A Guideline on the Analysis of Stochastic Interdependencies in Critical Infrastructures

Sandra König, Thomas Grafenauer,
Stefan Schauer, Manuel Warum
Austrian Institute of Technology GmbH
Digital Safety & Security Department
Vienna, Austria
{sandra.koenig, thomas.grafenauer,
stefan.schauer, manuel.warum}@ait.ac.at

Stefan Rass
Universität Klagenfurt
Institute of Applied Informatics
Klagenfurt, Austria
stefan.rass@aau.at

Abstract—**Protecting Critical Infrastructures (CIs) requires decisions made about systems with complex dynamics, which rarely admits accurate descriptions or precise predictions. For this reason, simulation models are often probabilistic, embodying known (physical) laws to the extent possible, but generally adding a random element to account for unexpected events that security management is primarily concerned with. One such complex element is the interplay between different components of a CI, i.e., the dynamic *inside* a CI, which is more than just the sum of the mutual dependencies. Another important factor is the interplay *between* CIs, which might be understood between two specific CIs but less well known when it comes to mutual impacts. Simulations can help assessing these dependencies, but only to the extent as they are accurately specifiable. This work addresses the practical issues of using a probabilistic model to simulate cascading effects in interdependent CIs by proposing methods to allow for specifications carrying subjective uncertainty. The description follows a running example of a fictitious water provider, where a stochastic simulation model of incident propagation is embedded into its existing risk management process. Our exposition runs up to the final question of the decision maker about where to take action and how to prioritize assets regarding their need for protection, but also their role in impact propagation. The final picture delivered by the method outlined here is meant as a support for risk management decisions, containing possible concrete scenarios in an aggregate form. The value for a decision maker is the revelation of previously unseen influences and impacts besides the known causes and threats being subject of the risk management. This paper demonstrates how a stochastic model of dependencies between CIs can be integrated in a standard risk management process and illustrates each step for the case of a fictitious water provider.**

Keywords-critical infrastructure; dependencies; risk management; water supply

I. Introduction

Critical Infrastructures (CIs) are essential pillars of today's society that relies on availability of water, power, health care and transportation but also on the availability of food. Due of the high impact of a failure of even one of these infrastructures on society, a lot of research focuses on investigation of CIs. Of particular interest are the various interdependencies between CIs as these increase in number and type. For example, a hospital nowadays does not only depend on electricity, water and food supply, and a well functioning transportation system but also on information and control systems for intensive care or surgery. Even worse, interdependencies may amplify consequences of an reduced availability of one CI due to cascading effects. Such effects need to be taken into account when conducting a risk analysis [1]. While the local protection of a CI is possible using respective domain knowledge, securing the compound of several interacting CIs requires cross-domain expertise that is hardly available to the expert(s) in charge. Thus, to understand wide-range impacts cascading over several interdependent CIs, simulations are an indispensable tool to discover scenarios that require an orchestrated defence involving a collaboration of security officers in several distinct CI. Our model meets this need by letting each domain expert describe its own local CI, and leaving the interdependencies between two CIs as a matter of two domain experts agreeing on how their individual CIs interact with and depend on one another.

Since risk analysis is only one step in a more comprehensive risk management process, we here illustrate how such an advanced risk analysis can be integrated in an existing risk management process with the aim of yielding more accurate results. Our analysis uses the simulation tool described in [2]. The paper gives a step-by-step description of how to integrate a mathematical model into risk management processes and illustrates the procedure with an example. Practical aspects such as the use of expert opinions are discussed.

Related Work

Interdependencies between CIs have increased during the past decade and turned formerly loosely dependent CIs into a complex and highly interconnected network of CIs. The increasing complexity gave rise to numerous models of the dynamics inside a CI and between CIs. Early methods to describe those dynamics include Hierarchical Holographic Modeling (HHM) [3], followed by a multi-graph model for random failures [4] or input-output models [5]. However, most of these methods do not pay enough attention to nowadays interdependencies that yield to manifold effects of a single incident. The unpredictability of consequences shifted the focus towards stochastic models. While Markov models are popular due to their simple structure, the applicability is often limited due to the exponentially growing state space. Models trying to cope with this issue allow for memory [6] but are challenging to put to practice due to their high complexity. The Interdependent Markov Chain (IDMC) model describes cascading failures in interdependent infrastructures in power systems [7], where every infrastructure is described by one discrete-time Markov chain where interdependencies between these chains are represented by dependent transition probabilities. A stochastic model allowing different degrees of failure while still being relatively simple to implement has been introduced in [8]. This paper extends previous work in the filed of cascading effects in interconnected networks [8], [9] but is also similar to approaches in IT security such as [10].

Incidents such as the disruption of electric power in California in 2001 [11], a power outage in Italy in 2003 [12] or failure of the nuclear plant in Fukushima, Japan, have demonstrated that interdependencies between various systems exist of which even experts were not aware. Intentional attacks such as the hacking of the Ukrainian power grid in 2015 [13], the Stuxnet worm [14] or the WannaCry ransomware (that hit hospitals particularly hard [15]) demonstrated the vulnerability of CIs due to the growing digitalization. Awareness of vulnerability due to cyberattacks has increased after recent attacks such as botnets [16] that nowadays also focus on critical information infrastructures [17]. Despite these well-known events, data is sparse and not sufficient to enable statistical analysis. Instead, simulation of such events and discussions with experts are needed to investigate consequences of incidents. The incidents of interest are both natural events (such as natural disasters) and man-made events, including unwanted interventions like cyberattacks or human error. Especially cyberattacks have recently moved into the center of attention and the EU Directive 2016/1148 on cyber security (also called the NIS-directive) describes regulations to increase the protection of CIs [18]. Simulating the consequences of an incident only requires knowledge about propagation dynamics and not about the type of incident (i.e., about the trigger) but even this information is typically difficult to get due to missing experience. Simulation methods are available to some extent, e.g., [19], and allow comparison of different models for specific situations. Motivated by the consequences of recent incidents, there is a growing interest in resilience of critical infrastructures [20]. An overview on models on interdependent CIs is presented in [21], while [22] gives an extensive review and comparison of different models of cascading effects in power systems. A general review on interdependencies between infrastructures with a focus on the different types of dependencies is given in [23]. The amount of research focusing on water supply and water providers seems to be more limited. A study focusing on security weaknesses of Industrial Control Systems (ICSs) and Supervisory Control and Data Acquisition (SCADA) systems and how to find good practices for water providers can be found in [24]. So far, there is only limited research focusing on incidents affecting a water provider. The impact of an Advanced Persistent Threat (APT) on a water provider has been investigated in [25] and [26].

In the following we demonstrate how to take interdependencies into account when analysing a critical infrastructure. The method is described for any critical infrastructure, as defined according to the European Commission [27].

> ‘Critical infrastructure’ means an asset, system or part thereof located in Member States which is essential for the maintenance of vital societal functions, health, safety, security, economic or social well-being of people, and the disruption or destruction of which would have a significant impact in a Member State as a result of the failure to maintain those functions.

The approach is illustrated focussing on a water provider.

Paper Outline

Since this article aims at illustrating the use of a theoretical model in practice, the remainder of this paper illustrates the various steps in detail for a example CI. Section II describes step by step how to integrate the analysis of stochastic dependencies between CIs in a risk management process. Section III then illustrates how the risk management process can be implemented for a fictitious water provider. Section IV provides concluding remarks and points out some directions of future work.

II. Incorporating Stochastic Dependencies in a Risk Management Process

This section illustrates how stochastic interdependencies between CIs can be incorporated in a standard risk management process. The process follows the ISO 31000 framework [28] and provides a step by step guide whose application will be illustrated by the analysis of a fictitious water provider in the following section. Our focus lies on the risk analysis (Step C below) as this is where the interdependencies between CIs have the biggest effect.

The upcoming analysis is based on modelling a CI as a set of interdependent *assets* (i.e., the relevant components of the CI) as a directed graph whose nodes represent the assets and edges represent the dependencies between them. Each asset carries the following information:

- *Criticality*: How important is the asset for the overall functionality of the CI?
- *Dependencies*: How critical is the asset for the functionality of other related assets?
- *Status indicator*: What is the level of functionality of each asset?

As for the last question, we here use different *states* on the scale $\{1, 2, 3\}$ to express increasing degrees of affection, ranging from status “working” (represented by value 1) up to the worst case status “outage” (represented by value 3), where the intermediate status level corresponds to limited functionality. More granular scales are possible but these three states are enough for illustration purposes.

Remark 1: It is important to note that we use the general term "asset" here, as in the standard risk management literature, when investigating risks *within* a CI. When interested in dependencies *between* CIs, e.g., the interplay between power providers, hospitals and water suppliers, a high-level perspective may see each CI as an asset and apply the model of stochastic dependencies. In this work, we focus on a single CI's situation and will thus hereafter use the term asset to describe a component of a CI.

A. Establishing the context

As a very first step it is necessary to understand the situation at hand. This includes both a description of how the water provider works internally as well as a deeper understanding of the overall context, i.e., dependencies to other CIs that provide input or require input for their part to ensure smooth operation. Dependencies between CIs are manifold and require a thorough analysis. In particular, dependencies are in no way limited to visible physical and known cyber connections but the analysis should also take into account logical interdependencies between different parts, as in the case of a control system.

A useful way to obtain an overview of the situation (and to discover potential missing dependencies) is visualization through a graph. To this end, a full list of components of the infrastructure as well as a full list of providers they depend on is represented in a network model. In a large network it may be useful (or even necessary) to classify dependencies according to their properties and only assign values to every class of connection, such as assigning one of the types "water", "communication" or "electricity" to every connection. This allows for distinction of different relationships but at the same time avoids an excessive amount of assessments.

Finally, context and focus of the risk management process are determined, defining which parts of the organization is covered by the analysis (which assets are relevant) and which criteria are used to evaluate the significance of risks, but also answering organizational questions such responsibilities and resources for the upcoming analysis.

B. Risk Identification

Next, it is important to identify the relevant risks. It is useful to distinguish the terms "threat", "vulnerability" and "risk" in the following: a *threat* is any factor or condition that can impact the correct functionality or security of a system. A *vulnerability* is any condition or property of a system that can lead to affection by a threat. A *risk* is the coincidence of threat and vulnerability. A security incident is then the physical event of a threat that hits some vulnerability and thereby causes an impact on the system. Quantitatively, risk is understood as the product of impact and the likelihood for the impact to occur.

In order to get a comprehensive overview, several sources need to be taken into account. General technical vulnerabilities are collected in databases such as the National Vulnerability Database (NVD) [29] while specific vulnerabilities of software and hardware components may be detected by use of automated vulnerability scanners such as Nessus [30] or OpenVAS [31]. Historic data help identify threats specific to the CI and discussions with experts form the field help understand which of these threats are actually risks. This step yields a list of the relevant risks that are analysed in the following.

C. Risk Analysis

This step is about getting a deeper understanding of the risks identified in the last step. Which assets are directly affected by each risk? What are the indirect consequences? How likely is each risk to occur?

Reports on past incidents are a good source to enhance understanding of risks. However, such data is not always available, either due to the rare occurrence or due to the fact that only necessary information is reported.

Consequences of an Incident

This step aims at estimating the consequences of an incident, i.e. of a realisation of each of the risks considered. To this end we apply the stochastic model introduced in [8] that also allows simulation. The simulation assumes that a certain incident has just occurred and directly affects one or more assets by putting them from functional into affected or even outage state. Based on the dependency information the status of related (dependent) assets is updated accordingly, where each asset may individually undergo different status changes, depending on the importance of the other asset (e.g., a mild affection may occur if the failed asset provides only a small part of the supply, or a severe affection may occur if an asset vitally depends on another yet failed asset). This reveals *cascading effects*, i.e., indirect impacts of a realisation of a risk scenario that are not evident at first sight. State transitions are supposed to happen probabilistically to cover cases of deterministic dependencies (e.g., a pump is fully relying on a continuous electricity supply) as well as probabilistic dependencies (e.g., water shortage can temporarily be overcome by backup water reservoirs). Application of the model [8] asks experts to provide assessments for the identified risks and to discuss the consequences of a realisation of each of these risks. While the model describes the propagation mechanisms and provides an estimate of the overall impact of an incident on a CI it requires knowledge about the effect of a failure of one single component on the ones directly depending on it. The core duty of the modeling then boils down into two major tasks:

1) Enumerate all assets and identify interdependencies between them as detailed as possible. In the following, we use the arrow notation $A \to B$ to denote a dependency of asset B on asset A (cf. Figure 1, e.g., where the pump B depends on the water A). A directed graph may be used to illustrate the situation.
2) Based on this information, specify probabilities for state changes in a dependent asset B, if the provider asset A is not working properly (i.e., is in state 2 or 3).

The first step typically collects information that is known and available to the CI operator. The challenging part is specification of the transition probabilities in the second step. This is a general issue in any probabilistic model (i.e., not specific to the one applied here), together with the occurring costs to the CI operator in terms of human resources.

The sought transition probabilities describe how likely it is that limited functionality or a complete shutdown of one component affects the dependent components. While historic data may provide some information on these transitions such data is rarely (publicly) available so that expert knowledge is the only remaining source. Despite the human resource cost to the CI this source is of high value since experienced employee often have a profound knowledge that is not available in written form. Still, experts may find it hard to provide precise estimates of likelihoods but rather have an idea about what is likely to happen next (based on past incidents). Aware of this problem, we avoid asking for precise numerical values but rather look for an assessment on a qualitative scale, as recommended in risk management (e.g., by the German Federal Office for Information Security (BSI) [32]). One way to address this problem is to ask for a qualitative prediction combined with a statement on the confidence of this estimate [1]. According to this approach we ask experts to answer the following two questions for each transmission probability t_{ij} that needs to be estimated.

> If the provider is in state i, how likely is it that this will put the dependent asset into state j?

Since this is usually hard to answer, we replace it by the following two simpler questions.

1) "If the provider is in state i, what is the most likely state j that the dependent asset will be in after this incident?"
2) "How certain are you about this prediction?"

The answers are to be chosen from a set of predefined values, namely from the set of states $\{1, \ldots, k\}$ for 1) and a set of possible confidence levels for 2). Even if the expert is unsure about the consequences, he typically still has a reasonable idea about the intensity of the consequences, i.e., if the expected consequences will be "close" to the predicted value. Because of this, the method assumes that in the case of uncertain assignments similar values as the predicted one are also likely to occur. In case an expert does not feel competent enough to make a prediction it is assumed that all possible values occur with the same likelihood. This intuition can be formalized and yields a probability distribution over the set of all possible states as shown in Table I.

TABLE I. Distribution over the CI's possible next state based on the expert's assignment

prediction	totally sure	somewhat unsure	totally unsure
1	(1,0,0)	(2/3, 1/3, 0)	(1/3,1/3,1/3)
2	(0,1,0)	(1/4, 2/4, 1/4)	(1/3,1/3,1/3)
3	(0,0,1)	(0, 2/3, 1/3)	(1/3,1/3,1/3)

Exact predictions may be difficult to provide even for very experienced people working in the respective domain. In order to capture this fact, the model can be extended to allow experts to be mistaken with a small probability ε even when they are sure about their prediction. We discussed this problem and a possible solution in [1] but will not go into detail here to keep the focus on the overall risk management process.

The input to the simulation is a network graph of connected assets of a critical infrastructure where each of these assets is in one specific state representing its functionality level. This graph resembles the picture in Figure 1 but additionally augments each node with matrices indicating the status change probabilities for each dependency. Dependencies may change over time, i.e., a short-term outage of a providers it typically less severe than a lasting reduced availability. This is particularly true for power supply (that is the backbone of every CI) where emergency power supply should prevent damage for a limited time period. The simulation will thus need a state transition probability matrix *per dependency* and *per time frame*.

The simulation prototype [2] distinguishes short, medium and long-term dependencies. For each of these time frames the probabilities $t_{ij} = \Pr(B \text{ is in state } j | A \text{ switches into state } i)$ describe the transition regimes. While the stochastic model allows a recovery (i.e., switching back into a better status), this is not yet implemented in the current version of the prototype.

The simulation starts when a risk scenario becomes reality, imitating probabilistic transitions (as for a Markov chain), and stops after a predefined time horizon (with the rationale that far-fetching forecasts become increasingly unreliable and hoping that after a decent amount of time some countermeasures can be taken). Each simulation run yields a time series of state transitions for all assets of the CI.

Given a number of simulations of a risk scenario (that may contain several risk, see [1]), the final states per assets can be averaged to get an estimate of the likelihood that this part of the CI is affected. For visualization it is helpful to apply color codes, ranging from green (symbolizing a working state) to red (symbolizing an outage), alerting about the criticality of the current condition. Numerically, the simulation results may be summarized as a table that lists the number of components which are on average in any of the possible states. The `OMNeT++` tool is used here to support the visualization and execution of our simulation, as described in [2]. Various additional outputs are possible, such as plots of time-lines relating to a single simulation run. This would display the times when a CI asset changes its state, and would show the temporal development of the cascading impacts.

Remark 2: It must be kept in mind that the simulation does *not* provide information about the likelihood for an incident to occur but starts from a given scenario that is assumed to have happened. The simulation then yields information on the likelihood of the consequences of this scenario.

D. Risk Evaluation

Risk evaluation is concerned with prioritizing the risks identified in Step II-B according to the criteria chosen in Step II-A. Classical risk management approaches use a cost-benefit analysis to decide on which risks are treated first. More advanced approaches based on game theory allow optimizing several goals simultaneously, thus also take into account non-monetary factors such as reputation or employees satisfaction [33].

As part of the overall risk management process, risk evaluation takes the results of the risk analysis step II-C into account. Based on this, it is possible to compare the different risks in terms of the impact they have on the CI (according to the considered goals). Ordering risks is doable in manifold ways, such as taking expectations (i.e., risk = impact $\times$ likelihood) or by lexicographic order on impacts. This is our choice in the following, based on the HyRiM stochastic ordering [34]. This approach corresponds to minimization of the likelihood of the worst possible damage.

E. Risk Treatment

Risk Treatment classically focuses on ways to mitigate risks by means of improving existing controls or implementing new controls in order to either reduce the likelihood of occurrence or the magnitude of the consequences and the selection of the controls to be implemented is often subjective. Advanced approaches apply game theory to find (a mix of) optimal controls [35] by considering various defence actions and selecting an optimal selection of these.

Based on the risk evaluation step II-D, risk treatment evaluates mitigation actions tailored to the specific risk scenario. In order to achieve that, it is necessary to have information on the *reasons* for a failure (root cause analysis), which simulations can deliver (we will illustrate this in Section III-E). Based on this information, we can proceed to find precautions and defence actions that provide optimal protection. Once these are implemented, risks can be reassessed (following the same steps but using updated assessment reflecting the new situation) to measure the effectiveness of the treatment.

Another way to treat the analysed risks is to identify mitigation strategies. Since resources are limited, an important task is selecting from a set of potential strategies. A method to solve this optimization problem is to consider the different risks as strategies of an attacker (nature in this case) and let the operator of the CI defend his system. At first sight, this might be an inappropriate model for two reasons. First, game theory assumes *rational* players that are able to predict the (best) responses of the other players to their actions. Still game theory is able to provide reasonable results when applying a zero sum game, as this type of game assumes that the attacker want to cause maximal damage. When the defender plays his optimal attack for this worst case, he will only be better off if the attacker deviates from his optimal strategy due to the characterization of a Nash equilibrium. The identified solution may not be as efficient as if we knew the attackers intentions but they yield an upper bound to the expected damage. The second issue with game theoretic models is the traditional assumption that payoffs are real valued. This assumption can however be generalized under very mild assumptions [36] such payoffs may be random. Thus the impact can be estimated through a simulation model as described above and payoffs of the game are the estimated distributions over all possible states. A Nash equilibrium can be computed numerically, e.g., in R [37]. Application of this game theoretic setting to CIs are illustrated in [35], [38].

III. Analysing a Water Provider

This section demonstrates how to put the process described in Section II into practice for a specific CI. In the following, we analyse risks faced by a fictitious water provider following the steps laid out above. All assessments and estimates presented in this paper are illustrative only, since any such data is sensitive and hence protected. However, the data used is based on discussions with experts of the field to be as realistic as possible. The main goal is to illustrate how to analyse the consequences of a risk scenario affecting a CI that is part of an entire network of interdependent CIs, i.e., depends on some CIs and in turn provides important input to other CIs.

In the following, we consider a European water provider of average size providing its services to a town with several hundred thousands inhabitants as well as some municipalities in the surrounding area. The water provider is responsible for planning, building and maintaining the entire water network, but his main focus is on the availability of the drinking water. In order to ensure a sustainable water quality, the provider supports water processing and sewage cleaning by using an ICS. The considered use case assumes various sources for water, namely a mountain spring, a river and a well. The two latter use pumps to lift the water above ground level. The water is further treated at the water plant to increase the quality (e.g., through removal of undesired chemicals or adding of minerals). Transportation paths are short due to the geography of the landscape and the number of necessary lines is correspondingly low. Several reservoirs are available to ensure water supply in case of high demand, e.g., to extinguish fire.

The water provider relies on the transportation system, in particular on roads, e.g., to be able to perform maintenance and manual checks, and as many other CI it crucially depends on electricity. In this scenario, we consider an internal power plant that contributes approximately 30% of the required energy while the remaining power comes from external providers. Redundancy in the system and an emergency power supply help to mitigate the criticality of this dependency. In case of a reduction or an interruption of electricity, the utility provider is still able to guarantee supply with drinking water up to three days due the precautionary measures.

On the other hand, the water provider is important for numerous other CIs in the region. It supplies drinking water to hospitals, schools and grocery stores but also cooling water for both hospitals and industrial companies. The actual relevance of each of these connections can only be assessed by the CIs depending on it, which requires discussions with the corresponding experts and thus goes beyond the scope of this use case that focuses on the risk management for a water utility. The remainder of this section presents an analysis of the effects stemming from the realisation of a risk using a qualitative risk assessment performed by experts from the water domain.

A. Establishing the context

The main input to context establishment is discussion with experts. In detail, we discussed the importance of each asset for the functionality of water supply which yielded a list as given in Table II. Besides the elements required for production (river and well pump), purification (water plant) and storage (water reservoir) this list also contains essential elements of water distribution, namely two parts of a distribution network. Functionality of this network is essential for the delivery of water to dependent CIs. In order to at least indicate the complexity of this distribution network we composed it of two different parts where the second part is located at a higher altitude than the first part such that a pump is required for transportation. Both the river pump and the well pump are abstract nodes that are supposed to describe the behaviour of all pumps of the corresponding type (if more than one exist). This abstraction allows us to model situations where a significant number of pumps fail so that it affects availability of water.

TABLE II. List of Infrastructures and Providers

Number	Object
1	Water Plant
2	Mountain Spring
3	Well
4	Well Pump
5	Water Reservoir
6	River
7	River Pump
8	Power Grid
9	Communication
10	SCADA server
11	Distribution network 1
12	Distribution Pump
13	Distribution network 2

Besides the assets themselves, emergency systems and backups need to be accounted for in the modeling process, e.g., we assume the existence of emergency power stations as required by (Austrian) law. Such components are indirectly taken into account when assessing the dependencies (in particular when setting up the transmission regime; shortage in power supply only affects other parts after a certain time) rather than being assets themselves.

The situation of the fictitious water provider is displayed in Figure 1 showing all relevant interdependencies between the assets. Initially, we assume that the everything is working smoothly, i.e., every asset is in states 1, until an incident happens. In our small example, we refrain from categorising the connections but rather assess every single connection.

B. Risk Identification

In order to identify the relevant risks we discussed potential threats with experts. The most significant ones for a water provider turned out to be the following ones.

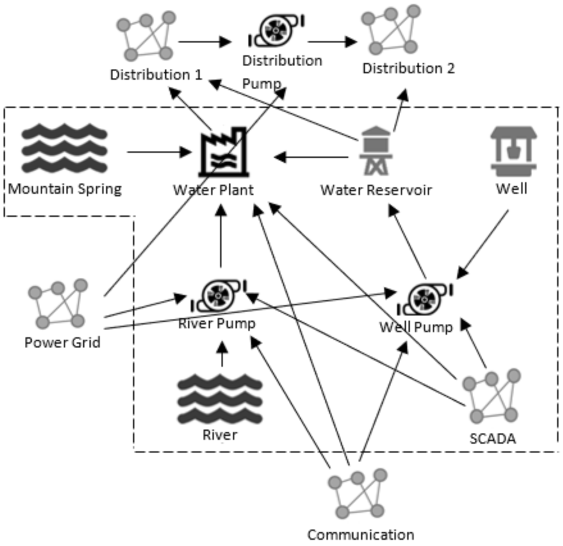

Fig. 1. Visualization of water use case

- R_1: flooding
- R_2: extreme weather conditions
- R_3: leakage of hazardous material (water contamination)
- R_4: cyberattack (e.g., on SCADA server)

Risk R_2 indicates the general vulnerability of water supply on the weather. For the upcoming analysis however, we need to be more specific in order to analyse the risk in the next step. Short-term heavy rain is not as a severe problem for a water provider (it certainly is for other CIs), since the main source of the water provider is groundwater. While it may cause smaller damage to the infrastructure, it will not interrupt water supply which is the core interest of our fictitious water provider. Some experts see flash floods an underestimated hazard and raise awareness [39]. Currently, droughts are considered more problematic for a water utility, in particular because they are likely to become more frequent in near future. Thus, we analyse the risk R_2 to be a *heat wave* in the following.

C. Risk Analysis

Knowing the relevant risks, it is necessary to understand each of these in detail. First, we identify the likelihood of each asset to be directly affected by a realisation of each of the risks. A flooding affects single sites (such as a well) but is typically not critical for the overall functionality for the water supply. Further, a realisation of risk R_1 may yield a limited operation of wells and springs as water may be contaminated by particles (germs, bacteria and others) induced by the flood. Depending on the degree of contamination, water may be boiled to make it drinkable or needs to be purified technically, which is a costly and time-consuming process. Recent floods such as the flooding in central Europe in 2002 [40] and 2013 [41], [42] indicate an increased likelihood of occurrence [43]. Concerning a realisation of R_2 we here focus on an extraordinarily dry period. Various water sources may dry up, in particular rivers or wells, but we assume that at least some sources like ground water remain available. A drought implies also an increased need for water and thus yields a peak in consumption which in turn challenges the infrastructure. Peaks will cause additional costs for the provider but are not considered here any further since this does not affect other parts of the system. The effect of leakage of hazardous material strongly depends on the circumstances of the leakage and to some extent on the material. The crucial factor is the extent of the leakage as this changes the impact significantly. A bounded contamination is not a severe issue as long as the water network is close-meshed (i.e., there is enough redundancy in the network). However, if groundwater or a large number of wells are affected, water purification may take several months. For our use case, we assume a realisation of R_3 that is a limited spreading, affecting mainly the river and only with small likelihood also affects the mountain spring or the well. Since contamination seriously affects the quality of drinking water [44], the probability that it switches to the worst possible state 3 is high. A realisation of R_4 is most challenging to investigate because data is typically (and luckily) sparse. Recent incidents in ports [45]–[47] and in particular prominent attacks such as Wannacry [15] or NotPetya [48] allow a heuristic estimation of consequences. In order to perform simulations it is necessary to make some basic assumptions about the spreading process (e.g., does it spread via email or not).

Understanding a risk includes identifying those assets that are directly affected by a realisation of the risk and estimating the expected impact. Tables III, IV, V and VI give the estimated probability distributions over the various states for short, medium and long-term of those assets that are directly affected (those not directly affected are omitted and the corresponding likelihoods set to zero). Note that the tables only contain the necessary information, i.e., the chance that the asset is not affected by the considered risk (stays in state 1) is such that the sum of the corresponding row is one.

TABLE III. Direct Impact of R_1 on Assets

Asset	Impact	Short	Medium	Long
Mountain Spring	limitation	0.3	0.4	0.5
	failure	0.2	0.3	0.4
Well	limitation	0.2	0.3	0.4
	failure	0.1	0.2	0.3

TABLE IV. Direct Impact of R_2 on Assets

Asset	Impact	Short	Medium	Long
Mountain Spring	limitation	0.1	0.2	0.3
	failure	0.0	0.2	0.3
Well	limitation	0.2	0.3	0.4
	failure	0.0	0.1	0.2
River	limitation	0.6	0.4	0.2
	failure	0.2	0.4	0.6

Identification of indirect consequences of a risk on each of the assets is supported by the simulation, as described in Section II-C and will be demonstrated in detail in the next paragraph.

TABLE V. Direct Impact of R_3 on Assets

Asset	Impact	Short	Medium	Long
Mountain Spring	limitation	0.0	0.0	1/3
	failure	1.0	1.0	2/3
Well	limitation	0.0	0.0	1/3
	failure	1.0	1.0	2/3
River	limitation	0.0	0.0	1/3
	failure	1.0	1.0	2/3

TABLE VI. Direct Impact of R_4 on Assets

Asset	Impact	Short	Medium	Long
Communication	limitation	0	$1/3$	$1/3$
	failure	1	$2/3$	$2/3$
SCADA	limitation	$1/2$	$1/3$	$1/3$
	failure	$1/4$	$1/3$	$2/3$

Finally, the likelihood of occurrence is estimated for each risk to add to the picture. Additionally, the likelihood of failure and impairment of the CI due to a realisation of each risk are rated as “negligible”, “low”, “medium”, “high” or “very high” by experts. Where available, public reports and statistics may complement such subjective assessment and yield refined estimates. In case of a cyberattack the estimates highly depend on the assumptions about the attacker, e.g., whether he plans a highly sophisticated APT or a more general malware attack. The values for this use case are given in Table VII.

TABLE VII. Overall Likelihood Assessment for Risks

Risk		Occurrence	Failure	Impairment
R_1:	flooding	medium	negligible	negligible
R_2:	heat wave	medium	negligible	medium
R_3:	contamination	low	negligible	medium
R_4:	cyberattack	medium	low	medium

Consequences of an Incident

The consequences of a realisation of a risk are estimated based on a simulation model, as described in Section II-C. The dependencies between the assets are assessed for three different time horizons, taking into account the dynamic nature of CIs. For example, if emergency power supply is available, the likelihood for a pump to switch to the outage state 3 given the electricity goes off is zero for the first couple of hours, and changes to 1 if the emergency generator runs out of fuel, unless the original power supply has recovered in the meantime. The assessments given in Tables VIII, IX and X are based on several discussion with domain experts. Impact was measured on the three-tier scale “negligible” (state 1), “medium” (state 2) and “high” (state 3) while the experts’ confidence in the provided prediction is described as “totally sure”, “somewhat unsure” or “totally unsure”. Note that these assessments need to be made for each specific connection and do neither contain information about potential substitutes (e.g., if several pumps are available) nor take into account the option of repair or recovery. The assessment is solely concerned about the nature of this specific dependency between the two assets.

The assessments from Tables VIII, IX and X are mapped to the corresponding transition matrices as described in Table

TABLE VIII. Short-term impact assessment

Link	Problem	Prediction	Confidence
River Pump →	limitation	negligible	totally sure
Water Plant	failure	negligible	totally sure
Mountain Spring →	limitation	negligible	totally sure
Water Plant	failure	negligible	totally sure
Communication →	limitation	medium	somewhat unsure
Water Plant	failure	negligible	totally sure
Water Reservoir →	limitation	negligible	totally sure
Water Plant	failure	negligible	totally sure
SCADA →	limitation	high	somewhat unsure
Water Plant	failure	high	somewhat unsure
Well →	limitation	negligible	totally sure
Well Pump	failure	negligible	somewhat unsure
Communication →	limitation	medium	somewhat unsure
Well Pump	failure	negligible	totally sure
Power Grid →	limitation	negligible	totally sure
Well Pump	failure	negligible	totally sure
SCADA →	limitation	medium	somewhat unsure
Well Pump	failure	high	somewhat unsure
River →	limitation	negligible	totally sure
River Pump	failure	negligible	somewhat unsure
Communication →	limitation	medium	somewhat unsure
River Pump	failure	negligible	totally sure
Power Grid →	limitation	negligible	totally sure
River Pump	failure	negligible	totally sure
SCADA →	limitation	medium	somewhat unsure
River Pump	failure	high	somewhat unsure
Well Pump →	limitation	negligible	totally sure
Water Reservoir	failure	negligible	totally sure
Water Plant →	limitation	negligible	totally sure
Distribution 1	failure	negligible	totally sure
Water Reservoir →	limitation	negligible	totally sure
Distribution 1	failure	negligible	totally sure
Water Reservoir →	limitation	negligible	totally sure
Distribution 2	failure	negligible	totally sure
Distribution Pump →	limitation	negligible	totally sure
Distribution 2	failure	negligible	somewhat unsure
Power Grid →	limitation	negligible	totally sure
Distribution Pump	failure	negligible	totally sure
Distribution 1 →	limitation	negligible	totally sure
Distribution Pump	failure	negligible	somewhat unsure

TABLE IX. Medium-term impact assessment

Link	Problem	Prediction	Confidence
Pump →	limitation	negligible	totally sure
Water Plant	failure	negligible	somewhat unsure
Mountain Spring →	limitation	negligible	totally sure
Water Plant	failure	negligible	somewhat unsure
Communication →	limitation	negligible	totally sure
Water Plant	failure	negligible	totally sure
Water Reservoir →	limitation	negligible	totally sure
Water Plant	failure	negligible	somewhat unsure
SCADA →	limitation	medium	somewhat unsure
Water Plant	failure	medium	somewhat unsure
Well →	limitation	medium	somewhat unsure
Well Pump	failure	high	somewhat unsure
Communication →	limitation	negligible	totally sure
Well Pump	failure	negligible	totally sure
Power Grid →	limitation	negligible	totally sure
Well Pump	failure	negligible	totally sure
SCADA →	limitation	medium	totally sure
Well Pump	failure	medium	totally sure
River →	limitation	medium	somewhat unsure
River Pump	failure	high	somewhat unsure
Communication →	limitation	negligible	totally sure
River Pump	failure	negligible	totally sure
Power Grid →	limitation	negligible	totally sure
River Pump	failure	negligible	totally sure
SCADA →	limitation	medium	totally sure
River Pump	failure	medium	totally sure
Well Pump →	limitation	negligible	totally sure
Water Reservoir	failure	negligible	somewhat unsure
Water Plant →	limitation	negligible	totally sure
Distribution 1	failure	negligible	somewhat unsure
Water Reservoir →	limitation	negligible	totally sure
Distribution 1	failure	negligible	somewhat unsure
Water Reservoir →	limitation	negligible	totally sure
Distribution 2	failure	negligible	somewhat unsure
Distribution Pump →	limitation	negligible	somewhat unsure
Distribution 2	failure	high	somewhat unsure
Power Grid →	limitation	medium	somewhat unsure
Distribution Pump	failure	negligible	somewhat unsure
Distribution 1 →	limitation	medium	somewhat unsure
Distribution Pump	failure	negligible	somewhat unsure

I, e.g., yielding transition regimes as shown in Table XI corresponding to the first few rows of Table VIII. The fact that dependencies may change over time is here captured by the three different time horizons ”short”, ”medium” and ”long”. The simulation requires definitions for all three time periods for each dependency. Table XII shows the definitions applied in this use case. Time starts at zero and the numbers given in each column are the upper bound of the corresponding range, e.g., for the connection River Pump → Water Plant short-term is a duration up to 12 hours, medium-term is between 12 an 48 hours and long is between 48 and 72 hours (when the simulation stops). For consistency, the lower limits are excluded and the upper ones are included, e.g., medium-term means $t \in (12, 48]$. The time definitions for the impacts are all equal and set to 1 hour for short-term, 12 hours for medium-term and 48 hours for long-term.

With this information the simulation can be run to estimate the impact of a realisation of the identified risks. Figure 2 shows the results of one run of the simulation (for a realisation of a contamination). The colouring of the nodes represents the state of the assets where dark grey indicates failure (state 3), medium grey indicates limited availability (state 2) and light grey means normal operation (state 1). Since the simulation contains stochastic elements, it is necessary to run a large number of repetitions. With the parameters specified above, we run the simulation 1000 times and estimate the probability distribution over the possible states for each asset. For a realisation of R_1, we considered the scenario of a flooding that lasts 4 days. The empirical probability distributions over the possible states for each asset are shown in Table XIII. For a realisation of R_2, we considered the scenario of heat wave lasting 7 weeks. The empirical probability distributions over the possible states for each asset are shown in Table XIV. For a realisation of R_3, we considered the scenario of a contamination that lasts 7 weeks. The empirical probability distributions over the possible states for each asset are shown in Table XV. For a realisation of R_4, we considered the scenario of a cyberattack that lasts 2 hours. The empirical probability distributions over the possible states for each asset are shown in Table XIII.

Note that the impact scale is influenced by the interests of the CI provider. The degree of damage could be measured in

TABLE X. Long-term impact assessment

Link	Problem	Prediction	Confidence
Pump → Water Plant	limitation	negligible	totally sure
	failure	medium	somewhat unsure
Mountain Spring → Water Plant	limitation	negligible	totally sure
	failure	medium	somewhat unsure
Communication → Water Plant	limitation	negligible	totally sure
	failure	negligible	totally sure
Water Reservoir → Water Plant	limitation	negligible	totally sure
	failure	medium	somewhat unsure
SCADA → Water Plant	limitation	medium	somewhat unsure
	failure	medium	somewhat unsure
Well → Well Pump	limitation	medium	somewhat unsure
	failure	high	totally sure
Communication → Well Pump	limitation	negligible	totally sure
	failure	negligible	totally sure
Power Grid → Well Pump	limitation	negligible	totally sure
	failure	high	totally sure
SCADA → Well Pump	limitation	medium	totally sure
	failure	medium	totally sure
River → River Pump	limitation	medium	somewhat unsure
	failure	high	totally sure
Communication → River Pump	limitation	negligible	totally sure
	failure	negligible	totally sure
Power Grid → River Pump	limitation	negligible	totally sure
	failure	high	totally sure
SCADA → River Pump	limitation	medium	totally sure
	failure	medium	totally sure
Well Pump → Water Reservoir	limitation	negligible	totally sure
	failure	medium	somewhat unsure
Water Plant → Distribution 1	limitation	negligible	somewhat unsure
	failure	medium	somewhat unsure
Water Reservoir → Distribution 1	limitation	negligible	totally sure
	failure	negligible	somewhat unsure
Water Reservoir → Distribution 2	limitation	negligible	totally sure
	failure	negligible	somewhat unsure
Distribution Pump → Distribution 2	limitation	medium	somewhat unsure
	failure	high	totally sure
Power Grid → Distribution Pump	limitation	medium	somewhat unsure
	failure	high	somewhat unsure
Distribution 1 → Distribution Pump	limitation	medium	somewhat unsure
	failure	high	totally sure

TABLE XI. Short-term transition probabilities

$$P^{\text{short}}_{\text{river pump} \to \text{water plant}} = \begin{array}{c} \\ s_{rp}=1 \\ s_{rp}=2 \\ s_{rp}=3 \end{array} \begin{array}{c} \begin{array}{ccc} s_{wp}=1 & s_{wp}=2 & s_{wp}=3 \end{array} \\ \left(\begin{array}{ccc} 1 & 0 & 0 \\ 1 & 0 & 0 \\ 1 & 0 & 0 \end{array} \right) \end{array}$$

$$P^{\text{short}}_{\text{mountain spring} \to \text{water plant}} = \begin{array}{c} \\ s_{ms}=1 \\ s_{ms}=2 \\ s_{ms}=3 \end{array} \begin{array}{c} \begin{array}{ccc} s_{wp}=1 & s_{wp}=2 & s_{wp}=3 \end{array} \\ \left(\begin{array}{ccc} 1 & 0 & 0 \\ 1 & 0 & 0 \\ 1 & 0 & 0 \end{array} \right) \end{array}$$

$$P^{\text{short}}_{\text{communication} \to \text{water plant}} = \begin{array}{c} \\ s_{com}=1 \\ s_{com}=2 \\ s_{com}=3 \end{array} \begin{array}{c} \begin{array}{ccc} s_{wp}=1 & s_{wp}=2 & s_{wp}=3 \end{array} \\ \left(\begin{array}{ccc} 1 & 0 & 0 \\ 1/4 & 2/4 & 1/4 \\ 1 & 0 & 0 \end{array} \right) \end{array}$$

TABLE XII. Time assessments for short-, medium- and long-term

Link	Short	Medium	Long
River Pump → Water Plant	12 hours	48 hours	72 hours
Mountain Spring → Water Plant	12 hours	48 hours	72 hours
Communication → Water Plant	12 hours	48 hours	72 hours
Water Reservoir → Water Plant	12 hours	48 hours	72 hours
SCADA → Water Plant	1 hour	2 hours	6 hours
Well → Well Pump	12 hours	48 hours	72 hours
Communication → Well Pump	12 hours	48 hours	72 hours
Power Grid → Well Pump	8 hours	48 hours	72 hours
SCADA → Well Pump	1 hour	2 hours	6 hours
River → River Pump	12 hours	48 hours	72 hours
Communication → River Pump	12 hours	48 hours	72 hours
Power Grid → River Pump	8 hours	48 hours	72 hours
SCADA → River Pump	1 hour	2 hours	6 hours
Well Pump → Water Reservoir	12 hours	48 hours	72 hours
Water Plant → Distribution1	12 hours	48 hours	72 hours
Water → Distribution1	12 hours	48 hours	72 hours
Water Reservoir → Distribution2	12 hours	48 hours	72 hours
Distribution Pump → Distribution2	24 hours	48 hours	72 hours
Power Grid → Distribution Pump	8 hours	48 hours	72 hours
Distribution1 → Distribution Pump	12 hours	48 hours	72 hours

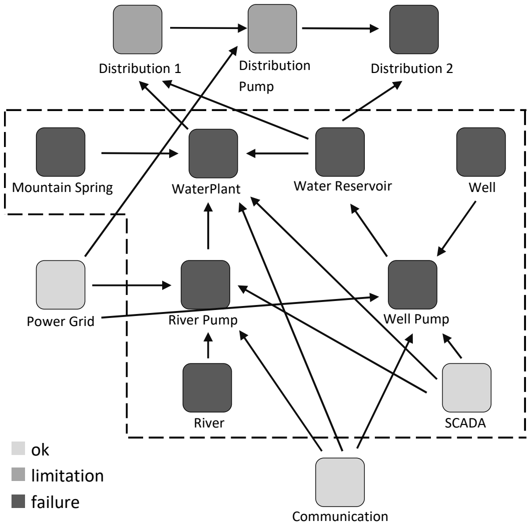

Fig. 2. Results of one run simulating scenario R_3

TABLE XIII. Estimated Impact of R_1

	state 1	state 2	state 3
Water Plant	0.384	0.422	0.194
Mountain Spring	0.011	0.325	0.664
Well	0.099	0.380	0.521
Well Pump	0.122	0.200	0.678
Water Reservoir	0.427	0.398	0.175
River	1.000	0.000	0.000
River Pump	1.000	0.000	0.000
Power Grid	1.000	0.000	0.000
Communication	1.000	0.000	0.000
SCADA server	1.000	0.000	0.000
Distribution network 1	0.672	0.290	0.038
Distribution Pump	0.690	0.148	0.162
Distribution network 2	0.670	0.137	0.193

terms of the number of affected customers, the time needed reassure availability of drinking water, the amount of resources necessary to overcome a shortage (in terms of money or person

TABLE XIV. Estimated Impact of R_2 on Assets

	state 1	state 2	state 3
Water Plant	0.144	0.547	0.309
Mountain Spring	0.228	0.344	0.428
Well	0.191	0.547	0.262
Well Pump	0.227	0.268	0.505
Water Reservoir	0.584	0.286	0.130
River	0.013	0.188	0.799
River Pump	0.023	0.093	0.884
Power Grid	1.000	0.000	0.000
Communication	1.000	0.000	0.000
SCADA server	1.000	0.000	0.000
Distribution network 1	0.524	0.394	0.082
Distribution Pump	0.551	0.196	0.253
Distribution network 2	0.565	0.150	0.285

TABLE XV. Estimated Impact of R_3 on Assets

	state 1	state 2	state 3
Water Plant	0.032	0.492	0.476
Mountain Spring	0.000	0.000	1.000
Well	0.000	0.000	1.000
Well Pump	0.000	0.000	1.000
Water Reservoir	0.145	0.600	0.255
River	0.000	0.000	1.000
River Pump	0.000	0.000	1.000
Power Grid	1.000	0.000	0.000
Communication	1.000	0.000	0.000
SCADA server	1.000	0.000	0.000
Distribution network 1	0.361	0.525	0.114
Distribution Pump	0.388	0.262	0.350
Distribution network 2	0.406	0.185	0.409

TABLE XVI. Estimated Impact of R_4 on Assets

	state 1	state 2	state 3
Water Plant	0.070	0.305	0.625
Mountain Spring	1.000	0.000	0.000
Well	1.000	0.000	0.000
Well Pump	0.230	0.553	0.217
Water Reservoir	0.811	0.139	0.050
River	1.000	0.000	0.000
River Pump	0.230	0.565	0.205
Power Grid	1.000	0.000	0.000
Communication	0.000	0.000	1.000
SCADA server	0.230	0.487	0.283
Distribution network 1	0.346	0.507	0.147
Distribution Pump	0.385	0.250	0.365
Distribution network 2	0.432	0.149	0.419

hours), the reputation damage due to the incident (e.g., in case of insufficient protection against cyberattacks) or many more. Further, assessing a criticality level to each asset is possible to represent the importance of each asset for the overall process. Such criticality levels may also have different meaning for individual scenarios; e.g., if a pump or water tower fails for one day, this is more critical than a water contamination, since in the latter case, households can be advised to boil the water before drinking it, whereas a failure of the pump may completely cut off the household from water supply.

D. Risk Evaluation

In order to evaluate the different risks, consequences of the risks are compared. As we use an ordinal scale to measure the impact of each risk, the stochastic ordering mentioned in Section II-D simplifies to a lexicographic ordering with following interpretation. If one risk has a lower likelihood of worst case impact (state 3) than the other, we prefer this one; in case these are equal, the likelihoods of the second worst impact (state 2) are compared, and so on (and we randomly decide on the ordering of two risks with identical distributions).

Risk evaluation focuses on an asset of special interest (e.g., due to its vital importance to the CI) and is here illustrated with a focus on the water plant, comparing the impacts the all four identified risks on this asset. Figures 3, 4, 5 and 6 show the estimated impacts of the corresponding risks on the water plant.

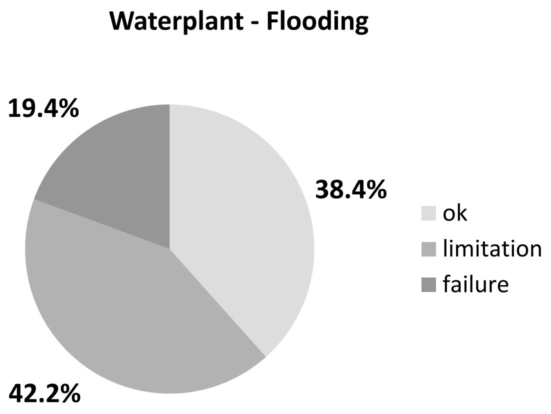

Fig. 3. Estimated Impact of R_1 on Water Plant

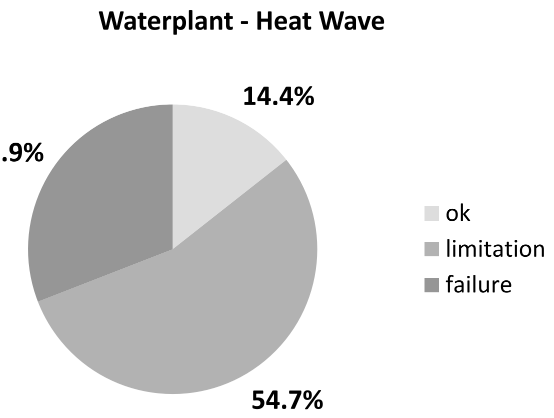

Fig. 4. Estimated Impact of R_2 on Water Plant

Comparing the probabilities we find that R_4 has the biggest chance of failure (62.5%), followed by R_3 (47.6%), R_2 (30.9%), and R_1 (19.4%). Thus, we think of R_4 as being the most dangerous one and R_1 as the one with the smallest chance of causing the most severe problems, i.e., we can write

$$R_1 \leq R_2 \leq R_3 \leq R_4.$$

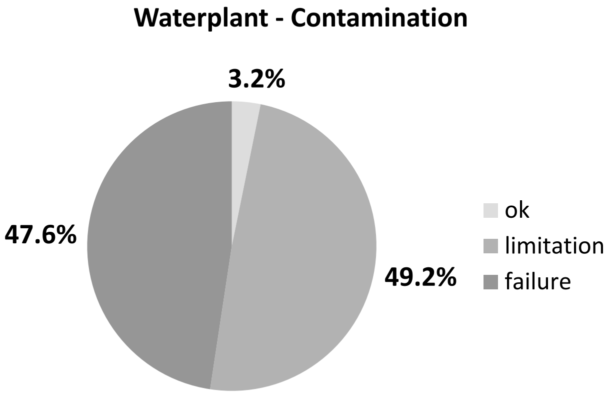

Fig. 5. Estimated Impact of R_3 on Water Plant

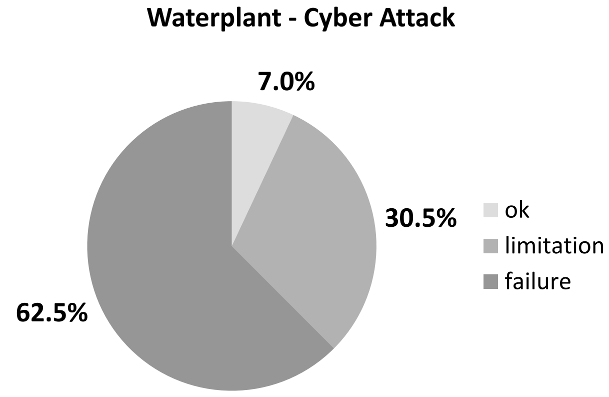

Fig. 6. Estimated Impact of R_4 on Water Plant

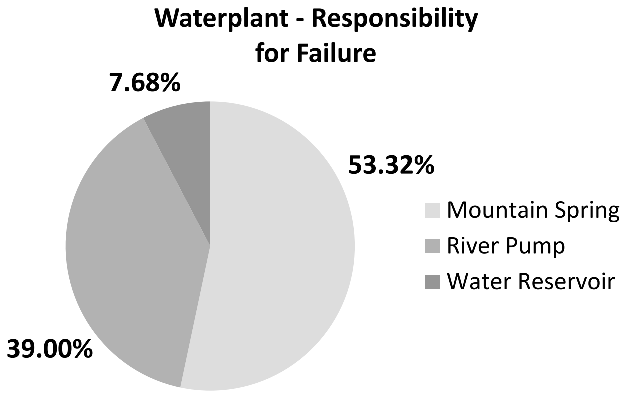

Fig. 7. Triggering factors for failure of the water plant during a contamination

E. Risk Treatment

As the last step of the risk management process, risk treatment deals with the understanding of the relevant risks and their effects on the CI can (and should) be used to identify ways to mitigate the risks. This is what risk treatment is about as the last step in the risk management process. Several methods may be used here to identify controls that need to be implemented or improved but in any case it is helpful to know the *trigger* of the failure. Figure 7 shows a pie chart illustrating which assets were responsible for failure of the water plant during a contamination (i.e., a realisation of R_3).

The most frequent cause for failure of the water plant is clearly the mountain spring which is not surprising during a contamination if it happens near the mountain spring. This problem is not easily fixed (purification is an expensive and long-lasting process). In such a case it is faster to substitute water, e.g., through an agreement with other water providers to help out in such a critical situation. However, the analysis provides more information, it shows that also limited operation of the river pump as well as the water reservoir may lead to a failure of the water plant. While both problems are very likely to be also due to the contamination, it might sometimes be a bit easier to fix these indirect issues and to reduce the overall likelihood of failure due to a risk. Similarly, Figure 8 shows the different triggers for a limitation of the water plant in case of a contamination. The triggers are the same as in the case of failure but the mountain spring is now clearly the main source of problem while the water reservoir does not significantly trigger a limitation of the water plant.

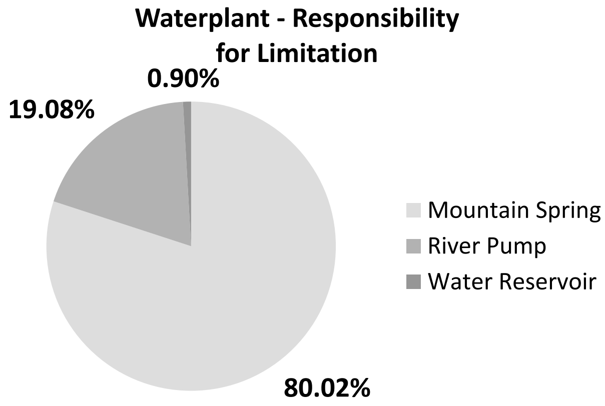

Fig. 8. Triggering factors for limitation of the Water Plant during a contamination

For a practitioner, the simulation's outcome is like a heat map, directly pointing out the most vulnerable spots in the network of critical infrastructures, which each CI domain expert can be informed about afterwards (see Figure 2 for an example scenario, from which each CI security officer can instantly see the degree of affection due to an incident). Towards a more fine-grained understanding of the affection's extent, the domain expert can continue by asking how likely an affection is to be medium or severe, which the pie charts (see Figures 3 to 6) directly tell. Given this knowledge about the local affection, the expert may then strive for a root cause analysis, which the pie chart in Figures 7 and 8 help with: here, the domain expert gets the information of who is the most relevant "neighboring" CI that has the strongest impact on one's own CI. That is, if some CI C has relations on two other CIs A and B, an incident at A may be more or less severe than an incident happening at B. For example, Figure 7 explains that a failure of the water plant in a contamination scenario is most likely due to a problem with the mountain spring, or possibly also due to a problem with the river pump, but least likely, the cause is found at the water reservoir. This can be a guidance for fixing the problem in practice. Similarly, Figure 8 would advise the expert, upon an incident, to first look at the mountain spring as an external trigger of the local issue, but only in rare cases, the water reservoir will have caused the trouble.

IV. Conclusion

Applying simulation is a straightforward proposal to extend the understanding of how security incidents propagate through and affect one or more critical infrastructures. Setting up proper simulation models, and using the information that these deliver is, however, a different story with its own challenges. This work used a hypothetical water provider in the background to describe a step by step method of

(i) how to specify interdependencies between critical infrastructures in a way that allows domain experts to include their subjective uncertainty,

(ii) how the data and specification for a concrete simulation model (chosen here for illustration) could look like, and

(iii) how the simulation model's results could be compiled into a digestible form to ease decision making by revealing previously unexpected roles of assets in incident propagation and loss estimation.

Open issues includes accuracy assessments of such a simulation (relative to reported incidents in real life), but equally important, a study of usability from domain experts perspectives. Having an accurate model is not enough, unless people outside the scientific realm and concerned with the practical things feel capable of using it. The "tutorial" style of this work shall be a step towards bridging this gap.

References

[1] S. König, T. Grafenauer, S. Rass, and S. Schauer, "Practical risk analysis in interdependent critical infrastructures - a How-To," in *The Twelfth International Conference on Emerging Security Information, Systems and Technologies (SECURWARE) 2018*. IARIA XPS Press, 2018, pp. 150–157.

[2] T. Grafenauer, S. König, S. Rass, and S. Schauer, "A simulation tool for cascading effects in interdependent critical infrastructures," in *International Workshop on Security Engineering for Cloud Computing (IWSECC 2018), collocated with the 13th International Conference on Availability, Reliability and Security (ARES 2018)*, 2018, (in press).

[3] Y. Y. Haimes, "Hierarchical Holographic Modeling," *IEEE Transactions on Systems, Man, and Cybernetics*, vol. 11, no. 9, pp. 606–617, 1981.

[4] N. K. Svendsen and S. D. Wolthusen, "Analysis and statistical properties of critical infrastructure interdependency multiflow models," in *2007 IEEE SMC Information Assurance and Security Workshop*, June 2007, pp. 247–254.

[5] R. Setola, S. D. Porcellinis, and M. Sforna, "Critical infrastructure dependency assessment using the input-output inoperability model," *International Journal of Critical Infrastructure Protection (IJCIP)*, vol. 2, pp. 170–178, 2009.

[6] S.-J. Wu and M. T. Chu, "Markov chains with memory, tensor formulation, and the dynamics of power iteration," vol. 303, pp. 226–239, 2017.

[7] M. Rahnamay-Naeini and M. M. Hayat, "Cascading failures in interdependent infrastructures: An interdependent markov-chain approach," *IEEE Transactions on Smart Grid*.

[8] S. König and S. Rass, "Stochastic dependencies between critical infrastructures," in *SECURWARE 2017: The Eleventh International Conference on Emerging Security Information, Systems and Technologies*. IARIA, 2017, pp. 106–110.

[9] S. König, S. Schauer, and S. Rass, "A stochastic framework for prediction of malware spreading in heterogeneous networks," in *Secure IT Systems: 21st Nordic Conference, NordSec 2016, Oulu, Finland, November 2-4, 2016. Proceedings*, B. B. Brumley and J. Röning, Eds. Springer International Publishing, 2016, pp. 67–81.

[10] A. V. Uzunov, E. B. Fernandez, and K. Falkner, "Securing distributed systems using patterns: A survey," *Computers & Security*, vol. 31, no. 5, pp. 681–703, 2012.

[11] S. Fletcher, "Electric power interruptions curtail California oil and gas production," *Oil Gas Journal*, 2001.

[12] M. Schmidthaler and J. Reichl, "Economic Valuation of Electricity Supply Security: Ad-hoc Cost Assessment Tool for Power Outages," *ELECTRA*, no. 276, pp. 10–15, 2014.

[13] J. Condliffe, "Ukraine's Power Grid Gets Hacked Again, a Worrying Sign for Infrastructure Attacks," 2016. [Online]. Available: https://www.technologyreview.com/s/603262/ukraines-power-grid-gets-hacked-again-a-worrying-sign-for-infrastructure-attacks/

[14] D. Kushner, "The real story of stuxnet," *IEEE Spectrum*, vol. 3, no. 50, pp. 48–53, 2013.

[15] National Audit Office, "Investigation: Wannacry cyber attack and the nhs," 2017. [Online]. Available: https://www.nao.org.uk/wp-content/uploads/2017/10/Investigation-WannaCry-cyber-attack-and-the-NHS.pdf

[16] C. Kolias, G. Kambourakis, A. Stavrou, and J. Voas, "DDoS in the IoT: Mirai and other botnets," *Computer*, vol. 50, no. 7, pp. 80–84, 2017.

[17] Z. Bederna and T. Szadeczky, "Cyber espionage through botnets," *Security Journal*, Sep 2019.

[18] European Parliament, "Directive (EU) 2016/1148 of the European Parliament and of the Council: concerning measures for a high common level of security of network and information systems across the Union," *Official Journal of the European Union*, 2016. [Online]. Available: https://eur-lex.europa.eu/legal-content/EN/TXT/HTML/?uri=CELEX:32016L1148&from=EN

[19] S. Rinaldi, J. Peerenboom, and T. Kelly, "Identifying, understanding, and analyzing critical infrastructure interdependencies," *IEEE Control Systems Magazine*, pp. 11–25, 2001.

[20] A. Gouglidis, B. Green, J. Busby, M. Rouncefield, D. Hutchison, and S. Schauer, *Threat awareness for critical infrastructures resilience*. IEEE, 9 2016.

[21] S. M. Rinaldi, "Modeling and simulating critical infrastructures and their interdependencies," in *37th Annual Hawaii International Conference on System Sciences, 2004. Proceedings of the*. IEEE, 2004, p. 8 pp.

[22] H. Guo, C. Zheng, H. H.-C. Iu, and T. Fernando, "A critical review of cascading failure analysis and modeling of power system," *Renewable and Sustainable Energy Reviews*, vol. 80, pp. 9–22, dec 2017.

[23] S. Saidi, L. Kattan, P. Jayasinghe, P. Hettiaratchi, and J. Taron, "Integrated infrastructure systems—A review," *Sustainable Cities and Society*, vol. 36, pp. 1–11, Jan 2018.

[24] E. Luiijf, M. Ali, and A. Zielstra, "Assessing and improving SCADA security in the Dutch drinking water sector," *International Journal of Critical Infrastructure Protection*, vol. 4, no. 3-4, pp. 124–134, 2011.

[25] A. Alshawish, M. A. Abid, H. de Meer, S. Schauer, S. König, A. Gouglidis, and D. Hutchison, "G-dps: A game-theoretical decision-making framework for physical surveillance games," in *Game Theory for Security and Risk Management: From Theory to Practice*. Cham: Springer International Publishing, 2018, pp. 129–156.

[26] A. Gouglidis, S. König, B. Green, K. Rossegger, and D. Hutchison, "Protecting water utility networks from advanced persistent threats: A case study," in *Game Theory for Security and Risk Management: From*

Theory to Practice. Cham: Springer International Publishing, 2018, pp. 313–333.

[27] European Comission, “COUNCIL DIRECTIVE 2008/114/EC of 8 December 2008 on the identification and designation of European critical infrastructures and the assessment of the need to improve their protection,” *Official Journal of the European Union*, no. L345, pp. 75–82, 2008.

[28] ISO International Organization for Standardization, *ISO 31000:2018 Risk management - Guidelines*. Geneva, Switzerland: ISO International Organization for Standardization, 2018.

[29] National Institute of Standarts and Technology (NIST), “National Vulnerability Database (NVD),” https://nvd.nist.gov/, accessed: 2019-12-05.

[30] “Nessus vulnerability scanner,” https://www.tenable.com/products/nessus-vulnerability-scanner, accessed: 2019-08-22.

[31] “OpenVAS: Open Vulnerability Assessment System,” http://www.openvas.org/, accessed: 2019-08-22.

[32] I. Münch, “Wege zur Risikobewertung,” in *DACH Security 2012*, P. Schartner and J. Taeger, Eds. syssec, 2012, pp. 326–337.

[33] S. Schauer, “A risk management approach for highly interconnected networks,” in *Game Theory for Security and Risk Management*. Springer International Publishing, 2018, pp. 285–311.

[34] S. Rass, S. König, and S. Schauer, “Decisions with uncertain consequences—a total ordering on loss-distributions,” vol. 11, no. 12, p. e0168583.

[35] ——, “Defending against advanced persistent threats using game-theory,” vol. 12, no. 1, p. e0168675. [Online]. Available: http://dx.plos.org/10.1371/journal.pone.0168675

[36] S. Rass, S. König, and S. Schauer, “Uncertainty in games: Using probability-distributions as payoffs,” in *Lecture Notes in Computer Science*. Springer International Publishing, 2015, pp. 346–357.

[37] S. Rass and S. König, “Package ’HyRiM’: Multicriteria risk management using zero-sum games with vector-valued payoffs that are probability distributions,” 2019. [Online]. Available: https://cran.r-project.org/web/packages/HyRiM/index.html

[38] S. König, “Choosing ways to increase resilience in critical infrastructures,” in *Proceedings of the 16th ISCRAM Conference – Valencia, Spain May 2019*, Valencia, 2019, pp. 1245–1251.

[39] Zurich, “Flash floods: The underestimated natural hazard,” Zurich Insurance Company Ltd, Tech. Rep., 2017, accessed: 2019-08-22.

[40] J. H. Christensen and O. B. Christensen, “Severe summertime flooding in europe,” *Nature*, vol. 421, no. 6925, pp. 805–806, 2003.

[41] A. H. Thieken, T. Bessel, S. Kienzler, H. Kreibich, M. Müller, S. Pisi, and K. Schröter, “The flood of june 2013 in Germany: how much do we know about its impacts?” *Natural Hazards and Earth System Sciences*, vol. 16, no. 6, pp. 1519–1540, 2016.

[42] K. Schröter, M. Kunz, F. Elmer, B. Mühr, and B. Merz, “What made the June 2013 flood in Germany an exceptional event? a hydro-meteorological evolution,” *Hydrology and Earth System Sciences*, vol. 19, pp. 309–327, 2015.

[43] Z. W. Kundzewicz, U. Ulbrich, T. brücher, D. Graczyk, A. Krüger, G. C. Leckebusch, L. Menzel, I. Pińskwar, M. Radziejewski, and M. Szwed, “Summer floods in central europe – climate change track?” *Natural Hazards*, vol. 36, no. 1-2, pp. 165–189, 2005.

[44] J. A. Brown and W. P. Darby, “Predicting the probability of contamination at groundwater based public drinking supplies,” vol. 11, pp. 1077–1082, 1988.

[45] C. Cimpanu. (2018) Port of San Diego suffers cyber-attack, second port in a week after Barcelona. [Online]. Available: https://www.zdnet.com/article/port-of-san-diego-suffers-cyber-attack-second-port-in-a-week-after-barcelona/

[46] (2018) 2018 highlights: Major cyber attacks reported in maritime industry. [Online]. Available: https://safety4sea.com/cm-2018-highlights-major-cyber-attacks-reported-in-maritime-industry/

[47] World Maritime News. (2018) COSCO shipping lines falls victim to cyber attack. [Online]. Available: https://worldmaritimenews.com/archives/257665/cosco-shipping-lines-falls-victim-to-cyber-attack/

[48] A. Greenerg. The untold story of NotPetya, the most devastating cyberattck in history. Accessed: 2019-08-23. [Online]. Available: https://www.wired.com/story/notpetya-cyberattack-ukraine-russia-code-crashed-the-world/