Practical Risk Analysis in Interdependent Critical Infrastructures - a How-To

Sandra König, Stefan Schauer
Austrian Institute of Technology GmbH
Digital Safety & Security Department
Vienna, Austria
{sandra.koenig,stefan.schauer}@ait.ac.at

Stefan Rass, Thomas Grafenauer
Universität Klagenfurt
Institute of Applied Informatics
Klagenfurt, Austria
{stefan.rass,thomas.grafenauer}@aau.at

Abstract—**Critical infrastructures (CIs) have become more and more interconnected in the recent past. Disturbances in one affect many others and consequences tend to become unpredictable due to manifold interdependencies and cascading effects. A decent amount of various stochastic models has been developed to capture this uncertainty and aid the management of security and risk. However, these models are not frequently used in practice, not to the least because many experts feel that there is a gap between theory and practice. In this article, we illustrate how to apply such a model by investigating the situation of a water provider that is part of an entire network of CIs step by step and describe the results of the analysis. While the data used is for illustration purpose only and describes the situation of a fictitious water provider, the assignments are based on several discussions with experts from the field. Besides pure damage prevention, simulations of incident propagation may be of independent interest for trust management and reputation.**

Keywords*–*critical infrastructure; dependencies; stochastic model; risk propagation; water supply.

I. INTRODUCTION

Critical infrastructures such as power or water providers, food systems, health care and transportation networks satisfy the basic needs of society. Each of them is crucial for the functionality of a society and significantly contributes to the economic welfare of people as well as their security. During the last years, mutual dependencies among CIs have become stronger; e.g., a hospital depends on electricity, water, food supply and working transportation lines. The increasing sensitivity of this network of connected CIs has been illustrated in the past by incidents such as the disruption of electric power in California in 2001 [1], the power outage in Italy in 2003 [2] or the hacking attack on the Ukrainian power grid in 2015 [3], only to name a few. The dependencies are getting more complex in nature, i.e., a water provider does not only need electricity for the pumps but also to keep the monitoring systems, e.g., Supervisory Control and Data Acquisition (SCADA) systems or Industrial Control Systems (ICSs), running. This increasing complexity makes it even harder to predict the consequences of a limited availability of one CI on other connected CIs. This is the main reason why we apply a stochastic model to investigate the consequences of interdependencies on the impact of a risk. Since electricity is a commonly fundamental provider for many CIs built on top, we pick the water supply as one example of these, to illustrate how incidents like the reported ones could affect a water provider *depending* on electricity (amongst others). More complex examples like hospitals are conceptually similar yet substantially more complex to describe, and are thus outside the scope of this current work.

Incidents of interest for simulation can be of various kind, including natural events, but also man-made unwanted interventions like cyber-attacks or human error. Especially cyber-attacks have recently (in 2016) been moved into the center of attention by the EU Directive 2016/1148 on cyber security [4]. The consequences of cyber incidents primarily relate to matters of privacy breaches and communication infrastructures, yet extend up to potential dangers of damaging infrastructures through cyber-attacks causing malicious configurations to vital parts of the system (such as the Stuxnet worm did). We stress that this kind of incident is its own kind of challenge to describe in the terms of the model that we study, yet no different in the simulation. To ease matters in the following, we thus confine ourselves to physical events and dependencies, leaving aspects of cyber-dependencies as straightforward adaptions.

Related Work

The increasing interest in interconnections and dependencies between CIs (and the effects on other utility providers) yields a growing number of publications investigating these dependencies. Various methods are used, including Hierarchical Holographic Modeling (HHM) [5], a multi-graph model for random failures [6] or input-output models [7]. Due to the unpredictability of consequences, stochastic models gained a lot of attention. A Interdependent Markov Chain (IDMC) model is used to describe cascading failures in interdependent infrastructures in power systems [8], where every infrastructure is described by one discrete-time Markov chain and the interdependencies between these chains are represented by dependencies between the corresponding transition probabilities.

A stochastic model that allows different degrees of failure while still being easy to implement is introduced in [9]. To some extent, simulation methods are available, e.g., [10], and allow comparing of different models for specific situations. Motivated by recent incidents, there is also a growing interest in the resilience of critical infrastructures [11]. An overview on models on interdependent CIs is presented in [12], while [13] gives an extensive overview on different models on cascading effects in power systems and presents a comparison of the various approaches.

When it comes to the domain of water supply and water providers as CIs, the amount of research seems to be more

ISBN: 978-1-61208-661-3

limited. In the context of the water sector, some research has been focusing on the security weaknesses of ICSs and SCADA systems and how to find good practices for water providers [14]. Further, effects of an Advanced Persistent Threat (APT) on a water utility provider have been investigated in [15] and [16] due to the increasing number of incidents based on such complex attack strategies. However, there is only little research specifically looking into the situation of a water provider depending on and influencing CIs in its vicinity.

Paper Outline

The remainder of this article is organized as follows: Section II describes the considered use case, Section III analyses the use case, which is further discussed in Section IV and Section V provides concluding remarks.

II. The Situation of a Water Provider

We describe the situation of a hypothetical water provider that we are going to analyze in the next section. Therefore, we are using information which is obtained from discussions with experts from a real-life water provider. The main goal is to illustrate how to analyze the consequences of a risk scenario affecting a CI that is part of a entire network of interdependent CIs. We investigate a utility organization that provides water to more than one hundred municipalities in its surrounding region. The main focus lies on availability of drinking water as well as on the water quality. In order to ensure a sustainable water quality, the provider supports water processing and sewage cleaning by an ICS. For our use case, we assume the existence of a well and a river head, each supported by a pump that conveys the water to the plant where it is further treated (e.g., undesired chemicals are removed or minerals added). A further source of water is a mountain spring nearby. Due to the geography of the landscape transportation paths are short and the number of necessary lines is low. A number of reservoirs are available to ensure supply with water needed to extinguish fire.

Further, the water provider depends on an transportation system, in particular on roads, e.g., to be able to check wells and springs. As any other CI, a water provider crucially depends on electricity (e.g., electric pumps). An internal power plant contributes approximately 30% of the required energy while the rest comes from external providers. Redundancy in the system and an existing emergency power supply help to mitigate this dependency on an electricity provider. In case of a (temporary) interruption of electricity, the utility provider is able to guarantee supply with drinking water up to three days due to available emergency power.

On the other hand, the water provider is important for a number of other infrastructures. In particular, it supplies drinking water to hospitals and grocery stores but also cooling water for hospitals and industrial companies. The actual importance of each of these connections can only be assessed by the CIs that depend on the water provider, which requires discussions with the corresponding experts and thus goes beyond the scope of our use case. A visualization of the use case is given in Figure 1.

Based on a desktop research and discussions with experts, the following risks have been identified as the most significant ones for a water provider:

- R_1: flooding
- R_2: extreme weather conditions
- R_3: leakage of hazardous material (water contamination)

In order to analyze the effects of a realization of one of these risks, we performed a qualitative risk assessment with experts from the water domain. The next section presents the results of this assessment together with a discussion on the consequences of such an incident.

III. Model-based Analysis of an Incident

The situation of the fictitious water provider described above will be analyzed in this section to illustrate how a practical risk analysis based on a theoretical model can be conducted. Based on the stochastic dependency model between CIs [9], consequences of an incident are simulated and the results are then visualized and discussed. All the assessments and estimates given in this paper are of illustrative use only, since it is not possible to disclose the water provider's original sensitive data. However, the data used is based on discussions with experts of the field to be as realistic as possible.

The model we apply is aligned with standard risk assessment methods like ISO31000, and considers a set of interdependent assets, being individual parts of a CI; a water-provider in our case. The water provider maintains a list of *assets*, each of which can be affected by a certain risk scenario. Each asset carries, among others, the following information:

- *Criticality*: How important is the asset for the overall function of the CI (a related question is that on the importance of the CI itself for other depending CIs or the society itself. Such assessments are outside the scope of this article, yet briefly sketched in Section IV to illustrate a possible post-processing of the simulation that we will describe later).
- *Dependencies*: How critical is the asset for the functionality of other related assets? E.g., how important is the mountain spring or well for the water plant (i.e., how much of the water supply is covered by the

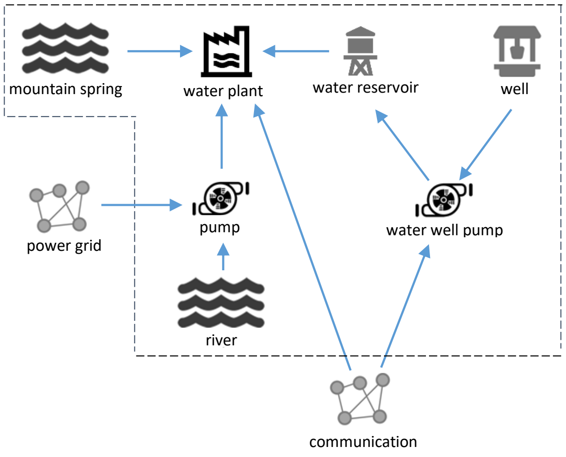

Figure 1. Visualization of Water Use Case

spring, how much is covered by the well, etc.)? How important (e.g., for control and signalling matters) would the company communication or office network be for the service as such, if an outage by a cyber incident or attack occurs?

- *Status indicator*: In normal operation, the assets would all be in working state, but can be in several other states, too (e.g., maintenance). For the risk assessment, the status can be related to the impact when the asset is affected. We shall use the scale $\{1, 2, 3\}$ to express increasing degrees of affection, ranging from status $1 =$"working" up to the worst case status $3 =$"outage", with the intermediate status level expressing anyhow limited functionality. More status levels are of course admissible, yet not used hereafter for the sake of simplicity.

Remark 1: It is important to stress that we use the general term "asset" as a link to standard risk management literature. As such, the term is appropriate for risk management *within* a CI. Adopting a more high-level perspective, such as national authorities may have, their view is on a whole network of CIs, such as power providers, hospitals and water suppliers, with those again depending on each other and so forth. From this high level perspective, a CI is itself an "asset" to the country/nation itself, and we can synonymously exchange the terms CI and asset. Since our focus in this work is on risk management from a single CI provider's perspective, we will hereafter use the term *assets*.

The simulation model will assume a certain incident to "just occur", which in the first place affects some assets by putting them from functional into affected or even outage state. The simulation then uses the dependency information to update the status of related (dependent) assets accordingly, where each asset may undergo individually 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 realization of a risk scenario.

The status transitions are generally probabilistic to cover cases of deterministic dependency (e.g., such as a pump continuously depending on electricity supply), and probabilistic dependencies (e.g., such as water supplies can temporarily be covered from backup water reservoirs). The main duty of the modeling then boils down into two major tasks:

1) Enumerate all assets and identify their interdependencies as detailed as possible. Hereafter, we let the arrow notation $A \to B$ denote a dependency of asset B on asset A (cf. Figure 1, e.g., where the pump B depends on the water A, and similar).
2) Use this information to specify probabilities for status changes in a dependent asset B, if the provider asset A has a status $\neq 1$ (i.e., any abnormal condition, not in normal working state).

The first of these two steps is typically a matter of compiling information that is already known and available to the CI provider. The actual difficulty is the specification of transition probabilities in step two of the above. We believe that this is a general issue in any probabilistic model (not only applying to [9] but also to many others of the references). Nonetheless, the remainder of this work will discuss both aspects in order of appearance.

A. *Identification of Dependencies*

In the beginning, it is necessary to identify all dependencies between the different components of the system. This is not limited to visible (physical) connections but also includes logical connections as in the case of a control system. During the upcoming analysis, it is necessary to assess every link between two components. If the network is large, it may be handy to classify dependencies according to their properties and assign values to every class of connection. In our small example, we refrain from categorizing the connections but rather assess every single connection.

B. *Expert Assessment of Risks*

Once the various components and the interdependencies have been identified, we focus on the assessment of the considered risks and its consequences of a realization in the network. The risk assessments are based on discussions with domain experts that rate each risk as "negligible", "low", "medium", "high" or "very high" while the recovery time is either rated as "short", "medium" or "long". The assessments are given in Table I.

TABLE I. OVERALL LIKELIHOOD ASSESSMENT FOR RISKS

Risk		Occurrence	Failure	Impairment
R_1:	flooding	medium	negligible	negligible
R_2:	extreme weather conditions	medium	negligible	medium
R_3:	leakage of hazardous material	low	negligible	medium

A flooding may affect single sites (e.g., a well), but is not critical for the overall functionality for the water supply as recent incidents like the flooding in central Europe in 2013 have shown. Still, single wells and springs may be used only partly as water may be contaminated by particles (germs, bacteria and others) induced by the flood. Depending on the degree of contamination, water can be boiled to make it drinkable. However, if this is not enough to ensure drinking water quality, the water needs to be purified technically which is a costly and time-consuming process. A realization of risk R_1 may thus yields a limited operation of wells and springs. The risk of an extreme weather situation needs to be considered in further detail based on the type of weather condition. Heavy rain is not a severe problem in our case, since the main source of the water provider is groundwater. It might cause smaller damage to the infrastructure, but will not interrupt water supply. As another extreme, droughts need to be considered, since they are likely to become more frequent in the future. Various sources may dry up, such as rivers or wells, so we may assume (here) that at least some sources like ground water remain available. The drought implies an increased water consumption and yields to peak consumptions that in turn challenge the infrastructure. The 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. As a realization of R_2, we assume an extraordinarily dry period, causing the well to produce only limited outcome while groundwater is still available; due to the drought, water consumption increases significantly at the same time. The realization of this risk may thus be similar as in the previous case which is why we combine the analysis with that of risk R_1.

The assessments related to leakage of hazardous material are challenging as the impact of such an event highly depends on the extent of the leakage. E.g., 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). Nevertheless, if groundwater or several wells are affected, water purification may take several months. Similarly as for the risk of flooding, the amount of hazardous material that has leaked matters a lot. For our use case, we assume that a limited amount affects some parts of the countryside used for water extraction so that a realization of risk R_3 affects the mountain spring. As contamination is a serious problem, we assume the spring switches into the worst state 3.

For our illustrative example, we here assume a scenario where communication is limited due to some internal problems. After some time, a realization of risk R_2 (an extremely dry period of time) or of risk R_3 (a contamination) yields to limited availability of the river source. In the remainder we model the consequences this event has on the other components of the water network. Note that the respective risks, say outages or resource shortages, may also be triggered by cyber-events, e.g., if a hacker switches off the pump or configures the systems towards reduced or zero supply volumes. As such, cyber events may constitute their own risks, but may also be reasons for risk scenarios to “kick in”.

C. Discussions of Consequences of an Incident

While the simulation is able to describe the propagation of the consequences of an incident, the analysis of the overall impact on a specific CI requires knowledge about the effect of a failure of one single component. In particular, it is necessary to estimate how likely it is that a problem or a failure in one component affects the dependent components. These values can be estimated from two sources of information: data from past incidents and expert knowledge. The first source is of limited use when working with critical infrastructures since only few data is known (and even less is publicly available). As for the second source, experts may struggle or be reluctant to estimate precise values, despite their profound knowledge about the infrastructure. Systematic approaches like the Delphi method can help with this issue [17].

Aware of this problem, we avoid asking for exact estimates but rather look for an assessment on a qualitative scale, as is typically recommended in risk management (e.g., by the German Federal Office for Information Security (BSI) [18]). However, this yields to the problem of estimating a whole distribution (namely, all the likelihoods of changing to any of the possible states) from a few qualitative values. In this section, we show one way to approach this problem without pretending an accuracy that cannot be achieved in real life.

In order to determine the transmission probability t_{ij}, a CI needs to answer the following question:

> If your provider is in state i, how likely is it that this will put you into state j?

Since this is usually hard to answer, we replace it by two simpler questions, namely

1) “If your provider is in state i, what is the most likely state j that you will end up with upon this incident?”
2) “How certain are you about your assessment?”

The answers can be chosen from a set of predefined values, namely the number of states $\{1,\ldots,k\}$ for 1) and a set of possible confidence levels for 2). If the expert is unsure about the consequences, we still assume that he has an idea about the intensity of the consequences, i.e., if the expected consequences will be very bad or close to negligible. Because of this, we assume that in the case of uncertain assignments similar values as the predicted one are also possible.

This additional assignment of an assurance value is of twofold benefit. First, it takes pressure form the expert and allows him to choose the answer “I don’t know” (represented by the statement that he is totally unsure about the prediction). Second, this information can be incorporated into the analysis by assigning some likelihood to neighboring values. We propose the following heuristic on an ordered scale of severity:

- If confidence is high (“totally sure”), assign all likelihood to the predicted value j from question 1 above.
- If confidence is medium (“somewhat unsure”), assign likelihood to direct neighbors $j-1$ and $j+1$ (as far as they exist on the scale) such that these are half as likely as the predicted value j.
- If confidence is low (“totally unsure”), assign the same likelihood to all possible values, i.e., choose a uniform distribution over all potential outcomes.

So, for the case of three possible states and the levels of assurance (i.e., the possible answers to question 2) form above) be “totally sure”, “somewhat unsure” and “totally unsure” we take the uncertainty about the assessment into account as follows: if the expert chooses “totally sure”, we assign the likelihood to the proposed status and all other states have a probability of zero. If he chooses “somewhat unsure”, we assign some likelihood to the two neighboring states (i.e., the next smaller and the next larger integer). If we can assume a symmetric situation where a deviation to both sides is equally likely, one approach is to assign to both neighbors half the likelihood of the predicted value. Finally, if the expert chooses “totally unsure”, we assume a uniform distribution over all possible states, representing the situation where we do not have any information at all. The described mapping from a predicted value and a level of uncertainty is explicitly given in Table II. In this table, a triple (p_1, p_2, p_3) represents the distribution over the three possible states, so state k is assumed with probability p_k $(k = 1, 2, 3)$. These estimated distributions then build up the rows of the transition matrices.

As it is quite difficult in practice to make predictions that are totally sure, we incorporate a small chance of an error even for these assessments. That is, we always assign a small probability ϵ to the states nearest to the predicted one, as exemplified in Table III. This makes the model more

TABLE II. 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)

realistic and takes some pressure from the experts performing the assessment.

TABLE III. DISTRIBUTION OVER POSSIBLE NEXT STATE WITH POTENTIAL ERROR

prediction	totally sure	somewhat unsure	totally unsure
1	$(1-\epsilon, \epsilon, 0)$	(2/3, 1/3, 0)	(1/3,1/3,1/3)
2	$(\epsilon/2, 1-\epsilon, \epsilon/2)$	(1/4, 2/4, 1/4)	(1/3,1/3,1/3)
3	$(0, \epsilon, 1-\epsilon)$	(0, 2/3, 1/3)	(1/3,1/3,1/3)

In the upcoming analysis we will consider the cases $\epsilon = 1\%$. We discussed several scenarios with experts from the field to understand the dependencies between the different assets. The assessments are given in Tables IV, V and VI. We measure the impact on a 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 are made for one specific connection and neither contain information about potential substitutes (e.g., if several pumps are available) nor the option of repair or recovery. It is only concerned about the nature of a specific dependence between two assets.

D. Simulation of Incidents

The input to the simulation is a network graph of connected critical infrastructures, where each component of the CI is in one specific state. This graph essentially resembles the picture in Figure 1, and augments each node with a matrix indicating the status change probabilities for each dependency and over time. The time aspect accounts for the fact that short-term outages of a provider may have different impact than long-term outages. E.g., if a power supply goes off, then emergency power supplies may cover for a limited time, thus causing no immediate service interruption. Consequently, the likelihood for a pump, having an emergency supply, to go into outage state 3 if 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 been fixed. However, the same pump is vitally dependent on its water source, and if this runs dry, the pump will immediately go into outage state 3. Therefore, the simulation will need a state transition probability matrix *per dependency* $A \to B$ and *depending on the time scale*.

TABLE IV. SHORT TERM IMPACT ASSESSMENT

Link	Problem	Prediction	Confidence
pump → water plant	limitation	negligible	totally sure
	failure	negligible	totally sure
mountain spring → water plant	limitation	negligible	totally sure
	failure	negligible	totally sure
communication → water plant	limitation	medium	somewhat unsure
	failure	negligible	totally sure
water reservoir → water plant	limitation	negligible	totally sure
	failure	negligible	totally sure
well → well pump	limitation	negligible	totally sure
	failure	negligible	somewhat unsure
communication → well pump	limitation	medium	somewhat unsure
	failure	negligible	totally sure
river → river pump	limitation	negligible	totally sure
	failure	negligible	somewhat unsure
power grid → river pump	limitation	negligible	totally sure
	failure	negligible	totally sure
river pump → water reservoir	limitation	negligible	totally sure
	failure	negligible	totally sure

TABLE V. MEDIUM TERM IMPACT ASSESSMENT

Link	Problem	Prediction	Confidence
pump → water plant	limitation	negligible	totally sure
	failure	negligible	somewhat unsure
mountain spring → water plant	limitation	negligible	totally sure
	failure	negligible	somewhat unsure
communication → water plant	limitation	negligible	totally sure
	failure	negligible	totally sure
water reservoir → water plant	limitation	negligible	totally sure
	failure	negligible	somewhat unsure
well → well pump	limitation	medium	somewhat unsure
	failure	high	somewhat unsure
communication → well pump	limitation	negligible	totally sure
	failure	negligible	totally sure
river → river pump	limitation	medium	somewhat unsure
	failure	high	somewhat unsure
power grid → river pump	limitation	negligible	totally sure
	failure	negligible	totally sure
river pump → water reservoir	limitation	negligible	totally sure
	failure	negligible	somewhat unsure

TABLE VI. 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
well → well pump	limitation	medium	somewhat unsure
	failure	high	totally sure
communication → well pump	limitation	negligible	totally sure
	failure	negligible	totally sure
river → river pump	limitation	medium	somewhat unsure
	failure	high	totally sure
power grid → river pump	limitation	negligible	totally sure
	failure	high	totally sure
river pump → water reservoir	limitation	negligible	totally sure
	failure	medium	somewhat unsure

The simulation prototype we developed [19] embodies this by taking three such matrices, one for short-term, one for medium-term and one for long-term effects in which the probabilities $t_{ij} = \Pr(B \text{ is in state } j | A \text{ switches into state } i)$ describe the transition regime.

While the general model allows a recovery (i.e., switching back into a better status), this is not yet implemented in the current version of the prototype.

IV. RESULTS OF THE ANALYSIS

In a nutshell, the simulation delivers at least three output artifacts:

1) Textual sequence of events with time stamps, and showing the status of all assets at the given time (such lists are usually extensive and are thus not presented here

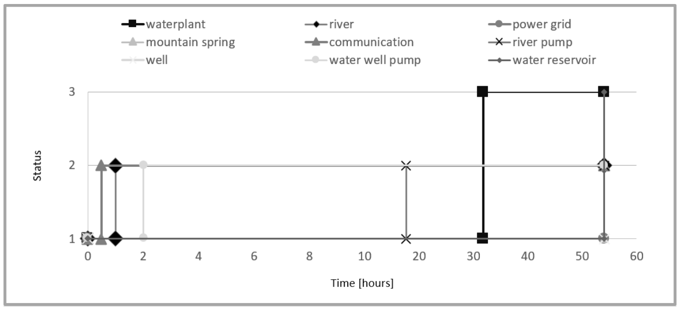

Figure 2. Example simulated time line for a water plant and its components

for space reasons). They are the basic data to compute further information for the risk management, such as the following:

2) Time-lines showing the evolution of the impact on different assets over time. Figure 2 shows an example for nine components in Figure 1.
3) Information about chances on when to expect status changes. Figures 3 and 4 show examples, with explanations to them and the preceding points following below.

Given a set of simulated scenarios, we can average the final states per asset to reflect the likelihood of this part of the CI (or CI network) to become affected (in a degree expressed by the state). For visualization, we 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 can be summarized as a table that lists the number of components which are on average in any of the possible states. We use `OMNeT++` as a tool to support the visualization and execution of our simulation.

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 "evolution" of the cascading impacts. Figure 2 shows an example result for one simulation run.

If numerous simulations are conducted, we can compile the resulting state transition times into an empirical distribution, to learn the expected, median, mode or other characteristic feature of the time when an asset goes into malfunctioning state. E.g., we can measure the expected time until an outage of an asset. Figures 3 and 4 display examples of such a simulation output. Based on this data, we can easily compute the average, i.e., expected, time for a transition from working (1) → outage (3), for the asset "water plant" to be slightly less than five days (with and without the uncertainty of ϵ artificially added to the expert assessment; cf. Table III). In our example, introduction of a small uncertainty yielded to a different empirical distribution of the transition times. If this difference is significant needs to be checked in detail and is beyond the scope of this work but it indicates that potential errors need to be taken into account (just as the concept of trembling hand equilibrium does for game theory) and should not be ignored when analyzing cascading effects.

Usually, the state itself is not exactly a measure of real impact, and needs conversion into a measurable number for management matters. The simulation output will thus in most cases undergo a post-processing that translates the status into a set of facts about what this status actually means, based on the criticality of the asset.

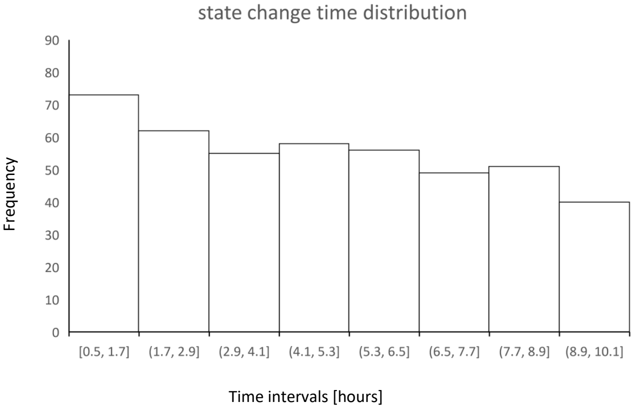

Figure 3. Simulated histogram of $1 \rightarrow 2$ state change times

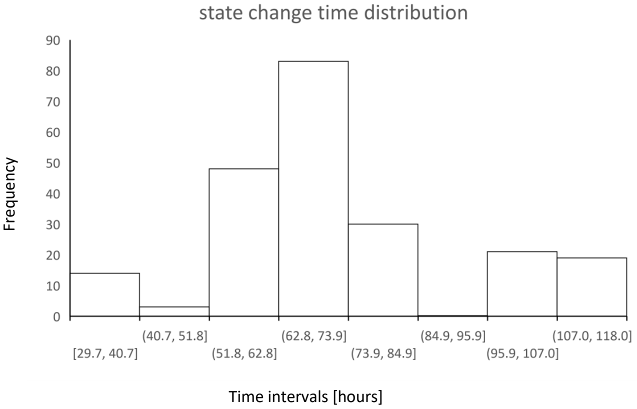

Figure 4. Simulated histogram of $1 \rightarrow 3$ state change times

As for the case of a water utility provider, the degree of damage could depend on the number of affected customers, the time needed to fix the issue, the amount of resources needed to cover the outage, and so forth. Table VII displays an example of such a classification using artificial numbers (for obvious reasons of real data's sensitivity, as already pointed out above) to characterize criticality levels in numeric ranks. In general, criticality levels may also have different meaning for individual scenarios; e.g., if a pump or water tower fails for one day, the criticality may be higher than if water is contaminated, since in the latter case, households can be advised to boil the water before drinking it, whereas if the pump fails, the household would be cut off from water supply completely.

Knowing which parts of the CI network fail at which times and for how long it is a simple matter to apply conditions as exemplified in Table VII to determine the criticality level for this *single* round of simulation.

Repeating this procedure for many times and recording the relative frequencies of occurrence for all criticality levels, we end up with probabilities for each criticality level as $p_i :=$

TABLE VII. DETAILED DESCRIPTIONS OF CRITICALITY LEVELS

	Incident scenario			
Criticality level 1	#1	#2	#3	...
No. of affected households	< 1000	1001...5000	...	...
duration of problem	< 1 day	1...7 days	...	...
costs to fix it (per hour)	100	150	...	...
⋮	⋮	⋮		⋱
Criticality level 2	1	2	3	...
No. of affected households				
duration of problem				
costs to fix it (per hour)				
⋮				

$\Pr($ criticality level $i)$. These likelihoods quantify the odds for running into a certain amount of trouble in a given scenario. Partitioning the range $[0, 1]$ into a fixed set of levels, say in thirds, we can convert these probabilities into *warning levels*. That is, if criticality level 2 occurred in a fraction of 60% of the simulated runs, $p_2 \approx 0.6$ falls into $1/3 < p_2 < 2/3$, giving middle warning level (e.g., yellow alert). Likewise, if criticality level 1 occurred in 90% of the simulation runs, then criticality level 1 has warning level 3 (red alert) in the final output.

It must be kept in mind that the simulation cannot provide any detailed information about the likelihood for an incident as such to occur; the simulation starts straight away from the given scenario that is assumed to have happened.

V. CONCLUSION

A major challenge in the simulation of critical infrastructures is the expert assessment of probabilities for a stochastic simulation. In this context and for the example given in this article, it is important to specify dependencies on a local level only, meaning that the opinion must be formed with consideration on only directly dependent assets, and *not* the overall CI, since this is the purpose of the simulation. We stress that these dependencies are not constrained in nature and physical and cyber-aspects of a CI can be unified under the same modeling framework. Thus, simulation methods like the described one aid even a holistic cyber-physical view on incident propagation in a CI, if dependencies between physical assets (e.g., a hospital) and cyber assets (e.g., the telecommunication network on which the hospital relies for emergency communication and signalling) are included in one model.

An independent difficulty lies in assessing the temporal aspects like the meaning of short-term, medium term and long-term impacts. Certainly, these need to be distinguished, but good heuristics or models to support experts in these regards are rarely available. Polling multiple experts here creates the additional challenge of unifying opinions from different domains, say from experts on the physical matter (like water), vs. people specialized in cyber-security (none of which is necessarily skilled in the other's domain). Aggregating such different assessments into a single value for a simulation is a matter of opinion pooling and subject of supplementary research related to ours (e.g., [20]–[23]). As for future research, it is thus required to develop models that help parameterizing other models. Matters of describing system dynamics are well understood, but helping experts cast their domain knowledge into reasonable figures for a simulation is a challenge on its own. The main contribution of this work is the almost complete picture of the work flow, not least to display the difficulties besides the potential of simulation-based risk analysis in critical infrastructures. While many sophisticated methods of modeling exist, matters of *using* such models have received significantly less attention. Our discussion, though based on a concrete example and method, covers issues of wider applicability. Extending and studying possibilities to make stochastic models more useful is, in our view, an important and promising direction of future research.

ACKNOWLEDGMENT

This work was done in the context of the project "Cross Sectoral Risk Management for Object Protection of Critical Infrastructures (CERBERUS)", supported by the Austrian Research Promotion Agency under grant no. 854766. We thank the experts from Linz AG for fruitful discussions and valuable insights.

REFERENCES

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

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

[3] 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/

[4] 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

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

[6] 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.

[7] 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, 2009, pp. 170–178.

[8] M. Rahnamay-Naeini and M. M. Hayat, "Cascading failures in interdependent infrastructures: An interdependent markov-chain approach," IEEE Transactions on Smart Grid, vol. 7, no. 4, jul 2016, pp. 1997–2006. [Online]. Available: https://doi.org/10.1109/tsg.2016.2539823

[9] 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.

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

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

[12] 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.

[13] 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, dec 2017, pp. 9–22.

[14] 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, 2011, pp. 124–134.

[15] A. Alshawish, M. A. Abid, H. de Meer, S. Schauer, S. König, A. Gouglidis, and D. Hutchison, “Protecting water utility networks from advanced persistent threats: A case study,” in HyRiM, S. Rass and S. Schauer, Eds. Springer International Publishing, 2018, ch. 6.

[16] A. Gouglidis, S. König, B. Green, S. Schauer, K. Rossegger, and D. Hutchison, “Advanced persistent threats in water utility networks: A case study,” in HyRiM, S. Rass and S. Schauer, Eds. Springer International Publishing, 2018, ch. 13.

[17] J. Rohrbaugh, “Improving the quality of group judgment: Social judgment analysis and the delphi technique,” Organizational Behavior and Human Performance, vol. 24, no. 1, 1979, pp. 73–92.

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

[19] T. Grafenauer, S. König, S. Rass, and S. Schauer, “A simulation tool for cascading effects in interdependent critical infrastructures,” in Proceedings of the 13th International Conference on Availability, Reliability and Security, ARES 2018, Hamburg, Germany, August 27-30, 2018, 2018, pp. 30:1–30:8.

[20] S. Rass and S. Schauer, Eds., Game Theory for Security and Risk Management: From Theory to Practice, ser. Static & dynamic game theory : foundations & applications. Cham, Switzerland: Birkhäuser, 2018.

[21] S. Rass, J. Wachter, S. Schauer, and S. König, “Subjektive Risikobewertung – Über Da-tenerhebung und Opinion Pooling,” in D-A-CH Security 2017, P. Schartner and A. Baumann, Eds. syssec, 2017, pp. 225–237.

[22] F. Dietrich and C. List, “Probabilistic opinion pooling generalized. Part one: General agendas,” Social Choice and Welfare, vol. 48, no. 4, 2017, pp. 747–786.

[23] J. Wachter, T. Grafenauer, and S. Rass, “Visual Risk Specification and Aggregation,” in SECURWARE 2017: The Eleventh International Conference on Emerging Security Information, Systems and Technologies, IARIA, Ed., 2017, pp. 93–98.