

Modeling Damage Paths and Repairing Objects in Critical Infrastructure Systems

Justin Burns, Brajendra Panda, and Thanh Bui
 Computer Science and Computer Engineering Department
 University of Arkansas
 Fayetteville, AR 72701 USA
 email: {jdb083, bpanda, tbui}@uark.edu

Abstract—Recently, critical infrastructure systems have become increasingly vulnerable to attacks on their data systems. If an attacker is successful in breaching a system’s defenses, it is imperative that operations are restored to the system as quickly as possible. This research focuses on damage assessment and recovery following an attack. We review work done in both database protection and critical infrastructure protection. Then, we propose a model using a graph construction to show the cascading affects within a system after an attack. We also present an algorithm that uses our graph to compute an optimal recovery plan that prioritizes the most important damaged components first so that the vital modules of the system become functional as soon as possible. This allows for the most critical operations of a system to resume while recovery for less important components is still being performed.

Keywords—critical infrastructure; damage assessment; recovery.

I. INTRODUCTION

Critical infrastructure systems are those that are considered extremely critical to the functioning of a government or a country. As described in [1], critical infrastructures are like the vital organs of a body that need to perform their own roles for the human body to function efficiently and painlessly. The US Department of Homeland Security [2] declares that such systems are “so vital to the United States that their incapacity or destruction would have a debilitating impact on our physical or economic security or public health or safety.” Therefore, the protection and smooth functioning of our nation’s critical infrastructures are indispensable and cannot be ignored.

These systems are becoming prime targets of attackers – primarily state actors – and a major attack on one can cripple the economy of the victim nation. These systems are also more likely to be connected to the internet now to provide benefits like cost reduction (where large systems can be remotely managed over the public network), increased capability (by providing sufficient computing resources for infrastructure hardware with less capability power), and improved efficiency and transaction speed. This connectivity unfortunately makes it easier for attackers to hack into these systems. Consider the New York Times report about the attack on Colonial Pipeline [3]. While the details of the attack are not yet disclosed, a group of cybercriminals were able to compromise data systems using the internet, which resulted in Colonial Pipeline shutting down their pipeline. This outage affected mass transit and other industries across the entire

U.S. East Coast and exposed a lack of preparation for such a crisis. This illustrates how an external system can have a relationship with a critical infrastructure system and how such relationships can be exploited to carry out an attack.

It is clear from past incidents and recent reports ([4]-[7]), to cite just a few) that attacks on critical infrastructures are occurring frequently, which indicates that prevention mechanisms are not enough to stop them. Thus, it is of utmost importance to aggressively prepare for post attack activities, which include damage assessment and recovery mechanisms that are critical to making the affected systems available at full functioning mode as soon as possible. This research aims at meeting this important goal.

We propose a framework that models damage spread within a set of data objects based on object dependencies and prioritizes making repairs to the most critical objects first. The framework is based on some of the models explored in critical infrastructure protection and uses a version of previously proposed repair methods that is modified to focus on meeting specific goals when determining the order in which repairs are made.

The rest of the paper is organized as follows. Section 2 offers some work performed in this area. Section 3 defines the problem that we aim to build our model for. We provide details on our model in section 4, which includes three subsections to explain our definitions, model description, and algorithm. Section 5 concludes our work.

II. RELATED WORKS

This paper aims to examine methods and frameworks used for database and critical infrastructure protection and apply it towards protecting a set of data objects. This section describes some of the publications that are relevant to our proposed framework. One of the major works on damage assessment and recovery within a database uses data dependency to find data affected by an attack to optimize recovery [8]. While this method relies on the direct relationships between data items, an alternate model to recover data from an attack instead uses the transaction log for assessment [9].

Kotzanikolaou et. al describe a model in [10] that assists in risk assessment for possible scenarios that can result in cascading failures within a CI system. For critical infrastructures with data-rich operations, the use of Cyber-Physical Systems can cause new vulnerabilities as described in [11]. Their model analyzes threats that can appear due to

these vulnerabilities and analyzes the potential cascading damage they can cause. System dynamics modeling can also be used to analyze disruptive events to characterize such disruptions to critical infrastructure by risk assessment and various impact factors as shown in [12].

Rehak et. al [13] model an infrastructure system as elements and linkages with different types of relationships establishing dependencies and interdependencies. They note that these elements can have varying criticality, causing some elements to cascade more damage into the system than others in the event of a failure. This work is important because by establishing criticality, they quantify damage within a system. We use this concept of criticality later in this paper to direct the optimal repair path of data objects.

We also consider models that assist with recovery during an attack. In [14], an algorithm is proposed to restore damaged element paths by recursively breaking down demand flows into simpler problems. They use a centrality metric to rank damaged nodes and determine which ones should be repaired first and expand on the use of centrality to make repair decisions in further work [15]. We use the concept of centrality to rank data objects in a case where two or more are equally critical. In our algorithm, we also utilize their method of simplifying damage paths to find the fastest route to restoring intermediate data objects. However, the novelty of our approach is twofold: we must repair all components within the system because data objects cannot have computations rerouted, unlike the network components in the work we have reviewed, and we aim to restore the most important components first so that their functions can be restored while repairs to the system are still ongoing.

III. PROBLEM DEFINITIONS

In the occasion when an adversary information attack succeeds, the victim must have the capability to degrade gracefully and recover damaged data and/or services in real-time if it is to survive. It is necessary to immediately carry out damage assessment and recovery process in order to bring the systems to working states. Otherwise, the damage would spread to other unaffected systems that are interconnected. This happens when a valid user or an unaffected system module reads a damaged object during its computation and updates another object based on the compromised value, causing the latter damaged as well. As time goes on, more and more objects become affected in this manner causing the spread of damage to fan out through the system quickly.

For damage assessment and recovery purpose, information about all processes that have been executed must be stored in the log (more on this presented later). This will help in determining the relationships among the processes, thus helping in establishing the damage trail. Moreover, during recovery, the operations of processes that have spread the damage have to be undone and then redone in order to produce correct states of affected objects. The problems with existing systems are: (1) They do not store process execution information in the log, and they purge the log periodically,

(2) Their recovery mechanisms are not designed to undo the effects of executed processes, (3) The size of the log, as it must not be purged, will make it almost impossible to continue the recovery process in real-time, and (4) During the damage assessment and recovery process, the system remains unavailable to users. This delay induces a denial-of-service attack, which is highly undesirable in time-critical applications that the critical infrastructures are designed to provide. Due to massive amount of data in the log that needs to be processed, the problem becomes even worse.

The goal of this research is to develop fast, accurate, and efficient damage assessment and recovery techniques so that critical information systems not only survive the attacks gracefully but will continue to operate providing as many vital services and functions as possible even before the system is fully recovered. In the next section, we explain how our model can accomplish this.

IV. THE MODEL

In this section, we describe our model in detail. The first subsection defines important graphs and metrics that we use for our model. In the next subsection, we describe how the model is built and is used to determine an optimal recovery plan. Finally, we describe the algorithm we use to implement our model.


A. Definitions

We first define the concept of information flow in a system. This also defines dependencies among various objects in the system and is used in our graph-based model.

Definition 1: Given two objects O_i and O_j in a system, if the value of O_j is calculated using the value of O_i , we say that there is information flow from O_i to O_j . Thus, O_j is said to be dependent on O_i and is denoted as $O_i \rightarrow O_j$.

The above definition helps in determining the spread of damage in the system. That is, if an object is damaged, then all its dependent objects will be considered damaged. During recovery, the parent (pre-cursor) object must be recovered before any of its dependent objects can be recovered.

Next, we define a graph containing the set of objects and all possible paths among them. We call it Possible Paths graph and it spans the entire system of objects and all dependency paths among them. An example of this graph is shown by Figure 1(a).



Definition 2: Consider a system containing the set of objects O . The **Possible Paths Graph (PPG)** is built by having a node N_i for each object in O . There exists an edge E_{ij} from N_i to N_j in the PPG if there is a possibility that information may flow from N_i to N_j , that is, N_j may be modified based on the value of N_i .

The purpose of building a PPG is that it will help during the damage assessment preparation phase. By assuming the point of attack one can identify the set of items that may be affected consequently. Thus, security officers can be prepared for different types of eventuality.

The second set of objects contains the actual paths that were used to make changes in the system within a specified time span, which for the purposes of the third graph that will be defined, is usually the time passed since an object has been damaged. This set is represented by the Active Paths Graph (APG), and all objects and dependencies in this set exist in the PPG. This graph will help in determining the damage flow in case of an attack. Given an initial attack point (an object), one can determine which objects in the system may be affected by the attack and which ones will not be. Therefore, the ability of the system to carry out its intended functions can be calculated. That is, during the recovery process, the set of damaged objects will be made unavailable while the rest can be made accessible. Knowing which objects will remain unaffected, one will be able to identify what services the system will be able to offer while the recovery continues.

Definition 3: The *Active Paths Graph (APG)* contains nodes N and edges E such that for every $N_i \in N$ and every $E_{ij} \in E$, both N_i and E_{ij} are also present in PPG, and E_{ij} illustrates an actual information flow; that is N_j was updated based on the value of N_i .




Figure 1b. The Active Paths Graph (APG)

Figure 1(b) provides an example of an Active Paths Graph and as can be seen it is a sub-graph of Figure 1(a). As discussed before, once an initial attack point is determined, the APG will help in accurately determining the damage flow and the set of objects affected by the attack. As discussed before, as time goes on, more and more objects will be affected as new objects will be updated based on the value of an affected object. Thus, to stop the spread of damage, all affected objects must be quickly identified and taken offline as soon as possible. This can be achieved by doing a flow assessment using the APG. This leads to the concept of actual damage spread path showing exactly which objects were affected by an attack. If a system is damaged, we represent the spread of damage as a third set of objects, the Damage Spread Graph (DSG). The set of objects and dependencies in this graph must exist within the APG, as damage spread occurs when objects make changes based on their dependencies. Like how the APG is a subsection of the PPG, the DSG is a subsection of the APG. Figure 1(c) is an example of what a damage path may look like. It is important to note that over time, a damaged object will always cascade its damage down to dependent nodes included in the APG. Definition 4 formally defines the DSG.

Definition 4: A *Damage Spread Graph (DSG)* contains nodes N and edges E such that for every $N_i \in N$ and every $E_{ij} \in E$, both N_i and E_{ij} are also present in APG and every node in N is damaged through an attack on the system.

Moreover, an edge E_{ij} depicts that N_i was damaged first and then N_j was damaged through the flow of information from N_i to N_j .




Figure 1c: The Damage Spread Graph (DSG)

Note that the edges between two objects may be bidirectional or recursive. For example, if an object O_j can have a dependency on object O_i and vice versa, then there will be a bidirectional edge between O_j and O_i . Similarly, if an object can be dependent on itself, it will result in a recursive graph. To clarify, let us consider an object “salary”. When an employee receives an increment that is based on a percentage of the current salary of the employee, it causes the new salary to be dependent on the old salary and is depicted by using an edge from salary to salary itself. However, it must be noted that, for simplicity, we use neither bidirectional nor recursive edges in APG or DSG. Rather, when an object is modified, we note that as a new version of the object, thus creating a new node for the object with the version number.

To minimize the time needed to restore the most important objects within a system of object dependencies, we also define criteria used to determine the order in which repairs are made:

Definition 5: The *criticality* of a node N is its predetermined level of importance to the system’s functions. This must be predetermined for the flexibility of the model to fit various systems and align the model with the goals of each specific system. For example, one system may need to prioritize certain components that other systems do not. The criticality of a component can be measured by various characteristics such as the intensity or scope of an impact caused by its failure as described in [13].

We assign a positive whole number to each node N to represent criticality. A lower assigned value indicates higher criticality. For example, a node N_i with a criticality of 2 would be considered more important than a node N_j with a criticality of 4. It is important to note that criticality values are not unique, meaning multiple nodes can have the same criticality value. When that happens, we use the following metric in the next definition to serve as a first “tiebreaker”.

Definition 6: Objects that have more damaged dependencies take longer to repair. Therefore, the *repair time* of a node N is defined as how many inward-flowing edges E^i it is receiving damage from.

When two or more objects are assigned the same importance, we choose to first repair the one that has a lower repair time. For example, consider two nodes N_i and N_j that are equally critical. If N_i needs 5 other nodes repaired to

repair it, and N_j needs 3 other nodes to repair it, then we will repair N_j first, because its operation can be restored more quickly than that of N_i .




Figure 2a. A parent node with high centrality



Figure 2b. A parent node with low centrality

Definition 7: The *centrality* of a node N is the number of outward-flowing edges E^o it has.

We use the above metric to decide the next object to repair when two or more are equal in both criticality and repair time. An object with a higher number of E^o will have higher centrality. Figure 2a and Figure 2b show two subsections of a DSG that highlights centrality. As shown in Figure 2a, N_4 has three nodes that are dependent on it: N_1 , N_2 , and N_3 , while as Figure 2b depicts, N_6 only has a single node N_5 dependent on it. Assume that the repair algorithm has repaired the parent node(s) of N_4 and that of N_6 . To clarify the situation, N_4 and N_6 need not have the same parents; it is just that both are in line to be repaired next. In this scenario, repairing N_4 before N_6 reduces the repair time for the three dependent nodes of N_4 instead of only one of N_6 , which can make future repairs be performed faster. Therefore, N_4 is considered to have a higher centrality than N_6 .

B. Model Description

The model uses the three graphs defined in the previous section to construct a representation of a given system and its sustained damage from the time of the initial attack. The PPG is a preprocessed map of all components and dependency paths within a system. We assume that we know how much time has passed since the initial attack and build the APG by including components and dependency paths that were used in a transaction log in that period. By knowing the component where the initial attack occurred, we build the DSG by tracing the damage through the transaction log. For damage to spread from one component to the next, it must follow two criteria: 1) there is a damaged node N_i that has an edge E_{ij} flowing from it to node N_j and 2) E_{ij} is used for a transaction while N_i is damaged. For the DSG to exist, the initial attack must occur within the APG, otherwise there is no cascading damage.

The goal of the model is to find the optimal sequence of repairs to restore the most important operations of a system as quickly as possible. We use the metrics defined in the previous section to decide which components should be repaired first. The first metric is criticality – the most critical components must be restored first to resume important operations. However, these components may also be dependent on other components that are damaged. These components must be repaired first before the base component can be repaired. At this point, the same problem is applied to the dependency components, and the most critical one is chosen first. If there is a tie, then components with a lower

repair time are picked first. For example, a component that has two damaged parent components will be prioritized over a component with three or more damaged parent components if both components are equally critical.




Figure 3. Recovery sequence decision

To clarify, let us consider the graph presented in Figure 3. As shown in the figure, nodes N_1 , N_2 , and N_3 are dependent on N_4 . Assume that the damage assessment method identified N_4 as damaged; thus, nodes N_1 , N_2 , and N_3 are also identified as damaged. During the recovery process, N_4 was recovered before the other three nodes. However, since it has three dependents all of which are damaged, the question is, which one should be repaired first. As our goal is to have the vital functions of the system to be made available before the other operations, our algorithm would choose the node among N_1 , N_2 , and N_3 having the most criticality.

C. The Algorithm

First, we discuss the primary objective of our work. Let us consider the notations used in the following table:

TABLE I. NOTATIONS

Notations	Descriptions
$P = (V, E)$	Possible Path Graph
$A = (V_A, E_A)$	Active Path Graph ($V_A \subseteq V, E_A \subseteq E$)
$D = (V_D, E_D)$	Damage Spread Graph ($V_D \subseteq V_A, E_D \subseteq E_A$)
$D = (V_C, E_C)$	Critical Node Graph ($V_C \subseteq V_D, E_C \subseteq E_D$)
δ_{ij}	Decision to fix edge i to j
δ_i	Decision to fix node i
t_i	Time to fix node i
c_i	Centrality of node i
P_{ij}	Dependency indicator of node i and j

Our objective is to find $\min \sum_{i \in V_D} t_i \delta_i$ subject to

$$\delta_i \sum_{j \in V_C} P_{ij} \leq \sum_{j \in V_C} P_{ij} \delta_j \quad \forall i, j \in V_C \quad (1)$$

$$\delta_i c_i \geq \sum_{(i,j) \in E_C} \delta_{ij} \quad \forall i \in V_C \quad (2)$$

$$P_{ij} \in \{0,1\} \quad \forall i, j \in V_C \quad (3)$$

$$\delta_i, \delta_{ij} \in \{0,1\} \quad \forall i \in V_C, (i,j) \in E_C \quad (4)$$

That is, the goal is to minimize the time required to fix all critical nodes subjected to conditional constraints of the system. To make sure that each preceding nodes of i are fixed before node i being processed, condition (1) is used. For example, if there is a node j connecting to i but in a prequel order, the sum product of all nodes j status and dependency indicator P_{ij} should be greater or equal than the product of sum of all dependency indicator P_{ij} with node i . To make sure that there would not be more out-going flows than the given capability of node i , equation (2) is imposed to make sure the total out-going edge would not surpass the centrality of node i . Conditions (3) and (4) were built to impose the

binary attribute of the dependency indicator P_{ij} , the decision whether to fix node i or edge from node i to node j .

The algorithms provided in this section use the model described in the previous section to compute the optimal order of repairs to restore the most important functions of a system first. When an attack occurs, we expect an Intrusion Detection System (IDS) to identify the attack and provide the initial point of damage. The working principles of IDSs are not within the scope of this work and so, not described here.

After receiving notification from an IDS, a precise damage assessment is performed. If the damage assessment process is unable to make accurate assessment, i.e., in case a damaged node is not correctly identified, it and its dependent nodes, which are also damaged, will remain unrecovered. This will result in valid users or procedures reading them and spreading damage by updating other objects, as discussed earlier. For a detailed discussion on damage assessment, one may review [8] and [9], which were developed particularly for database systems. However, the methods are still applicable to critical infrastructure systems. Below we provide a basic mechanism to carry out the assessment.

Damage assessment begins with the APG, which shows the actual dependency relationships among the objects in the system (Note that the APG can be built as transactions are executed and dependencies are established among various nodes of the PPG). Given the initial attack point, the corresponding node is then marked as damaged. This is the starting node of the DSG. Then by scanning the log from the corresponding location of the attack point, transactions that read the marked node are identified. Any objects written by those transactions are then marked as damaged in the APG. This process continues until the end of the log. Finally, all unmarked nodes and the edges showing their dependencies are removed. The resulting graph is the completed DSG.

Once damage assessment is carried out, recovery procedure must begin immediately in order to make the system operational quickly. We use Algorithm 1 as the main procedure to initialize an object set for repairs. The algorithm starts by initializing the set of damaged objects O . Each node N within O consists of a system component and its relationships with other nodes in O . As mentioned previously under Definition 4, some system components may have recursive or bidirectional dependencies between each other. Therefore, system components can have repeat nodes within O to represent their different versions. Each node is assigned values for criticality, repair time, and centrality. Using those metrics, the algorithm determines an initial target node N_0 based on criticality. If there are two or more nodes with the highest criticality, then the node with the lower repair time is selected. In the event of another tie, the node with higher centrality is selected. Further ties are broken by random selection. N_0 , along with O and the repair queue Q , are used to make the first call to the recursive function Algorithm 1.1 at step 4.5. Algorithm 1 proceeds until O is completely empty, and then the repair queue is finalized, and Q is printed.

As previously discussed, a node must have its parent nodes repaired before it can be considered eligible for repairs. Algorithm 1.1 ensures that nodes are scheduled for repairs in the proper order while still adhering to the rules set for determining priority. It does this by using a while loop to check the currently selected node N for repair eligibility. If N is eligible for repairs, then it is removed from O and Q is updated, then returned. If N is not eligible, then O' , a subsection of O made up of all dependency paths above the currently selected node is created and used to find the next highest priority node N' within O' . Algorithm 1.1 is recursively called using N' and O' , which can either result in the node's repair or another node being selected for repair again. The recursive nature of this algorithm ensures that each time a decision needs to be made on which node needs to be repaired next, it will prioritize criticality and efficiency among all the nodes that can be repaired at any given step. In this way, the bulk of the work done by the algorithm is choosing the next object for repair within each iteration. Each function call will result in one object being repaired and $n - 1$ additional function calls, where n is the number of nodes within the set of nodes being passed. Since repaired objects need to be removed from the DSG, function calls will need to update and return the global DSG and Q .

Algorithm 1: Initialization for object set repair

Result: Queue of objects ordered by repair priority

- 1 Initialize set of damaged objects O
- 2 Preprocess object priority using criticality, repair time, and centrality
- 3 Initialize repair queue Q
- 4 while O has damaged nodes remaining
 - 4.1 Select the highest critical node(s) N within O
 - 4.2 if Two or more nodes are tied for highest criticality
 - 4.2.1 Select the node(s) N with the lowest repair time R within O
 - 4.3 if Two or more nodes are tied for lowest repair time
 - 4.3.1 Select the node(s) N with the highest centrality within O
 - 4.4 if Two or more nodes are tied for highest centrality
 - 4.4.1 Select a single node at random from those still tied
 - 4.5 Update repair queue(N_0, O, Q) $\rightarrow Q$
- 5 Print Q

Algorithm 1.1: Recursive repair function

Result: Schedules a node N for repairs and returns the updated repair queue Q

- 1 Update repair queue(*Selected node N , object set O , repair queue Q*):
- 2 while *Current object has unrepaired dependencies*:
 - 2.1 Create subset of damaged nodes O' of all nodes N' and edges E' that N is dependent on
 - 2.2 Select the highest critical node(s) N' within O'
 - 2.3 if Two or more nodes are tied for highest criticality
 - 2.3.1 Select the node(s) N' with the lowest

- repair time R within O
- 2.4 if Two or more nodes are tied for lowest repair time
 - 2.4.1 Select the node(s) N' with the highest centrality within O
 - 2.5 if Two or more nodes are tied for highest centrality
 - 2.5.1 Select a single node at random from those still tied
 - 2.6 Update repair queue(N', O', Q) $\rightarrow Q$
 - 2.7 Remove the most recent object in repair queue from O
- 3 Repair N
 - 4 Add N to Q
 - 5 Return Q

The algorithm produces a list of system nodes in the order in which they should be repaired. Recovery procedure then continues to the next step to begin repairs on the system. It is important to note that while repairs are simulated by the algorithm, the process for repairing the actual components of the system is not within the scope of this work.

V. CONCLUSION

In this research, we have presented a method to repair data objects that prioritizes quick recovery for the most important components of a system. This allows for the partial restoration of functions during the recovery process with an emphasis on restoring service to the most necessary functions. This was first done by building out three graphs to represent the entire system, what changes the system made after an attack, and the cascading damage as a result of those changes. Next, we developed an algorithm to optimally schedule repairs by using those graphs to find damage paths that affect the most critical nodes of a system and calculate the fastest repair order to fully restore those nodes. Our work is most applicable to protecting critical infrastructure systems where services need to be restored as quickly as possible to avoid economic or societal disruptions.

Further work includes considering the frequency at which an object is used to update its dependencies. Objects that are updated at a higher frequency would be prioritized as more important. A method to select the order of repairs for non-critical objects after all critical objects have been repaired is also needed. Finally, a performance analysis of this model is required to be carried out to evaluate the model under various conditions.

ACKNOWLEDGEMENT

This work has been supported in part by grant H98230-20-1-0419 issued by the National Security Agency as part of the National Centers of Academic Excellence in Cybersecurity's mission to expand cybersecurity research and education for the Nation.

REFERENCES

- [1] E. Viganò, M. Loi, and E. Yaghmaei, "Cybersecurity of Critical Infrastructure", In Christen M., Gordijn B., Loi M.

- (eds), *The Ethics of Cybersecurity*, The International Library of Ethics, Law and Technology, vol 21, Springer
- [2] *Critical Infrastructure Security*: <https://www.dhs.gov/topic/critical-infrastructure-security>. [retrieved: October 2021]
- [3] D. E. Sanger and N. Perloth, (2021, May 14). "Pipeline Attack Yields Urgent Lessons About U.S. Cybersecurity", <https://www.nytimes.com/2021/05/14/us/politics/pipeline-hack.html>. [retrieved: October 2021]
- [4] A. Anastasios, "Is the Electric Grid Ready to Respond to Increased Cyber Threats?", <https://www.tripwire.com/state-of-security/ics-security/electric-grid-ready-increased-cyber-threats/>. [retrieved: October 2021]
- [5] B. Barrett, "An Unprecedented Cyberattack Hit US Power Utilities", <https://www.wired.com/story/power-grid-cyberattack-facebook-phone-numbers-security-news/>. [retrieved: October 2021]
- [6] K. O'Flaherty, "U.S. Government Issues Powerful Cyberattack Warning As Gas Pipeline Forced Into Two Day Shut Down <https://www.forbes.com/sites/kateoflahertyuk/2020/02/19/us-government-issues-powerful-cyberattack-warning-as-gas-pipeline-forced-into-two-day-shut-down/#5f3061645a95>". [retrieved: October 2021]
- [7] M. Lewis, "Cyberattack Forces Gas Pipeline Shutdown", <https://www.jdsupra.com/legalnews/cyberattack-forces-gas-pipeline-shutdown-76217/> [retrieved: October 2021]
- [8] B. Panda and J. Giordano, (1999) Reconstructing the Database After Electronic Attacks. In: Jajodia S. (eds) *Database Security XII. IFIP — The International Federation for Information Processing*, vol 14. Springer, Boston, MA.
- [9] S. Patnaik and B. Panda, (2003). Transaction-Relationship Oriented Log Division for Data Recovery from Information Attacks. *Journal of Database Management*, 14(2), pp. 27-41.
- [10] P. Kotzanikolaou, M. Theoharidou, and D. Gritzalis, (2013) Cascading Effects of Common-Cause Failures in Critical Infrastructures. In: J. Butts and S. Sheno (eds) *Critical Infrastructure Protection VII. ICCIP 2013. IFIP Advances in Information and Communication Technology*, vol 417. Springer, Berlin, Heidelberg.
- [11] J. Ding, Y. Atif, S. Andler, B. Lindström, and M. Jeusfeld, (2017). CPS-based Threat Modeling for Critical Infrastructure Protection. *ACM SIGMETRICS Performance Evaluation Review*. 45. pp. 129-132. 10.1145/3152042.3152080.
- [12] E. Canzani, H. Kaufmann, and U. Lechner, (2016). Characterising Disruptive Events to Model Cascade Failures in Critical Infrastructures. 10.14236/ewic/ICS2016.11.
- [13] D. Rehak, J. Markuci, M. Hromada, and K. Barcova, "Quantitative evaluation of the synergistic effects of failures in a critical infrastructure system", *International Journal of Critical Infrastructure Protection*, Volume 14, 2016, pp. 3-17, ISSN 1874-5482
- [14] N. Bartolini, S. Ciavarella, T. F. La Porta, and S. Silvestri, "Network Recovery After Massive Failures," 2016 46th Annual IEEE/IFIP International Conference on Dependable Systems and Networks (DSN), 2016, pp. 97-108
- [15] S. Ciavarella, N. Bartolini, H. Khamfroush, and T. Porta, (2017). "Progressive damage assessment and network recovery after massive failures," *IEEE INFOCOM 2017 – IEEE Conference on Computer Communications*, 2017, pp. 1-9.