BUCEPHALUS: a BUsiness CEntric cybersecurity Platform for proActive anaLysis Using visual analyticS

Marco Angelini, Graziano Blasilli, Silvia Bonomi, Simone Lenti, Alessia Palleschi, Giuseppe Santucci*
Sapienza University of Rome
Emiliano De Paoli†
MBDA Italia

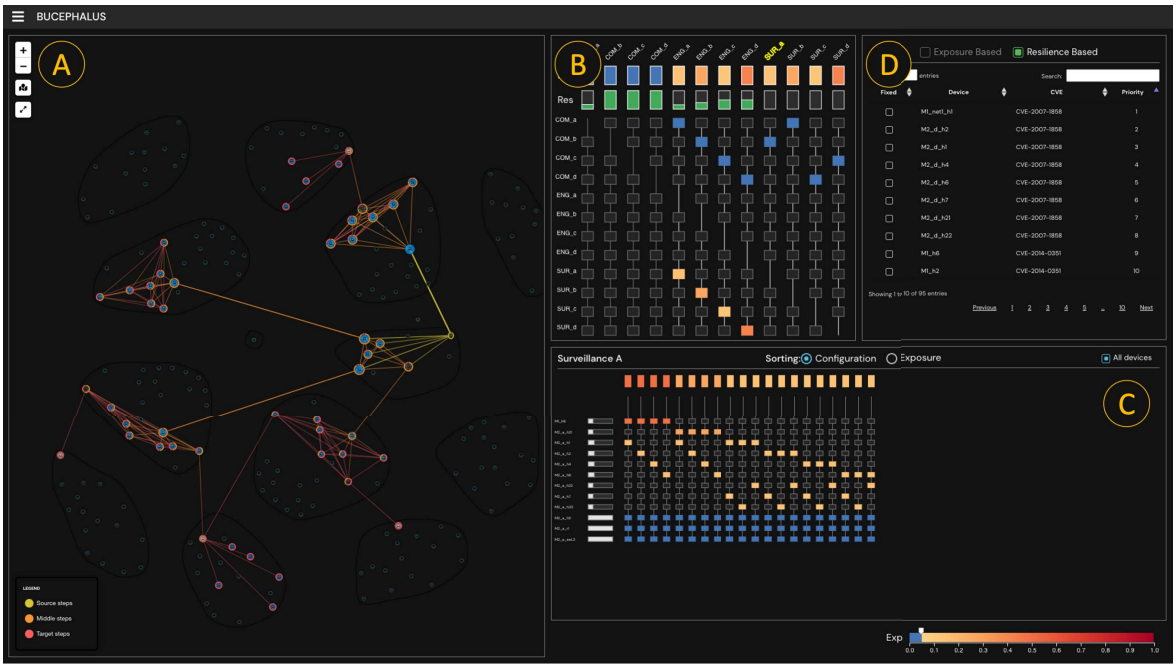

Figure 1: BUCEPHALUS is composed of four different panes, giving the security operator the possibility to investigate the relationship between the devices in the monitored network and the supported business functions. The *Network Pane* (A) shows the attack graph, highlighting in yellow the sources of attack, in orange the intermediate attack steps, and in red the steps toward a target device. At the same time, the *Business Pane* (B) gives the overview of the business functions' status and their inter-dependencies, allowing to explore the exposure and resilience levels. Choosing a business function to investigate ("Surveillance A", *SUR_a*), the *Dependencies Pane* (C) shows in a matrix-like view the relations between devices (rows) and equivalent configurations (columns), representing how the cyber-exposure affects the business functions working state. Finally, the *What-if analysis Pane* (D) proposes two strategies for mitigating the exposure (*Attack Paths Based* and *Resilience Based*), listing the pair device-vulnerability in accordance with the chosen strategy. The security operator can conduct a what-if analysis by simulating mitigations, obtaining a mitigation plan that raises the resilience of the business functions of the organization.

Abstract

Analyzing and mitigating the threats that cyber-attacks pose on the services of a critical infrastructure is not a trivial activity. Research solutions have been developed using data about the devices used for implementing the services, services dependencies, network topology, and the vulnerabilities that can be exploited to attack the network. However, most of the proposed solutions fail to consider these aspects in an integrated fashion, allowing the user to understand global dependencies and weaknesses. This paper contributes this issue with BUCEPHALUS, a Visual Analytics solution providing a) a visual overview of the existing relationships among business functions, devices, and vulnerabilities, and b) a what-if analysis scenario, in which the user is supported on making decisions on which vulnerabilities are more appropriate to fix. BUCEPHALUS has been developed and validated within a user-centered design process involving security professionals.

Index Terms: Cybersecurity—Business Impact Analysis—Network Hardening—Attack graphProactive analysis; Visual Analytics—What-if analysis

1 Introduction

The management of the security risks due to cyber-attacks or failures is gaining increasing attention: a non-exhaustive list of research topics include the identification of business dependencies on supporting systems, the analysis of the network vulnerabilities and the associated threats, the automated assessment of business risk, and

*e-mail: {surname}@diag.uniroma1.it
†e-mail: emiliano.de-paoli@mbda.it

2639-4332/21/$31.00 ©2021 IEEE
DOI 10.1109/VizSec53666.2021.00007

the support of the hardening of the system through proactive actions. Even if business functions continuity and cyber-exposure are the key issues to deal with while hardening the cyber-posture of an organization, their interconnection is not a trivial topic to consider, where critical vulnerabilities could have very different impacts on business continuity for different organizations, while even a single low critical one could impair a whole business function. Dependencies between business functions and resilience of a business function, i.e., the capability of a business function to provide a good service level even if under a cyber-attack, make this scenario even more complex. Most of the proposed solutions in literature fail to consider those aspects in an integrated fashion, as reported in Section 2.

This paper introduces BUCEPHALUS, a Visual Analytics solution under development within the collaboration between Sapienza University of Rome and the MBDA company, for monitoring and hardening an organization cyber-posture driven by business dependencies and continuity. MBDA, a world leader in the military aviation sector, is a multinational company with thousands of employees working in Europe and the United States, producing missiles and missile systems.

The proposal presented in this paper relies on the modeling of the three critical aspects associated with the security analysis and hardening processes (see Section 3), and provides a contribution in linking those perspectives. The system business functions are the core of the analysis, and the adopted model allows for describing dependencies among functions and linking each function to the devices it relies on. The second aspect is associated with the cyber-threats modeled through a topological attack graph that provides the means for computing the exploitation likelihoods it poses on the devices used by the business functions. Finally, the third aspect is the resulting business quality, modeled in terms of function exposition to failure and resilience. The Visual Analytics system presented in Section 4, designed through a user-centered process with MDBA security experts, builds on such models to visually inform the users about the dependencies among business functions and between business functions and devices, clearly showing the threats the actual attack graph poses on supporting devices and how this affects the quality (business functions exposure and resilience) of the organization. Moreover, the system supports hardening activities by computing strategies for fixing the attack graph vulnerabilities; in this way, the security operator can focus on the most relevant vulnerabilities, i.e., the vulnerabilities that have the highest impact on the network exposure or the organization's business. Finally, acknowledging that some business constraints could prevent the straightforward application of the optimal fixing order, the system provides the user with a what-if environment in which she can simulate fixing one or more vulnerabilities, observing the business quality improvement.

Summarizing, the contributions of the paper are the following:

- a user-centered design of a Visual Analytics solution for the joint analysis of cyber-ecposure and business dependencies, and the resulting set of general requirements;
- a Visual Analytics solution, BUCEPHALUS, designed with MBDA security experts, that implements those requirements;
- a proactive what-if functionality supporting the user in exploring both the optimal vulnerability fixing sequence and sub-optimal strategies associated with business constraints, allowing to recommend a mitigation strategy to prioritize the former, the latter, or a combination of the two;
- two concrete usage scenarios that help in validating the efficacy of the proposed solution.

2 Related work

Looking at previous proposals for monitoring the state of operation of an enterprise with respect to cybersecurity, several contributions exist. They focused on the monitoring and analysis of the cyber-exposure level, where most of them are based on the attack graph visualization and analysis [3, 6, 12, 15, 30, 31, 35, 38]. All those approaches have in common that they focus only on the network attack surface perspective, without considering how it can affect the business layer of an organization.

Several researchers have conducted activities with the goal of linking cyber-exposure data to other perspectives in an attempt to raise situational awareness. Pike et al. [33] correlate network anomalies and attacks to real-world social or geopolitical events. Ferebee et al. [16] apply similar considerations, providing requirements for security visualization and business impact analysis visualization by building on previously gained knowledge on understanding weather maps used in meteorology. DAGGER by Peterson [32] is a modeling and visual framework for representing knowledge and information from network security data for decision-makers. In our work, the Business Centric view can be looked at from three different perspectives: (1) the service level that each business function provides during operations, (2) the exposure to cyber-threats that each of the functions can potentially suffer, and (3) the resilience that each of the functions shows with respect to the cyber-exposure.

Starting from the first perspective, Motzek et al. [26] propose a business dependency model normalization and matching approach by exploiting structures and dependencies of business resources to model the dependencies between the business functions of an organization. Matthes et al. [25] provide a three-phase method to systematically identify dependencies between business capabilities and other elements of an Enterprise Architecture. Bouchaala et al. [9] developed DAT, a dependency analysis tool for Business processes expressed in Business Process Model and Notation (BPMN). A recent survey from Stein Dani et al. [36] explores the different methods for visualizing information coming from those models. With respect to their six classes classification (Augmentation of existing elements, Creation of new elements, Exploration of the 3D space, Information visualization, Visual feedback concerning problems detected in process models and Perspectives), we position our proposal into the "Visual feedback concerning problems detected in process models". Overall those works focus only on representing the dependencies and not on linking them to cyber-exposure data as our approach does.

With respect to the link between the cyber-exposure of a device and the supported business function, Motzek and Möller [28] provide a formal, mathematical model for bias and context-free mission impact assessment, eventually applied to a cybersecurity scenario [27]. Chen et al. [11] introduce a new business process impact assessment method that measures the impact of an attack towards a business-process-support enterprise network. The impact scores for business processes result from the severity of the vulnerabilities and the relations between vulnerabilities and business processes. Those papers propose only automatic models without any visualization or support for visual exploration and decision-making. There exist instead solutions that provide visual support and actionability to those models. Goodall et al. [18] propose Camus, a system to automatically map cyber assets to the users who depend on them, to the missions they support, and to the services they provide. Tannian [37] proposes a design study on visualizing the effects of cyber-attacks on business continuity, conducted with seven IT professionals from the Des Moines metropolitan area. However, this work focuses only on the reactive aspects, with an ongoing attack, and does not focus on the resilience of the business functions or their exposure (proactive analysis). Angelini and Santucci [4] propose a visual metaphor (Corruption of area) to represent the degradation of service level for critical infrastructure

business functions superimposed on a geographic map. This work focuses only on the service level and does not consider the resilience to cyber-threats of business functions. The authors extended this work by considering high-level management personnel to support during the review of the operational status of an enterprise [5]. However, this solution provides only an overview, aggregated at the lowest level of detail, and does not allow actionability for any countermeasure. On the contrary, Jajodia et al. [22] present Cauldron, a solution that provides visualization of attack paths, with automatically generated mitigation recommendations, along with analysis of mission impact from attacks. While not visually sophisticated as our solution, their paper focuses only on the impact on service level and not on business resilience. Finally, CyGraph is a system by Noel et al. [29] that links together assets to mission and dependencies among mission requirements; results are explorable in a set of task-driven visualizations. Gonzalez-Granadillo et al. [17] propose a proactive and reactive management system that evaluates a cyber-threat scenario, considering the likelihood of success, the induced impact, the cost of the possible responses, and the negative side-effects of a response. However, no visual environment nor capability for the user to explore the results are provided.

With respect to the last perspective, Horn and D'Amico [21] present an initial effort to use visual analytics to support the modeling of the computer network defense (CND) decision process of an organization and tracing relationships between decision goals, sub-goals, and data sources, like IDS alerts, asset management, and network flows. At the top, there is the one overarching goal to capture an organization's mission from observations of its practices. This overarching goal can be decomposed into sub-goals. However, different from our approach, they do not use an explicit representation of the business assets and functions and do not exploit modeling of original data sources like attack graphs. Finally, their design is based on superimposing this information over a node-link hierarchical structure. Still, D'amico and Sals [14] discuss a 3D representation of information security breaches, assets involved, and their support to mission-critical aspects. One proposal similar to our approach is the work by Hao et al. [19]. The authors introduce VisImpact, a visualization technique that represents operational business data into valuable information reducing data complexity and abstracting the most critical factors, called impact factors, which influence business operations. While the authors propose a case study on fraud-analysis, the focus is on the business flow-graph. It considers in a limited form the cyber-exposure of the organization, as our contribution does. Creese et al. [13] present CyberVis, a 3D visual system that combines traditional network diagram icons with BPMN, a risk-propagation logic that connects the network and business-process and task layer, and a flexible alert input schema able to support intrusion alerts from any third-party sensor. CyberVis abstracts the visuals to show only noteworthy information about attack data and indicates potential impact across the network and enterprise tasks. Different from our approach, they do not consider the resilience level of a business function and how far it could be from being degraded, but only relations between exposure and service level.
Finally, some visualization works exist that coped with the concept of the resilience of an organization (e.g., [10, 39]). However, those works target resilience to a phenomenon not necessarily tied to cybersecurity, like natural disasters or physical security.

3 Business Exposure model

This section provides details about how BUCEPHALUS models the relationships between cyber-exposure, service level, and resilience of a business function. These models are implemented in an automatic module in the system for their computations. The section first introduces the cyber-exposure model that describes the cyber-exposure level of devices inside a network organization. Then it moves on to describe the business dependency model, which illustrates the relations and inter-dependencies existing between devices and supported business functions, and among business functions themselves. Finally, it introduces a linkage between the two, which we define as the Business Exposure model, that allows for describing the effect of cyber-exposure both on a business function service level and its resilience.

3.1 Attack Graph model

An *Attack Graph* (AG) represents possible ways via which a potential attacker can intrude into a computer network by exploiting a series of vulnerabilities on various hosts and gaining certain privileges at each step. Many different AG models have been defined in the literature depending on the specific semantics assigned to nodes and edges of the graph [23].

In this paper, we will focus on *host-based Attack Graphs* where a node represents a specific level of privilege gained by the attacker on a specific host (e.g., *None*, *User* or *Root* on the host h_i) while an edge between node p_s and node p_t represents the possibility to exploit a vulnerability on the destination host h_t gaining a privilege p_t stating from the privilege p_s earned on the source host h_s.

Given an attack graph $\mathcal{AG}$ and two hosts h_s and h_t, it is possible to compute all the existing *Attack Paths* between two hosts[1] h_s and h_t simply by computing all the possible paths existing over $\mathcal{AG}$ connecting any privilege existing on h_s with any level of privilege gained on h_t. The result is a collection of alternate sequences of nodes (i.e., level of privilege over a host) and exploitable vulnerabilities where each path has the form $p_s, vul_j, p_j, vul_k, p_k, \ldots, vul_t, p_t$ and is called *multi-step attack path*.

3.2 Business Dependency Model

Failure or compromising of elements in the Information and Communications Technology (ICT) network may have a strong impact on the ability of a company to provide its services correctly. Several studies exist trying to relate elements characterizing the ICT network layer (e.g., host or other devices connected to the network) with the business processes supported by the ICT infrastructure. Bahşi et al. [7] provide a systematic literature review of existing frameworks for assessing the impact of cyber actions on missions or business processes up to 2018. Among all the existing models for representing dependencies between business processes and network devices, we decided to use a general, simple, and flexible model similar to those used by Gonzalez-Granadillo et al. [17]. In particular, we will consider a model representing dependencies as direct relationships between dependency nodes. A dependency node could be any of the following:

- *Business process or function*: it represents a functional process needed to support the company's mission (e.g., environmental monitoring for a company working in cultural heritage or billing sub-system for a generic service provider);
- *Host/device involved in services provisioning*: it represents an element of the ICT network that contributes to the implementation and support one or more business process (or functions).

Given two dependency nodes n_i and n_j (either two business nodes, a business and a host node or two host nodes), we say that n_i *depends on* n_j (i.e., there exists a dependency between n_i and n_j) if *a failure or compromising of* n_j *impacts the correct functioning of* n_i.

This model allows the extraction of "equivalent configurations" of devices supporting a business function, allowing to define a degree of redundancy for a business function. Nominally, if n different equivalent configurations support a business function, it means that

[1] For ease of explanation, we just considered here one source and one target host. However, attack paths can be computed between any set of source and any set of target hosts by simply iterating.

just one of them is needed to be operational to support the business function to the desired service level. This implies, at least nominally, that the added $n-1$ redundant configurations make more resilient the business function. For $n=1$, the redundancy is zero, and the business function is the least resilient possible.

3.3 Linking the two worlds: the Business Exposure model

Let us note that even if attack graphs and business dependency models can be defined and studied independently of each other, in that way each of them represents only a partial view of the resilience posture of the organization. Matching them instead allows modeling the effect of cyber-exposures on a business function and its correct working state. The correct working state of a business function depends on a set of conditions that have to be satisfied (related through a logical AND). Each condition corresponds to the correct working state of another business function, a single device supporting the business function, or a group of redundant devices (grouped by a logical OR) where only one of them has to behave for the corresponding business function to operate correctly. The presence of logical ORs in the resulting dependency tree generates multiple configurations, i.e., "equivalent configurations" (meaning that only one of them is needed to work for the supported business function working correctly), supporting the correct working state of a business function despite the potential impairment of a subset of its supporting devices. The execution of this joint model, called the Business Exposure model, gives the capability to:

- analyze the direct effect on business functions caused by exposure to a specific set of attack paths computed by using the attack graph (**Effect on service level**);
- evaluate the effects that cyber-exposure has on the business dependencies themselves, where some of them could be very resilient and guaranteed at their nominal value defined in the business dependency model, while others could results weaker, or worse already compromised due to high exposure of their supporting devices (**Identification of weak dependencies**);
- weighting "equivalent configurations" on attack paths, it is possible to compute the real level of redundancy and the resilience of a business function. For example, if a business function has three equivalent configurations, $\langle c_1, c_2, c_3 \rangle$ but it exists in the attack graph an attack path that includes a device from c_2 and a device from c_3, the real redundancy will get lowered from 2 to one, expressing a less resilient business function (**More accurate evaluation of resilience**);
- connected to the previous point, capability to automatically suggest a mitigation plan that is driven by business function resilience (**Resilience-driven mitigation plan**).

Exposure The dependencies of a business function f_i can be expressed as $f_i \rightarrow (f_1 \wedge f_2 \wedge \cdots) \wedge (c_1 \vee c_2 \vee \cdots)$, which is the logical AND among all its functions dependencies, and the logical OR among all its equivalent configurations. The exposure to attacks E is defined for devices, equivalent configurations, and business functions. From the attack graph model, each attack path has associated a likelihood l expressing the probability that the path will be instantiated during an attack. The exposure of a device is defined as the maximum likelihood among all the attack paths that involve that device. The exposure of an equivalent configuration $E(c_1)$ is defined as the maximum exposure of the devices involved in the equivalent configuration. The exposure E of a business function f is defined as:

$$E(f) = \max \Big[\max[E(f_1),\, E(f_2),\, \cdots] \,,\, \min[E(c_1),\, E(c_2),\, \cdots] \Big]$$

Resilience For a business function f, the set of its equivalent configurations is $\mathcal{C}_f = \{c_1,\ c_2,\ \cdots\}$. Given an exposure threshold t, we can assume that if $E(c_1) < t$ this particular equivalent configuration has a low probability of being compromised. The resilience R of a business function f expresses the proportion of how many equivalent configurations have a low probability ($< t$) of being compromised.

$$R(f,t) = \frac{\| \{ \forall c_i \in \mathcal{C}_f \,,\, E(c_i) < t \} \|}{\| \mathcal{C}_f \|}$$

4 Visual support to Business Exposure Model

This section describes how we designed a Visual Analytics environment supporting the Business Exposure model. We first introduce the requirements collection process, intertwined with the principal design decisions and intermediate results that led our process. Following this, the description of the final version of the BUCEPHALUS environment is provided in Section 4.2.

4.1 Requirements collection

To design the proposed solution, we worked in conjunction with the MBDA company, which has relevant needs for monitoring the cyber-exposure of their products, not only in terms of degraded services but also in resilience to possible cyber-threats and proactive prevention of cyber-attacks. Inside the weapon installations of the company, the workstations allow the operator to interact and control the system. The design and development of the human-machine interface is a crucial aspect of the quality of the entire product. Five key-personnel figures were involved during the design process: one Administrator, one Technical officer, one operative, all experts in managing cyber-exposure and business continuity monitoring and analysis, and two experts from the area dedicated to Human Factor studies, which the main objective is to provide and guarantee customers a product conceived and built with the user in mind. The design activities spanned one year and started reasoning on an existing initial solution for the visual analysis of pure cyber-exposure of an enterprise network, MAD [2]. This solution proved very good in representing the cyber-exposure status (proactive and reactive) for the organization's devices. However, it did not help relate those data to business function exposure and business function resilience.

Through a set of five think-aloud sessions (two initial sessions of brainstorming, three following meetings with mock-ups and prototypes), lasting on average from 1.5 to 3 hours each, we designed the solution presented in Section 4.2. During the first two meetings, we set up the initial goals for a new system capable of managing and representing the structure and dependencies among the business functions of the organization (**Requirement RQ1: Capability to see the overview of business functions structure and interdependencies**) and their operational level (**Requirement RQ2: Capability to see the overview of business functions operational level**). This led to the proposal of the first mock-up of the visual interface that would support the analysis of those data. Limitations were reported in terms of the inability to relate the business functions' service levels to the originating cause of problems. Additionally, it was reported that operators tend to consider both perspectives (i.e., business functionality and cyber-exposure of the enterprise devices) simultaneously. This led us to consider in the first revision of our design two additional requirements: **Requirement RQ3, Capability to see the cyber-exposure level of the monitored environment**, and **Requirement RQ4, Capability to proactively analyze the resilience of business functions with respect to cyber-exposure of their supporting devices**. While RQ3 was the direct consequence of what our stakeholders reported, RQ4 derived from the considerations that the link between cyber-exposure of devices and degradation of their supported business function(s) is not the only perspective that can be considered in a proactive analysis. Even the "distance" of a business function from its possible degradation is very useful in managing correctly the cybersecurity posture of an

organization. The more the core business functions are distant from their degraded state, the more the organization will be resilient to cyber-attacks.

Figure 2: Intermediate prototype (cold mock-up) produced during user-centered design iterations. It is visible how dependencies between devices and business functions are represented in a matrix-style view, while the inter-dependencies among business functions are detached and moved on top of the topology view.

This time a cold mock-up was produced with some functionalities running for the cyber-exposure representation inspired by previous work on pure visual representation of cyber-exposure [2, 8]. At the same time, we added a new part concerning the representation of the relations between business functions and devices' cyber-exposure visible in Figure 2. It used a matrix-like visualization, where rows represent devices and columns the business functions. Each cell of this sparse matrix could be colored with respect to cyber-exposure, contributing to understanding its impact on the business function. Business functions could contribute to the functionality of other business functions. The hierarchical representation on the top part of Figure 2 captured this behavior, where bottom layer functions contribute to higher-level functions. The use of horizontal bars represented the current exposure for each function.

During a new face-to-face meeting, it was noted that the association between devices and business functions in a direct form was a good overview, but details about equivalent configurations of devices that still support the business functions should be actionable on-demand. Additionally, it was noted that the visualization of inter-dependencies among business functions needed a higher level of detail. Combined, they contributed to the formulation of a new requirement: **Requirement RQ5, Usage of a top-down approach for the whole visual environment**. This requirement follows the classic visual information seeking mantra ("Overview first, zoom and filter, details on demand") [34] and was explicitly required by the stakeholders to respect the common way in which business and security analysts use visual systems.

This requirement, followed even during the previous design phases, was considered a hard requirement from this moment on for all the remaining design aspects. Interestingly, what at first could seem a classic requirement for a visual environment, was coupled with two additional requirements coming directly from the security operators' workflow: the first one, **Requirement RQ6**, requires to have the **capability to reduce the analysis only on the devices and/or business functions that present problems in terms of exposure and/or resilience**. The additional requirement, **Requirement RQ7**, asked instead for **positional stability of visual elements for the main visualizations, where the user needs to find the same information in the same part of the screen all the times she wants to access them**. The union of RQ6 and RQ7 asked for careful visual design choices that are discussed in Section 4.2, and led to the first release of the system.

By a new meeting in which the system was presented, arose the final requirement, **Requirement RQ8, capability for the system to suggest possible mitigation plans, with the decision-maker having the final word on which actions to perform**. This requirement also includes the capability to conduct what-if analysis scenarios where the operator can simulate the effects of mitigation action and eventually confirm it for real execution. This requirement was implemented and contributed to the final design of the system presented in the following section. Table 1 lists the requirements collected during the design process.

Table 1: Requirements description

ID	Requirement description
RQ1	Capability to see the overview of business functions structure and inter-dependencies.
RQ2	Capability to see the overview of business functions operational level.
RQ3	Capability to see the cyber-exposure level of the monitored environment.
RQ4	Capability to proactively analyze the resilience of business functions with respect to cyber-exposure of their supporting devices.
RQ5	Usage of a top-down approach for the whole visual environment.
RQ6	Capability to reduce the analysis only on the devices and/or business functions that present problems in terms of exposure and/or resilience.
RQ7	Positional stability of visual elements for the main visualizations, where the user needs to find the same information in the same part of the screen all the times she wants to access them.
RQ8	Capability for the system to suggest possible mitigation plans, with the decision-maker having the final word on which actions to perform.

4.2 The BUCEPHALUS Visual Analytics environment

The user-centered design presented in the previous section led to a Visual Analytics environment subdivided into four panes (see Figure 1) supporting the different perspectives of the analysis (**RQ7**).

The experts' need to monitor the cyber-exposure of the enterprise devices (**RQ3**) requires an explicit representation of the monitored network that is visible in the *Network Pane* (Figure 1A).

This pane shows a node-link representation of the network topology on which the attack graph is projected. We choose the node-link representation to display complex information at the node level and support tasks like following path(s) and users' habits. Homer et al. [20] present methodologies that can automatically identify the less interesting portions of an attack graph and group similar attack steps as virtual nodes, to immediately increase the understandability of the data. We follow a similar approach for visualizing aggregations of attack paths, showing the different roles that the nodes play in the paths according to the encoding presented in Blasilli et al. [8] (see Figure 3).

The background color of a node represents the higher privilege reached by all the attack paths involving that device: gray, blue, and purple stand for none, user and root privileges, respectively. The attack path proportions on nodes are shown with the internal donut chart: red color identifies the final step of an attack, yellow is used to identify attack paths source nodes, and orange represents every intermediate step. The external donut chart represents the proportion of vulnerabilities of the node used by the current attack paths, in gray, while the subset used for performing privilege escalation, in blue. The color-coding is the same presented in [8] conforming to the one usually adopted by the experts; in addition, the system allows to choose also colorblind safe colors.

The general requirement of a top-down approach for the analysis (**RQ5**), coupled with the need to relate the business functions and the devices that support them, calls for a hierarchical visualization that provides a high-level overview of the business functions attributes and gives the possibility to analyze on-demand their relations with the functioning of devices. The design relies on the matrix-based representations presented in Section 4. This requirement combined with the need for visual stability (**RQ7**) of the main visualizations led us to the creation of two different panes: the *Business Pane* (see Figure 1B) and the *Dependencies Pane* (see Figure 1C).

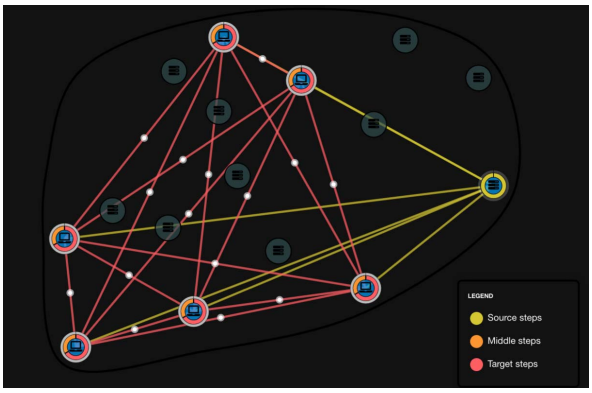

Figure 3: Detail of the *Network Pane* showing the attack graph projected on a node-link representation of the network topology. The color of the edges represents the type of attack step. The nodes encode in the same way this information in a donuts chart showing the cardinality of attack paths in which the node is included.

The *Business Pane* (see Figure 4) adopts a matrix-based representation in which the rows and the columns are the business functions. The matrix has two additional rows (the first two rows) that encode the exposure E (**RQ2**) and the resilience R (**RQ4**) of each business function, as defined in Section 3.3. In the first row, the color encodes the exposure of the function; the element is encoded in blue when the exposure is below a configurable threshold (e.g., 0.05), otherwise it is colored according to a yellow-red color scale. The exposure of the functions does not consider their resilience to cyber-attacks. The second row shows the resilience level of each function through a bar-chart encoding to convey this information. The height of the bar is proportional to the number of equivalent configurations that have an exposure below the threshold.

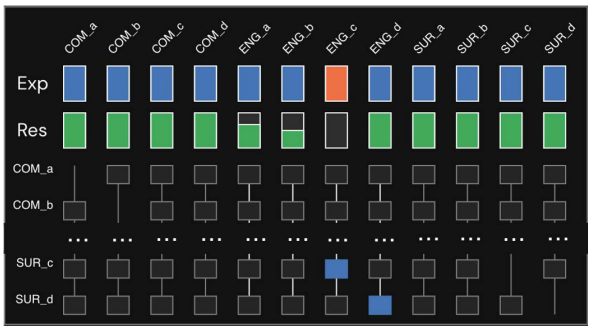

Figure 4: The *Business Pane* shows the exposure and resilience levels of the business functions and the dependencies between them. The first row shows the exposure *Exp* of each function: blue if it is below a threshold (e.g., 0.05), while a yellow-red color scale encodes greater values. The second row shows the resilience *Res* encoding the height of the bar proportionally to the number of equivalent configurations that have an exposure below the threshold. The rest of the matrix represents the dependencies between the functions.

While this part of the visualization provides an overview of the business functions status, it does not give any details on the relationships between them (**RQ1**). The underlying matrix encodes the dependencies between the functions; for each column of the matrix, the cells of the supporting functions (connected through a logical AND) are colored according to their exposure. The visualization is enriched with vertical lines loosely inspired by UpSet [24] recalling how to interpret the dependency matrix: a line represents a logical AND among the elements that it traverses. While this pane provides an overview of the business functions and their relations, it still not describes the dependency of the functions from their supporting devices (**RQ4**). By selecting a function from this pane, the *Dependencies Pane* is updated showing the dependencies of the function chosen from the correct operation of the devices. Remembering how the Business Exposure model works (see Section 3), we represented the dependencies and equivalent configurations in a matrix-like view.

The matrix rows represent the devices that contribute to supporting the business function, while each column represents one equivalent configuration (see Figure 5). An additional row is added

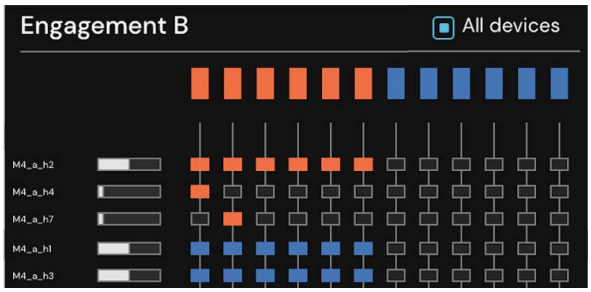

Figure 5: The *Dependencies Pane* showing the equivalent configurations supporting the correct working state of the selected business function, i.e., "Engagement B". In the matrix, rows represent devices and column equivalent configurations. Each cell represents the dependency of the configuration from the device, with the corresponding exposure level encoded. For each device, a horizontal bar shows the number of configurations on which it takes part.

on top of the matrix, which encodes the maximum exposure of the devices' configuration. The configurations are computed considering the availability of a single device for each OR group; their number is thus equal to the product of the cardinalities of the OR groups potentially leading to a high number of configurations. The need to reduce the analysis only to the elements that present problems in terms of exposure (**RQ6**) is supported by a slider that allows to exlcude from the analysis the devices and the configurations with exposure below a threshold.

The matrix analysis (which has dense areas for non-redundant devices and sparse areas for redundant ones) is aided by a horizontal bar-chart aligned with the list of devices that encode the number of configurations in which a device is present. Furthermore, the matrix rows are sortable according to the number of configurations on which the device occurs or to its level of exposure.

The selection of the devices in this pane is synced with the selection in the *Network Pane*; the analyst can, thus, identify devices of interest in one analysis (network-driven or business-driven) and see their role in the other one easily switching between them.

The presence of multiple configurations and the need to include multiple functions in the analysis can make it difficult to prioritize the devices according to their exposure or their contribution to the correct working of the function(s) (**RQ8**). The *What-if Analysis Pane* supports this task by presenting the list of device-vulnerability pairs that are present in the Attack Graph. The list is sortable according to two different strategies:

- **Attack Paths Based**: this mitigation strategy has been defined in VULNUS [1], as *AG Environmental Strategy*. It aims at reducing the number of attack paths considering only topological information. It considers the role that each device-vulnerability pair has in the attack paths. Pairs are ordered according to their number of occurrences in the attack paths. The first proposed vulnerability is the one that allows to interrupt the greatest number of attack

paths, and so on.

- **Resilience Based**: this strategy prioritizes device-vulnerability pairs whose exploits play central roles for the resilience level of the business functions. This strategy aims to increase the resilience level of either all the business functions, or a subset of them which the user chooses. By considering the topological and business information of the functions, the strategy suggests fixes that improve the resilience of the functions.

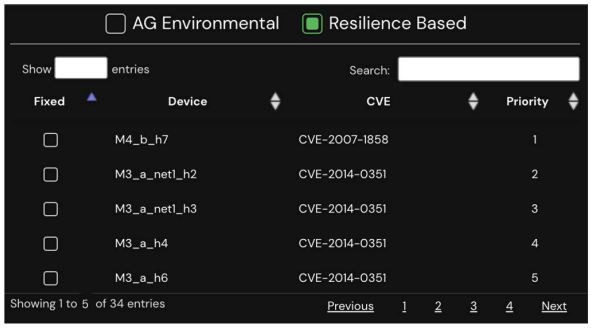

Figure 6: The *What-if Analysis Pane* supports the analysis of the proposed fixing strategies. The analyst simulates the application of a fix by clicking on the toggle next to the vulnerability to fix. The other panes update showing the state that corresponds to the fix.

The analyst can analyze the impact of fixing one or more vulnerabilities by selecting them from the table and looking at their effect in the other panes. Requirements coverage is summerized in Table 2. A video demonstration of `BUCEPHALUS` showing the described functionalities is available at `https://aware-diag-sapienza.github.io/BUCEPHALUS`

Table 2: `BUCEPHALUS` requirements coverage.

	Network Pane	Business Pane	Dependencies Pane	What-if Analysis Pane
RQ1		•		
RQ2		•		
RQ3	•			
RQ4		•	•	
RQ5	•	•	•	
RQ6			•	
RQ7	•	•	•	•
RQ8				•

5 USAGE SCENARIOS: THE MBDA ORGANIZATION

To show the added capabilities that visually assisted business-centric analysis provides with respect to classic cyber-exposure analysis and to understand whether the proposed solution effectively covers the collected requirements, we tested our system on two usage scenarios. Those scenarios have been conducted with MBDA personnel (security analysts) on one of their product, a weapon system (complex installation constituted by multiple interconnected devices) consisting of 242 devices, 52 distinct vulnerabilities, and 12 business functions (belonging to three main classes, Surveillance, Engagement, Communication). In the first scenario, the proactive analysis of business functions resilience is the security analyst's core activity. In contrast, the second scenario concerns the difference between a global mitigation plan that considers only cyber-exposure aspects with respect to a mitigation plan that focuses on overall business functions resilience. Our system generates both plans, and the security analyst conducts a what-if analysis by testing several alternatives from them.

5.1 Scenario 1: Analysis of the resilience of a business function

The first usage scenario aims to allow the security operator to explore the business functions' service level and their resilience with respect to potential cyber-threats. For service level, we mean the business function exposure to cyber-threats, namely how probable is that the business function will be degraded if a cyber-threat occurs. She begins the analysis by looking at the top part of the *Business Pane* that reports information about aggregated exposure per business function and aggregated resilience level per business function. She is interested in identifying three types of conditions, presented in decreasing order of importance for inspecting anomalies:

- **Resilient business functions**: those are business functions that are working at the desired service level, and that present a good level of resilience, meaning that not only they guarantee now the desired service level, but they are resilient (i.e., they exhibit redundancy) that helps them in guarantee the same service level even if under the effect of a cyber-threat;
- **Working business functions**: those are business functions that are guaranteeing the nominal service level, but that at the same time do not present an adequate resilience, meaning that a single cyber-threat can lower the desired service level;
- **Degraded business functions**: those are functions that present a degraded service level due to their exposure and lack of resilience, that lower their service level under a threshold defined by the Administrator. In this case, it becomes crucial to eliminate the causes for this degradation first and then reason about resilience level afterward. The resilience level can be variable, depending on the number of equivalent configurations compromised by the cyber-exposure of the supporting devices.

The security operator spots all three cases (three resilient business functions, one working business function and eight degraded business functions), as visible in Figure 7a for the first class, Figure 7b for the second, and Figure 7c for the third class.

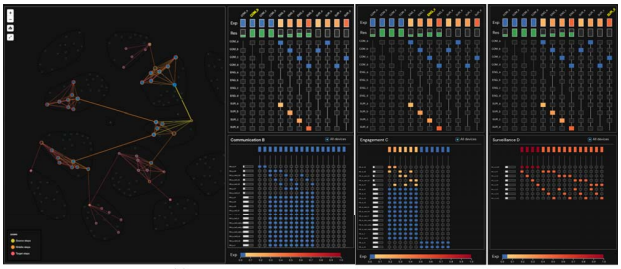

Figure 7: Three different business functions, each of them representing an instance of the three classes defined in scenario 1. (a) Business function COM_b: good service level and maximum resilience. (b) Business function ENG_c: degraded service level, with a potential residual resilience to exploit. (c) Business function ENG_c: degraded service level, with a potential residual resilience to exploit

She focuses on the function *SUR_d* (Surveillance d), which seems the most degraded and exposed at the same time. She first looks at inter-dependencies among business functions in the *Business Pane*. It is visible that the *SUR_d* function depends on the *COM_c* function. Given that the *COM_c* function operates at its nominal conditions and presents its maximum resilience, the security operator does not look at it as the cause of the degradation. She proceeds to explore the dependencies between the *SUR_d* function and its supporting devices in the *Dependencies Pane*. She spots (looking at the color-coding) that the device `M2_d_h20` plays a strong part in

degrading the first four equivalent configurations. She then clicks on its label, and the *Network Pane* gets updated accordingly to show the portion of the attack graph that includes this device. The device presents a high number of vulnerabilities and cannot be restored easily. Additionally, its restoration cannot be enough to recover any equivalent configuration, given that it works in conjunction (AND rule) with `M2_d_h1`, `M2_d_h2`, `M2_d_h4` and `M2_d_h6` devices.

The security operator then proceeds to other devices in order of their effect on resilience. She spots that the device `M2_d_h1` can be easily fixed, along with `M2_d_h7`, `M2_d_h21`, and `M2_d_h22`, given that they are all affected by CVE-2007-1858. Their combined fixes result in three equivalent configurations restored, meaning that the business function operates at a higher service level (due to lower exposure) and gains a slighter higher resilience. The reason for a little gain in resilience is that the analyst restores three different equivalent configurations, but unfortunately all of them depends from the same device `M2_d_h1`. From the *What-if analysis Pane*, she simulates the fixes one by one, in an incremental way, to check the effects they have on the *SUR_d* exposure and resilience. The result is visible in Figure 8(top), where the function *SUR_d* shows a reduced exposure (light orange versus initial strong orange), but resilience is still not present. Considering promising the identified devices, the security operator continues inspecting them, solving an additional vulnerability (CVE-2014-0351). At this time, the *SUR_d* function has recovered the correct service level (reducing its exposure), and it presents a slight degree of resilience (different equivalent configurations can support its nominal service level), as visible in Figure 8(bottom). She iterates on this workflow, exploiting the Business Exposure model information to preserve the business functions service levels (reducing the exposure driven by business functions requirements) and increasing their resilience.

Figure 8: On top is reported the intermediate fixing of business function *SUR_d*, displayed with the horizontal layout. It shows how exposure is reduced but resilience is still not present. On bottom instead is reported the result at the end of the work. This time the business function has recovered its nominal working state and it presents a slight amount of resilience.

5.2 Scenario 2: Business driven mitigation strategies

The workflow described in usage scenario 1, while targeted at specific business functions behavior (i.e., in the case in which the operator clearly identified a subset of business functions to work on) can be a long activity to conduct. Otherwise, if she is interested in global optimization of the business exposure, she can rely on the what-if analysis capabilities provided by `BUCEPHALUS`. In this scenario, the analyst's goal is to identify a suitable mitigation strategy to improve the security level of the installation. The high number of vulnerabilities and their spread in the network make it difficult to prioritize them. The analyst can thus be guided in the analysis by the strategies proposed in the *What-if Analysis Pane*. The first strategy, exposure-driven, focuses on reducing the exposure surface; it thus aims to reduce the number of attack paths with the minimum number of vulnerability fixings. The vulnerabilities are thus ordered according to the number of attack paths they enable. This strategy effectively reduces the number of attack paths: fixing the first five vulnerabilities [2] drastically reduces the attack paths and contributes to improving the overall level of exposure and resilience (see Figure 9 center). This strategy is effective in reducing the exposure and improving the resilience of six business functions, i.e., *ENG_a*, *ENG_b*, *ENG_c*, *SUR_a*, *SUR_b*, *SUR_c*. However, this strategy does not impact the security level of the two business functions with the highest exposure, i.e., *ENG_d* and *SUR_d*. This is mainly because the contributions of the devices to the configurations that support the business functions are not taken into account. Vulnerabilities on devices that are compromised in few attack paths but that are essential to provide resilience are therefore overlooked in favor of those that enable several attack paths, regardless of their impact on the overall exposure and resilience.

Figure 9: Details of the *Business Pane* showing the functions exposure and resilience levels in the initial scenario (top), after fixing the first 5 vulnerabilities proposed by the attack paths driven strategy (center), and by the resilience-driven strategy (bottom).

The second strategy, resilience-driven, conversely prioritizes the vulnerabilities according to their impact on resilience. By analyzing the first five vulnerabilities, we can see that they play an essential role in the functions, and their fixing is highly effective in improving the overall exposure and resilience levels (see Figure 9 bottom).

By mitigating two vulnerabilities on `M2_d_h21` and one vulnerability on `M1_h1`, `M1_h2`, `M1_h6`, five functions (*ENG_c*, *SUR_b*, *SUR_c*, and the two most exposed, *ENG_d* and *SUR_d*) recover the correct service level and their full resilience capabilities. *ENG_b* recovers the correct service level and its resilience significantly increases. *ENG_a* presents a slight decrease of the exposure and a slight increase of the resilience while *SUR_a* is not impacted.

The proposed fixing strategy results not directly applicable due to an external constraint on the `M1_h1` that does not allow to stop it to apply the necessary hot-fix. The analyst explores alternative solutions by selecting sub-optimal choices and evaluating their effectiveness. In this scenario, she evaluates the effectiveness of alternative plans, evaluating the inclusion of a patch for vulnerability CVE-2007-1858 on `M1_h4`. This strategy has a significant impact on the business: it effectively restores the full capabilities of the two most exposed functions with slightly less impact on the others (*ENG_b* and *SUR_b*,

[2] Vulnerability CVE-2014-0351 in device `M1_h6`, CVE-2007-1858 in device `M1_h6`, CVE-2016-0494 in device `M2_a_h21`, CVE-2015-0395 in device `M2_d_h20`, and CVE-2015-0412 in device `M2_d_h20`.

in particular, do not recover the correct service level), see Figure 10. The analyst approves this fixing strategy and pass it to operations for deployment.

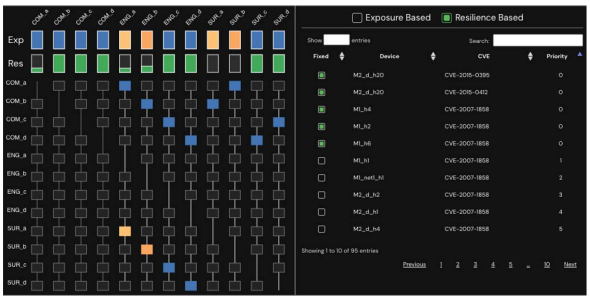

Figure 10: Exposure and resilience levels after the application of a sub-optimal strategy driven by the resilience having a slightly lesser increase in the overall security level with respect to the optimal one.

6 Discussion

This paper explored the possibility of conducting proactive analysis of an organization's network driven by the business functions' service level and their resilience. This approach has the advantage to not only consider the effects that eventual cyber-threats can have on business functions but also introduce the capability to plan what additional actions can "move away" a business function from its possible degraded state, making it more resilient with respect to cyber-exposures. We achieve those analysis capabilities through the conjunction of two models, the classic attack graph for network exposure modeling, and the business dependency model for representing dependencies among business functions, obtaining the mapping of relations between the cyber-exposure status and the business functions. Through the use of Visual Analytics techniques we allow an analyst to explore the results of this new model (see Usage scenario 1), and to construct on top of it a recommender algorithm capable of computing effective planning for mitigating the exposure and raising the resilience and service level of business functions (see Usage scenario 2). This algorithm can be exploited to conduct what-if analysis.

We collected feedback from the experts who tested the environment and identified three promising research directions:

Comparison and evaluation of used features: As reported in Section 3, we used a classic network attack graph and a classic business dependency model to build the approach developed for BUCEPHALUS. Interestingly, using more sophisticated versions of those two models could potentially lead to additional parameters and derived features that could inform the business-centric analysis. More research could be conducted even on correlating those features in different situations and see which of them tend to go in accordance for both cyber-exposure and business-centric views, and which are more biased toward one of those perspectives;

Granularity of mitigation plans: The computed mitigation plans used in the *What-if analysis Pane* are computed at the highest possible granularity, namely a couple $< nodeID, vulnerability >$ according to the classic definition of a network attack graph. While this information is correct and helps in achieving the presented results, during our analysis we discovered that it could exist a second way of modeling this problem based on attack paths (ordered sequences of couples $< nodeID, vulnerability >$). We plan to add this functionality to BUCEPHALUS;

Exploration of attack paths based information: Apart from the computation of mitigation plans, even more interesting is the representation of this information in the *Business and Dependencies panes*, allowing the security operator to inspect causes of coupling equivalent configurations for a business function. We coped with this problem in the last part of this work and designed a visual solution integrated with the existing visual encoding, presented in Figure 11. This design can be integrated directly into the *Dependencies pane* , where equivalent configurations (c_i) are represented. It exploits the spaces existing between equivalent configurations to represent the degree of coupling that attack paths model. A column represents an equivalent configuration, while horizontal segments encode the number of attack paths that include devices coming from different configurations, effectively coupling them. The color encodes in both cases the degree of exposure. Looking at the different business functions, it is visible as f_3 is the most resilient function (it does not exist any attack path that includes devices from its equivalent configurations), f_4 is quite resilient but not perfect (it has exposure on its c_1, f_1 has effect from attack paths of length 2 that couples (c_1, c_2) and (c_2, c_3). Finally, for f_2 there exist also attack paths that couple all the equivalent configurations (c_1, c_2, c_3), meaning that if one of those attack paths effectively occurs the business function will be for sure degraded, without any resilience. By interacting with this chart (e.g., selecting one or more horizontal segments), the security operator could obtain the set of attack paths that, if mitigated, decouple two equivalent configurations, making the function more resilient. We are currently implementing this design in the system.

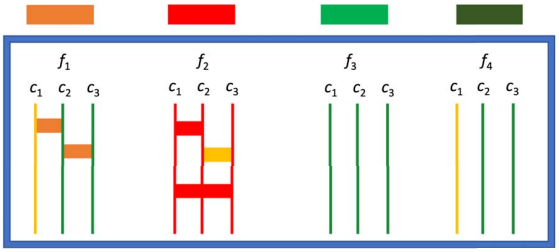

Figure 11: Exploring the relationship between attack paths and equivalent configurations of a function. Vertical lines represent equivalent configurations (c), while horizontal segments encode the number of attack paths that include nodes coming from different equivalent configurations, effectively coupling them. The color encodes in both cases the degree of exposure.

7 Conclusions and Future Work

This paper presented BUCEPHALUS, a Visual Analytics system that eases the analysis of the relationships among business functions, devices, and network vulnerabilities, visually providing an overview of dependencies and weaknesses. BUCEPHALUS supports the proactive hardening of the network through a what-if analysis scenario, in which the user is presented with an optimized order of vulnerability fixing, exploring the effect of sub-optimal strategies that satisfy business constraints. The system has been implemented through a user-centered design with MBDA professionals, producing eight requirements whose visual implementation has been validated and tuned by the feedback provided by the experts involved in the process. Moreover, two usage scenarios provided a step-by-step validation of the implemented functionalities. As future work, we plan to conduct further evaluation activities that, starting from usage scenarios proposed, collect efficacy and efficiency metrics to quantify the advantages of the proposed approach. We also plan to extend this approach to reactive actions, i.e., to model the consequences of suitable mitigation actions in terms of business continuity and quality. Finally, we will investigate an automatic extension of BUCEPHALUS able to operate during a cyber-attacks balancing the continuity of the business functions with the halting of the attack.

Acknowledgments The authors wish to thank Daniele Buonadonna for his initial efforts on the topic.

REFERENCES

[1] M. Angelini, G. Blasilli, T. Catarci, S. Lenti, and G. Santucci. Vulnus: Visual vulnerability analysis for network security. *IEEE Transactions on Visualization and Computer Graphics*, 25(1):183–192, Jan 2019. doi: 10.1109/TVCG.2018.2865028

[2] M. Angelini, S. Bonomi, S. Lenti, G. Santucci, and S. Taggi. Mad: A visual analytics solution for multi-step cyber attacks detection. *Journal of Computer Languages*, 52:10–24, 2019. doi: 10.1016/j.cola.2018.12.007

[3] M. Angelini, N. Prigent, and G. Santucci. Percival: proactive and reactive attack and response assessment for cyber incidents using visual analytics. In *2015 IEEE Symposium on Visualization for Cyber Security (VizSec)*, pp. 1–8, 2015. doi: 10.1109/VIZSEC.2015.7312764

[4] M. Angelini and G. Santucci. Visual cyber situational awareness for critical infrastructures. In *Proceedings of the 8th International Symposium on Visual Information Communication and Interaction*, VINCI '15, p. 83–92. Association for Computing Machinery, New York, NY, USA, 2015. doi: 10.1145/2801040.2801052

[5] M. Angelini and G. Santucci. Cyber situational awareness: from geographical alerts to high-level management. *Journal of Visualization*, 20(3):453–459, 2017.

[6] D. L. Arendt, R. Burtner, D. M. Best, N. D. Bos, J. R. Gersh, C. D. Piatko, and C. L. Paul. Ocelot: user-centered design of a decision support visualization for network quarantine. In *2015 IEEE Symposium on Visualization for Cyber Security (VizSec)*, pp. 1–8, 2015. doi: 10.1109/VIZSEC.2015.7312763

[7] H. Bahşi, C. J. Udokwu, U. Tatar, and A. Norta. Impact assessment of cyber actions on missions or business processes: A systematic literature review. In *ICCWS 2018 13th International Conference on Cyber Warfare and Security*, p. 11. Academic Conferences and publishing limited, 2018.

[8] G. Blasilli, E. D. Paoli, S. Lenti, and S. Picca. Lessons learned while supporting Cyber Situational Awareness. In K. Vrotsou and J. Bernard, eds., *EuroVis Workshop on Visual Analytics (EuroVA)*. The Eurographics Association, 2021. doi: 10.2312/eurova.20211093

[9] O. Bouchaala, M. Yangui, S. Tata, and M. Jmaiel. Dat: Dependency analysis tool for service based business processes. In *2014 IEEE 28th International Conference on Advanced Information Networking and Applications*, pp. 621–628, 2014. doi: 10.1109/AINA.2014.76

[10] S. Buckman, M. Arquero de Alarcon, and J. Maigret. Tracing shoreline flooding: Using visualization approaches to inform resilience planning for small great lakes communities. *Applied Geography*, 113:102097, 2019. doi: 10.1016/j.apgeog.2019.102097

[11] C. Cao, L.-P. Yuan, A. Singhal, P. Liu, X. Sun, and S. Zhu. Assessing attack impact on business processes by interconnecting attack graphs and entity dependency graphs. In F. Kerschbaum and S. Paraboschi, eds., *Data and Applications Security and Privacy XXXII*, pp. 330–348. Springer International Publishing, Cham, 2018.

[12] M. Chu, K. Ingols, R. Lippmann, S. Webster, and S. Boyer. Visualizing attack graphs, reachability, and trust relationships with navigator. In *Proceedings of the Seventh International Symposium on Visualization for Cyber Security*, VizSec '10, p. 22–33. Association for Computing Machinery, New York, NY, USA, 2010. doi: 10.1145/1850795.1850798

[13] S. Creese, M. Goldsmith, N. Moffat, J. Happa, and I. Agrafiotis. Cybervis: Visualizing the potential impact of cyber attacks on the wider enterprise. In *2013 IEEE International Conference on Technologies for Homeland Security (HST)*, pp. 73–79, 2013. doi: 10.1109/THS.2013.6698979

[14] A. D'Amico and S. Salas. Visualization as an aid for assessing the mission impact of information security breaches'. In *Proceedings DARPA Information Survivability Conference and Exposition*, vol. 2, pp. 190–195 vol.2, 2003. doi: 10.1109/DISCEX.2003.1194964

[15] R. F. Erbacher. Visualization design for immediate high-level situational assessment. In *Proceedings of the Ninth International Symposium on Visualization for Cyber Security*, VizSec '12, p. 17–24. Association for Computing Machinery, New York, NY, USA, 2012. doi: 10.1145/2379690.2379693

[16] D. Ferebee, D. Dasgupta, M. Schmidt, and Q. Wu. Security visualization: Cyber security storm map and event correlation. In *2011 IEEE Symposium on Computational Intelligence in Cyber Security (CICS)*, pp. 171–178, 2011. doi: 10.1109/CICYBS.2011.5949412

[17] G. Gonzalez-Granadillo, S. Dubus, A. Motzek, J. Garcia-Alfaro, E. Alvarez, M. Merialdo, S. Papillon, and H. Debar. Dynamic risk management response system to handle cyber threats. *Future Generation Computer Systems*, 83:535–552, 2018. doi: 10.1016/j.future.2017.05.043

[18] J. R. Goodall, A. D'Amico, and J. K. Kopylec. Camus: Automatically mapping cyber assets to missions and users. In *MILCOM 2009 - 2009 IEEE Military Communications Conference*, pp. 1–7, 2009. doi: 10.1109/MILCOM.2009.5380096

[19] M. C. Hao, D. A. Keim, U. Dayal, and J. Schneidewind. Business process impact visualization and anomaly detection. *Information Visualization*, 5(1):15–27, 2006. doi: 10.1057/palgrave.ivs.9500115

[20] J. Homer, A. Varikuti, X. Ou, and M. A. McQueen. Improving attack graph visualization through data reduction and attack grouping. In J. R. Goodall, G. Conti, and K.-L. Ma, eds., *Visualization for Computer Security*, pp. 68–79. Springer Berlin Heidelberg, Berlin, Heidelberg, 2008.

[21] C. Horn and A. D'Amico. Visual analysis of goal-directed network defense decisions. In *Proceedings of the 8th International Symposium on Visualization for Cyber Security*, VizSec '11. Association for Computing Machinery, New York, NY, USA, 2011. doi: 10.1145/2016904.2016909

[22] S. Jajodia, S. Noel, P. Kalapa, M. Albanese, and J. Williams. Cauldron mission-centric cyber situational awareness with defense in depth. In *2011 - MILCOM 2011 Military Communications Conference*, pp. 1339–1344, 2011. doi: 10.1109/MILCOM.2011.6127490

[23] K. Kaynar. A taxonomy for attack graph generation and usage in network security. *J. Inf. Secur. Appl.*, 29(C):27–56, Aug. 2016. doi: 10.1016/j.jisa.2016.02.001

[24] A. Lex, N. Gehlenborg, H. Strobelt, R. Vuillemot, and H. Pfister. Upset: Visualization of intersecting sets. *IEEE Transactions on Visualization and Computer Graphics*, 20(12):1983–1992, 2014. doi: 10.1109/TVCG.2014.2346248

[25] F. Matthes., A. Nowobilska., C. Schulz., and A. Freitag. A method for business capability dependency analysis. In *Proceedings of the Second International Conference on Innovative Developments in ICT - INNOV,*, pp. 11–20. INSTICC, SciTePress, 2011. doi: 10.5220/0004471100110020

[26] A. Motzek, C. Geick, and R. Möller. Semantic normalization and matching of business dependency models. In *2016 IEEE 18th Conference on Business Informatics (CBI)*, vol. 01, pp. 7–15, 2016. doi: 10.1109/CBI.2016.10

[27] A. Motzek and R. Möller. Probabilistic mission defense and assurance. In *NATO IST-148: Symposium on Cyber Defence Situation Awareness, STO-MP-IST-148, Sofia, Bulgaria*, pp. 4–1, 2016.

[28] A. Motzek and R. Möller. Context- and bias-free probabilistic mission impact assessment. *Computers & Security*, 65:166–186, 2017. doi: 10.1016/j.cose.2016.11.005

[29] S. Noel, E. Harley, K. Tam, M. Limiero, and M. Share. Chapter 4 - cygraph: Graph-based analytics and visualization for cybersecurity. In V. N. Gudivada, V. V. Raghavan, V. Govindaraju, and C. Rao, eds., *Cognitive Computing: Theory and Applications*, vol. 35 of *Handbook of Statistics*, pp. 117–167. Elsevier, 2016. doi: 10.1016/bs.host.2016.07.001

[30] S. Noel, M. Jacobs, P. Kalapa, and S. Jajodia. Multiple coordinated views for network attack graphs. In *IEEE Workshop on Visualization for Computer Security, 2005. (VizSEC 05).*, pp. 99–106, 2005. doi: 10.1109/VIZSEC.2005.1532071

[31] S. O'Hare, S. Noel, and K. Prole. A graph-theoretic visualization approach to network risk analysis. In J. R. Goodall, G. Conti, and K.-L. Ma, eds., *Visualization for Computer Security*, pp. 60–67. Springer Berlin Heidelberg, Berlin, Heidelberg, 2008.

[32] E. Peterson. Dagger: Modeling and visualization for mission impact situation awareness. In *MILCOM 2016 - 2016 IEEE Military Communications Conference*, pp. 25–30, 2016. doi: 10.1109/MILCOM.2016.7795296

[33] W. A. Pike, C. Scherrer, and S. Zabriskie. *Putting Security in Context:*

Visual Correlation of Network Activity with Real-World Information, pp. 203–220. Springer Berlin Heidelberg, Berlin, Heidelberg, 2008. doi: 10.1007/978-3-540-78243-8_14

[34] B. Shneiderman. The eyes have it: A task by data type taxonomy for information visualizations. In B. B. BEDERSON and B. SHNEIDERMAN, eds., *The Craft of Information Visualization*, Interactive Technologies, pp. 364–371. Morgan Kaufmann, San Francisco, 2003. doi: 10.1016/B978-155860915-0/50046-9

[35] R. Skowyra, S. R. Gomez, D. Bigelow, J. Landry, and H. Okhravi. Quasar: Quantitative attack space analysis and reasoning. In *Proceedings of the 33rd Annual Computer Security Applications Conference*, ACSAC 2017, p. 68–78. Association for Computing Machinery, New York, NY, USA, 2017. doi: 10.1145/3134600.3134633

[36] V. Stein Dani, C. M. Dal Sasso Freitas, and L. H. Thom. Ten years of visualization of business process models: A systematic literature review. *Computer Standards & Interfaces*, 66:103347, 2019. doi: 10.1016/j.csi.2019.04.006

[37] M. F. Tannian. Business impact visualization for information security and compliance events. In *Graduate Theses and Dissertations. 13050.*, 2013.

[38] L. Williams, R. Lippmann, and K. Ingols. Garnet: A graphical attack graph and reachability network evaluation tool. In J. R. Goodall, G. Conti, and K.-L. Ma, eds., *Visualization for Computer Security*, pp. 44–59. Springer Berlin Heidelberg, Berlin, Heidelberg, 2008.

[39] C. W. Zobel. Comparative visualization of predicted disaster resilience. In *Proceedings of the 7th International ISCRAM Conference*, pp. 1–5. ISCRAM, 2010.