Simulator

The Open Distribution System Simulator (OpenDSS) is a comprehensive electrical system simulation tool for electric utility distribution systems [5]. EPRI has made this program freely available to spur the advancement of grid modernization efforts by providing researchers with a capable tool.

The program supports nearly all rms steady-state (i.e., frequency domain) analyses commonly performed for utility distribution system planning. In addition, it supports types of analyses that are designed to meet future needs, many of which are being driven by the development of the Smart Grid. The program was originally intended to support the needs of distributed generation analysis. Other features support energy efficiency analysis of power delivery and harmonics analysis. It is somewhat unique for a distribution planning tool, being designed to simulate discrete events.

B. Co-Simulation

The problem scenario is analyzed using co-simulation between a communications simulator (ns-2) and a utility power distribution system simulator (OpenDSS). As the solar PV output ramps down, the storage controller will attempt to

dispatch the distributed storage (DS) elements once per second to maintain a smooth voltage. ns-2 will be used to simulate the arrival of the messages at the storage units, which are assumed to respond instantly. The timings of the dispatch commands are fed back into the OpenDSS as a script. The response of the system voltages and power flow is then computed. Fig. 5 depicts the interaction between OpenDSS and ns-2.




Fig. 5. Co-simulation flow chart

The OpenDSS environment provides output data representing the time of the PV ramp event, the topological coordinates of the DS nodes, and the power output (load) profile for the DS nodes. Scripts parse the OpenDSS output to configure ns-2 with the node topology.

ns-2 simulates the arrival of the dispatching messages at each of the 84 storage units. The arrival times are then sorted and used to create OpenDSS scripts that are fed into the OpenDSS engine. The OpenDSS then performs a sequence of power flow solutions at the specified time assuming the storage elements respond without delay. In future simulations, this assumption will be adjusted, once the actual behavior of the storage elements is better understood.

The control that monitors the solar PV output is assumed to sample the output once per second. The control is assumed to send another batch of dispatch messages with each sample. It is assumed the communications latency is identical for each batch of messages. The simulation continues, repeating the sampling and communications latency, re-dispatching the storage elements until there is no more remaining capacity. Since there is only 2.1 MW of storage to compensate for 2.5 MW of PV output it cannot be expected to perfectly compensate for the assumed solar ramp.

The control signal in this simulation is always 1.0 s behind the solar ramp without considering the communications latency. This naturally creates a sawtooth characteristic in the voltages. In future work, when the simulators are more closely linked, we intend to consider proportional controllers that send dispatch messages based on the percentage change. This will likely create more messages at random intervals, which raises the possibility of increased communications congestion.

OpenDSS is an interactive discrete event simulator that is

driven by a script or another program. A typical snippet of the OpenDSS script used here is shown in Fig. 6. This snippet shows the dispatch of the first 2 storage units. The time is set to the arrival time of the dispatch message computed by ns-2. Then the state of the target storage unit is set to be discharging at the designated rate, 11.9 % in this case. This compensates for the first 10 % drop in the solar PV output, 250 kW. Next, the new circuit condition is solved (Solve) and the results captured in the monitors (Sample). This is repeated for each storage element at each step of the simulation.

```
Set sec = 22.020834372 ! Unit 1
storage.jo0235001304.state=discharging %discharge=11.9
Solve
Sample
Set sec = 22.022028115 ! Unit 2
storage.jo0235000257.state=discharging %discharge=11.9
Solve
Sample
...etc...
```

Fig. 6. Snippet of OpenDSS script for simulating storage unit dispatch

C. The 802.11 Simulation Model in ns-2

The wireless network is based on 802.11, since mature simulation models have been developed in ns-2. Specifically, we are using the 802.11Ext model which is distributed as a part of ns-2 [6] from version 2.34. The 802.11Ext model was developed by Mercedes Benz Research and Development North America and the University of Karlsruhe [7].

The model is configured to simulate 802.11 operating in the 915 MHz Industrial, Scientific and Medical (ISM) band. This band is not currently standardized by IEEE 802.11, although a Study Group has been formed with that objective. The band is 26 MHz wide, (902—928) MHz, so it only supports a single 20 MHz channel. However, IEEE 802.11-2007 [8] contains mechanisms for operating the Orthogonal Frequency-Division Multiplexing (OFDM) PHY in 5MHz and 10 MHz channels (first introduced with the 802.11j amendment). Our model is configured to simulate a 5 MHz channel operating at 910 MHz. Based on the 5 MHz channel parameters specified in IEEE 802.11-2007 Table 17-1, Table 17-4, Table 17-15 and clause 17.3.8.6, we are using the Physical Layer (PHY) and Medium Access Control (MAC) parameters shown in Table I.

TABLE I
PHY AND MAC PARAMETERS

Frequency	910 MHz
Channel Bandwidth	5.0 MHz
Noise Floor	2.51189e-13
Path Loss Exponent	2.5
Data Rate	1.5 Mb/s
Slot Time	21 μ s
Short Inter-Frame Space (SIFS)	64 μ s
Symbol Duration	16 μ s

A 100 byte payload carried by the MAC Service Data Unit (MSDU) frame is selected as sufficient to support a small command with User Datagram Protocol (UDP), Internet

Protocol Security (IPsec), and Internet Protocol Version 6 (IPv6) overhead.

D. Range and RF Power

For the Smart Grid scenario under consideration the distance distribution between the Storage Controller and the DS devices is shown in Fig. 7.




Fig. 7. Number of nodes located in distance zones

The simulation sends message frames in the order of node number, which is not related to position or distance. All 84 frames are queued for transmission at $t = 10$ ms (a fixed delay factor to allow for the reception of the PV message and processing delay). Each frame is sent sequentially, and each either succeeds or is discarded if the maximum number of retry attempts is exceeded.

The RF power is varied to simulate wireless impairments such as propagation loss, fading, and interference. At lower power, the receiver has an increased probability of failing to receive a frame, and thus require retries. As power levels continue to drop, frames are dropped due to retry timeouts. Since the frames are sent sequentially, plotting the arrival time is a useful way to visualize the system performance. Fig. 8 shows the arrival of frames plotted against time for different power levels.




Fig. 8. Frame arrival time

In the 400 mW case, all 84 frames arrive in 118 ms. The full simulation trace shows only a few frames required retransmissions. At 80 mW, more frames require retries, and the delay increases to 206 ms. At 30 mW, the total delay is 753 ms, and the curve does not reach 100 % because 3 frames are dropped due to retry timeout. The 10 mW power level is not usable; as 23 frames are dropped due to retry timeout.




Fig. 9. Summary of time and packet loss vs. RF power

Fig. 9 summarizes the wireless performance. For RF power levels over 100 mW, the wireless communication is fast and robust, and leaves a good margin for interference and fading. The inflection point for frame dropping due to retry timeout is at 30 mW.

V. CO-SIMULATION RESULT

In the base case, the output of the 2.5 MW PV installation begins to ramp down by 10 % per second at 5 s, as seen in Fig. 10. The solar ramp ends at $t = 15$ s and the voltage stabilizes on all three phases. The average voltage drop on the phases due to the solar dropout is 1.85 %.




Fig. 10. PV voltage magnitude with solar ramp start at 5 s, with no distributed storage

When the output from the solar panel begins to drop, the storage controller is notified to trigger discharge of the DS units to mitigate the drop in voltage. As the output continues to drop, the controller is notified each second to keep increasing the output of the DS units until the drop in output ceases or the storage units are fully dispatched.

A. Case 1: Transmission RF Power of 400 mW

The output of the solar panel is sampled once a second and a 10 ms delay is assumed between sensing the drop and sending the messages from the storage controller to the DS units. The solar dropout begins at 5 s. Therefore, the first set of messages is sent from the controller at 6.01 s. In this case a transmission RF power, P_r , of 400 mW is used. All 84 messages reach their destinations at the DS units, with a maximum communications delay of 118 ms.

The first round of dispatch sets all DS units to discharge at a rate of 11.9 %. As the voltage continues to drop a new set of dispatch messages is sent every second, and the discharge rate is increased by 11.9 %, until 100 % discharge rate is reached. By $t = 15$ s, after 9 dispatch requests, all units are at 100 % discharge rate. Fig. 11 shows the voltage at the PV as

the DS units compensate for the drop in PV output power.

With the use of distributed storage the average voltage drop on the phases is 0.74 % compared to 1.85 % without storage and the steepness of the drop is lessened. Note that the voltage drop in each phase is different than in the base case, due to the distribution of the storage units among the phases.




Fig. 11. Voltage magnitude without communication and with communication at transmitted RF power of 400 mW

B. Case 2: Transmission RF Power of 30 mW

In Case 2 messages are sent from the storage controller with transmission RF power of $P_t = 30$ mW, with the first round of messages again being sent at 6.01 s. With decreased RF power, only 96 % of the dispatch messages are successfully received by the DS units and only 81 units begin discharging. The maximum communications delay is 753 ms and the effect of the DS units coming online is more spread out than in Case 1. The PV voltage response can be seen in Fig. 12.




Fig. 12. Voltage magnitude without communication and with communication at transmitted RF power of 30 mW

Because only 81 DS units are dispatched at each interval, the voltage drop is not mitigated the same as in the previous cases. The average voltage drop on the phases is 0.77 % which is slightly higher than the drop in Cases 1 and 2, but still significantly less than the 1.85 % drop without storage.

C. Discussion

The DS units can compensate for about half of the voltage drop. This is sufficient to prevent the regulators from tapping and the rate of change of voltage is significantly reduced. This simulation shows that the location of storage elements on the system definitely matters, which demonstrates the value of this kind of simulation. To best compensate for the

PV power ramp, a similar amount of storage capacity would have to be located near the PV array site.

The latency in the control sampling and communications produces the sawtooth shape of the voltage characteristic. It is not likely that this will cause many complaints because the voltage changes are quite small. However, there may be sensitive load equipment that reacts adversely to this wavering voltage. The co-simulation technique would be valuable in determining a control and communications scheme that might function more acceptably.

VI. CONCLUSIONS AND FUTURE WORK

We are working to achieve a much tighter linking of ns-2 and OpenDSS so that we can investigate a wide range of Smart Grid issues where communications latency can adversely impact the expected behavior of the power system. A linking of a communications simulator and a power systems simulator has been accomplished and applied to a practical problem. The results of the simulation demonstrate the value of this capability for investigating Smart Grid applications.

ACKNOWLEDGMENT

The example distribution system was inspired by an actual feeder being studied in an EPRI Smart Grid Demonstration project with American Electric Power Co., Inc. (AEP), Columbus, Ohio USA, that features AEP's Community Energy Storage (CES) concept.

REFERENCES

- [1] S. Steffel, "Distribution Grid Considerations for Large Scale Solar and Wind Installations", 2010 IEEE PES Transmission and Distribution Conference and Exposition, New Orleans, April 2010.
- [2] Rappaport, T.S., Wireless communications: principles and practice, Prentice Hall PTR New Jersey, 2002.
- [3] Bianchi, G., "Performance analysis of the IEEE 802.11 distributed coordination function", IEEE Journal on selected areas in communications, vol. 18, no. 3, pp 535—547, 2000.
- [4] Zhai, H., Kwon, Y., and Fang, Y., "Performance analysis of IEEE 802.11 MAC protocols in wireless LANs," Wireless communications and mobile computing, vol. 4, no. 8, pp. 917—931, 2004.
- [5] OpenDSS Users Manual, Sourceforge.net website, <http://electricdss.svn.sourceforge.net/viewvc/electricdss/doc>
- [6] "Network Simulator ns-2," <http://www.isi.edu/nsnam/ns/>
- [7] Qi Chen, Felix Schmidt-Eisenlohr, Daniel Jiang, Marc Torrent-Moreno, Luca Delgrossi, Hannes Hartenstein, "Overhaul of IEEE 802.11 Modeling and Simulation in NS-2 (802.11Ext)", in http://dsn.tm.uni-karlsruhe.de/medien/downloads_old/Documentation-NS-2-80211Ext-2008-02-22.pdf
- [8] "IEEE Std. 802.11-2007, Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) specifications," IEEE Std. 802.11, 2007

DISCLAIMER: The full description of the procedures used in this paper requires the identification of certain commercial products and their suppliers. The inclusion of such information should in no way be construed as indicating that such products or suppliers are endorsed by NIST or EPRI, or are recommended by NIST or EPRI, or that they are necessarily the best materials, instruments, software or suppliers for the purposes described.