15th Annual Conference on Systems Engineering Research
Disciplinary Convergence: Implications for
Systems Engineering Research

Eds.: Azad M. Madni, Barry Boehm
Daniel A. Erwin, Roger Ghanem; University of Southern California
Marilee J. Wheaton, The Aerospace Corporation
Redondo Beach, CA, March 23-25, 2017

Using systems engineering to create a survivable communications system that will operate in the presence of "Black Sky" hazards

Neil Siegel, Ph.D.[a]

[a]University of Southern California, siegel.neil@gmail.com

Abstract

Many studies have shown (and incidents like super-storm Sandy have demonstrated) that existing U.S. electronic communications systems (e.g., land-line phones, cell phones, the internet, etc.) would not be available in the event of a large-scale, long-term power outage, whether that outage were to be triggered by terrorism, by war, or even by certain types of natural events.

However, several of the critical tasks involved in getting the power back on after such an event – such as synchronizing the re-start of large-scale generators and users of electricity, allocating and planning the movement of spare parts and personnel with critical skills, and so forth – cannot be accomplished without some level of real-time, wide-area electronic communications.

Because of this dependency on electronic communications, and the fact that all such electronic communications systems would fail just a few hours into such an outage, in a scenario created by a group of industry experts, they found that – in the absence of some survivable wide-area emergency communications system – the duration of an outage would be measured in months or years. Casualties would be very high. When the same experts modelled this scenario while positing the existence of a survivable, wide-area emergency communications system, the duration of the outage was measured in weeks, and casualties were orders of magnitude less. It appears that the existence of such a survivable, wide-area emergency communications system is an important aspect of emergency preparations.

This paper uses systems-engineering methods to examine the question of how best to provide reliable electronic communications under such conditions at the required scale, data rates, and reliability.

Keywords: emergency response, electric outages, recovery from large-scale outages, recovery from large-scale disasters, networked low-frequency radios, storable emergency power sources

1. Statement of the problem

Many studies have shown (and incidents like super-storm Sandy have demonstrated) that existing U.S. electronic communications systems (e.g., land-line phones, cell phones, the internet, communications satellites, etc.) would not be available in the event of a large-scale, long-term power outage, whether that outage were to be triggered by terrorism, by war, or even by certain types of natural events[i].

However, several of the critical tasks involved in getting the power back on after such an event – such

as synchronizing the re-start of large-scale generators and users of electricity, allocating and planning the movement of spare parts and personnel with critical skills, and so forth – cannot be accomplished without some level of real-time, wide-area electronic communications.

Because of this dependency on electronic communications, and the fact that all such electronic communications systems would fail just a few hours into such an outage, in a scenario created by a group of industry experts [in which the power outage was triggered by an electromagnetic-pulse attack event], they found that – in the absence of some survivable wide-area emergency communications system – the duration of an outage could be measured in months or years – and restoration could come in increments, rather than all at once. Casualties would be very high, comparable to large-scale nuclear war (these casualties would result primarily from sewage-borne disease). When the same experts modelled this scenario while positing the existence of a survivable, wide-area emergency communications system, the duration of the outage was measured in weeks, and casualties were orders of magnitude less. It appears that the existence of such a survivable, wide-area emergency communications system is an important aspect of emergency preparations.

The author has been leading a team that is studying this problem[ii]. Drawing upon the emerging results of this study, this paper uses systems-engineering methods to examine the question of how best to provide reliable electronic communications under such conditions at the required scale, data rates, and reliability[iii].

To state the problem another way: If a power outage is "big" enough (for example, all of western Europe, or the entire eastern half of the United States – and now, there are ways to cause such a large-scale outage), there is at present essentially no way to recover from this outage before casualties become very high. No organization or country has ever practiced "turning off and re-starting" a large portion of their power grid, so there is no actual experience to serve as a credible guide to the recovery effort from a power outage of this scale.

The apparent root cause for the long endurance and difficulty in recovery from such a large-scale power outage is that if a power outage is "big" enough, because of the number of generation and transmission / distribution sites involved, crews cannot get to all sites quickly, and therefore the outage will last more than a few hours. But after some period of time – certainly less than 24 hours – the back-up power systems (mostly, on-site batteries) for our electronic communications systems (e.g., the internet, the land-line phone system, the cell-phone system, the ground segments of satellite communications systems, and so forth) run out of power, and those systems then shut-down[iv]. Experience from periodic adverse weather events seem to indicate that (perhaps due to poor maintenance of these batteries), these back-up power systems in fact run out of power in less than 4 hours.

Once these batteries run down, the electronic communications systems they support stop operating (actual experience from weather-induced emergencies actually indicates that some of these systems fail ***before*** their battery back-up systems run out of power, due to their being unable to cope with the large spike of usage that occurs as people try to "phone home" after the onset of the emergency). As noted above, there are many key steps in restoring electric service after a large-scale outage that are fundamentally different than restoring power after a local outage; as indicated below, these steps require real-time communications across distance, and therefore, need access to some sort of electronic communications system. This creates a "threshold" effect; an outage that is large enough creates conditions that make recovery fundamentally different than the city-wide or state-wide / region-wide outages that we have occasionally experienced, and the recovery mechanisms that are in place to recover from those city / state / regional types of outages will not suffice to recover from these larger outages.

To understand this aspect in more detail, recall that in general, we do not store electricity; it is used immediately upon generation. This requires that generation capacity and offered electric load be *kept in synchronicity*. Of course, electricity-consuming devices are constantly being turned on and off, but their effect on the overall power consumption within a segment of the grid is fairly small (the total use of electricity in the U.S. is around 5 quadrillion – 5,000,000,000,000,000 – watts per year[v]), and because the fluctuations

caused by individual actions are relatively small, during steady-state power-grid operations electric use can be monitored passively at the generators, by measuring fluctuations in voltage or frequency. Generation is then matched to the real-time demand by adjusting the speed (and hence output) of generators so as to keep voltage and frequency within very small bounds of their target values[vi].

Such passive monitoring, however, will *not* work for re-start operations over a large area; instead, the generation station must *coordinate* in real-time with the large-scale user of electricity (for example, a water-pumping station) that is ready to come back on-line, so that the amount of electricity generation about to be re-started *exactly* matches the power consumption of the system about to be brought back on line. Failure to do this correctly will result in significant fluctuations in voltage and/or frequency, which can damage equipment. Then we are worse off than before; not only are things off, but they are damaged. Once the small number of very large-scale users of electricity (such as water pumping, sewage treatment, natural gas pumping – which together, account for about 20% of all electricity use in the United States) are back on-line, bringing on small groups of residential and business users will again individually result in small (and hence, manageable) fluctuations, which can be dealt with by the normal passive-monitoring method. But one needs some significant steady-state, pre-coordinated load to create a "denominator" of predictable electric load before one can start to introduce the random electric load of residences and businesses; without a big "denominator" (what matters is the *percentage* of fluctuation), the fluctuations will be too large and too fast for the passive method to control.

So, restarting the power grid after a large-scale failure requires real-time communications . . . but all of the electronic communications systems will be down. Analysis appears to indicate that it is not economically feasible to increase the battery back-up capacity at every point along these regular communications systems; analysis shows that we need to provide 30-60 days (not a few hours!) of back-up power to deal with the sort of large-scale power outage considered herein. Doing that at all of the millions of locations where there is electrically-powered equipment for the internet, the phone system, the cell-phone system, etc. would cost between 20 to 50 times what it would cost to provide a separate emergency communication system, which we estimate would require only about 100,000 points of presence across the U.S. (in contrast to the several hundred million points of presence of the regular electronic communications systems).

So it appears that we need some sort of emergency communications system that can operator on its own stored power for 30 to 60 days. In the rest of this paper, I use systems engineering methods to determine the necessary communications links and capacity, the necessary participants, the optimal configurations, and the appropriate re-start sequencing / procedures for such an emergency communications system.

2. Constraints on potential solutions, and the resulting systems-engineering trade-space for candidate solutions

Not included in this paper, but included in a longer report[vii], the study team has developed what we call a "social architecture" for such an emergency communications system. This social architecture allows us to identify the users and customers for such an emergency communications system, determine how they define value within their operational context, and in general, to capture the information necessary to create a system that is both (in the terminology used in the U.S. Federal Acquisition Regulations, or FAR[viii]) "effective and suitable" for their mission.

The following are a small sub-set of the goals and considerations from that social architecture:

- A trained operator should be able to set up the emergency communications suite at a location (e.g., to effect the transition from the long-term storage configuration into the operating configuration) in 4 hours or less. This limitation of setup time will drive decisions about how the emergency communications equipment is configured for long-term storage.
- The emergency communications system should be designed so that – unless there has been physical damage sustained at a particular node – 99.9% of the locations should be able to operate (in at least a partial capacity) after being brought out of long-term storage. Achieving this level of reliability and availability will require some on-site spare parts at every site.

- The emergency communications system should be designed so that if at least 90% of the nodes in a region are brought into operation, AND also at least 75% of the nodes in adjacent regions are brought into operation, then at least 75% of the emergency communications system nodes within that region should *automatically* discover a route to their regional reliability coordinator (e.g., the local emergency communications system hub). Additional nodes can achieve connectivity to the local emergency communications system hub through manual operator actions. Achieving this high level of reliability and connectivity will require the use of multiple data paths (the technical term is "communications path diversity") for most of the point-to-point linkages needs by the emergency personnel. This in turn implies a need for multiple communications devices of different types at most emergency communications system sites.
- "Connectivity" in the paragraphs above means "achieve at least push-to-talk voice and a modest data" capability.
- Push-to-talk voice communication serves a high portion of the emergency communications system use-cases. The next most important capability is the ability to send a digital photograph from one location to the hub. More generalized data service is a still-lower priority. These priorities will be used to establish dynamic priority-of-service within the emergency communications system.
- When we reach full deployment, there will be about 100 thousand emergency communications system installations across the United States (a full European deployment would be somewhat larger), so cost per site becomes a design consideration. An initial look suggests that the 30-day to 60-day power requirement could become the driving per-site cost element, so as the technical design evolves, taking design actions that reduce the power requirement, and creating viable strategies for lowering per-site power-unit cost (but without decreasing operational availability!) will be an area of concentration.
- Although extensive pre-emergency planning (down to the level of checklists for individuals), will be required, no plan will survive the first hour of an actual event without requiring modification, perhaps extensive modification. Collecting status from the field is a key emergency communications system role, as this is the information required for the regional reliability coordinator (and other responsible parties at higher echelons) to adapt the plan to the actual situation on the ground, and to communicate the altered and adapted plan to those in the field who will actually execute it.
- Even the relatively small number of "tier-1" restoration priorities will involve a complex web of sites, equipment, relationships, and supply-chains. For example, it is probably not sufficient to get power to the stations that pump natural gas in the big main pipelines that service electric generation stations; in addition, we will probably need to include the capability to provide power (perhaps by emergency batteries or generators) to whatever control station remotely plans and operates the valves and pumps on those big natural-gas pipelines. Each item included in tier-1 (and the other tiers, too, of course) will have a similar web of such relations and a tiered supply chain that need to be thought through: people, power, spare parts, procedures, and so forth.
- At present, neither the United States nor any European country has such an emergency communications system; therefore, building consensus that one is required is an essential step. We believe that voluntary actions are vital to getting the process started of building such consensus underway, but at some point, we believe that we will need federal statutes &/or regulations to enforce compliance, and help convince utility-oversight boards that it is legitimate to recover emergency communications system costs through utility rates, and so forth. We point out that seat belts and car pollution control were not widely adopted until government starting creating mandates (in those cases, first the state of California, and then the federal government). The way the cyber protection story is playing out reinforces our belief that mandates will eventually be required in that field, too, but that voluntary efforts will be needed in order to create the environment necessary to get those mandates enacted, and to figure out what the correct form and content of those mandates ought to be. We believe that the process of building a societal consensus to build such an emergency communications system will have to go through a similar process.

We next mapped the goals, requirements, and considerations developed in the emergency

communications system social architecture into candidate solutions. This process was undertaken herein through several steps:

- Identify candidate communications technologies
- Create a list of key issues / risk areas
- Using that list of key issues / risk areas, identify a set of key technical trade studies
- Create a set of candidate designs, together with methods and metrics for selecting among those candidates designs. Make the initial design selection, provide the rationale, and make a preliminary assessment of the feasibility and performance of the selected design.

These are discussed in the following sections.

2.1 Identify candidate communications technologies

Based on our experience implementing many sorts of high-reliability communications systems[ix], we selected the following technical methods of communications as the candidates for consideration within this emergency communications system:

- HF radio, supplemented with some sort of networking
- VHF radio, supplemented with some sort of networking
- Higher-frequency radio, supplemented with some sort of networking
- Meteor-burst radios, supplemented with some sort of networking
- Use of power lines to carry communications signals
- So-called “dark fiber”, that is, fiber optic communications cables that are installed but not in active service
- Commercial satellites (low-Earth orbit)
- Commercial satellites (geo-stationary orbit)
- Various combinations of the above

2.2 Create a list of key issues / risk areas

Given the results of the social architecture, and the above list of candidate technologies, we identified the following as key risks that need to be addressed through the technical trade study process for the emergency communications system:

- How to provide power for the emergency communications system at each location for the specified 30-day to 60-day period?
- What spectrum (RF frequencies) would be available for the emergency communications system during emergency operations? Not all of this spectrum allocation need be available for use during nominal (non-emergency) use.
- What techniques and materials would allow the emergency communications system equipment to be stored for long periods of time (years or decades), yet still be periodically tested, maintained, and support periodic training?
- How to allow the emergency communications system to adapt, ideally almost automatically, to the likely differences between the anticipated emergency conditions and those that actually come to pass?

Many other risks were identified, but there were the ones that through our assessment process represented a combination of likelihood and potential impact that stood out as potential key system disablers if not properly addressed and mitigated.

2.3 Using that list of key issues / risk areas, identify a set of key technical trade studies

Given the above list of key risk areas, we then identified the following as the key trade studies that we needed to perform:

- How to provide 30-day self-contained power, with a long storage life, reasonable maintenance requirements, and high availability at need
- Spectrum availability during emergencies, and based on that, selection of the actual communications mechanisms, and eventually, the actual devices
- Approaches and materials to achieve effective long-term storage of the emergency communications system equipment
- Techniques to support self-adaptation of the emergency communications system

The principle results of these trade studies led to a candidate solution, presented below. Reasonable technical approaches were found to exist to mitigate all risk areas. We were not provided with an explicit "design-to" per-site cost, but we believe that most of the study outcomes will allow options to be selected that are deemed reasonable in cost. One area that still needs cost-per-site optimization is the 30-day to 60-day self-power requirement; we have excellent options, but the specific mix of technologies will probably vary significantly with the specifics of the site (e.g., is the use of solar panels as augmentation at this site feasible and attractive?), and therefore it is too early to roll up this portion of the cost to a system-wide total; actual regional site-surveys will probably be required in order to arrive at credible system cost estimates.

3. Candidate solution, arising from the systems engineering studies

In this section, we describe the candidate design, and the initial work that has been undertaken in order to validate the efficacy of that design.

In our work to-date, we have gone as far as selecting many key technical parameters (such as radio frequencies and transmission polarizations), but not all technical parameters (e.g., we have not yet selected exact RF power levels, antenna sizes, and manufacturers; that latter level of detail will come later). But we have progressed far enough to demonstrate the technical feasibility of the proposed design, and to estimate some of the key system parameters, such as availability, network connectivity rates, and storage life.

The following are the components selected for use within the emergency communications system:

- The following devices are at each emergency communications system location:
    - HF NVIS radios, with small magnetic antennas
    - UHF radios
    - Packet routers
    - Software to implement and control the above functions
    - A power sub-system, based on vanadium redox flow batteries, with control circuitry to enhance reliability, supplemented at many locations by solar panels, and perhaps supplemented at a small number of sites by wind-powered generators.
- The following device is located at each regional reliability center:
    - A packet router equipped with special, mission-specific software agents
- The emergency communications system can optionally include the following components:
    - Mobile emergency communications system nodes (e.g., trucks or other vehicles that are equipped with emergency communications equipment, a packet router, and appropriate power equipment).
    - Mobile emergency power systems. These are trucks that have battery-power sub-systems that can be driven to power an emergency communications system location whose batteries have failed. They can also be used, of course, to power other types of equipment in case of need.

3.1 The emergency communications system design story

These devices are combined to form the emergency communications system as follows:

1. Site-to-site communications are provided by the HF radios and the UHF radios. The disadvantage of large size traditionally associated with HF radio antennas is corrected through the use of magnetic antennas. UHF provides very high-quality service, but at a shorter range than HF, while

HF radios can operate beyond line-of-sight. Hence we include ***both*** frequencies in the emergency communications system design.

2. 30 to 60 days of stand-alone power is provided at each site. In order to be affordable, the size of the stand-alone power array will vary from site to site, driven by a cautious estimate of emergency communications system power requirements for 30 to 60 days. At many sites, the size (and cost) of the battery array is decreased by the addition of solar panels. There is the potential to add wind power at a few, suitably-situated, sites.
3. Vehicle-mounted emergency communications system configurations are possible. The UHF component could be configured to operate on-the-move; the HF NVIS component will likely be configured only to operate at-the-pause (e.g., when the vehicle is parked); this limitation is due to the need for the UHF antenna to be 15 to 20 feet above ground level, implying that the extendable mast must be erected for HF NVIS coverage to work. Power for the emergency communications system equipment on these vehicles will be provided by a combination of batteries and enhanced vehicle alternators.
4. Truck-mounted emergency communications system portable battery configurations are possible, providing a portable power source that can be moved from site to site during an emergency. This could power emergency communications system equipment at a site where the battery has been damaged, for example, but could also be used at any site to provide temporary emergency power.
5. The single-hop, direct site-to-site communications success-rate is improved through the use of error correction coding, and other higher-level communications protocols. These are implemented in the packet router located at each emergency communications system site.
6. At each site, there are 2 independent radios, on different frequency bands, utilizing different modulations. This provides a basic type of *communications path diversity*, and thereby improves system reliability. The router at each site determines which radio to use for each transmission attempt (whether voice or data), based on its radio "visibility" to adjoining sites. No manual action is required by the emergency communications system user to select the best radio for each transmission; this is accomplished for them automatically; we do not expect the emergency personnel to be radio-frequency-propagation experts!
7. The packet router also uses the same visibility information to implement multi-hop communications for both voice and data: a communications link need not be "direct"; I can talk to you through a set of intermediate nodes, e.g., the data are in fact routed through other emergency communications system sites. The finding and utilization of such paths is automatically accomplished by the packet routers; no manual action is required by the emergency personnel to find and implement such multi-hop paths.
8. Frequency selection is based on time-of-day, atmospheric conditions, and other factors. An "intelligent director" (a packet router equipped with mission-specific software agents) controls and coordinates this process, providing direction to the packet routers, which in turn command the radios to use the appropriate frequencies and other radio settings. No manual action is required by the emergency personnel to account for day / night frequency preferences.
9. The over-all policy for frequency utilization must be coordinated with regional and national civil officials. This also is implemented in the intelligent director; no manual action is required by the emergency personnel user to comply with the ever-changing radio-frequency policy.

Figure 3-1 (below) depicts the methodology used to provide a preliminary validation of some of the key technologies selected for emergency communications system (ECOM). On the left side of Figure 3-1, you can see that we draw upon the requirements and goals for the emergency communications system (as developed and described in the social architecture section, above) as the criteria against which we measure candidate designs. You can also see that we use the specified performance for the products selected as part of the inputs to our emergency communications system performance model.

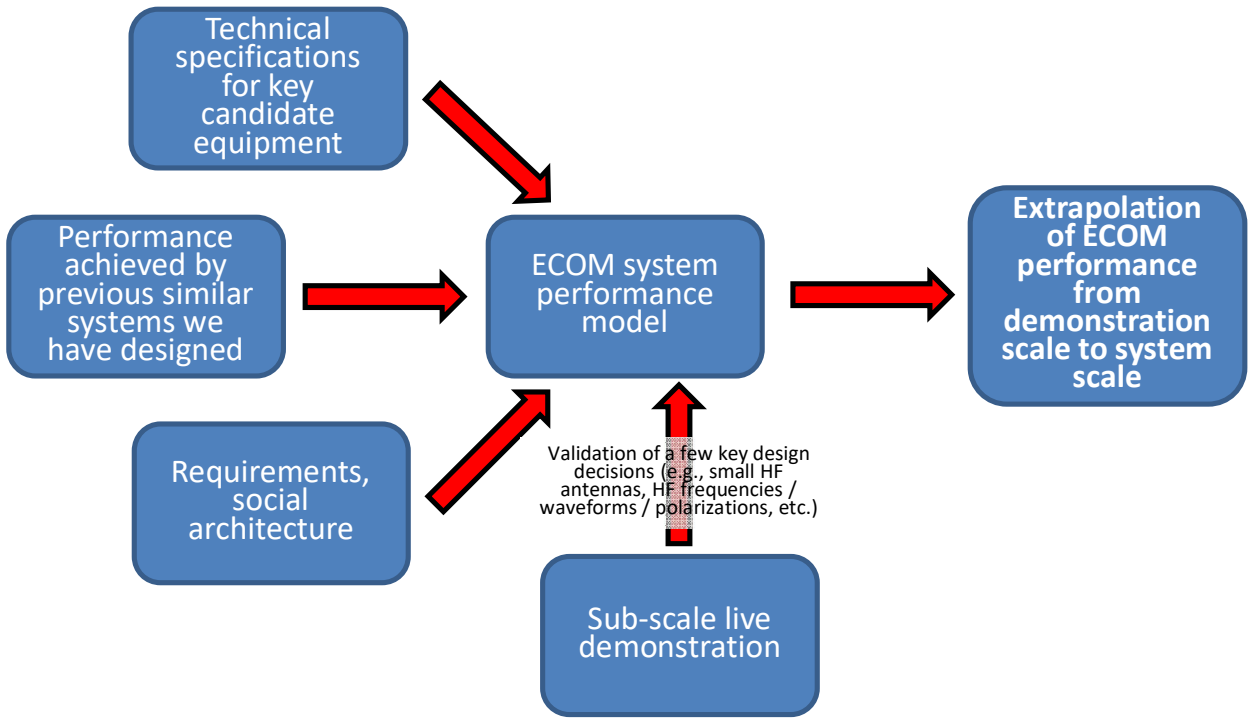

Figure 3-1. Process to achieve preliminary validation of the selected design for the emergency communications system.

A system performance model makes various assumptions about certain aspects of the system it is modeling. By drawing upon actual performance achieved by similar systems, we can create a model that is of high credibility.

The author was the program manager for the U.S. Army "Blue-Force Tracker"[x]. This system uses VHF and UHF radios that are interconnected via local routers to implement its communications network, and therefore, is in some critical ways similar to the communications systems proposed herein for ECOM; since the Army system was fielded in 1999[xi], and has been in successful and continuous operational use since that time[xii], this means that many of the key technical and architectural features proposed for ECOM have been validated through their use on that system.

The motivation for the design of this Army system was similar to the goals for this emergency communications system: very high reliability was required, and the design solution was to achieve this via *communications path diversity*: multiple radios, on different frequencies, using different waveforms, with the best path determined in real-time by the attached routers[xiii]. In one cardinal way, the Army system was much more difficult than this emergency communications system: most of its nodes have to operate while *on the move*, and therefore, had to deal with continuously-changing line-of-sight interruptions, interruptions caused by terrain masking and foliage masking, and so forth. Most of the operating locations for this emergency communications system are, in contrast, at fixed sites, and therefore this system does not face these difficulties. On the other hand, the sites for this emergency communications system are on average farther apart than the units for the Army system; this is what led to the selection of HF radios for use within ECOM, as HF can achieve longer single-hop communications distances than the VHF and UHF radios used in the Army system.

At the bottom of Figure 3-1, you see that we actually conducted a sub-scale live demonstration of certain critical aspects of our candidate design; specifically, the HF radios and their magnetic antennas. We selected these components for a sub-scale live demonstration because they are the principle component-

level difference from the Army system mentioned above. This demonstration validated that our assumptions about the performance of HF radio with the magnetic antennas were correct. We also did a live performance-measurement field test with a pair of UHF radios.

Quantitative data about HF NVIS radio performance were also collected. Given the ranges and packet-completion rates provided therein for HF NVIS, the measurement results with the UHF radios [together with the knowledge we have from the Army system about comparable radio performance data for UHF], we made estimates for the performance of ECOM at the system level, by using system-modelling methodologies validated through our work on the Army system (whose architecture is similar). These results show that the candidate emergency communications system design will perform very well, indeed.

4. Summary, portions of the problem remaining unsolved, candidate next steps

The combination of social architecture, operational mission threads, and other analyses have led to insight regarding both what will make the subject emergency communications system both "effective" and "suitable", using the terms from the Federal Acquisition Regulations. The work described herein to capture the data needed by the emergency personnel will allow the development of an effective system, whereas the social architecture (and the implications derived from it) will allow the development of a suitable system. As noted in the discussion, we believe that the technical challenge of creating a *suitable* system is in the case of this emergency communications system the *harder* part of the problem, and therefore, we have placed more emphasis on that portion. A set of specific design features – role-based processing, remote authentication, automatic management and configuration of the radio network, and so forth – have been identified as high leverage for the emergency communications system user community: both those *coordinating* the recovery operation, and those actually *implementing* the recovery operation.

Of course, the emergency communications system design, development, and deployment effort is just beginning. Consensus (and funding) must be acquired across a wide range of stakeholder communities. Prototyping – in order to create tangible artifacts showing what the emergency communications system would look like, and how it would operate – is probably essential in order to help build this consensus. The following provides a list of candidate next steps for the emergency communications system:

- Implement a pilot (e.g., smaller scale) implementation of the emergency communications system. This would create a tangible artifact that could support demonstrations. Experience suggests that this is very important in terms of building support for the implementation of the system, in addition to the technical value that would result from the learning that would come from building the pilot system. Use the technical lessons-learned from the pilot to finalize an actual scaleable detailed design.
- Build and exercise the emergency communications system system-level performance model, described in section 3.
- Continued work on the concepts for implementing the organizational relationships and coordination mechanisms required for using the emergency communications system.
- Continue work for the manning plan: e.g., training regular utility-company employees, and also creating a surge-capacity concept that can bring in needed personnel from other industries.
- Build and utilize an emergency communications system system-level cost model that allows for estimates about number of sites, hours of utilization per day, solar availability, and other factors to be adjusted, so as to create parametric estimate curves for emergency communications system cost. Estimate the cost and time of implementation. Suggest phasing and incremental capabilities.
- Investigate concepts for funding (acquisition, training, maintenance) and cost-recovery, e.g., who would pay, how might cost be shared, how could costs borne by the operators and owners of critical infrastructure recover those costs through utility rates, and so forth.

- Develop ideas for artifacts that would help other stakeholders understand the problem and the proposed solution (so as to help build a constituency that would support the implementation of such a national emergency communications system).
- Consider the creation of some use-cases that would use some of the emergency communications system equipment in non-emergency or lesser-emergency situations, in order to help reinforce the business case

Acknowledgements

This work was funded in part by the Electric Infrastructure Security Council, a non-profit research entity established by the electric utilities and their major customers, whose support is appreciated.

Appendix A. About the author

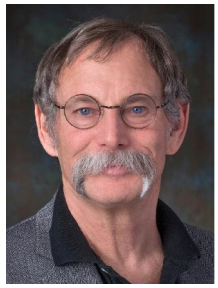

Neil Siegel is the IBM Professor of Engineering Management in the Department of Industrial and Systems Engineering at USC. Previously, he was for 15 years Sector Vice-President & Chief Technology Officer at Northrop Grumman, and before that, he held a variety of senior leadership positions within that company. He holds more than 20 patents, and his inventions are used in a billion devices worldwide. He has been elected to the U.S. National Academy of Engineering, received the Simon Ramo medal for systems engineering and systems science, is a fellow of the IEEE and an INCOSE-certified Expert Systems Engineering Practitioner, and has received a variety of other awards and honors.

References

[i] Note that the U.S National Academy of Science and the U.S. agency NASA now estimate that the probability of an adverse space weather event (e.g., solar-induced electronic storms severe enough to cause continent-scale electric power outages) in the next decade is around 12%. See http://sites.nationalacademies.org/cs/groups/ssbsite/documents/webpage/ssb_153147.pdf and https://science.nasa.gov/science-news/science-at-nasa/2014/23jul_superstorm.

[ii] Some of the preliminary results of this study are available at http://www.eiscouncil.com/Library; see the item “BSX White Paper” on that web page.

[iii] This paper is one of four that is part of a coordinated CSER conference session that is considering the topic of “Protecting Electric Power Sources”. This panel was chaired by the author of this paper.

[iv] Generally, the back-up power devices for these systems are designed to provide between 4 and 8 hours of operation.

[v] For example, see information from the U.S. Energy Information Administration (http://www.eia.gov/electricity/annual/).

[vi] In the U.S., the target values are the familiar 120 volts of alternating current, alternating at a rate of 60 cycles per second. In Europe, 220 volts and 50 cycles per second.

[vii] “Emergency Communications System (ECOM), A Technical Report for the Electric Infrastructure Security Council”, by Neil Siegel, Ph.D. and Bran Ferren, 8 August 2016.

[viii] Available on-line at https://www.acquisition.gov/?q=browsefar.

[ix] Including the emergency communications system in the City of New York, and the U.S. Army’s “tactical internet” and Blue-Force Tracker. See “The Digital Battlefield: A Behind-the-Scenes Look from a Systems Perspective”, Neil G. Siegel and Azad M. Madni, Elsevier, 2014.

[x] Formally known as “Force XXI Battle Command Brigade-and-Below”, or FBCB2.

[xi] “Force Tracking – Experience in Combat”, Neil Siegel, 2005

[xii] “The Digital Battlefield: A Behind-the-Scenes Look from a Systems Perspective”, Neil Siegel and Azad Madni, Elsevier, 2014.

[xiii] “Digitizing the Battlefield”, Neil Siegel, a chapter in the book *Fateful Lighting*, information technology association of America, 2002