Chapter 1

CRITICAL INFRASTRUCTURE ANALYSIS: A METHODOLOGY FOR INTEGRATED MODELING AND SIMULATION

William J. Tolone, Seok-Won Lee, Wei-Ning Xiang, Joshua Blackwell, Cody Yeager, Andrew Schumpert and E. Wray Johnson

Abstract Integral to effective critical infrastructure analysis is the assessment of infrastructure vulnerabilities, which seeks to provide insights into potential disruptions, insights that may increase the efficacy of protection plans as well as operations for response and recovery. Effective critical infrastructures analysis, however, must account for both the multi-dimensional, highly complex characteristics within each infrastructure as well as the high level of dependency among infrastructures. In this paper we present a new methodology for integrated modeling and simulation that supports such analysis. In addition, we present our Integrated Modeling Environment, which embodies this new methodology.

Keywords: Modeling and Simulation, Critical Infrastructure Protection, Critical Infrastructure Analysis, Geospatial Analysis, Ontological Analysis

1. Introduction

Critical infrastructures, by definition, are those infrastructures that, if disrupted, can undermine a nation's security, economy, public health, and/or way of life [10]. Numerous recent incidents, e.g., the blackout in the northeast United States and southeast Canada in 2003 and the hurricane damage in Louisiana and Texas in 2005, demonstrate the potentially catastrophic impacts of critical infrastructure disruptions. While it is unlikely that disruptions can be prevented, an effective practice of critical infrastructure analysis can reduce their frequency or at least min-

To Appear in Proceedings of the 2008 International Conference on Critical Infrastructure Protection

imize their impacts by improving vulnerability assessments, protection planning, and strategies for response and recovery.

Critical infrastructure analysis seeks to provide insights into infrastructure behaviors and potential disruptions, insights that may increase the efficacy of protection plans as well as operations for response and recovery. The U.S. government has identified thirteen (13) critical infrastructure sectors [10] (e.g., energy, communications, and banking and finance). Each sector involves multi-dimensional, highly complex collections of technologies, information, processes, and people - i.e., each is a mission-critical, socio-technical system. Moreover, all sectors are highly interdependent where disruptions within one sector often cascade and escalate across other sectors [13].

Effective critical infrastructure analysis, therefore, must account for these characteristics - leading to two important requirements. First, critical infrastructure analysis must emphasize not only the engineering properties, but also the behavioral properties of each infrastructure. The engineering properties primarily describe the technical characteristics of an infrastructure in terms of the underlying physics-based properties of the inanimate objects that shape and constrain the operation of that infrastructure. Behavioral properties describe the relational properties that emerge from business processes, decision points, human interventions, information availability, reliability, and consistency, etc. in addition to the engineering properties of the infrastructure.

Second, critical infrastructure analysis must be conducted *in situ*, i.e., in context. Critical infrastructures operate in place and time. Examining infrastructures in isolation improperly ignores the complex dependencies that exist among infrastructures as well as contextual factors that shape and constrain infrastructure behavior. Suchman argues that context gives meaning to action [14] - that one cannot separate actions from the context in which they are performed without losing the meanings or implications of those actions. Consequently, examining critical infrastructures outside of place and time can lead to a loss in the meanings or implications of infrastructure behaviors, which results in vulnerability assessments that are at best incomplete and at worst invalid.

These two requirements must be the foundation for any comprehensive, holistic, and systemic analysis of critical infrastructures. This is particularly true for modeling and simulation solutions that enable analysis in support of critical infrastructure protection. In this paper, we present a new methodology for critical infrastructure modeling and simulation and demonstrate how this methodology addresses the above requirements. First, we present an overview of related work. Second, we describe our methodology and its realization within the Integrated

Modeling Environment (IME), our critical infrastructure modeling and simulation solution. Third, we illustrate the application of the methodology and the IME to a critical infrastructure analysis problem. Finally, we offer some concluding thoughts and discuss future work.

2. Related Work

Modeling and simulation are important techniques to facilitate the exploration and analysis of complex phenomena. In fact, for many phenomena, modeling and simulation may be the only viable means for exploration and analysis. This is particularly true for phenomena that are characterized by organic collections of events that occur within open systems - systems that may include social, economic, technical, civic, environmental, informational, and geographic context. Effective modeling and simulation of such phenomena often require a system of systems approach that recognizes the various dimensions of the phenomena as well as the relations among these dimensions. These characteristics are particularly true for critical infrastructure modeling and simulation.

A comprehensive survey of critical infrastructure modeling and simulation solutions can be found in [11]. This survey highlights several methods to critical infrastructure modeling and simulation. Several solutions decompose analysis to the exploration of individual infrastructures. Many useful single infrastructure solutions exist, e.g., [1, 12]. However, decomposition methods fail to recognize the importance to critical infrastructure analysis of the behavioral properties of each infrastructure as well as the complex dependencies that exist among infrastructures. Furthermore, these solutions cannot necessarily be generalized due to the unique characteristics of each infrastructure.

Other solutions focus on the interdependencies among infrastructures, e.g., [4, 6]. These solutions attempt to recognize the *in situ* requirement and model the complex behavior that emerges from the dependencies among participating infrastructures. However, these solutions do not adequately incorporate the unique behavior of the underlying infrastructures. While dependencies among critical infrastructures can lead to cascading and escalating effects [13], such effects emerge specifically from the interplay of these dependencies and the individual behavior of each infrastructure. By eliminating individual infrastructure behavior from the model, the fidelity of the model is greatly reduced.

Still other solutions attempt to build comprehensive models of critical infrastructures (e.g., [3, 5, 9, 13, 15]). However, comprehensive models are not necessarily tractable due to the unique characteristics of each

infrastructure. As a result, these models typically emphasize higher levels of analysis while deemphasizing detailed analysis.

Recently, there have been efforts to develop hybrid solutions for critical infrastructure modeling and simulation, e.g., [2, 16]. Pederson et. al. describes these efforts as a coupled modeling approach [11]. Under this approach, individual infrastructure models are integrated in a generalized way with models of infrastructure dependencies to enable system of system analysis - thus, coupling the fidelity of individual infrastructure models with the requirement for analysis occurring *in situ*. The modeling and simulation solution presented in this paper is properly described as a coupled modeling approach.

3. Methodology Foundation

Drawing upon our previous work [16] and recognizing the importance of the above mentioned requirements, we are developing a new methodology for integrated modeling and simulation. The foundation for the methodology is grounded in our framework for leveraging existing infrastructure models and our representation of context and behavior.

3.1 Integration Framework

Our methodology for integrated modeling and simulation is based, in part, on our ability to integrate separate infrastructure models under a single modeling and simulation framework. This framework is designed around a service-oriented architecture supported by a common service provider API (see Figure 1). Under this framework, each infrastructure model is integrated by implementing framework connector, which realizes the common service provider API, and registering the connector with the framework's service registry. These models, then, are leveraged during analysis via the connectors by the Integrated Modeling Environment (see Section 5), which functions as a service requester. Interaction between the service requester and service providers is event-driven. Thus, our methodology enables discrete simulations in support of analysis activities. Individual infrastructure models, however, may or may not be event-based. For example, to integrate continuous simulation models, one must implement a framework connector that adapts continuous simulations to discrete simulations.

3.2 Representing Context and Behavior

As context gives meaning to action [14],examining the behavior of critical infrastructures in isolation and outside of place and time can lead to a loss in the meaning or implication of infrastructure behavior. To rep-

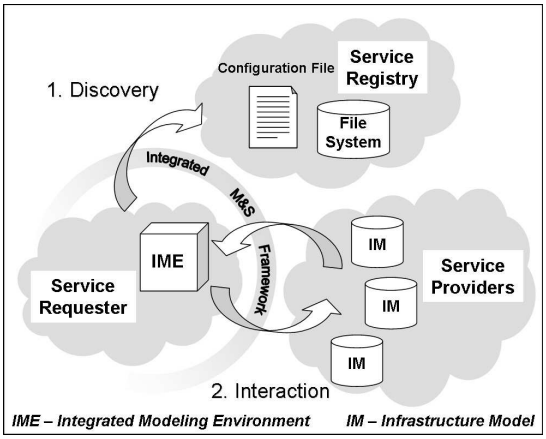

Figure 1. Integrated Modeling and Simulation Framework.

resent context and the meaning it embodies, we draw on John Locke's definition of *knowledge*. Locke describes knowledge as the ability to distinguish concepts or ideas [8]. In other words, knowledge emerges from the relationships among concepts. As such, our representation leverages this definition of knowledge while drawing upon ontological modeling principles and the notion of a relation to provide a representation of context and behavior. Our methodology uses relations to support the specification of contextual and behavioral properties along three distinct dimensions: function, time, and space. These dimensions situate infrastructure features and their collective behavior by answering the questions how?, when?, and where? features are related. In this context, an infrastructure feature is any modeled component of an infrastructure.

3.2.1 Functional Relations. Under our methodology, each infrastructure feature may be associated functionally with other infrastructure features. We define our functional relations according to a specified commodity and relational rule, and by leveraging a provider/subscriber paradigm. Commodities are tangible or intangible goods or services that may be generated, transported, and/or consumed by infrastructure features. Relational rules further restrict the relation by constraining the set of valid origin features that may provide the specified commodity to the specified destination feature. Most relational rules constrain this behavior according to provider/subscriber proximity. Represented by the following tuple, $(origin \times commodity \times destination \times relational\_rule)$, a functional relation under our methodology states that infrastructure feature *origin* provides the specified *commodity* to infrastructure feature

destination according to the specified *relational_rule*. Given that the collective critical infrastructures for a given region may have in excess of 20,000 features, it is not feasible that every functional relation be individually specified. As such, we allow functional relations to be specified at both a type/subtype level and an instance level using selection sets. A selection set is a specification that resolves to a set of features according to a specified criterion. For example, our methodology enables the specification of functional relations that state that infrastructure features of type *origin_type* provide a specified *commodity* to infrastructure features of type *destination_type* according to the *relational_rule*.

3.2.2 Temporal Relations. Under our methodology, each infrastructure feature may be associated with temporal latencies for enabling or disabling the feature. Represented by the following tuple, $(feature \times commodity \times effect \times duration)$, a temporal relation under our methodology states that when a specified infrastructure $feature$ losses or gains access to a specified *commodity*, the specified $effect$ (i.e., disable or enable, respectively) is delayed a specified *duration*. For example, if an infrastructure feature losses access to the essential commodity *electricity*, the disabling effect of losing that commodity is delayed until a specified latency has passed. This latency may model a limited alternative commodity source (e.g., battery backup). Similarly, once an infrastructure feature gains access to its essential commodities, the enabling effect is delayed until a specified latency has passed. This latency may model the startup time required to enable the feature. If access to an essential commodity is restored before the disablement latency has expired, then the disable event is discarded. Similar to functional relations, temporal relations for infrastructure features may be specified at either a type/subtype level or an instance level.

3.2.3 Spatial Relations. Finally, our methodology recognizes that as physical objects, infrastructure features are spatially tangible. As such, each infrastructure feature may associated with a location in the geographical space. Its location and spatial relationships with other infrastructure features are represented by geographic coordinates and further, as in many geographic information systems, by topological relationships [7]. Represented by the following tuple, $(feature \times location)$, a spatial relation under our methodology states that infrastructure $feature$ is located at *location* in geographic space. Spatial relations of infrastructure features are used in numerous ways, including for proximity analysis according to relational rules (e.g., nearest provider within a specified radius), spatial correlations (e.g., map overlays), and geo-visualizations.

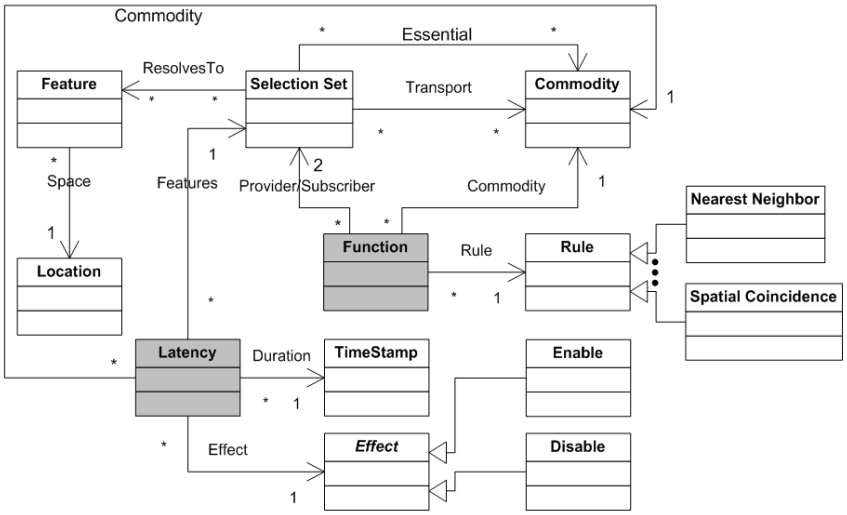

Figure 2. Infrastructure Context and Behavior Ontology.

3.2.4 Infrastructure Context and Behavior Ontology.

Integrating functional, temporal, and spatial relations leads to the following representation, i.e., ontology, for modeling infrastructure context and behavior (see Figure 2). An ontology models the well-defined dimensions of a domain in terms of objects, attributes, and relations. An ontology also enables the construction of a common understanding through a common language and representation for analytical discourse. In our ontology, functional and temporal relations are represented by the objects in grey. Spatial relations are modeled by the "Space" association between the "Feature" object and the "Location" object.

4. The Integrated Methodology

Leveraging our integration framework and our context and behavior ontology we now presented our methodology. This new methodology is comprised of five (5) key steps.

1. *Infrastructure model identification and development* - Infrastructure models are realized by using 3^{rd} party products (e.g., [1, 12]) or by instantiating generic infrastructure models that are built into the integration framework (i.e., utility, transport, and channel networks).
2. *Connector development according to the integration framework* - Each infrastructure model must instantiate a connector in order for the model to participate in the integration framework. The framework defines a simple connector API to support connector development.

3. *Infrastructure model import* - The IME, as a service requester, requires from each infrastructure model a representation of the infrastructure features for the model in order for those features to participate in the context and behavior ontology.
4. *Integrated model development* (i.e., ontology instantiation) - Functional, temporal, and spatial relations are specified. From these specifications, relationships are instantiated.
5. *Integrated modeling and simulation* - Models are explored, simulations are executed and analyzed, models are validated, analysis products are constructed.

The relationship among these steps is not strictly sequential. Rather, each step remains ongoing as analysis questions change, infrastructure models evolve (due to data acquisition, verification, and validation), and the integrated model evolves (due to model evolution, verification, and validation). Analysis, therefore, is organic activity that is seamlessly integrated with infrastructure model development, integrated model development, and verification and validation.

5. The Integrated Modeling Environment

The Integrated Modeling Environment (IME) is a modeling and simulation solution that facilitates system-of-systems analysis by enabling the horizontal fusion of zero or more infrastructure models. System-of-systems analysis seeks to explore and understand the collective behaviors of integrated systems. In the context of critical infrastructure protection, one might integrate separate models for the electric power, telecommunications, natural gas distribution, and transportation infrastructures for a given geographic region. Then, using the IME, analysts may conduct integrated, multi-model analysis via simulations to explore and understand the collective behaviors of these integrated models.

To illustrate, we provide a brief introduction to the IME and an example of the analysis that it supports. The primary interface for the analyst includes a multi-tab palette and a geo-visualization of a given region. Figure 3 depicts the analyst interface palette. Included in this palette are three tabs. The first tab, the "Objectives" tab, allows analysts to specify desired and undesired effects aggregated under a named objective. Here, an effect represents the disablement of a specified domain model element, i.e., an infrastructure feature. Objectives may be specified either from a red-team or blue-team perspective.

The second tab allows analysts to specify sequences of scheduled events (i.e., courses of action), where each scheduled event represents

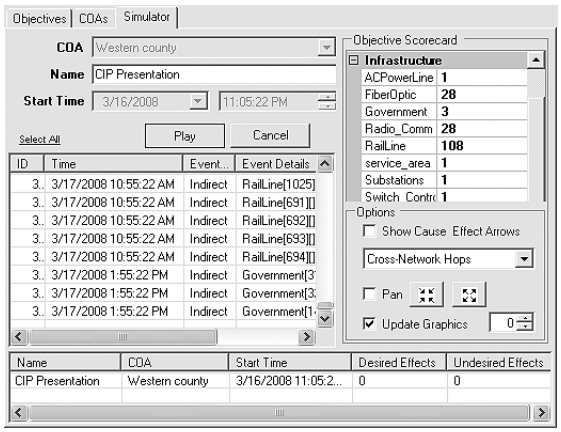

Figure 3. Analyst Interface - Simulator Tab.

the enabling or disabling of specified infrastructure features at a given time, represented as a delta from the start of a simulation.

The third tab (visible in Figure 3) is the "Simulator" tab. This tab allows analysts to select a course of action, specify a start time and initiate a simulation. As the simulation executes, the analyst sees on the left side of this tab an event stream capturing infrastructure feature enable and disable events. Features are enabled/disabled as a function of individual infrastructure model behavior and as a function of the relations specified in the infrastructure context and behavior ontology. Each simulation event includes a timestamp. The right side of the simulator tab contains a scorecard that aggregates simulation event stream data along various dimensions (e.g., time, infrastructure, feature type). Saved simulations are listed at the bottom of this tab. As the simulation executes, analysts can observe the effects in the geo-visualization. Dynamic changes in feature symbology reflect domain model state changes (i.e., the enabling and disabling of features).

To conduct meaningful analysis, however, the context and behavior ontology must be specified. This activity is supported by a separate "model builder" palette. This palette includes, among other tabs, interfaces for the specification of commodities, relationships, latencies, and connectors. The relationship tab (visible in Figure 4) provides model builders a means to specify and manage the functional relations for the given domain models. The latencies tab provides model builders with a means to manage the temporal relations that specify infrastructure feature enabling and disabling latencies. Finally, the connectors tab provides model builders with a means to manage the participating in-

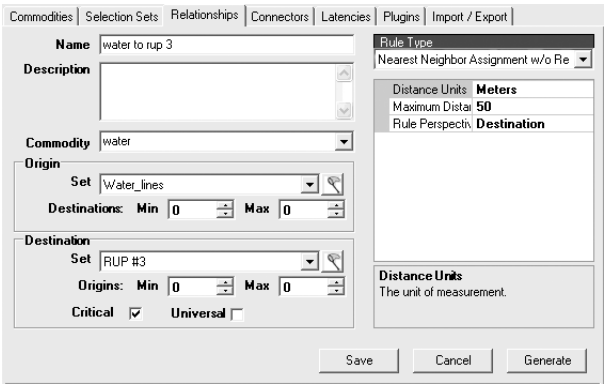

Figure 4. Model Builder Interface - Relationships Tab.

frastructure models. Several infrastructure models have been integrated into the IME via the connector framework (e.g., [1, 12]). In addition, the IME, by default, provides the aforementioned built-in models - i.e., utility, transport, and channel networks.

6. Critical Infrastructure Analysis Illustration

In this section we provide a simple illustration of our methodology. This illustration focuses on a small geographic region with several buildings and critical infrastructures. In particular, this illustration includes infrastructure models for natural gas, steam, and water. Figure 5 depicts each infrastructure as well as a layered view of all infrastructures.

To support integrated modeling and simulation across these infrastructures, one first geo-codes relevant infrastructure features to establish spatial context. Next, one identifies the commodities that are essential to the operation of the infrastructures in question. In this illustration, the following commodities are identified: steam, gas, and water. Next, one specifies temporal latencies for infrastructure features to establish temporal context. For this illustration, there is only one temporal latency specification (see Table 1). Finally, the model builder specifies the relevant functional relations among infrastructure features. For this illustration, three functional relations are specified (see Table 2).

To conduct analysis, analysts utilize the previously described interface (see Figure 3) to specify objectives and courses of action, and to execute and explore simulations. For this illustration, the objective is to maintain the operation of Buildings #2 and #3 (i.e., the analyst specifies an objective with the undesired effects of disabling these buildings). The course of action is initiated by a gas line fracture due to ongoing

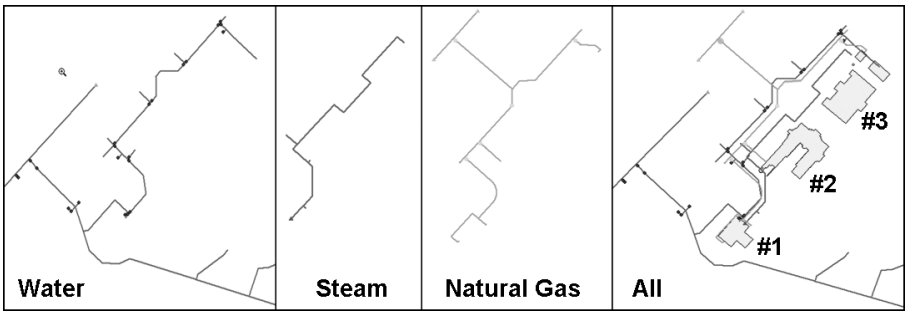

Figure 5. Illustrative Infrastructure Models.

Table 1. Temporal Relations.

Selection Set	Commodity	Effect	Duration
Steam Source	Steam	Disable	1.00:00:00 (d.h:m:s)

Table 2. Functional Relations.

Origin	Commodity	Destination	Rule
Water Line	Water	Building #1	Nearest Neighbor†
Building #1	Steam	Steam Source	Nearest Neighbor†
Steam Line	Steam	Building #2 & #3	Nearest Neighbor†
Gas Line	Gas	Building #1	Nearest Neighbor†

construction. Subsequent to the fracture, downstream gas is lost (Figure 6, panel 1). To contain the leak, a gas valve is scheduled to be closed, as part of the course of action, one hour into the simulation. This results in the loss of the gas commodity to Building #1 (Figure 6, panel 2). The loss of gas to Building #1 halts the production of steam (Figure 6, panel 3). After a twenty-four hour delay, Buildings #2 and #3 can no longer function due to a loss of requisite heat (Figure 6, panel 4). The integrated modeling and simulation behavior as demonstrated by this simulation is realized by the behaviors of the individual infrastructure models and the temporal, spatial, and functional relations in the IME context and behavior ontology.

Once simulations complete, they may be explored, replayed, and saved for further analysis. Using the scorecard interface analysts can examine

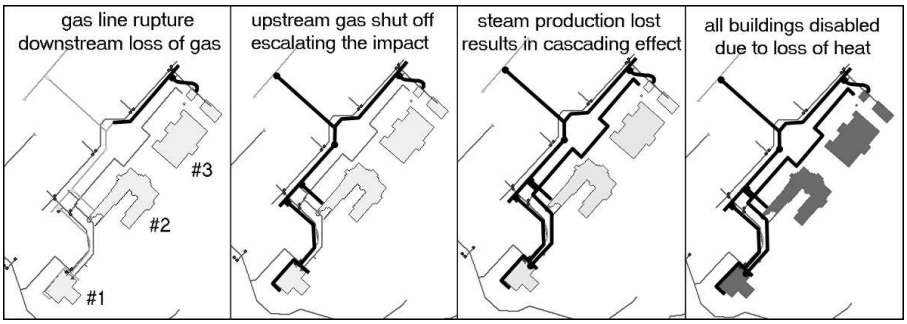

Figure 6. Example Simulation - Disabled Features in Bold.

the time-sequenced order-of-impact of simulation events as well as the plausible impact to each critical infrastructure. In addition, analysts can examine the event trace to understand and/or validate the event chain that led to a (un)desired effect. During analysis, analysts may refine the ontology, e.g., by adding/deleting/modifying commodities, functional relations, and/or temporal latencies, to explore "what-if" scenarios.

7. Conclusions and Future Work

The methodology presented in this paper reflects nearly five years of research and development. Over this time, the importance of context in all of its richness (e.g., function, space, and time) to analysis constantly reappeared. The research embodied in our methodology has been successfully translated into practice and in the process uncovered numerous further research questions.

Evaluation of our research remains an ongoing priority and research question. We have explored this issue from several perspectives [17, 18]. As a result, verification and validation of our methodology has become inherent to the practice of using the IME and is further enabled by the underlying principle of transparency that is embodied in our methodology. In practice, our methodology, as realized in the IME, is actively being used by analysts to explore and analyze critical infrastructures for large scale (>100,000 km^2) geographic regions. In addition, we have developed an integrated model for an urban region, the extent of which is >500 mi^2 with a population that exceeds 800,000. The critical infrastructures in this integrated model include: electric power, natural gas distribution, water distribution, telecommunications, and transportation. Furthermore, we have demonstrated the IME on a corporate IT infrastructure model (for a Fortune 500 company) that integrates models

for IT hardware, system software, business applications, business processes, and business units. Finally, we have developed models for an urban neighborhood with an extent of roughly 1000 contiguous acres that serves a population of over 20,000. Further verification and validation is enabled by our adherence to the underlying principle of transparency. All analysis enabled by our ontology is completely transparent to the analyst. Event traces can be explored and questioned by subject matter experts. In fact, this practice is encouraged by our methodology and regularly utilized by its practitioners. The result is an ongoing seamless activity of analysis with verification and validation, the impact of which improves the underlying ontology as well as the resulting analysis.

At the same time, there are several limitations to our work. First, all simulations are currently deterministic. For our user community, this is considered an advantage and a disadvantage. It is an advantage because many of the analysts that we encounter want to retain the responsibility to assess the level of certainty in the analysis. In addition, many non-deterministic analysis techniques are based on prior probabilities that are unavailable for the models in question. On the other hand, it is a disadvantage because the modeled phenomena frequently contain high levels of uncertainty and techniques for non-deterministic analysis may be able to expose more easily the range of plausible outcomes. Research focused on the targeted introduction of non-determinism into our methodology is ongoing. Second, the IME is currently unable to represent infrastructure degradation. Research to extend the methodology to account for degradation, however, is also ongoing as the issues of degradation and non-determinism appear intertwined. Third, we continue to explore ways to expand the expressiveness of the IME ontology. Finally, the IME faces several visualization limitations. Methods must be developed to visualize in an integrative and effective manner multiple infrastructure models along their functional, spatial, and temporal dimensions. These questions are also being actively pursued by our team.

References

[1] ArcGIS Network Analyst, http://www.esri.com/software/arcgis/extensions/networkanalyst/index.html.

[2] E. Casalicchio, E. Galli, S. Tucci. Federated agent-based modeling and simulation approach to study interdependencies in IT critical infrastructures, *11th IEEE Symposium on Distributed Simulation and Real-Time Applications*, IEEE Computer Society, 2007.

[3] A. Chaturvedi, A society of simulation approach to dynamic integration of simulations, *Proc., Winter Simulation Conference*, 2006.

[4] D.D. Dudenhoeffer, M.R. Permann, M. Manic, CIMS: a framework for infrastructure interdependency modeling and analysis, *Proc., Winter Simulation Conference*, 2006.

[5] F. Flentge, U. Beyer, The ISE metamodel for critical infrastructures, *Critical Infrastructure Protection*, Goetz and Shenoi (Eds.), Springer, pp 323-336, 2007.

[6] O. Gursesli, A.A. Desrochers, Modeling infrastructure interdependencies using petri nets, *IEEE International Conference on Systems, Man and Cybernetics*, 2003.

[7] C.P. Lo, A.K.W. Yeung, Concepts and Techniques of Geographic Information Systems, Upper Saddle River, NJ: Prentice Hall, 2007.

[8] J. Locke, *An Essay Concerning Human Understanding*, 1690.

[9] J.R. Marti, J.A. Hollman, C. Ventrua, J. Jatskevich, Design for survival real-time infrastructures coordination, *Proc., Int'l Workshop on Complex Network and Infrastructure Protection*, March 2006.

[10] National Strategy for Homeland Security, 2002.

[11] P. Pederson, D. Dudenhoeffer, S. Hartley, M. Permann, Critical infrastructure interdependency modeling: a survey of U.S. and international research, Report No. INL/EXT-06-11464, Critical Infrastructure Protection Division, Idaho National Laboratory, 2006.

[12] PowerWorld Simulator, http://www.powerworld.com/products/simulator.asp.

[13] S.M. Rinaldi, J.P. Peerenboom, T.K. Kelly, Identifying,understanding, and analyzing critical infrastructure interdependencies, *IEEE Control Systems Magazine*, Dec. 2001.

[14] L.A. Suchman, Plans and Situated Actions: the Problem of Human-Machine Communication, Cambridge University Press, 1987.

[15] N. Svendsen, S. Wolthusen, Multigraph dependency models for heterogeneous critical infrastructures, *Critical Infrastructure Protection*, Goetz and Shenoi (Eds.), Springer, pp 337-350, 2007.

[16] W.J. Tolone, D. Wilson, A. Raja, W.N. Xiang, H. Hao, S. Phelps, E.W. Johnson, Critical infrastructure integration modeling and simulation, *Proc., Second Symposium on Intelligence and Security Informatics (ISI-2004)*, LNCS #3073, Springer, June 2004.

[17] A.J. Weeks, An assessment of validation methods for critical infrastructure protection modeling and simulation, MA Thesis, UNC Charlotte, 2006.

[18] A.J. Weeks, A. Schumpert, S.W. Lee, W.J. Tolone, W.N. Xiang, A new approach to V&V in CIP modeling and simulation,*Proc., ESRI International User Conference*, San Diego, CA, Aug. 2006.