Simulation-driven Risk Model for Interdependent Critical Infrastructures

Stefan Schauer*
AIT Austrian Institute of Technology†
Stefan.Schauer@ait.ac.at

Stefan Rass
University of Klagenfurt ‡
Stefan.Rass@aau.at

Sandra König
AIT Austrian Institute of Technology
Sandra.Koenig@ait.ac.at

ABSTRACT

Critical infrastructures (CIs) in urban areas or municipalities have evolved into strongly interdependent and highly complex networks. To assess risks in this sophisticated environment, classical risk management approaches require extensions to reflect such interdependencies and include the consequences of cascading effects into the assessment. In this paper, we present a concept for a risk model specifically tailored to the requirements of interdependent CIs. We will show how interdependencies among CI can be reflected in the risk model on different levels of abstraction in a generic way. Furthermore, we will highlight how the simulation of cascading effects can be directly integrated into the risk model to consistently assess the evolution of complex CI systems over time. In this way, the model supports municipalities' decision makers in improving their risk and resilience management of the CIs under their administration.

Keywords

Risk model, risk assessment, interdependent critical infrastructures, cross-domain simulation, cascading effects.

INTRODUCTION

Over the last decade, critical infrastructures (CIs) have become more and more interconnected with and interdependent on each other and have evolved into highly complex networks. In metropolitan areas in particular, where thousands or even millions of people live in a geographically narrow space, the uninterrupted and failure-free operation of CIs is vital to maintain essential services and the social life. As historic events show us, if some incident occurs in one of those CIs, several other CIs can be affected directly or indirectly. For example, in 2019 in Venezuela, an interruption of the main power line caused a major blackout in the capital Caracas (Dube and Castro 2019). As a consequence, not only telecommunication but also the supply with fresh drinking water, medical care and financial services were disrupted.

Risk and resilience management are important tools to tackle such cascading effects and to identify threats, their potential consequences and how to make CIs more robust. Still, risk and resilience management are often carried out with the focus on individual infrastructures, although the interdependencies among CIs been given particular attention over the last decade, starting with the first approaches to characterize them even back in 2001 (Rinaldi et al. 2001). A number of concepts and approaches have already been developed, which focus on the overall CI network and the interdependencies among CIs (see also the next Section for more details), but most of the concepts and frameworks used by CI operators, governmental bodies or city administration (such as ISO 27005 (International Organization for Standardization (ISO) 2018a), NIST SP800-30 (National Institute of Standards and Technology (NIST) 2012) or COBIT 5 for Risk (Information Systems Audit and Control Association (ISACA) 2013)) still focus on individual infrastructures.

*corresponding author
†https://www.ait.ac.at
‡https://www.aau.at

In this paper, we describe a conceptual model for risk assessment and risk management of interdependent CIs. This risk model is developed as part of the research project ODYSSEUS (running from 10/2019 until 09/2021), which focuses on risk and resilience aspects of CI networks and supply networks located within a metropolitan area. One major goal is to support not only CI operators but also governmental bodies and city administration to improve the risk management activities and resilience management in the metropolitan area. Therefore, the proposed risk model is closely aligned with the individual steps of the standard risk management process, as defined in ISO 31000 (International Organization for Standardization (ISO) 2018b), and can easily be integrated into already existing risk management activities. Additionally, the risk model has a strong focus on capturing the complex interdependencies among the CIs and assessing the potential cascading effects stemming from these interdependencies. To achieve that, the model integrates a cross-domain simulation approach which can assess the consequences of cascading effects, hence incorporating them in the risk assessment.

CI PROTECTION AND RISK MANAGEMENT

Over the last years, the protection of CIs has become a prevalent topic all over the world. In particular in Europe, the European Union (EU) has formalized the topic by the Network and Information Security (NIS) directive (European Comission 2016) from 2016, which has been implemented in national law by all EU member states. A revision to the NIS directive, called NIS 2.0, has been proposed in December 2020 (European Comission 2020a) and shifts the focus from individual infrastructures towards the security of supply chains by requiring individual companies to address cyber security risks therein. Additionally, also the European Program for Critical Infrastructure Protection (EPCIP) from 2008 (European Commission 2008) has been revised and extended in 2020 (European Comission 2020b) with a stronger focus on an all-hazards approach and on dependencies among "critical entities" (which can be understood as CIs or subsystems thereof) to improve their resilience in all member states. This way, the EU has responded to incidents that have threatened and disrupted CIs over the last years, such as the hacking of the Ukrainian power grid (Zetter 2016), the WannaCry and (Not-)Petya ransomware attacks in 2017 (Bill 2017; Fox-Brewster 2017) with their global effects on supply chains (PTI 2017), the large-scale power blackout in Venezuela in 2019 (Dube and Castro 2019) and of course the COVID-19 pandemic, which still has wide-ranging effects on the health system and supply chains.

The importance of risk and resilience management concepts as well as the interdependencies among CIs have already been covered in the literature of critical infrastructure protection for almost two decades now. Rinaldi et al. started to explicitly look into the dependencies among CIs in 2001 (Rinaldi et al. 2001), proposing the first versions of interdependency graphs. These approaches got extended, among others, by the Hierarchical Coordinated Bayesian Model (HCBM) (Yan et al. 2006), which provided a more fine-grained description of those interdependency graphs. To estimate the cascading effects stemming from those interdependencies, methods like the Cross Impact Analysis and Interpretative Structural Model (CIA-ISM) (Bañuls and Turoff 2011) provided first insights on how (small) incidents on individual CIs can cause wide-ranging consequences among the entire CI network. However, in their framework, it was not possible to sufficiently model the complex dynamics in the CI network. Therefore, stochastic processes such as percolation theory (Salathé and Jones 2010), Bayesian networks (Schaberreiter et al. 2013) and Interdependent Markov Chains (IDMC) (Wang et al. 2012; Rahnamay-Naeini and Hayat 2016) were applied to model the propagation of cascading effects in the CI network. Recently, a stochastic model (König et al. 2019) as well as a simulation tool (Schauer, Grafenauer, et al. 2020) have been developed. Both the model and the tool describe CIs together with their operational states, analyzes their evolution after an incident has happened and assesses the respective consequences for the overall CI network.

Nowadays, risk managers of CIs are aware of their importance for society and have implemented standard risk management frameworks (e.g., ISO 31000 (International Organization for Standardization (ISO) 2018b), ISO 27005 (International Organization for Standardization (ISO) 2018a), NIST SP800-30 (National Institute of Standards and Technology (NIST) 2012) or COBIT 5 for Risk (Information Systems Audit and Control Association (ISACA) 2013)). Although concepts for integrating the analysis of interdependencies and cascading effects in standardized processes have already been proposed (Schauer 2018), the assessment of cascading effects is often neglected by the CI operators. New stakeholders, such as city administrations or governmental bodies are approaching the topic from an elevated point of view, with a specific focus on the overall CI network and its resilience. In this context, the national research project ODYSSEUS tackles the problem by developing a cross-domain simulation approach (Schauer and Rass 2020) to assess the cascading effects among CIs in a metropolitan area and estimate their risk and resilience.

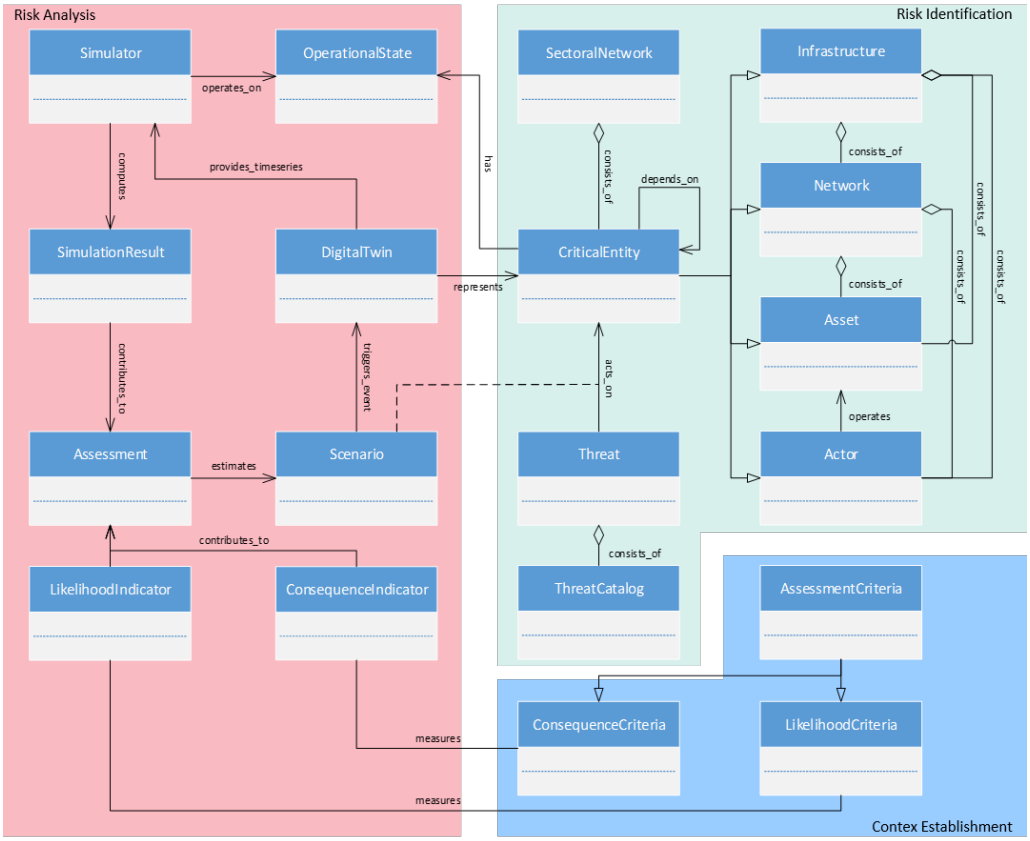

Figure 1. Illustration of the ODYSSEUS Risk Model as UML static diagram
(attributes of classes omitted here to simplify the presentation)

ODYSSEUS RISK MODEL

Risk Model Definition

When scrutinizing the risk assessment for interdependent CIs, three core aspects need to be taken into account: first, a generic format to describe the CI network, i.e., the *individual CIs* (even including their critical subsystems, if necessary) and their interdependencies; second, a *structure* to represent the dynamic aspects of the cascading effects among the CIs; third, an approach to *assess the consequences* of cascading effects according to different categories. The ODYSSEUS risk model we propose here covers all three aspects and integrates the relevant concepts. In Figure 1, we are using a simplified UML class diagram notation for describing these concepts and connect them in one consistent model.

Regarding the individual CIs and their interdependencies, we are starting off with the generic notion of a *CriticalEntity*, which is the core concept of the ODYSSEUS risk model. It represents each (critical) object that is of relevance for the risk assessment. Depending on the level of abstraction used, a *CriticalEntity* can stand for either an entire critical *Infrastructure* (e.g., an energy provider, telecom operator or a hospital), a *Network* inside a CI (e.g., the SCADA or ICT network) or a critical *Asset* (e.g., an important substation, pump or server), that itself be located either within the *Infrastructure* or the *Network*. Additionally, we also included *Actors* in the concept of the *CriticalEntity*, because specific employees, such as IT network administrators, maintenance personnel or others can also be considered as critical for an infrastructure. Each *CriticalEntity* can have multiple *OperationalStates*, which describe the current status of the *CriticalEntity*. The individual states like "normal operation", "limited capacities" or "complete breakdown", are used later on for the cross-domain simulation.

The structure of a *CriticalEntity* is chose in this way to make the ODYSSEUS risk model easily adaptable to many different situations within a CI. Furthermore, the definition of a *CriticalEntity* and its connected concepts is extensible with additional information if required. For example, data on the lifespan of critical *Assets* within a CI can be integrated into the model based on manufacturer specifications, or historic data collected by the CI. In this way, the model would also allow to specify a hazard function indicating different phases in the lifespan of an asset (and thus also indicate a change of the *OperationalState* of a *CriticalEntity*). This can be done, for example, using a

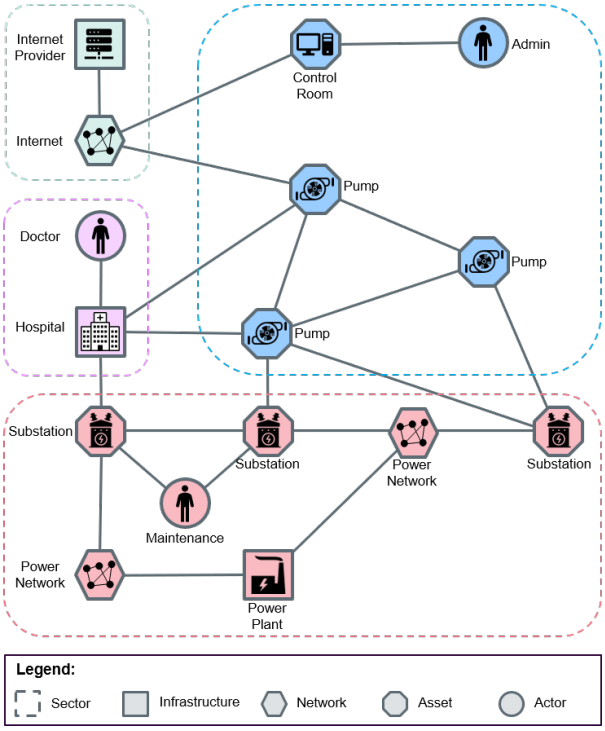

Figure 2. High-level example of an interdependency graph using the concepts from the risk model.

Weibull distribution function, which is popular for modelling reliability. Additionally, other metrics such as Mean Time to Failure (MTTF), Mean Time Between Failures (MTBF) or Mean Time to Repair (MTTR) can be used to describe the behaviour of the *CriticalEntities* over time.

To capture the interdependencies among all the *CriticalEntities*, we explicitly included the *depends_on* relation in the risk model. This allows us to directly build the interdependency graph later on, which serves as the main structure to simulate and analyse the cascading effects. As we are focusing in ODYSSEUS on several interdependent networks within an metropolitan area in ODYSSEUS, the risk model also allows to define *SectorNetworks*. They consist of *CriticalEntities*, i.e., each *SectorNetwork* can be built from individual *Infrastructures*, entire (sub-) *Networks*, critical *Assets* and *Actors*. The *CriticalEntity* and the related classes up to *SectorNetwork* mainly contribute to the first step in the risk management Process, commonly referred to as “Establishing the Context”.

To illustrate how an interdependency graph is constructible based on *CriticalEntities* and the related concepts, Fig. 2 provides a small example. It shows four *SectoralNetworks* as dashed rectangles, covering electric power supply (red), water supply (blue), health services (violet) and telecommunication services (green). The Power Plant, Hospital and Internet Provider are modelled as *Infrastructures* (rectangles), as they represent complex entities. For their assessment, no deep functional details of the behaviour are needed, and an abstract description at high level is sufficient. Although the networks operated in the sectors, e.g., the Power Network or the Internet, are highly complex systems, they are also represented as single entities of type *Network* (hexagons). However, in the case of water supply, the network is of particular interest and therefore modelled at a finer level of detail up to specific important *Assets* (octagons) in the network, including several Pumps as well as a Control Room. Finally, some *Actors* (circles) appear in this example graph, e.g., the Admin of the Control Room or the Maintenance personnel responsible for the Substations. Edges among the *CriticalEntities* represent *depends_on*-relations as identified in the risk model. They establish the connection not only among entities within one sector but also across *SectoralNetworks*. In this way, the ODYSSEUS risk model allows to compile diverse entities from different sectors, as relevant for an assessment, into a unified graph representation. This graph then serves as the basis for a subsequent simulation of cascading effects. The graph in Fig. 2 is hereby intentionally kept small for illustrative reasons, and is therefore not complete (some dependencies are omitted to simplify the picture, e.g., the Internet Provider has no connection to the power grid); a real life interdependency graph would be much more complex.

Before turning to the dynamic aspects of cascading effects, we introduce the concept of *Threats*. In the ODYSSEUS risk model, *Threats* are collected in a *ThreatCatalog*, which can represent any real-life threat catalog. This provides the ability to freely define the scope of the risk model, as only those threats relevant for the overall assessment are

modeled as *Threats* and integrated in the *ThreatCatalog*. Further, a *Threat* can *act_on* one or more *CriticalEntities* as part of a *Scenario*. In this way, the risk model implements the general relation between entities, threats and scenarios that is usually found in standard risk management frameworks. Using scenarios makes it also easier to specify the time frame, location and magnitude of a threat taking place.

The analysis of the effects of a specific threat in the context of a scenario is covered in the risk model by the *Assessment* class. As usual in risk management, to assess the effects, the likelihood and the consequences need to be estimated. Therefore, the risk model is following a semi-quantitative approach and uses a set of *AssessmentCriteria*, consisting of *LikelihoodCriteria* (setting the scales for describing the likelihood) and *ConsequenceCriteria* (setting the scales for describing the consequences), which are usually defined in the course of the "Context Establishment" step of the risk management process (see next section for details). In the context of CIs, assessing the consequences of an incident can be particularly challenging, since the consequences have multiple dimensions that need individual quantifications and measurement: these may include, but are not limited to, monetary costs, ecological impacts, lives affected (e.g., injured), and many others. To reflect this multitude also in the proposed risk model, several *ConsequenceIndicators* are definable for an *Assessment* to account for consequences in different aspects. For example, the likelihood is often a scalar quantity in ordinal units, such as “low”, “medium” or “high” probability that some scenario happens (in a subjective assessment), or as a relative number of occurrences for a given time frame (conditional on the availability of enough information from historic events, e.g., how many times the event happened or will happen in one year). For extended flexibility, our model allows to define multiple *LikelihoodIndicators* to describe the occurrence of events in an *Assessment* (and also to be consistent with the estimation of the consequences). Such indicators can describe, for example, the plausibility of an incident to occur or the capabilities, motivation and resources of an adversary. These indicators can be particularly helpful when dealing with intentional threats, where the likelihood cannot be estimated well based on historic events but more specific information on the adversary needs to be considered. Examples for *LikelihoodIndicators* and *ConsequenceIndicators* will be given below in Figures 5 and 6.

The quantification for both the consequences and the likelihood is herein always in terms of categories that define the meaning of natural-language terms like ”low” or ”high” relative to the application context. This assures that the framework is flexibly applicable to different domains, at the cost of implying that categorizations of the same name but referring to different domains are generally not comparable or even incompatible. That is, a ”high” likelihood for system X can mean something entirely different from a ”high” likelihood for system Y. Similarly, may impact rankings, say ”medium”, may have incomparable interpretations when they refer to system X or system Y.

To include the dynamics of cascading effects into the assessment of a scenario, we are using a cross-domain simulation approach (Schauer and Rass 2020), which simulates the cascading effects within the CI network and highlights the critical entities affected by a specific incident. In the risk model, this approach is represented by the *DigitalTwin* and *Simulator* classes. The *DigitalTwin* serves as a representation of a *CriticalEntity* (technically, this is a neural network, see further details in next section) and produces time series which describe the behavior of the *CriticalEntity* in a specific scenario. The *Simulator* then takes these time series and simulates the cascading effects in the overall CI network (the graph structure is coming from the *depends_on* realtions among the *CriticalEntities*). As a result, a list of *CriticalEntities* together with their *OperationalStates* is produced, which is modeled as the *SimulationResult* class. Those results directly contribute to the risk assessment and thus are connected to the *Assessment* class. In detail, the *SimulationResults* represent the consequences of the cascading effects a *Scenario* has on the CI network and thus are measured according to the *ConsequenceIndicators*.

Integration into the Risk Management Process

One core requirement for the ODYSSEUS risk model is its applicability and easy integration into the risk management processes currently implemented within city administration and other governmental bodies. Therefore, the risk model is strongly aligned with the overall risk management process described in ISO 31000 (International Organization for Standardization (ISO) 2018b), which has been established as the standard framework for risk management. It has been adopted in several other frameworks (e.g., ISO 27005 (International Organization for Standardization (ISO) 2018a), NIST SP800-30 (National Institute of Standards and Technology (NIST) 2012) and others) and therefore is already applied in multiple organizations. The process itself is divided into five main steps, "Establishing the Context", "Risk Identification", "Risk Analysis", "Risk Evaluation", and "Risk Treatment", together with two concurrent activities, "Communication and Consultation" and "Monitoring and Review" (see also Figure 3. However, when it comes to interdependent CIs, there are several issues that need particular attention, e.g., the analysis of the consequences of cascading effects, since they need to be treated differently compared to looking at CIs individually. With the ODYSSEUS risk model particularly addressing the core steps "Risk Identification", "Risk Analysis" and "Risk Evaluation", it is specifically tailored to supporting a risk assessment for

Figure 3. Illustration of the generic ISO 31000 risk management process.

the complex environment of interdependent CIs. In the following, we will briefly address the different steps of the risk management process and highlight where the important aspects when considering interdependent CIs.

During *Context Establishment*, often many of the topics influencing the context are already given or described in some documentation. When looking at the external context, political and legal environment is usually given by national security strategies and law; with regards to CIs in Europe and EU countries, the NIS Directive (European Comission 2016) and its national laws are part of the external context which influences the risk assessment. When focusing on interdependent CIs, it is a little more difficult to define the internal context, since each individual CI in the CI network has its own organization culture, goals and processes but also operates different technologies and manages different infrastructure assets. Thus, an important step in the *Context Establishment* is create a consistent structure to describe the interdependencies among the carious CIs. ODYSSEUS risk model particularly addresses this aspect with the concept of *CriticalEntities* and the *depends_on* relation to model the interdependency graph structure.

The core aspect of *Risk Identification* is to compile a list of threats that might act on the CI network as a whole or the CIs individually. As the CIs in the network usually are of various type and operate in different industry sectors, they are facing a broad variety of different threats. Additionally, the threats can also be rather technical and thus not applicable to all (maybe just to a few) CIs. Hence, it is necessary create a more abstract view on threats when considering interdependent CIs and therefore use rather generic threat catalogues. Such catalogues are provided, for example, from national risk assessments in different countries, which gives a very detailed and usually a good fit to cover a large variety of threats. For Switzerland, one such catalogue is available from the Bundesamt für Bevölkerungsschutz (BABS) 2019. Using the *ThreatCatalog* class in the risk model, these catalogues can be directly covered by the risk model.

The step *Risk Analysis* mainly deals with assessing the likelihood and consequences of the identified threats. As already pointed out above, this is one of the most difficult tasks in risk management, since it is not trivial how to assess the likelihood and consequences when it comes to CIs. In general, we cannot expect a lot of data available to perform a completely quantitative assessment and thus the semi-quantitative approach covered in the ODYSSEUS risk model perfectly fits to the given situation. By using the concepts of *LikelihoodCriteria* and *ConsequenceCriteria* together with the respective *LikelihoodIndicators* and *ConsequenceIndicators*, it makes it easier to integrate more subjective assessment methods like expert workshops to obtain a good estimate of an incident's likelihood and consequences.

When taking interdependent CIs into account, it is particularly important to assess the influences among the CIs and evaluate the cascading effects, i.e., how a specific incident affects not only one CI itself but the entire CI network. This is a task that is only partly addressed in classical risk management frameworks, since interdependencies are hardly taken into account therein. With the concepts of an entity's *OperationalStates* and of the *DigitalTwin* as digital representation of an entity, the dynamic aspects are completely integrated into the ODYSSEUS risk model. As this is one of the main aspectsof the ODYSSEUS risk model, we will describe the interplay between the risk model and the cross-domain simulation approach in further detail in the next section.

In *Risk Evaluation*, the combination of the likelihood and the consequences into a single risk score is taking place, this is usually done by multiplying the numbers representing the likelihood and the consequences and the visualizing it in a risk matrix. However, since we are dealing with *LikelihoodIndicators* and *ConsequenceIndicators*, this becomes much more difficult. One solution is to break down the individual categories onto one single number (e.g., by using an arithmetic mean) or to illustrate the indicators in different risk matrices. The ODYSSEUS risk model provides the freedom to use any method computing and visualizing a risk score and therefore can be adapted to many application areas.

Based on the results from *Risk Evaluation* the *Risk Treatment* evaluates which of the identified risks should be treated first and what could be countermeasures to reduce the risk. Common activities are to reduce either the likelihood or the consequences or even get rid of the system which causes the threat (which is usually not an option when dealing with interdependent CIs).

RISK ASSESSMENT

Cross-Domain Simulation

The ODYSSEUS risk model is specifically focusing on the interdependencies among CIs and the potential cascading effects, which result from an incident taking place in one of those CIs. The main tool to assess those cascading effects is a cross-domain simulation approach (Schauer and Rass 2020), which is directly integrated in the structure of the risk model.

A general technical challenge is the integration of arbitrarily heterogeneous simulation systems in a joint co-simulation environment. We can hardly expect a simulation system for one CI, e.g., water, to be technically compatible to a different simulation environment for another CI, e.g., traffic. An additional dimension of concern is the sensitivity of the information embodied in simulation environments: even if two infrastructure providers maintain simulation models for their own (internal) risk assessment, they may be reluctant to opening up their environments to others for purposes of co-simulation, simply because this may leak out highly sensitive information. To overcome both difficulties, the technical compatibility as well as to respect data protection concerns, the ODYSSEUS risk model pursues the use of *DigitalTwins* that emulate the simulation environments of real-life CIs, while hiding their internal structures and details at the same time. Using machine learning, specifically artificial neural networks (ANN), we gain a mutually compatible representation of different simulations (compatible up to the consistency of physical parameters exchanged between CIs for the co-simulation), and avoid the need of access to a confidential simulation system. Overall, the digital twin is just a trained ANN that "mimics" the output of a simulation, based on training that can be done in a preparatory phase and under strong organizational and technical protection of the infrastructure provider.

The integration of these domain-specific simulations into an overall co-simulation model is based on the SAURON simulation model (Schauer, Grafenauer, et al. 2020), using probabilistic Mealy automata to describe the operational states of CIs, and invokes the individual digital twins of the corresponding CIs to determine their dynamic incident response behavior for the co-simulation, as would be delivered (approximately in the same way) by the real simulator or the real system. This method explicitly addresses the different domains that are covered in the context of a metropolitan area on a level that is separate to the interdependency level. Thus, it is useful to distinguish the Inter-Domain view, i.e., the co-simulation level, from the Intra-Domain view, i.e., the domain simulation level (see Figure 4.

The Inter-Domain view strongly relies on the SAURON model (König et al. 2019), i.e., it consists of a graph structure, where the nodes represent the critical entities and the edges represent the interdependencies among them. The basic concept of the Inter-domain model (and also of the SAURON model) is to represent the behavior of a node (i.e., a critical entity or CI) in the graph by using different abstract operational states (e.g., from "normal operation" to "complete breakdown") and to describe how the node changes from one operational state to another. Therefore, each node is modeled as a Mealy automaton, where the internal state represent the node's operational states and probabilistic transitions between those states, which represent the evolution from one operational state to another. One way that those state transitions can be triggered is by an event happening to the node itself, e.g., an incident such as a cyber attack). A second way to trigger a state transition is that a direct neighbor of the node in the graph (i.e., another critical entity or CI on, which provides some service the node is depending on) changes its operational state.

In the cross domain simulation approach, the Intra-Domain view extends the Inter-Domain view by improving the description of the CIs as well as the domains they are located in, using the aforementioned ANNs. Those ANNs emulate the behavior of individual nodes and in this way the probabilistic state transitions of the SAURON model

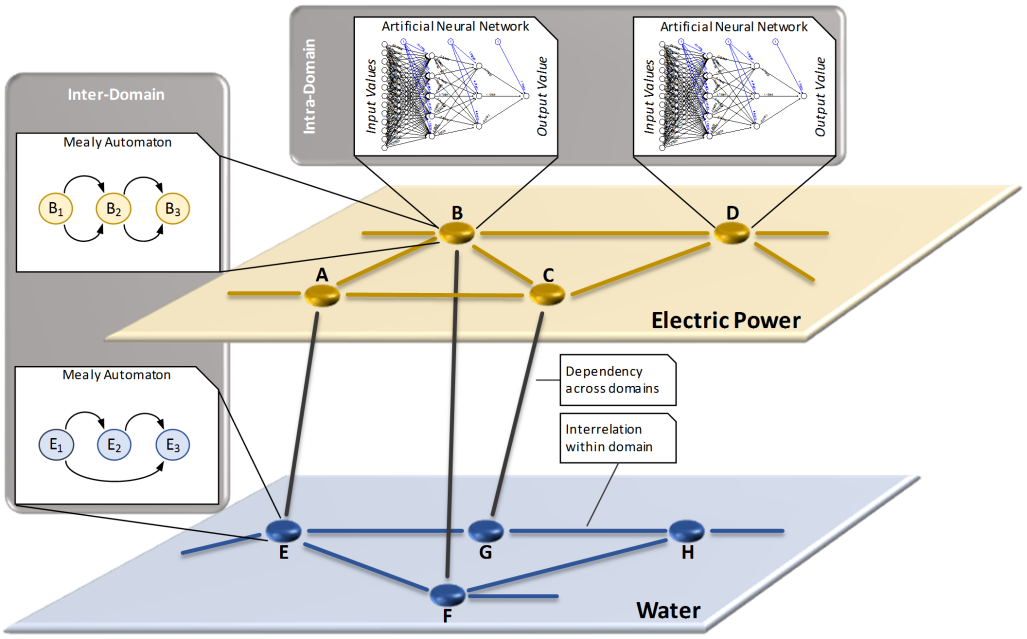

Figure 4. Illustration of the interplay between the Inter-Domain and the Intra-Domain model.

are extended, i.e. controlled, by the calculations of probabilistic ANNs (i.e., deterministic ANNs taking auxiliary random inputs). An event leading to a temporally complex behavior of the CI is here emulated using an ANN by generating a corresponding time series by the ANN, which in turn controls the state transitions of the SAURON node. To do this, an ANN is added to a SAURON node as shown in Figure 4. This ANN is specifically modeled and trained to represent the CI in the SAURON model, based on the temporal behavior of the CI corresponding to the SAURON node, or part(s) thereof. The tasks of an ANN can range from complex simulations to simple replications of ON/OFF dynamics.

For ODYSSEUS, a special form of artificial neural networks is used, namely recurrent neural networks (RNNs). These are particularly suitable for working with time series. Time dependencies are relevant for ODYSSEUS, since different domains (e.g., the traffic network or the power network of a city) are modeled and analyzed. A RNN can infer from the observed temporal relationships and can thus predict the behavior of the system in the next time unit based on a given time series. In the same way, repair times can also be emulated if an RNN is "started" by a failure event and generates a time series from which the current repair time can be read out over time, and the end can be predicted using learned distributions for repair times.

Risk Assessment

As mentioned during the description of the risk model in the previous section, the risk assessment is a crucial step in the risk management process and thus also a crucial part of the proposed risk model. The difficulties in this aspect evolve mainly around the estimation of the likelihood and the consequences of a specific incident scenario. In the context of CIs and in particular when looking at the interdependencies among them and the potential cascading effects stemming from those interdependencies, a large amount of uncertainty is introduced into the risk assessment. First of all, as CIs are complex systems, operating various machinery together with numerous monitoring and control systems as part of sophisticated processes. Therefore, it is not possible to model a CI in full detail; too many of the processes and internal functionality cannot be captured explicitly. Additionally, the interdependencies with other CIs introduce additional uncertainty, as the dynamics among CIs often are not explicitly known, i.e., usually they are not exposed before an incident is happening. Although simulation frameworks like the Cross-domain Simulation described above support the estimation of the cascading effects, the estimation of the consequences is still subject to uncertainty. Therefore, the proposed ODYSSEUS risk model follows a semi-quantitative approach for estimating the likelihood and consequences of an incident scenario by employing *LikelihoodCriteria* and *ConsequenceCriteria* together with the respective *LikelihoodIndicators* and *ConsequenceIndicators*. In this way, the ODYSSEUS risk model provides a multidimensional approach to tackle this uncertainty.

The *LikelihoodIndicators* and *ConsequenceIndicators* allow to capture different aspects when assessing the likelihood and impact of a specific incident, which leads to a better description of both elements compared to classical

Effects on People		
Value	Category	Description
5	Catastrophic	More than 25 slightly injured more than 15 severly injured more than 3 casualties
4	Significant	Up to 25 people slightly injured up to 15 people severely injured up to 3 casualties
3	Moderate	Up to 15 people slightly injured up to 5 people severely injured no casualties
2	Limited	Up to 5 people slighly injured, no severe injuries, no casualties
1	Insignificant	One person slightly injured, no severe injuries, no casualties

Effects on Assets		
Value	Category	Description
5	Catastrophic	Damage caused of more than € 300.000
4	Significant	Damage caused between € 150.000 and € 300.000
3	Moderate	Damage caused between € 50.000 and € 150.000
2	Limited	Damage caused between € 10.000 and € 50.000
1	Insignificant	Damage caused up to € 10.000

Figure 5. Example for the definition of consequence indicators for the effects on people and on assets.

approaches in risk assessment that put all information into a single scalar quantity (whether categorical or numeric). Regarding the assessment of likelihood, usually one scale is used, which indicates the relative frequency of an incident to happen in a specific time frame (e.g., one year). This can be misleading, on the one hand, when too little data is available on a specific threat, i.e., an extremely rare event, such that a relative frequency cannot be estimated based on that data without including a large amount of uncertainty. Using additional *LikelihoodIndicators* like the plausibility of that event to happen or experts' opinions on the current situation, this uncertainty can be addressed and reduced. On the other hand, when looking at complex threats like attacks, the intention behind the adversary highly influences the likelihood of such a threat to happen. In that case, indicators capturing, for example, the motivation, capabilities and resources of the adversary provide a better estimation (and thus reduces the uncertainty stemming from simply "guessing" the adversary's intention).

Due to the wide-spread effects an incident within a CI might have on society, also the consequences of this incident can only be insufficiently measured in one quantity (e.g., monetary value, which is often used in risk assessment of companies). This becomes even more valid when the interdependencies among CIs also need to be taken into account and the assessed values are subject to uncertainty. However, when several *ConsequenceIndicators* are defined to cover multiple societal effects as in the proposed ODYSSEUS risk model, different aspects of the consequences can be captured much better, i.e., focusing on economic and ecological impacts, effects on the health of individuals as well as the well-being and functioning of the society as a whole. In this way, the uncertainty regarding potential consequences can be reduced by this variety of indicators, providing a better overview on the consequences of an incident.

To provide a definition for the ODYSSEUS risk model, we partly adopt a guideline for risk management in crisis management, which was issued by the Austrian Ministry of the Interior (Bundesministerium für Inneres 2018). In detail, we use a set of several consequence indicators, reflecting the effects on people (physical and psychological damage), on assets (damage to infrastructure), on the economy in general (damage on supply chains), on the ecology (damage to ecosystems) and on the political and social system (political controversies or insecurity in the society). In this way, effects in almost all relevant aspects of society can be modeled. For the assessment of each individual indicator, we use a five-tier scale, ranging from "insignificant" to "critical" (see Figure 5).

Regarding the likelihood, we diverge a little from the guideline by using the *LikelihoodIndicators*. In this way, the proposed risk model allows to estimate the likelihood according to various aspects, e.g., historic events (how often a similar incident took place in the past), the exposure of a critical entity to a given incident, or the motivation and capabilities of an attacker (which is highly relevant for any sort of intentional attack). For sake of accordance with the consequence indicators, we also use a five-tier scale to describe each indicator (see Figure 6).

By implementing the concepts of *LikelihoodIndicators* and *ConsequenceIndicators*, a risk assessment carried out according to the ODYSSEUS risk model contains much more information than, for example, an assessment in a classical risk management process. Together with the cross-domain simulation approach, the effects on the entire CI network can be estimated based on multiple simulation runs. Hence, this additional amount of information and data on the effects of a particular incident on the different aspects of social life will provide an improved overview on the consequences with regards to the overall CI network. This will also have a direct influence on the selection of mitigation measures to reduce the risk level.

Frequency of Ocurrence		
Value	Category	Description
5	Highly likely	Happens at least once per year
4	Likely	Happens once in roughly 10 years
3	Probable	Happens once in roughly 100 years
2	Unlikely	Happens once in roughly 300 years
1	Highly unlikely	Happens less than once in roughly 300 years

Adversary's Capabilities		
Value	Category	Description
5	Very high	very sophisticated level of expertise, well-resourced, can generate opportunities
4	High	sophisticated level of expertise, significant resources and opportunities
3	Moderate	moderate resources, expertise, and opportunities
2	Low	limited resources, expertise, and opportunities
1	Very low	very limited resources, expertise, and opportunities

Figure 6. Example for the definition of likelihood indicators for the frequency of occurrence and the adversary's capabilities.

APPLICATIONS AND LIMITATIONS

A general limitation of the co-simulation by integration of emulations is the required amount of training data to craft a reasonably accurate digital twin. While neural networks exhibit nice "interpolating" behavior in some sense, for it to be an accurate substitute of a real simulation, it takes vast amounts of training data. This number further increases with the lot of parameters that can be adjusted, since a domain simulator may come with dozens of levers to push and pull, while the ANN will at some point need to be fixed with a given number of inputs that it can work with. This number determines how much training information is needed to approximate the simulator to a reasonable precision.

At this point, the semi-quantitative nature of the risk management comes in handy, since we actually may not need an arbitrarily accurate approximation of the real-life behavior, as long as the impact and likelihood categories are correctly determined. That is, if a category is defined by an interval $[a, b]$, say, describing monetary losses between a \$ and less than b \$, it makes no difference if the simulation would predict a loss of c \$, while the digital twin produces a loss of d \$ $\neq c$, as long as both fall into the same category $c, d \in [a, b]$. A second advantage of the digital twin is speed: simulations may take long time, and are thus not necessarily feasible to run in reality for a risk assessment. Contrary to this, the digital twin implementing an ANN can provide the result in an instant, making it feasible for many repetitions, in particular when systems exhibit stochastic behavior (which we most likely can expect).

Besides the data required to build the digital twins, a lot of qualitative information about the general structure and interdependencies among the CIs is necessary to build a realistic model of the CI network. A standard method to collect this information is through questionnaires and workshops with experts of the individual CIs to obtain a concise and consistent view on the network. This method is of course quite resource-intensive but can also build on already existing information gathered by CI operators and governmental bodies. Nevertheless, the CI operators need to be willing to provide this information, which might be not always the case due to the sensitive nature of the data. However, national legislation like the NIS law can provide a framework for the collaboration between CI operators and governmental bodies.

CONCLUSION

In this paper, we presented a concept for a risk model that is specifically tailored to the risk assessment requirements for interdependent CIs within a municipality. Therefore, this risk model allows to explicitly define the dependencies among "critical entities" (i.e., CIs or critical parts thereof) and thus to build an interdependency graph of the CI network. Additionally, we showed how the core aspects of a cross-domain simulation approach are integrated into the risk model such that the simulation and assessment of cascading effects is directly supported. To assure the applicability and practical relevance of the risk model, it is closely aligned to the generic ISO 31000 risk management process. As part of the national research project ODYSSEUS, the risk model is the basis for a framework to support the decision makers within municipalities or governmental bodies in assessing various threats and thus improving the resilience of the CIs under their administration. In this way, the risk model provides a structured way to integrate simulation-driven assessments in today's risk management for interdependent CIs.

In the course of the ODYSSEUS project, an artificial city already has been built and a set of use cases is formulated, which describe several incident scenarios acting on different parts of the city. As a next step, an abstract model

of the city is created with the critical entities therein being described according to the proposed risk model. By assessing these use cases, the risk model will be evaluated under realistic conditions. Further, the assessment process as well as the results from the assessment will be discussed by risk and security experts from CIs and municipalities in Austria. Their feedback will be used to improve the risk model to adjust it more closely to realistic risk assessment processes.

ACKNOWLEDGEMENTS

This work was supported by the research Project ODYSSEUS ("Simulation und Analyse kritischer Netzwerk-Infrastrukturen in Städten") funded by the Austrian Research Promotion Agency under Grant No. 873539.

REFERENCES

Bañuls, V. A. and Turoff, M. (2011). "Scenario construction via Delphi and cross-impact analysis". In: *Technological Forecasting and Social Change*. The Delphi technique: Past, present, and future prospects 78.9, pp. 1579–1602.

Bill, B. (2017). *WannaCry: the ransomware worm that didn't arrive on a phishing hook*. Tech. rep. Sophos Ltd.

Bundesamt für Bevölkerungsschutz (BABS) (2019). *Katalog der Gefährdungen. Katastrophen und Notlagen Schweiz*. Tech. rep. Bern, Switzerland: Bundesamt für Bevölkerungsschutz (BABS), pp. 1–44.

Bundesministerium für Inneres (2018). *Risikomanagement im Katastrophenmanagement - Leitfaden*. Tech. rep. Vienna, Austria: Bundesministerium für Inneres.

Dube, R. and Castro, M. (Mar. 2019). "Venezuela Blackout Plunges Millions Into Darkness". en-US. In: *Wall Street Journal*.

European Comission (2016). "DIRECTIVE (EU) 2016/1148 OF THE EUROPEAN PARLIAMENT AND OF THE COUNCIL of 6 July 2016 concerning measures for a high common level of security of network and information systems across the Union". In: *Official Journal of the European Union*, pp. L 194/1.

European Comission (2020a). *Proposal for a DIRECTIVE OF THE EUROPEAN PARLIAMENT AND OF THE COUNCIL on measures for a high common level of cybersecurity across the Union, repealing Directive (EU) 2016/1148*. Tech. rep. 2020/0359(COD). Brussels, Belgium, pp. 1–108.

European Comission (2020b). *Proposal for a DIRECTIVE OF THE EUROPEAN PARLIAMENT AND OF THE COUNCIL on the resilience of critical entities*. Tech. rep. 2020/0365 (COD). Brussels, Belgium, pp. 1–56.

European Commission (2008). "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". In: *Official Journal of the European Union* L345, pp. 75–82.

Fox-Brewster, T. (2017). *Petya Or NotPetya: Why The Latest Ransomware Is Deadlier Than WannaCry*.

Information Systems Audit and Control Association (ISACA) (2013). *COBIT 5 for Risk*. Rolling Meadows, USA: ISA.

International Organization for Standardization (ISO) (2018a). *ISO/IEC 27005:2018 Information technology - Security techniques - Information security risk management*. en. Geneva, Switzerland.

International Organization for Standardization (ISO) (2018b). *ISO/IEC 31000:2018 Risk Management - Principles and Guidelines*. englisch. Geneva, Switzerland.

König, S., Rass, S., Rainer, B., and Schauer, S. (2019). "Hybrid Dependencies Between Cyber and Physical Systems". en. In: *Intelligent Computing*. Ed. by K. Arai, R. Bhatia, and S. Kapoor. Advances in Intelligent Systems and Computing. Springer International Publishing, pp. 550–565.

National Institute of Standards and Technology (NIST) (2012). *NIST SP800-30 Rev. 1 Guide for Conducting Risk Assessments*. Gaithersburg, USA.

PTI (2017). *New malware hits JNPT operations as APM Terminals hacked globally*.

Rahnamay-Naeini, M. and Hayat, M. M. (2016). "Cascading Failures in Interdependent Infrastructures: An Interdependent Markov-Chain Approach". In: *IEEE Transactions on Smart Grid* 7.4, pp. 1997–2006.

Rinaldi, S. M., Peerenboom, J. P., and Kelly, T. K. (2001). "Identifying, understanding, and analyzing critical infrastructure interdependencies". In: *IEEE Control Systems* 21.6, pp. 11–25.

Salathé, M. and Jones, J. H. (2010). "Dynamics and Control of Diseases in Networks with Community Structure". In: *PLOS Computational Biology* 6.4, e1000736.

Schaberreiter, T., Bouvry, P., Röning, J., and Khadraoui, D. (2013). "A Bayesian Network Based Critical Infrastructure Risk Model". In: *EVOLVE - A Bridge between Probability, Set Oriented Numerics, and Evolutionary Computation II*. Ed. by O. Schütze, C. A. Coello Coello, A.-A. Tantar, E. Tantar, P. Bouvry, P. Del Moral, and P. Legrand. Vol. 175. Berlin, Heidelberg: Springer Berlin Heidelberg, pp. 207–218.

Schauer, S. (2018). "A Risk Management Approach for Highly Interconnected Networks". In: *Game Theory for Security and Risk Management*. Ed. by S. Rass and S. Schauer. Cham: Springer International Publishing, pp. 285–311.

Schauer, S., Grafenauer, T., König, S., Warum, M., Rass, S., and Rass, S. (2020). "Estimating Cascading Effects in Cyber-Physical Critical Infrastructures". en. In: *Critical Information Infrastructures Security*. Lecture Notes in Computer Science. Linköping, Sweden: Springer International Publishing, pp. 43–56.

Schauer, S. and Rass, S. (2020). "Creating a Cross-Domain Simulation Framework for Risk Analyses of Cities". en. In: *Critical Infrastructure Protection XIV*. Ed. by J. Staggs and S. Shenoi. IFIP Advances in Information and Communication Technology. Cham: Springer International Publishing, pp. 307–323.

Wang, Z., Scaglione, A., and Thomas, R. J. (2012). "A Markov-Transition Model for Cascading Failures in Power Grids". In: *2012 45th Hawaii International Conference on System Sciences*, pp. 2115–2124.

Yan, Z., Haimes, Y. Y., and Wallner, M. G. (2006). *Hierarchical coordinated Bayesian model for risk analysis with sparse data*. Baltimore, USA.

Zetter, K. (2016). *Everything We Know About Ukraine's Power Plant Hack | WIRED*.