

Modeling Critical Infrastructure Requirements

Stephen D. Wolthusen

Abstract—Critical infrastructures in industrialized nations form a highly interdependent network that must be protected against both intrinsic defects and active attacks. This requires local as well as joint situational awareness based on current, accurate, and semantically unambiguous data as well as simulations, particularly of attack scenarios, necessitating in turn automated information sharing measures that can span transitive dependency networks.

Since the infrastructure elements are frequently civilian-owned, providing provable assertions on the precise nature of the data shared and the extent of dissemination is crucial. In this paper, a layered graph-theoretical modeling technique is used; at a lower layer, a standards-based ontological model is described in which resources and interactions are formed into a common exchange format. From this, a simple dependency model amenable to combinatorial optimization and simulation is described, which is then also used as the foundation for the application of the schematic protection model by Sandhu to the information sharing problem.

I. INTRODUCTION

THE protection of critical infrastructures such as energy, financial services, health care, public services, and transportation [1]¹ has moved from being primarily driven by safety and engineering concerns to also incorporating elements of security, particularly from external hostile actions, but also including sabotage from within.

Critical infrastructures in most industrialized countries share several characteristics that make their protection, particularly as a cohesive system, difficult.

Foremost among these is that the majority of the infrastructure is owned by commercial or semi-commercial interests (e.g. municipal utilities) that must operate competitively, with limited capital investment and operational expenses. That, in itself, is already constraining decisions regarding the safety and security of infrastructure elements beyond what is either required by regulatory authority and legal requirements or by providing an immediate competitive advantage [3].

From the perspective of the overall or national critical infrastructures, respectively, however, another issue resulting from civilian ownership is that the very information and exchange of information required to maintain either interdependent elements of the infrastructure or the entire

infrastructure can put the civilian infrastructure owners at a disadvantage. This can e.g. occur when required information sharing in a given sector exposes business intelligence (e.g. cost structures, capabilities) to competitors in the same or another sector. Another potential impediment to information sharing is particularly prominent in case of natural monopolies (e.g. certain utilities). Here, disclosure of mishaps or even simply potential vulnerabilities can be correlated immediately with its most likely origin, potentially resulting in decreases in the valuation of the infrastructure owner.

Therefore, while information gathering and sharing are critical elements in both the prevention of harm to critical infrastructure and the timely and efficient remediation of any problems that occur within the network of interdependent infrastructure components, the flow of information itself must be closely monitorable and controllable if infrastructure owners and operators are to engage in it. This is e.g. clearly reflected in the U.S. approach to CIP following the catastrophic terrorist attacks of 2001 [4], [5].

Moreover, the information collected and exchanged in the interest of improving the robustness, availability, and overall assurance of infrastructure elements must also be protected against malicious outside interest and influence. While much of the focus in infrastructure protection (CIP) is traditionally provided by the safety engineering community, the possibility of deliberate attacks introduces a number of new failure modes that either have been considered impossible or at the very least highly improbable and have therefore not been given adequate consideration.

This paper therefore outlines a mechanism for modeling infrastructure elements both at the level of data collection and exchange and also for modeling and simulation with particular emphasis on protection for information collection and information sharing.

The remainder of the paper is structured as follows: Section II identifies key requirements for modeling from both the perspective of overall infrastructure protection and the individual infrastructure component owners and operators. Subsequently, section III provides a high-level outline of the underlying modeling mechanisms while section IV details the security controls and policies defined over the model. Finally, section V briefly discusses prior and related work.

S. Wolthusen is with the Security Technology Department, Fraunhofer-IGD, Darmstadt, Germany. E-mail: wolt@igd.fhg.de.

¹While there exists a consensus on the sectors considered to be part of the infrastructures beginning with [1], the precise elaboration and granularity of sectors differs in the various national approaches [2].

II. REQUIREMENTS

In addition to individual infrastructure component robustness and survivability [6], [7], [8], which is beyond the scope of this paper, CIP rests on understanding and dynamically adapting component configurations in such a way that overall objectives (e.g. power plant output to the electrical grid) are met.

This description is scale-free in its applicability; however, some issues become relevant only if geographical and organizational boundaries are crossed in the process.

At the largest scale (i.e. national and multinational structures), information sharing and analytical capabilities among independent entities must be considered the primary feasible mechanism for improving infrastructure reliability and survivability (see section V); for civilian infrastructure owners and operators to conduct the prerequisite information collection and storage activities, CIP activities should be concomitant with other cost benefits.

Consequently, the acceptability of CIP models may be enhanced significantly if, in addition to its primary objectives, it also assists in providing information that is internally useful to an organization, e.g. in identifying inefficiencies and redundancies beyond what is required for protection purposes.

Infrastructure dependencies are, in all but the most trivial cases, multilateral relations (among intra-organizational infrastructure elements, infrastructure providers, and with government).

However, the dependencies themselves are both dynamic (including feedback loops) and insufficiently characterized by a simple relation. The former observation implies that the efficacy of an infrastructure model is significantly influenced by its coupling with the underlying system (i.e. it is insufficient to operate on data collected by temporally isolated snapshots), while the latter implies a requirement for detailed annotation of relations.

Since individual infrastructure owners generally already operate information systems containing all or significant portions of the data required for CIP; however, direct harmonization among such entities for information sharing and exchange is both impractical given the scope of data contained in such databases, and infeasible given the cost sensitivities of infrastructure owners.

As a result, a key requirement for information exchange is the use of an interoperable intermediate format of sufficient generality to contain not only the data elements but also the underlying ontological structures. The latter requirement not only results from the need to translate data points and relational tuples, but also from the fact that infrastructure elements evolve over long time scales – during which the semantics of individual data points and metrics are likely to change.

To retain the ability to perform analyses over such changing data points, ontological information must be retained

for use in renormalization. Moreover, an additional dimension for data points (particularly those not founded in ground measurements but rather in derivations and inferences) that frequently arises is the need for associating a confidence level with the data point that can subsequently be used by belief revision techniques [9], [10], [11], [12], [13].

The information assurance requirements for both ground and derived data points (confidentiality, integrity, correctness, availability, timeliness, and non-repudiability) must be characterized in this way for both those derived from within a system under consideration (e.g. a single power plant) and when modeling interconnected infrastructure elements since feedback mechanisms can also impede proper modeling at the local level (e.g. in manipulating sensor data transmitted over insufficiently secure links as in chemical or power plants).

Even though the ultimate goal of any integrated CIP model is information sharing, access, including read-only provisions to information must be constrained by the least privilege principle [14] with well-defined information flows based on need-to-know and at the same time full auditability of any transfer. As noted in section I, some of these constraints are direct results from the fiduciary duties of civilian infrastructure owners and operators to equity holders, protecting information system assets and against competitive intelligence. Moreover, the information to be protected extends to indirect effects such as public confidence in the infrastructure operator, resulting in a requirement for protecting even indirect information flows where sources could easily be inferred (e.g. in case of a local power utility as the indirect causative agent for a failure at a water treatment facility).

At the same time, emergency access requirements must be defined in such a way that the transition from normal operation to emergencies is well defined and can be invoked by all authorized parties to the extent necessary (e.g. for providing electrical power to critical areas such as hospitals and air traffic control facilities).

Moreover, reliability models commonly used generally assume stochastic processes in assessing the likelihood of a malfunction or hazard; such assumptions may no longer be made safely since adversaries (particularly terrorists) may deliberately target interdependent networks, thereby inducing otherwise highly improbable fault chains. Since resources for protecting infrastructures are necessarily finite and the number of such scenarios may be assumed to be transfinite, comprehensive analytical methods are unlikely to yield usable results. The ability to perform simulations and case studies of such attack scenarios is therefore of particular pragmatic interest.

III. MODEL STRUCTURE

The following sections provide a high-level overview of the structures used for representation and reasoning over

critical infrastructure data. Section III-A represents the highest abstraction level at which entities and dependencies are represented in the dependency model; the primary focus of this sub-model is scale invariance and efficient representation for use in computation. Moreover, this model also provides the foundation for the control mechanisms discussed in section IV.

Section III-B provides a sketch for the the intermediate ontological model and exchange mechanism data format. The primary goals for this model (in addition to a straightforward bijection onto the graph-theoretic dependency model) were the provision of a relatively simple common abstract data format for critical infrastructure data, the ability to perform high-level computations over the represented entities as well as the dynamism required for real-time analysis.

The final layer within this model, namely the mapping of existing databases, sensor data, and other interfaces onto the ontological model (e.g. for geospatial models and databases, wiring topologies, etc.) is beyond the scope of this paper since such interfaces are necessarily product- and typically implementation-specific since most infrastructure owners and operators rely on considerable internal development to provide adequate cataloging and analytical capabilities.

A. Dependency Model

To satisfy the requirements outlined in section II, a dependency model based on multigraphs provides a powerful and general mechanism with a sound mathematical foundation.

Definition 1: Infrastructure components are separated into entities \mathcal{E} ($\mathcal{E} = \{e_1, \dots, e_k\}$) represented as vertices and dependencies \mathcal{D} ($\mathcal{D} = \{d_1^1, \dots, d_n^m\}$) among entities represented as directed edges where the set of edges is partitioned into m *dependency types*, resulting in a graph $\mathcal{G} = (\mathcal{E}, \mathcal{D})$. \mathcal{G} may contain parallel edges, but may not contain self-loops.

Edges between two given vertices e_a, e_b are not uniquely identified by the 2-tuple (e_a, e_b) as is the case in simple graphs since they may differ in their dependency type:

Definition 2: For two given vertices e_a, e_b within \mathcal{G} , the set of edges must not contain two edges of the same dependency type.

The set of all dependencies between given vertices e_a, e_b is denoted as (e_a, e_b) and abbreviated (a, b) . By collecting edges of different dependency type, a directed simple graph \mathcal{G}^s is produced and referred to as the *aggregate dependency graph*.

For a given dependency type t , a t -*dependency path* is a sequence $P = \{e_1, d_1^t, e_2, d_2^t, \dots, d_{i-1}^t, e_i\}$ of alternating vertices and edges such that for $1 \leq j \leq i$, d_j^t is incident with e_j and e_{j+1} .

Two paths are t -*edge disjoint* if they do not have an edge of type t in common.

Dependency paths and connectivity properties are preserved by the aggregate dependency graph, edge disjointness is defined analogously to t -edge disjointness.

For a given graph with edges e_k and dependency types t_j , a relation $(e_k \times t_j) \leftrightarrow \mathbb{N}$ is defined. The range of this relation is referred to as the *dependency strength* and denoted s_{e_k} for a given edge e_k .

Given a dependency graph, the graph can be partitioned into vertex subsets $\mathcal{E} = \mathcal{E}_1 \cup \dots \cup \mathcal{E}_k$ (where $k \leq |\mathcal{E}|$) called *partitions* (\mathcal{P}_i). For observing dependencies at higher levels of abstraction, theorem 1 provides a justification for coalescing graphs.

Theorem 1: For a given dependency graph $\mathcal{G} = (\mathcal{E}, \mathcal{D})$ and a partitioning over the vertices $\mathcal{E} = \mathcal{E}_1 \cup \dots \cup \mathcal{E}_k$, each partition (\mathcal{P}_i) can be substituted by a single *coalesced vertex*.

Proof (sketch): Without loss of generality select a partition \mathcal{P}_i . For each unique dependency type t_j in \mathcal{P}_i (determined by enumerating the edges of the subgraph induced by \mathcal{P}_i), insert a vertex $e_j^{t_j}$ into \mathcal{P}_i such that all dependency paths with edges of type t_j in \mathcal{P}_i for which a vertex lies outside of \mathcal{P}_i are incident with $e_j^{t_j}$.

For each separate inbound (outbound) dependency path now incident with $e_j^{t_j}$, form the sum over the dependency strengths.

All vertices except the $e_j^{t_j}$ can now be eliminated since no dependency path is cut by the removal.

Remove any edge from the graph that forms a self-loop for a given $e_j^{t_j}$.

Insert a vertex e_i into \mathcal{P}_i and extend the inbound (outbound) dependency paths such that the edges of the $e_j^{t_j}$ are incident with e_i .

The $e_j^{t_j}$ can now also be removed as above, and the partition \mathcal{P}_i is coalesced into a single vertex e_i . \square

Similar to vertex coalescion, *edge coalescion* can also be of interest; in this case two or more edges with different types t_i and t_j incident with vertices e_k, e_l are coalesced by forming the set union over the types with the derived type $t_{i,j}$. The coalesced edges are then removed from the graph; no self-loops can occur in this step.

If all edges are coalesced, the result is a *typeless dependency graph*. The dependency strength of a coalesced edge created from k individual edges is trivially defined as

$$\frac{1}{k} \sum_{1 \leq i \leq k} s_{e_i}$$

if a normalized dependency strength is used (see section III-A.3).

A.1 Transitive Dependency Strength

Given a dependency graph annotated with dependency strengths, it is of particular interest to identify critical tran-

sitive dependencies (including for edge and vertex aggregations). This can be determined by interpreting the dependency strengths applying e.g. the maximum flow algorithm by Ford and Fulkerson [15], [16]. Given a bound for dependency strength s_{max} , the algorithm determines identifies the maximum transitive dependency in $\mathcal{O}(|\mathcal{D}| s_{max})$.

This algorithm provides adequate results in all cases given the formulation using integral dependency strengths [17]; in dense dependency graphs, more efficient algorithms such as preflow-push algorithms can be applied [18]; using the dependency graph construction above, a maximum flow can be computed deterministically using $\mathcal{O}(n^{3/2}m^{1/2} + n^2(\log s_{max})^{1/2} + \log s_{max})$ flow operations and $\mathcal{O}(\min\{nm, n^3/\log n\} + n^2(\log s_{max})^{1/2} + \log s_{max})$ time by applying the algorithm by Cheriyan *et al.* [19], [20].

A.2 Multiple Dependency Paths

Another example of an important question for a dependency model is to determine the number of distinct dependency paths, particularly of vertex disjoint dependency paths (e.g. in analyzing requirements for establishing redundant systems); see figure 1 for one such instance.




Fig. 1. Identification of multiple t -dependency paths

Using the theorem of Menger [21], such a set of depen-

dependency paths can be found constructively with complexity $\mathcal{O}(\mathcal{E}^{\frac{1}{2}}\mathcal{D})$.

However, whether vertex pairs can be disjointly connected is a \mathcal{NP} -complete problem in the general case [22], [23], [24], requiring either the use of randomized algorithms or suitably constrained problem variants.

A.3 Interpretation

The interpretation of entities of the basic dependency graph corresponds directly to that of entities discussed in section III-B; in case of aggregate dependency graphs, the partitioning and aggregation must follow the semantics of the underlying model (e.g. coalescing vertices within a logical grouping or, at a higher level of abstraction, within a single organizational entity).

The interpretation of the dependency types is that of a specific category (e.g. electrical power, voice communication link, water supply) with edge coalescion providing aggregate dependencies between entities.

Several graph-theoretic algorithms useful in calculating properties of the dependency graph require a bounded edge valuation; dependency strength is therefore required to be individually bounded and re-normalized in case of edge coalescion.

Moreover, to retain the ability to apply certain combinatorial optimization algorithms to the dependency graph, dependency strength must be expressed by values $\in \mathbb{N}$.

Such numerical valuations are frequently not possible immediately; it is therefore frequently necessary to perform a translation from qualitative assessments onto a fixed but arbitrary scale (which must be used consistently throughout the dependency graph).

B. Ontological Model

As noted in sections III and II, the ontological model must provide a common abstraction layer for the plenitude of underlying data formats.

Data in this format must have well-defined semantics that can be retained over changes in underlying representations and storage and be archivable. This represents a particular challenge since the lifetimes of many infrastructure components encompasses a large number of information system generations (e.g. in excess of 100 years in case of some water and sewage conduits).

Moreover, in many cases the full semantics is not fully contained in data repositories but only accessible through interpretative logic layers. Given the cost associated with re-acquiring all data (in addition to direct sensor measurements) as well as the danger of inconsistencies among such parallel repositories, the use of an interpretative intermediate layer appears prudent and economical.

Such an intermediate layer can be provided based on open, interoperable standards defined by the World Wide

Web Consortium (W3C) in the form of the Resource Description Framework (RDF).

The RDF represents predicates (e.g. data points) over entities as a directed graph with vertices representing entities and edges annotated with properties and property values [25], [26], [27].

A basic RDF graph can therefore be considered a superset of the dependency graphs discussed in section III-A; several syntactical features such as RDF containers (bags, sequences) can be normalized and decomposed into regular directed graphs for this purpose.

Within RDF, both entities (vertices) and properties (edges) are represented by Uniform Resource Identifiers (URI); while a basic descriptive format exists, this can be extended arbitrarily using the RDF schema mechanism including RDF reification [28]; this definition includes a semi-rigorous model-theoretic definition of the formal semantics of RDF [29].

It should be noted that the use of URIs for representing underlying representations provides a natural solution for satisfying the requirement for real-time data access and mediation to existing data repositories; this mechanism also permits natural interaction e.g. with web service-based architectures such as those found in geographical information systems [30], [31], [32].

The actual ontological representation [33] can also be accomplished using open standards, in this case using the W3C Web Ontology Language (OWL) [34], [35], [36], [37], which can be considered a syntactical and semantic extension of RDF. These standards define descriptions of classes, properties and their instances and, more importantly, semantic entailments which can be used for reasoning within the ontological model.

For the purposes of the critical infrastructure model, a sublanguage of the full OWL language is selected, namely OWL DL; this constrains the expressiveness of the ontological model to those representable by description logics [38]. However, this rather severe constraint is required to retain computability of entailments. The actual ontology schema definition is beyond the scope of this paper.

IV. SECURITY CONTROLS

As noted in sections I and II, ensuring precise control on a need-to-know basis specified by infrastructure owners is a critical requirement to ensure information sharing that is not coerced by governmental measures.

Unfortunately, most commonly used access control mechanisms, particularly including access control lists used e.g. in network security devices and commercial off the shelf (COTS) operating systems, are already sufficiently expressive to ensure that the leaking of access rights (and hence access to information, potentially violating confidentiality and integrity requirements in particular), i.e. the security of a given protection system configuration cannot be proven

in the general case [39], [40].

However, one of the most salient characteristics of an critical infrastructure modeling, simulation, and retrieval system is that pre-determining the characteristics of entities to whom access to resources is granted cannot be effectively determined a priori. This is particularly caused by the need to include transitive dependencies and interactions in a number of computations both at the level of the dependency and the ontological model.

While some models commonly implemented in COTS systems, particularly operating systems used in defense and intelligence applications (e.g. the lattice-based model of Bell and LaPadula [41], [42], [40] or role-based access control models [43], [44]), can provide controls that prevent leaking of rights, the granularity levels feasible in such models require a coarse *a priori* stratification and, moreover, frequently lead to a “mushroom” configuration in which entities must be assigned high classification levels to permit access at operationally required levels.

For the purposes of the model discussed here, a model that combines more flexibility for the expression of (access) rights transfers but whose security can still be decided is, however, highly desirable.

One such model is the schematic protection model (SPM) developed by Sandhu based on a capability-oriented protection mechanism proposed by Minsky [45], [46], [47].

The basic SPM uses a strong type system [48], [49], [50] over entities within the protection system, further subdivided into type families for subjects and objects, and also distinguishing between rights that alter the protection state (control rights) and those leaving it invariant (inert rights) to represent confidentiality model similar in expressiveness to monotonic access matrix models.

The rights associated with an entity (which may ultimately be considered a capability list) are referred to as the domain of an entity in the SPM, whereas a single right descriptor is referred to as a ticket E/r where E denotes the entity to which the right r is to be applied. The model requires that any transfer of rights between subjects occur only if a predicate over each two fixed but arbitrary subject entities is valid per definition 3.

Such rights can be associated trivially both with the edges of the dependency model from section III-A and the ontological model from section III-B, directly mapping the dependency type in case of the dependency model.

Definition 3: Let X, Y be subjects and $dom(X)$ be the set of rights of X and let r be a control right. A local link predicate $link_i(X, Y)$ with formal parameters X, Y is defined as a conjunction or disjunction of the following atomic terms:

1. $X/r \in dom(X)$,
2. $X/r \in dom(Y)$,
3. $Y/r \in dom(X)$,
4. $Y/r \in dom(Y)$, or

5. true

If $\text{link}_i(A, B)$ evaluates to truth, this is referred to as a link_i between A and B .

In addition to this predicate, which can be considered the connection relation over a rights transfer graph, the SPM further constrains the copying of rights by way of filter functions for given rights:

Definition 4: Let T be the set of all types, R the set of all rights, T_S the set of types for subjects, and a filter function f_i be a function $f_i : T_S \times T_S \rightarrow 2^{T \times R}$.

For given subjects A, B, X of type a, b and x , respectively, the right $X/r : c$ can be copied if and only there exists an i such that $X/rc \in \text{dom}(A)$, $\text{link}_i(A, B)$, and $x/r : c \in f_i(a, b)$.

Given these preliminary definitions, a protection scheme is then given by definition 5:

Definition 5: A protection scheme consists of the following elements:

1. A finite set of entity types T , partitioned into subject entity types T_S and object entity types T_O ,
2. A finite set of rights symbols R , partitioned into inert rights R_I and control rights R_C ,
3. A finite set of local link predicates $\{\text{link}_i\}_{1 \leq i \leq n}$ ($n \in \mathbb{N}$),
4. A filter function $f_i : T_S \times T_S \rightarrow 2^{T \times R}$ whose domain covers the link_i ,
5. A demand function $d : T_S \rightarrow 2^{T \times R}$ authorizing a subject to demand a right from another entity,
6. A can-create function $cc : T_S \rightarrow 2^T$, $cc \subseteq T_S \times T$ such that subjects of type a can create entities of type b if and only if $cc(a, b)$ holds true.
7. A create-rule cr for each 2-tuple in cc such that for given entities A, B with types a, b , $cr(a, b)$ contains the rights to B placed in $\text{dom}(A)$ if $b \in T_O$, and which rights to A are placed in $\text{dom}(B)$ if $b \in T_S$.

For safety analysis, the function cr is of particular interest. In case of object creation, rights assignment is simply given as $cr(a, b) \subseteq \{b/r : c \bullet r : c \in R\}$, i.e. given a subject A of type a and an object B of type b , A obtains the right $B/r : c$ if and only if $b/r : c \in cr(a, b)$. Protection systems in which rights are not amplified (attenuating systems) are of particular interest; for a given state s the flow function of rights can be computed in polynomial ($\mathcal{O}(|T_S^s|^3)$) time in the number of subjects [48] by forming the transitive closure for each right, permitting the determination of a maximum state.

The general safety problem for SPM is also undecidable; however, by further restricting the SPM to acyclic attenuating instances, a decidable subset of instances can be obtained [51].

To this end, the dependency model itself must be transformed into an acyclical derivative instance, by removing the inbound edge adjacent to a fixed but arbitrary vertex in a given t -cycle.

The above description provides the means to verifying that no unintended rights transfers can occur within a given model instance. However, a separate category of rights is also required for emergency access.

Emergency access can occur in several variants; the simplest case is that of an intra-organizational emergency in which regular processes for rights transfers are revoked. Such limited access may be modeled by predetermining rights types in the schematic protection model instance and considering these rights in a separate step; a similar case can be modeled for multilateral emergency access situations (e.g. based on a predetermined contractual regulation that grants a party formal rights to invoke an emergency situation). Even so, the additional rights granted must be selected judiciously so as not to render the security analysis obsolete.

A third category, however, cannot be captured adequately from within the model, namely the exercise of sovereign power in the national interest in case of large-scale emergencies. In this case new rights would be injected into an existing model at will, rendering an *a priori* security analysis moot. Such rights injection mechanisms must therefore be modeled separately (e.g. by showing that adequate safeguards exist to prevent illegitimate entities from causing the invocation of such a governmental emergency state).

V. RELATED WORK

Complex information and control systems integrating sensors and actuators dispersed over large geographic scales, particularly ones robust to deliberate attacks on their assessment capabilities, have been developed in the context of nuclear command and control such as the SIDAC², with many of the problems encountered then still facing critical infrastructure protection – but with only a single entity in control of all assets [52].

A number of approaches have been proposed for modeling and simulation of critical infrastructures [53], [54] and vary considerably in the level of detail considered, ranging from simple dependency analyses to elaborate models containing continuous physical submodels (e.g. for pipelines and electrical grid systems) as well as behavioral models.

Among the earliest and most widespread is the application of a control systems approach [55] including hybrid mechanisms [56]. Particularly for behavioral modeling, agent-based systems have also been investigated in detail [57], [58].

In addition, metaanalyses have been conducted using techniques from reliability analyses and game theory; of particular interest (see section II) are situations where adversaries deliberately target infrastructure networks [59].

Several of the dependency graph problems outlined in section III-A can be traced back immediately to multiframe

²Single Integrated Damage Assessment Capability

problems in graph theory; for an example of an approach including time-variable flows see e.g. [60].

VI. CONCLUSIONS

The present paper has presented a brief summary of ongoing research in modeling and simulation for critical infrastructure protection. In this modeling approach, the precise specification of information sharing mechanisms regarding integrity and confidentiality was given particular attention with the goal of providing (civilian) infrastructure owners with well-defined controls over the dissemination of potentially sensitive information.

The graph-theoretical high level model along with a graph-based interoperable lower level model and information exchange format based on open standards presented provides a foundation particularly for automated analysis and preprocessing without requiring extensive modifications to data repositories and acquisition mechanisms at individual infrastructure components, thereby lowering the cost of adopting the proposed model.

Future work to be performed includes the formulation of a precise ontological model and representation schema for multiple domains and investigations into including expert knowledge of interactions and operations (e.g. for power grids) into such models.

A particular challenge, particularly for large-scale integrated models lies in constraining the model in such a way that it is suitably amenable to combinatorial optimization techniques, particularly with large number of simultaneous constraints (e.g. using interior point algorithms [61], [62], [63]).

Moreover, the inclusion of time-variable dependencies provides a particular challenge for both modeling and simulation but may not be amenable to exact computational approaches; in this case probabilistic modeling may need to be taken into consideration.

ACKNOWLEDGMENTS

The author would like to thank S. Ritter and W. Stein of the BSI (German Information Security Agency) for discussions and comments.

REFERENCES

- [1] R. T. Marsh, ed., *Critical Infrastructures: Protecting America's Infrastructures*. Washington D.C., USA: United States Government Printing Office, 1997. Report of the President's Commission on Critical Infrastructure Protection.
- [2] J. Brömmelhörster, S. Fabry, and N. Wirtz, eds., *Internationale Aktivitäten zum Schutz Kritischer Infrastrukturen*. Bonn, Germany: Bundesamt für Sicherheit in der Informationstechnik, 2004.
- [3] S. M. Rinaldi, J. P. Peerenboom, and T. K. Kelly, "Identifying, Understanding, and Analyzing Critical Infrastructure Dependencies," *IEEE Control Systems Magazine*, vol. 21, pp. 11–25, Dec. 2001.
- [4] 107th Congress of the United States of America, "Homeland Security Act of 2002." H.R. 5005., Jan. 2002.
- [5] J. Mintz, "U.S. to Keep Key Data On Infrastructure Secret." *Washington Post*, Feb. 19, 2004. Page A21.
- [6] N. G. Leveson, *Safeware: System Safety and Computers*. Reading, MA, USA: Addison-Wesley, 1995.
- [7] G. M. Koob and C. G. Lau, eds., *Foundations of Dependable Computing: Models and Frameworks for Dependable Systems*. Dordrecht, The Netherlands: Kluwer Academic Publishers, 1994.
- [8] G. M. Koob and C. G. Lau, eds., *Foundations of Dependable Computing: Paradigms for Dependable Applications*. Dordrecht, The Netherlands: Kluwer Academic Publishers, 1994.
- [9] C. E. Alchourrón, P. Gärdenfors, and D. Makinson, "On the Logic of Theory Change: Partial Meet Contraction and Revision Functions," *Journal of Symbolic Logic*, vol. 50, pp. 510–530, June 1985.
- [10] D. W. Etherington, *Reasoning with Incomplete Information*. Research Notes in Artificial Intelligence, London, UK: Pitman, 1988.
- [11] P. Gärdenfors, ed., *Belief Revision*, vol. 29 of *Cambridge Tracts in Theoretical Computer Science*. Cambridge, UK: Cambridge University Press, 1992.
- [12] G. Brewka, *Nonmonotonic Reasoning: Logical Foundations of Commonsense*, vol. 12 of *Cambridge Tracts in Theoretical Computer Science*. Cambridge, UK: Cambridge University Press, 1990.
- [13] G. Brewka, J. Dix, and K. Konolige, *Nonmonotonic Reasoning: An Overview*. Chicago, IL, USA: University of Chicago Press, 1997.
- [14] J. H. Saltzer and M. D. Schroeder, "The Protection of Information in Computer Systems," *Proceedings of the IEEE*, vol. 63, pp. 1278–1308, Sept. 1975.
- [15] L. R. Ford and D. Fulkerson, "A Simple Algorithm for Finding Maximal Network Flows," *Canadian Journal of Mathematics*, vol. 9, pp. 210–218, 1957.
- [16] L. R. Ford and D. Fulkerson, *Flows in Networks*. Princeton, NJ, USA: Princeton University Press, 1962.
- [17] V. Chvátal, *Linear Programming*. New York, NY, USA: W.H. Freeman & Co., 1983.
- [18] J. Cheriyan and S. N. Maheshwari, "Analysis of Preflow Push Algorithms for Maximum Network Flow," *SIAM Journal on Computing*, vol. 18, no. 6, pp. 1057–1086, 1989.
- [19] J. Cheriyan, T. Hagerup, and K. Mehlhorn, "An $O(n^3)$ -Time Algorithm Maximun-Flow Algorithm," *SIAM Journal on Computing*, vol. 25, no. 6, pp. 1144–1170, 1996.
- [20] A. V. Goldberg and S. Rao, "Beyond the Flow Decomposition Barrier," *Journal of the ACM*, vol. 45, no. 5, pp. 783–797, 1998.
- [21] K. Menger, "Zur allgemeinen Kurventheorie," *Fundamenta Mathematicae*, vol. 10, pp. 96–115, 1927.
- [22] R. Karp, "Reducibility Among Combinatorial Problems," in *Complexity of Computer Computations* (R. E. Miller and J. W. Thatcher, eds.), pp. 85–103, New York, NY, USA: Plenum Press, 1972.
- [23] M. Middendorf and F. Pfeiffer, "On the Complexity of the Disjoint Paths Problem," *Combinatorica*, vol. 13, no. 1, pp. 97–107, 1993.
- [24] J. Vygen, " \mathcal{NP} -Completeness of Some Edge-Disjoint Paths Problems," *Discrete Applied Mathematics*, vol. 61, no. 1, pp. 83–90, 1995.
- [25] F. Manola and E. Miller, "RDF Primer." W3C Recommendation, Feb. 2004.
- [26] G. Klyne and J. J. Carroll, "Resource Description Framework (RDF): Concepts and Abstract Syntax." W3C Recommendation, Feb. 2004.
- [27] D. Beckett, "RDF/XML Syntax Specification (Revised)." W3C Recommendation, Feb. 2004.
- [28] D. Brickley and R. V. Guha, "RDF Vocabulary Description Language 1.0: RDF Schema." W3C Recommendation, Feb. 2004.
- [29] P. Hayes, "RDF Semantics." W3C Recommendation, Feb. 2004.
- [30] J. de la Beaujardière, "Web Map Service Implementation Specification," Tech. Rep. OGC 01-068r2, Open GIS Consortium, Wayland, MA, USA, Nov. 2001.
- [31] J. D. Evans, "Web Coverage Service," Tech. Rep. OGC 03-065r6, Open GIS Consortium, Wayland, MA, USA, Aug. 2003.

- [32] P. A. Vretanos, "Web Feature Service Implementation Specification," Tech. Rep. OGC 02-058, Open GIS Consortium, Wayland, MA, USA, Sept. 2003.
- [33] J. F. Sowa, *Knowledge Representation: Logical, Philosophical, and Computational Foundations*. Pacific Grove, CA, USA: Brooks Cole Publishing, 2000.
- [34] D. L. McGuinness and F. van Harmelen, "OWL Web Ontology Language Overview." W3C Recommendation, Feb. 2004.
- [35] M. K. Smith, C. Welty, and D. L. McGuinness, "OWL Web Ontology Language Guide." W3C Recommendation, Feb. 2004.
- [36] S. Bechhofer, F. van Harmelen, J. Hendler, I. Horrocks, D. L. McGuinness, P. F. Patel-Schneider, and L. A. Stein, "OWL Web Ontology Language Reference." W3C Recommendation, Feb. 2004.
- [37] P. F. Patel-Schneider, P. Hayes, and I. Horrocks, "OWL Web Ontology Language Semantics and Abstract Syntax." W3C Recommendation, Feb. 2004.
- [38] F. Baader, D. Calvanese, D. McGuinness, D. Nardi, and P. Patel-Schneider, eds., *The Description Logic Handbook*. Cambridge, UK: Cambridge University Press, 2003.
- [39] M. A. Harrison, W. L. Ruzzo, and J. D. Ullman, "Protection in Operating Systems," *Communications of the Association for Computing Machinery*, vol. 19, pp. 461–471, Aug. 1976.
- [40] D. E. Denning, *Cryptography and Data Security*. Reading, MA, USA: Addison-Wesley, 1983.
- [41] D. E. Bell and L. J. LaPadula, "Secure Computer Systems: A Mathematical Model," tech. rep., The MITRE Corporation, Bedford, MA, USA, Nov. 1973. Volume II of MTR-2547 (also NTIS AD-771543). Performed for the Air Force Electronic Systems Division (AFSC).
- [42] D. E. Bell and L. J. LaPadula, "Secure Computer Systems: A Refinement of the Mathematical Model," tech. rep., The MITRE Corporation, Bedford, MA, USA, Apr. 1974. Volume III of MTR-2547 (also NTIS AP-780528). Performed for the Air Force Electronic Systems Division (AFSC).
- [43] D. F. Ferraiolo and D. R. Kuhn, "Role-Based Access Control," in *Proceedings 15th NIST-NCSC National Computer Security Conference*, (Baltimore, MD, USA), pp. 554–563, Oct. 1992.
- [44] D. F. Ferraiolo, R. S. Sandhu, S. Gavrila, D. R. Kuhn, and R. Chandramouli, "Proposed NIST Standard for Role-Based Access Control," *ACM Transactions on Information and System Security*, vol. 4, pp. 224–274, Aug. 2001.
- [45] N. H. Minsky, "An Operation-Control Scheme for Authorization in Computer Systems," *International Journal of Computer and Information Sciences*, vol. 7, pp. 157–191, June 1978.
- [46] N. H. Minsky, "The Principle of Attenuation of Privileges and its Ramifications," in *Foundations of Secure Computing* (R. A. DeMillo, D. P. Dobkins, A. K. Jones, and R. J. Lipton, eds.), pp. 255–276, New York, NY, USA: Academic Press, 1978.
- [47] N. H. Minsky, "Synergistic Authorization in Database Systems," in *Proceedings of the Seventh International Conference on Very Large Data Bases*, (Cannes, France), pp. 543–552, IEEE Computer Society, Sept. 1981.
- [48] R. S. Sandhu, *Design and Analysis of Protection Schemes Based on the Send-Receive Transport Mechanism*. PhD thesis, Rutgers University, New Brunswick, NJ, USA, Apr. 1983. Also available as technical report DCS-TR-130.
- [49] R. S. Sandhu, "Analysis of Acyclic Attenuation Systems for the SSR Protection Model," in *Proceedings of the 1985 IEEE Symposium on Security and Privacy (SOSP '85)*, (Oakland, CA, USA), pp. 197–206, IEEE Computer Society, Apr. 1985.
- [50] J. R. Hindley, *Basic Simple Type Theory*. Cambridge Tracts in Theoretical Computer Science, Cambridge, UK: Cambridge University Press, 1997.
- [51] R. S. Sandhu, "The Schematic Protection Model: Its Definitions and Analysis for Acyclic Attenuating Schemes," *Journal of the Association for Computing Machinery*, vol. 35, pp. 404–432, Apr. 1988.
- [52] P. Bracken, *The Command and Control of Nuclear Forces*. New Haven, CT, USA: Yale University Press, 1983.
- [53] M. Amin, "Toward Self-Healing Infrastructure Systems," *IEEE Computer*, vol. 33, pp. 44–53, Aug. 2000.
- [54] S. M. Rinaldi, "Modeling and Simulating Critical Infrastructures and Their Interdependencies," in *Proceedings of the 37th Annual Hawaii International Conference on System Sciences (HICSS'04)*, (Big Island, HI, USA), IEEE Computer Society Press, Jan. 2004.
- [55] K. Sullivan, J. C. Knight, X. Du, and S. Geist, "Information Survivability Control Systems," in *Proceedings of the 21st International Conference on Software Engineering*, (Los Angeles, CA, USA), pp. 184–192, IEEE Computer Society Press, May 1999.
- [56] J. James and F. Mabry, "Building Trustworthy Systems: Guided State Estimation as a Feasible Approach for Interpretation, Decision and Action Based on Sensor Data," in *Proceedings of the 37th Annual Hawaii International Conference on System Sciences (HICSS'04)*, (Big Island, HI, USA), IEEE Computer Society Press, Jan. 2004.
- [57] M. J. North, "Agent-Based Modeling of Complex Infrastructures," in *Proceedings of the Simulation of Social Agents: Architectures and Institutions Workshop*, (Chicago, IL, USA), pp. 239–250, Oct. 2000.
- [58] W. H. Thomas, M. J. North, C. M. Macal, and J. P. Peerenboom, "From Physics to Finances: Complex Adaptive Systems Representation of Infrastructure Interdependencies," tech. rep., Naval Surface Warfare Center Technical Digest, Dahlgren, VA, USA, 2003.
- [59] V. M. Bier and V. Abhichandani, "Optimal Allocation of Resources for Defense of Simple Series and Parallel Systems from Determined Adversaries," in *Risk-Based Decisionmaking in Water Resources X* (Y. Y. Haimes, D. A. Moser, E. Z. Stakhiv, G. I. Zisk, D. Dirickson, and B. I. Zisk, eds.), (Santa Barbara, CA), pp. 59–76, American Society of Civil Engineers, ASCE, Nov. 2002.
- [60] L. Fleischer and M. Skutella, "The Quickest Multicommodity Flow Problem," in *Proceedings of the 9th Conference on Integer Programming and Combinatorial Optimization (IPCO'02)* (W. J. Cook and A. S. Schulz, eds.), vol. 2337 of *Lecture Notes in Computer Science*, (Cambridge, MA, USA), pp. 36–53, Springer-Verlag, May 2002.
- [61] Y. Ye, *Interior Point Algorithms: Theory and Analysis*. New York, NY, USA: John Wiley & Sons, Inc., 1997.
- [62] P. M. Pardalos and M. G. C. Resende, eds., *Handbook of Applied Optimization*. Oxford, UK: Oxford University Press, 2002.
- [63] M. G. C. Resende and G. Veiga, "An Annotated Bibliography of Network Interior Point Methods," Tech. Rep. TD-5JBHXX, AT & T Laboratories Internet and Network Systems Research Center, Florham Park, NJ, USA, Jan. 2003.