Wireless Network Deployment in the Smart Grid: Design and Evaluation Issues

Camillo Gentile, David Griffith, and Michael Souryal,
National Institute of Standards and Technology

Abstract

The projected use of the power grid for "smart" applications such as advanced metering and distributed automation will require highly reliable, secure, well designed and managed communication networks. While many of the benefits of wireless communications, such as untethered access to information, support for mobility, and reduced infrastructure, would be available to the grid, there are still a number of unanswered questions regarding network performance, suitability, and security. In this article, we introduce the communication requirements that have been established for these applications thus far. Throughout, we highlight the implications of wireless deployment as it relates specifically to the smart grid and we identify key issues that must be considered when evaluating a wireless technology against the communication requirements.

The use of wireless technologies in the grid for applications such as system monitoring, metering, and data gathering goes back several decades, but they have the potential to be used to a much greater extent as utilities deploy the infrastructure for the smart grid. Wireless communications offer many potential benefits to utilities, including untethered data access by mobile repair crews and the rapid expansion of communications infrastructure to areas where building out a fiber or coax backbone is not economically or technically feasible. For distribution systems, which are by nature dispersed over a large geographic area, rapid detection of outages and repair scheduling is important, especially in severe weather conditions. Since communications lines tend to falter in these conditions, wireless can also serve as readily deployable on-demand backup systems. Wireless can also be used to support applications in which wired communications are not feasible, such as gas metering, where isolating the pipe from electric wires is a matter of safety.

At the same time, there are a number of challenges associated with using wireless communications in the smart grid. The most important issues involve network performance and security. Can existing networking technologies and, more specifically, wireless communication systems support the communication and reliability requirements specific to the smart grid? How does one evaluate different technologies with very different characteristics? Are there any interference issues, and what is their impact on system performance and reliability?

In this article, we describe how smart grid data traffic is being specified and how these specifications can be used as a basis for developing smart grid traffic models. We reference the ongoing effort by the Open smart grid (OpenSG) subcommittee of the UCA International Users Group (UCAIug) to catalog the applications and their performance requirements. We also reference a methodology described in our previous work to convert these specifications into a format that is useful to network designers. After giving a brief overview of the available wireless technologies, we identify key issues that must be considered when evaluating candidate technologies against the requirements. Among the factors that will differentiate smart grid wireless deployments from commercial deployments are the predominance of fixed devices, relatively low data rates at the network edge during normal operations, deterministic traffic patterns for some applications, and the fairly high robustness and reliability requirements.

Certain commercial equipment, instruments, or materials are identified in this article to foster understanding. Such identification does not imply recommendation or endorsement by the National Institute of Standards and Technology, nor does it imply that the materials or equipment identified are necessarily the best available for the purpose.

Smart Grid Communication Requirements

The NIST Framework and Roadmap for smart grid (SG) Interoperability Standards [1] is a working document prepared by the National Institute for Standards and Technology (NIST) in partial fulfillment of its mandate under the Energy Independence and Security Act of 2007 (EISA). It provides a catalog of SG applications over the various grid domains, operations, generation, transmission, distribution, markets, service providers, and customers. Network entities, known as *actors*, are the endpoints for communications flows associated with smart grid applications. For example, the smart meter (SM) actor at the customer premises must be able to respond to an on-demand read request from the customer information systems (CIS/Billing) actor in operations. In order to enable a thorough analysis of the grid communication network, the OpenSG/SG-Network (SG-NET) task force defined a suite of application communication requirements in [2].

The application communications requirements are defined in terms of the maximum latency and minimum reliability of data messages between a pair of designated source and sink actors. Latency is the time interval required for the successful transmission and reception of a message, and reliability is the

0890-8044/12/$25.00 © 2012 IEEE

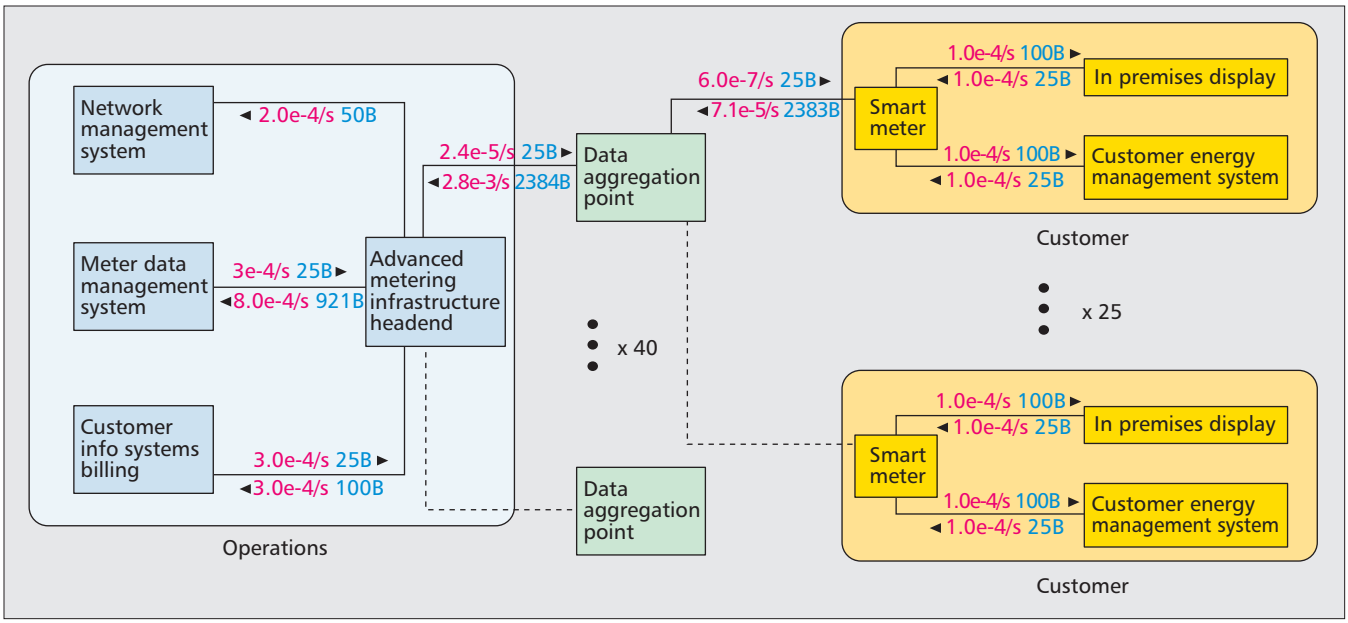

Figure 1. *Traffic characteristics for the meter reading applications defined in [2].*

probability that the latency lies within a given time interval. Most of the applications require 98 percent reliability within a delay of 5 s. Some applications, such as on-demand meter reading and switch of service, have higher priority, requiring 99.5 percent reliability within 1 s. Other less stringent requirements, such as scheduled meter reading, require 99.5 percent reliability within four hours.

The data messages are specified in terms of their arrival rate (number of transactions per day), the time interval of the day in which they may occur, and the message size (bytes).[1] At present, the requirements do not include information about the distribution of the message interarrival times; this is an important topic for study. The message load on the network, which can be expressed in bits per second, is computed by multiplying the message arrival rate by the mean message size. The arrival rate and message size are often given as a range of values. It follows that the peak (minimum) load corresponds to the maximum (minimum) arrival rate and the largest (smallest) message size. We define average load as the mean of the minimum and peak loads.

In order to translate these message specifications into communication requirements on a given link, it is necessary to create a network topological model of the physical links between the actors in the network. The SG-NET requirements provide some guidelines for the link topology of the grid, but they are loosely defined. In fact, no direct link between source and sink actors may even exist in the network, implying an interface through other actors in the network. Rather, the link topology is left flexible to the needs and preferences of the service providers. For example, consider the advanced metering infrastructure (AMI) deployment for the in-premises display and customer energy management system applications illustrated in Fig. 1. The actors in the Operations domain are connected to the actors in the Customer domain via the AMI head-end and the SMs, both serving as gateways to other domains. A utility has one or more AMI head-ends which collectively serve from several thousand to several million customers. Customer traffic to the AMI head-end is often collected by Data Aggregation Points (DAPs) that operate closer to the network edge. According to the SG-NET requirements, there can be up to 50,000 meters per DAP in automated metering networks. The actual number will likely vary significantly in a given deployment.

Once the network topology has been determined, one can develop a model for the traffic on a given link. One approach to developing a link traffic model appropriate for analytical performance models aggregates the flows from all applications utilizing the link, using the specifications in [2] for each flow and making simplifying assumptions about the message interarrival time to compute a mean arrival rate and message size on the link [3]. Event-driven simulations of link performance, on the other hand, may model each application flow on the link individually based directly on the requirements. Detailed guidelines on how to translate smart grid application flows into link traffic characteristics were documented by the NIST/smart grid Interoperability Panel Priority Action Plan for Wireless Communications [4].

An example of link traffic characteristics for advanced metering is presented in Fig. 1. In this example, the AMI head-end serves a total of 1000 customers through 40 DAPs. The arrival rate (pink) and the mean message size (blue) for average load are shown on each link. To generate the arrival rates, we considered only residential electrical meters, and we assumed that the rate of meter read request failures (resulting in the transmission of an error message to the request originator) is 1/1000. We used the maximum arrival rates for all applications; for example, on-demand meter read requests that originate at the IPD occur 1–10 times per day at each residence; we used 10 events per day. Since we considered daytime operations, we ignored bulk read events, which take place between 6:00 p.m. and 6:00 a.m.

One observes that the network is very lightly loaded: the heaviest loads occur within Operations, with a maximum of 104 kb/s from the meter data management system (MDMS) to CIS/Billing. Most links upstream of the AMI head-end in the Operations domain will likely be wired, as high bandwidth dedicated lines will be available to its limited number of actors. Downstream of the AMI head-end, the connections to the Customer domain may be overwired, wireless, or power-line communications links. The average load and in particular the peak load, which can be three orders of magnitude larger,

[1] *The specification of the message size is at the application layer. Additional protocol and security overhead from the lower layers of the communications stack will increase the actual length of the messages.*

Environment	Rural	Suburban	Urban
n_0	2.1	2.7	3.6
n_1	7.5	N/A	N/A
d_1 (m)	650	N/A	N/A
PL_0 (dB)	38.3	40	21.3
σ (dB)	2.2	7.4	7.4
R (m)	738	159	113
ρ (SM per km^2)	10	800	2000
SM in coverage area πR_2	17	64	80

Table 1. *Characteristics of the rural, suburban, and urban environments.*

across any link must be considered when determining whether wireless or powerline can meet the communications requirements. This is precisely the analysis we conduct in the following section for wireless.

Wireless Deployment Considerations

The unique characteristics of the expected smart grid communications network will make certain aspects of a wireless deployment easier than in other contexts. The previous section highlighted the low expected traffic loads associated with meter reading applications, particularly at the edge of the network on the link between the meter and the DAP. Furthermore, with the exception of home area networks, communication actors will be owned and operated by the utility. Therefore, the actors in the network can be managed and controlled (e.g., through scheduling) to efficiently utilize network resources during peak conditions such that they can operate in a non-capacity-limited mode, regardless of the medium access control (MAC) layer that is being used. The physical (PHY) layer will have a greater effect on the system performance, and the choice of the factors such as carrier frequency, transmit power, and modulation and coding scheme will be significant.

Another differentiating characteristic is that most communications devices will be fixed, not mobile, which will reduce the overhead needed for channel estimation, power control, and link adaptation.[2] Cellular deployments will require less support for handover, so links can operate with lower margins, and frequency allocation can be planned more efficiently. Mesh deployments will require relatively infrequent route updates, reducing control message overhead. Moreover, the use of directional antennas or beamforming to mitigate interference will be more feasible because the links are fixed.

Despite these mitigating factors, the key challenges to achieving the goals of the smart grid using wireless communications will be meeting the stringent coverage and reliability requirements in a variety of deployment environments and interference conditions. As the network evolves and new applications are introduced, network capacity will need to be expanded. All the while, security and confidentiality of user information must be maintained. This section details some of these considerations with a view toward how they can be modeled and using quantitative examples where appropriate. While these examples are based on specific environments and specific wireless technologies, the trends observed and the conclusions drawn generally apply to many other wireless technologies and deployment scenarios.

Radio Frequency Environment

The nature of the terrain and ground clutter affects the propagation of a signal on a wireless link. In a flat open area, for example, a signal will typically be attenuated less and hence propagate over longer distances than in an area with many buildings, mountains, and foliage. The attenuation is represented through a channel model for path loss as a function of link distance *dd*, such as

$$PL(d)(dB) = PL_0 + \begin{cases} 10n_0 \log_{10}(d/d_0), & d_0 \le d \le d_1 \\ 10n_0 \log_{10}(d_1/d_0) + 10n_1 \log_{10}(d/d_1) & d > d_1 \end{cases}, \quad (1)$$

where PL_0 is some value at reference distance d_0. The two exponents n_0 and n_1 allow for different degrees of attenuation after the breakpoint distance d_1, typically between line-of-sight (LOS) and non-LOS conditions. Other models such as the well-known Hata-Okumura model, map specific deployment characteristics, such as carrier frequency, antenna heights, and deployment environment, to specific values for the reference path loss and exponents. The variability in the attenuation from the predicted path loss — a phenomenon known as shadowing — is modeled as a lognormal random variable with standard deviations.

We examined the DAP-to-SM link (Fig. 1) in a neighborhood area network (NAN) in three representative environments that varied in terms of signal reach and variability. Table 1 shows the parameter values for these environments [5, 6] corresponding to the channel model in Eq. 1. The table also lists representative SM densities for the environments and the corresponding number of SMs within each coverage area. Assuming an IEEE 802.11 model [7] for the link, the coverage ranges were computed from the model specifications (effective isotropic radiated power of 25 dBm, required signal-to-interference-plus-noise ratio of 10.4 dB) and from the channel parameters in each environment for an outage probability of 0.1 percent. The coverage ranges R, R, are shown in the table. The rural environment is relatively benign, featuring shallow path loss and low-variance shadowing, translating into the largest coverage range. The suburban and urban environments feature steeper path loss and more severe shadowing.

Using the model in [7], we show in Fig. 2 MAC layer reliability and latency performance curves (plotted in blue and red, respectively) versus the offered load at the SM for the urban and rural environments. We assumed Poisson arrivals for the traffic which yields conservative bounds on network performance. The dotted blue horizontal line shows the lowest acceptable reliability (98 percent) and the dashed pink line shows the maximum latency (5 s) from the OpenSG requirements. Average and peak loads associated with the metering use case are indicated by the solid vertical lines. In both environments, reliability and mean latency are strongly correlated. The rural environment can support higher loads; in the urban environment the greater meter density results in a larger user population even though the DAP range is smaller. Because the average offered load is small, we have good performance

[2] *Link quality will vary with precipitation and the movement of objects in the paths of the signals, but these variations will occur on a much longer timescale and with lower frequency than in a mobile environment. Although the signal component is likely to be relatively stable, RF interference can still vary widely.*

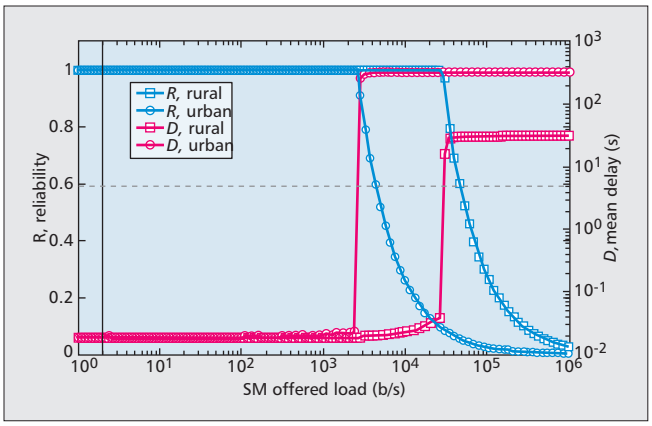

Figure 2. *Smart meter performance vs. offered load for a WLAN deployment.*

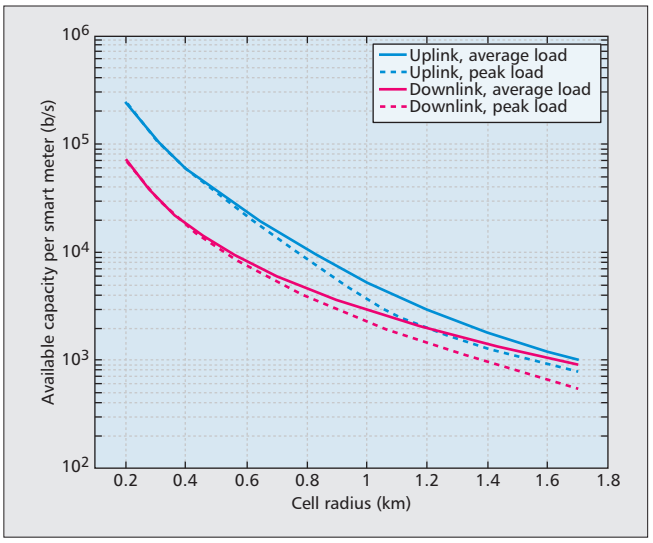

Figure 3. *Smart meter capacity vs. cell radius for a wide-area cellular network urban deployment.*

in both environments; at peak load, meters in the more benign rural environment (with less severe path loss and shadowing, and lower user density) perform well, while performance in the urban environment at peak load is unacceptable. This result suggests that additional resources (access points, in this case), beyond those needed simply for coverage, would be required to handle peak load in the urban environment.

The Coverage-Capacity Trade-off

The previous example highlighted the fact that a network can be coverage-limited in one situation (e.g., in a rural environment) and capacity-limited in another (e.g., in an urban environment under peak load). In most wireless networks, there is a tradeoff between coverage and capacity: the network can be optimized to maximize coverage (the geographic area served by a base station or access point) at the expense of capacity, or to maximize capacity (the number of devices or the aggregate load that can be served by a base station or access point) at the expense of coverage. In the latter case, since the coverage of each infrastructure element is reduced, additional infrastructure may be needed to serve the territory.

We illustrate an example of the trade-off between coverage and capacity through an analysis of a wide area cellular network connecting an AMI head-end with DAPs. Figure 3 shows a plot of the information capacity per smart meter on the WAN link versus the cell radius in an urban environment. The technology assumed in this analysis is 3rd Generation Partnership Project (3GPP) Long Term Evolution (LTE) operating at 700 MHz [9]. The figure shows capacity on both the uplink and downlink,[3] and for both average and peak load conditions. The greater level of intercell interference under peak load reduces the capacity somewhat, more so at larger cell sizes for which a given load utilizes more resources. In general, however, we see that capacity decreases as cell size increases. One can use this type of analysis to predict the infrastructure needed to support current and future loads. Initial deployments are likely to be coverage-limited and operate at the rightmost end of these curves (a few large cells); as load increases, the network operating point may move to the left (many smaller cells).

Extending Coverage though Multi-Hop Routing

When wireless deployments are coverage-limited rather than capacity-limited, multi-hop routing offers a means to fully utilize network resources by extending coverage. In the case of the DAP-SM link, the DAP can reach meters beyond its range by routing through other meters within its range acting as relays. This would be practical for an SM deployment with a high degree of connectivity. Reliability can also be improved through route diversity, in the case where a smart meter malfunctions, and transmitting shorter distances using lower power mitigates interference. Even though multihop routing has greater overhead due to route discovery than single hops alone, it can still be managed on the links that are lightly loaded, such as the DAP-SM link. However a consequence of multi-hop routing is the buildup of forwarded traffic at the relay, especially at relays with a direct link to the DAP. As the number of hops increases, so does their forwarded traffic, causing them to reach saturation before all others. Eventually they reach full capacity, at which point network resources are fully utilized.

Figure 4 shows the reliability of the multi-hop links between the DAP and the SMs as a function of the deployment range of the meters, defined as the maximum reach of the DAP with multihop. The results were generated from the IEEE 802.11 mesh model in [8]. In this example, the SMs operate at peak load. The coverage ranges in the three environments from Table 1 are marked by the vertical dashed-dotted lines. The horizontal dashed line indicates the reliability threshold value of 0.98 for acceptable performance. The networks in the urban and suburban environments are shown to reach saturation before the deployment range exceeds the respective coverage ranges (i.e., for single-hop communication), as indicated by the drop in reliability. This means that the networks are capacity-limited, not coverage-limited, so extending the deployment range would not prove beneficial. The drop is particularly sharp in the urban environment, which has the highest density of meters; here, incrementing the deployment range increases the forwarded traffic on the relays the most. When the deployment range reaches 2150 m, the throughput is virtually zero for both environments. Note the discontinuities on the abscissa in the plot.

In the rural environment, the network reaches saturation at a deployment range beyond 2150 m, at nearly triple the coverage range. The corresponding number of SMs that the DAP can support is 145 — much higher than the 17 in the coverage area alone. Note the gradual drop-off due to the low density of meters. So even at peak offered load, the rural environment is a suitable candidate for deployment extension, especially because wireless infrastructure in rural areas tends to be

[3] *Downlink SM capacity is lower due to asymmetric message sizes (Fig. 1). The assumed 42-byte overhead per SM message consumes a much larger percentage of downlink capacity than uplink capacity.*

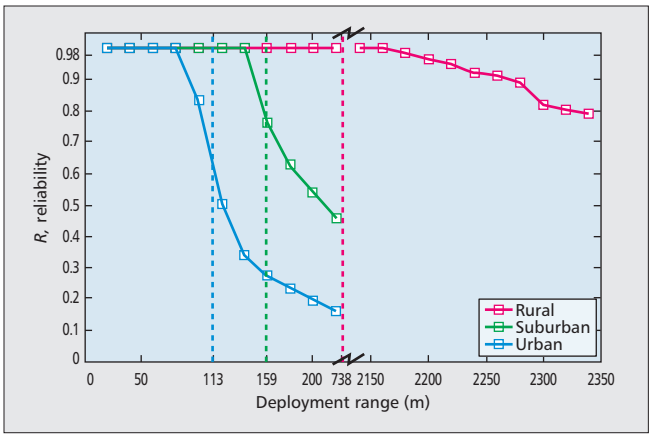

Figure 4. *Reliability of multihop links versus deployment range for the three environments.*

scarce compared to urban and suburban.

Although not shown, the MAC delay has the same saturation points as the reliability metric, where the delay rises sharply from zero, similar to the behavior in Fig. 2.

Interference

The particular sources of interference in an SG network depend on the frequency band that the utility ultimately chooses. Utilities' access to licensed spectrum has been limited as they have been required to share bandwidth with other users or to migrate out of bands where they previously operated, such as the 2 GHz band [10]. The 900 MHz and 2.4 GHz ISM bands are attractive because they are unlicensed, but they are heavily used by consumer devices. For example, a home area network (HAN) using the unlicensed 2.4 GHz band may interfere with the household's existing wireless LAN. The degree of interference will depend on many factors, such as the placement of LAN devices and HAN devices relative to each other, and the offered load associated with each network. If, as currently indicated, the amount of traffic generated by HAN devices is very low, interference may be less of an issue. At a lower level, immunity to electromagnetic interference needs to be designed into any smart grid device with an RF transceiver. For example, the microprocessor responsible for measuring electrical usage in a wireless-enabled smart meter must be adequately shielded from the meter's RF transmitter.

To illustrate the impact of interference on performance, the following example considers the effect of a constant level of ambient interference on the link. The source of this interference might be users from other cells using the same channel or other wireless users in the case of an unlicensed band. In Fig. 5, we plot reliability and mean MAC layer delay versus the ratio of the interference power spectral density, I_0, to the noise power spectral density, N_0. Recall that the minimum reliability and maximum delay are 0.98 and 5 s, respectively, as shown by the dashed horizontal lines in the figures. We considered four load scenarios: average (2 b/s) and peak (4 kb/s) loads, and higher loads of 10 kb/s and 100 kb/s. We used the rural environment parameters from Table 1. The figures clearly indicate the critical value of I_0/N_0, beyond which performance deteriorates. In this example, the largest acceptable value for x is about 30 dB in terms of the reliability from Fig. 5a, although we can see from Fig. 5b that the delay performance begins to degrade at the much lower I_0/N_0 value of around 15 dB. The performance degradation affects the system regardless of the load; adding resources will not solve the problem; the utility will have to use interference mitigation techniques such as directional antennas, higher transmit power, or more robust modulation and coding schemes to improve the link budget. We also note that at saturation loads, we have poor performance even at low interference-to-noise ratio values. The solution in this case is to deploy additional resources to reduce congestion.

End-to-End Performance

It is important to remember that analyzing link performance is only one part of the assessment process; it is equally important to consider the end-to-end application performance for which the requirements are defined. The link performance assessment provides us with useful bounds based on the application requirements. For example, if we have an application with a maximum acceptable end-to-end delay of 5 s, which is indicated by the horizontal dashed line in Fig. 2b, we know that loads which produce delays above this limit are clearly unacceptable because the end-to-end delay will be too great, even if the node and link delays along the rest of the path are essentially zero.

If we can treat the various network elements on a communications path as independent, we can estimate the end-to-end delay and reliability as the sum of the individual link delays and the product of the individual link reliabilities, respectively. If the components are not statistically independent, we must examine the behavior of the entire path in aggregate. In this case, one can sometimes obtain performance bounds by focusing on a bottleneck link or network element (e.g. a router) whose reliability or delay dominates the path. For example, a path that consists of a series of high-speed links and one slow link will be limited to the best performance of the slow link. This approach can facilitate rapid analysis of application performance in more complex networks, allowing more targeted simulation and testing.

Security

Wireless networks are more vulnerable than their wired counterparts due to the potential for direct access to the transport medium. Strategically placed rogue servers with seemingly higher data speeds can lure users away from intended servers; sniffers can determine the MAC address of an intended server and use it to mask a rogue server. This lets adversaries steal sensitive information such as usernames and passwords or inject malicious software into the network. If automated accounting and billing systems are built into the smart grid, any compromised interface could translate into the unauthorized disclosure of financial information on a massive scale.

Hence, security must be considered at every layer of the protocol stack. The SG-NET requirements, which are defined at the application layer, require additional protocol and traffic events to support security signaling, as in the case of authentication and authorization, and adding bits to existing payloads to achieve encryption. As a first step toward this goal, the SG-NET requirements list the security objectives of confidentiality, integrity, and availability (CIA) for each event. As a second step, a mapping between CIA levels (low/moderate/high) and the security protocols available at the various layers will be needed to adequately address security.

Summary

In this article, we describe the smart grid communication requirements that have been established to date by the OpenSG/SG-Network task force, namely the maximum reliability and minimum latency that application data messages between designated grid entities can tolerate, as well as arrival rates and sizes of those messages. In order to design a network model from the requirements, we point out assumptions

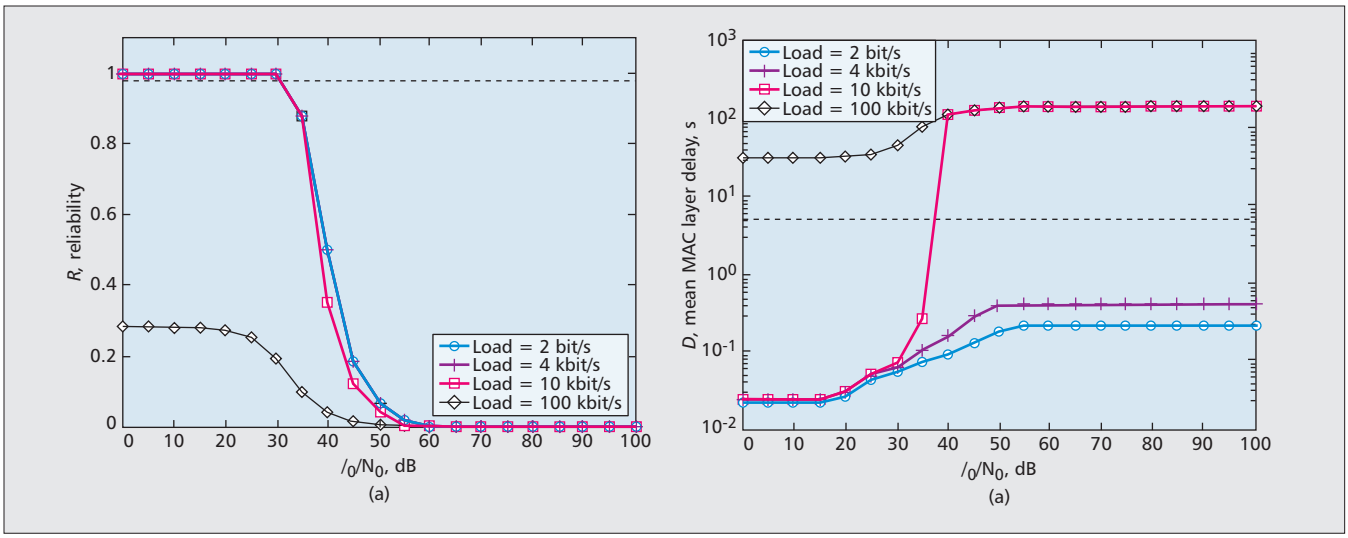

Figure 5. *Performance vs. interference-to-noise ratio for various offered loads: a) Reliability and b) delay.*

that must be made about the network topology and traffic model, since this information is not provided in the requirements. Using a network model, we highlight the implications of wireless deployment as they relate specifically to the smart grid through numerical examples. One observation is that most of the entities in the grid will be fixed, greatly simplifying radio resource management due to quasi-static channels and enabling enhanced interference mitigation techniques. Also, since the anticipated loads to the smart meters will be light (average load 2 b/s), the network will operate in a non-capacity-limited mode, so considering the MAC layer of a specific wireless technology will be less important than comparing the characteristics of the radio frequency environment. Finally, we explain why smart meters are a strong candidate for multihop routing to extend coverage and reliability through route diversity.

References

[1] Office of the National Coordinator for smart grid Interoperability, "The NIST Framework and Roadmap for smart grid Interoperability Standards," Release 1.0, NIST Special Publication 1108.

[2] "SG Network System Requirements Specification," ver. 4.0, OpenSG Users Group, July 2010, available: http://osgug.ucaiug.org/UtiliComm/Shared%20Documents/Latest_Release_Deliverables/SG%20Network%20System%20Requirements%20Specification%20v4.0.xls.

[3] M. Souryal *et al.*, "A Methodology to Evaluate Wireless Technologies for the smart grid," *Proc. IEEE Conf. Smart Grid Commun.*, Oct. 2010, pp. 356–61.

[4] "Guidelines for Assessing Wireless Standards for Smart Grid Applications," ver. 1.0, NIST IR 7761, to be published, draft available: http://collaborate.nist.gov/twiki-sggrid/bin/view/SmartGrid/PAP02Objective3.

[5] D. Laselva *et al.*, "Empirical Models and Parameters for Rural and Indoor Wideband Radio Channels at 2.45 and 5.25 GHz," *IEEE Int'l. Symp. Personal, Indoor, and Mobile Radio Commun.*, Sept. 2010, pp. 654–58.

[6] H. H. Xia *et al.*, "Microcellular Propagation Characteristics for Personal Communications in Urban and Suburban Environments," *IEEE Trans. Vehic. Tech.*, vol. 43, no. 3, Aug. 1994, pp. 743–52.

[7] D. Griffith *et al.*, "An Integrated PHY and MAC Layer Model for Half-duplex IEEE 802.11 Networks," *Proc. IEEE MILCOM*, Nov. 2010, pp. 1354–659.

[8] C. Gentile *et al.*, "Throughput and Delay Analysis of Half-Duplex IEEE 802.11 Mesh Networks," *Proc. IEEE ICC*, June 2011.

[9] M. R. Souryal and N. Golmie, "Analysis of Advanced Metering Over A Wide Area Cellular Network," submitted to the *IEEE Conf. smart grid Commun.*, Oct. 2011.

[10] B. Kilbourne and K. Bender, "Spectrum for smart grid: Policy Recommendations Enabling Current and Future Applications," *Proc. IEEE Conf. smart grid Commun.*, Oct. 2010, pp. 578–82.

Biographies

CAMILLO GENTILE (camillo.gentile@nist.gov) received B.S. and M.S. degrees from Drexel University, Philadelphia, Pennsylvania and Ph.D. degree from the Pennsylvania State University, University Park, all in electrical engineering. He has been a researcher in the Advanced Network Technologies Division NIST, Gaithersburg, Maryland since 2001. His current interests include RF channel modeling, smart grid, LTE, and millimeter-wave telecommunications.

DAVID GRIFFITH received his Ph.D. in electrical engineering from the University of Delaware. He worked on satellite communications systems at Stanford Telecommunications and Raytheon, and is currently with the Information Technology Laboratory at NIST. His research interests include mathematical modeling and simulation of wireless communications networks, including public safety broadband networks and smart grid.

MICHAEL SOURYAL is with NIST conducting research in wireless communication systems. He received his D.Sc. in electrical engineering from the George Washington University, M.S. in information networking from Carnegie Mellon University, and B.S. in electrical engineering from Cornell University. He was awarded an NRC Postdoctoral Fellowship at NIST in 2004. He holds an adjunct appointment as Professorial Lecturer at the George Washington University.