

NORIA-O: An Ontology for Anomaly Detection and Incident Management in ICT Systems

Lionel Tailhardat ^{a,*}, Yoan Chabot ^a and Raphaël Troncy ^b

^a *Orange, France*

E-mails: lionel.tailhardat@orange.com, yoan.chabot@orange.com

^b *EURECOM, Sophia Antipolis, France*

E-mail: raphael.troncy@eurecom.fr

Abstract. Large-scale Information and Communications Technology (ICT) systems give rise to difficult situations such as handling cascading failures across multiple platforms and detecting complex malicious activities occurring on multiple services and network layers. For network administrators and supervision teams, managing these situations while ensuring the high-standard quality of service and security of networks requires a comprehensive view on how communication devices are interconnected and are performing. However, the relevant information is spread across heterogeneous log sources and databases which triggers information integration challenges. There are several efforts to propose data models representing computing resources and how they are allocated for hosting services. However, to date, there is no model to describe the multiple interdependencies between the structural, dynamic, and functional aspects of a network infrastructure. In this paper, we propose the NORIA ontology that re-uses and extends well-known ontologies such as SEAS, FOLIO, UCO, ORG, BOT and BBO. NORIA-O has been developed together with network and cybersecurity experts in order to describe a network infrastructure, its events (user login, network route priority reconfiguration), diagnosis and repair actions (connectivity check, firmware upgrade) that are performed during incident management. A use case describing a failure on a fictitious network shows how this ontology can model complex ICT system situations and serve as a basis for anomaly detection and root cause analysis.

Keywords: Ontology, Incident, Network Infrastructure, NORIA

1. Introduction

An IT network is a set of computers, routers, and other devices connected and configured to allow data processing and sharing. Internet and corporate networks are typical examples of IT networks. An IT service is the usage of this data processing and sharing capability for specific purposes, from the most trivial ones (entertainment, ticket booking, home automation) to more challenging ones (stock exchange, road lights, or nuclear plant management). User expectations on IT service continuity and performance involve continuous monitoring and improvement of the IT network assets and configuration. Incidents (unplanned interruption or quality reduction of an IT service due to spontaneous hardware fault, cybersecurity attacks, or side problems) and daily network operations (system reconfiguration, hardware/software upgrade) directly affect the user experience depending on the IT network characteristics (device redundancy, IT resource mutualization, network size, and interconnections). Maintaining high-standard

*Corresponding author. E-mail: lionel.tailhardat@orange.com.

1 quality of service is a challenging task on large-scale networks due to varied Information and Communications 1
2 Technology (ICT) systems characteristics and behaviors. Tackling this situation with Artificial Intelligence (AI) re- 2
3 quires to handle data from heterogeneous sources (network devices, Network Monitoring Systems - NMSs, Security 3
4 Information and Event Management - SIEM systems) and apply reasoning for situation awareness on dynamic data 4
5 with both current and past knowledge (network characteristics, state and known errors). 5

6 Our main contribution is the so-called NORIA Ontology (NORIA-O thereafter) for representing network in- 6
7 frastructures, incidents and maintenance operations on networks. This ontology re-uses and extends well-known 7
8 ontologies such as SEAS [1, 2], FOLIO [3], UCO [4], ORG [5], BOT [6] and BBO [7] for describing technological 8
9 systems and activities from the business process modeling perspective. NORIA-O is also enriched with controlled 9
10 vocabularies allowing to handle heterogeneous data from varied ICT systems and incident situations through a 10
11 small set of centrally defined and shareable definitions. NORIA-O has been developed within the Orange¹ com- 11
12 pany, a leading international network infrastructure and service provider. Its long-standing experience on complex 12
13 network design and management allows us to back NORIA-O with insightful details from domain experts and to 13
14 evaluate the model with real-world data. The ontology, controlled vocabularies and their associated documentation 14
15 are available at <https://w3id.org/noria>. 15

16 The remainder of this paper is organized as follows. In Section 2, we review RDF-based and non-RDF based 16
17 data models enabling to represent network infrastructures and incidents. In Section 3, we describe the methodology 17
18 we follow to design NORIA-O, starting from eliciting Competency Questions (CQs) that capture the knowledge 18
19 of network and security experts. We introduce an overview of the ontology alongside four facets of discourse: 19
20 structural, functional, dynamic and procedural. We also present some key design choices such as re-using existing 20
21 models and vocabularies to capitalize on the community's efforts in order to facilitate the reuse of NORIA-O. In 21
22 Section 4, we take a deep dive into the different concepts of the ontology implemented in [RDFS/OWL-2](#) as well 22
23 as into the associated vocabularies. We evaluate the ontology with respect to our requirements and competency 23
24 questions (CQs) in Section 5. We exemplify how NORIA-O is used for the supervision of a realistic network 24
25 infrastructure in Section 6. Finally, we conclude and outline some future work in Section 7. 25

26 2. Related Work 26

27 27
28 28
29 29
30 Previous works have demonstrated that the use of semantic modeling is of interest for network infrastructure 30
31 monitoring (e.g. INDL [8], CRATELO [9], UCO [4], ToCo [10], ACCTP [11], DevOpsInfra [12]). Various tools 31
32 for the construction and exploitation of knowledge graphs have also been proposed for data integration (e.g. 32
33 RMLMapper [13], for log parsing and semantization (e.g. SLOGERT[14]), for vocabulary reconciliation (e.g. 33
34 String2Vocabulary [15]), and for visualization (e.g. Gephi [16]). We posit that these works partly cover the knowl- 34
35 edge domains required for describing ICT systems and related management activities (e.g. incident management, 35
36 cybersecurity risk evaluation). For example, the combination of the SEAS and PEP [1, 2] models is useful for 36
37 describing technological systems, commands and observed values from probing devices. However, SEAS mostly 37
38 targets the Internet of Things (IoT) domain and end-user devices, and the semantics of PEP relates to computer 38
39 process. The DevOpsInfra [12] ontology describes sets of computing resources and how they are allocated for hosting 39
40 services. However, concepts are missing for a finer grain description of the network topology. Moreover, the ontol- 40
41 ogy mostly focuses on the provisioning activity and is not aligned with other well-known models such as SOSA [17] 41
42 and the TMForum Open API² for interoperable definitions of states and operations. The CRATELO [9] model en- 42
43 ables describing and reasoning on cyber operations. Used in combination with the PACO [18] model, reasoning on 43
44 network traffic from the defenders' and attackers' perspective is possible. However, concepts for network topology 44
45 and operations are missing for contextualizing network traffic sessions within the network topology itself and the 45
46 day-to-day operations. 46

47 In this paper, we aim to fill this gap in developing a [comprehensive semantic model](#) for describing and reasoning 47
48 on the combination of network infrastructure characteristics (e.g. device type, links), network activity (e.g. user 48
49

50 ¹<https://www.orange.com/>

51 ²<https://github.com/tmforum-apis>

login, interface operational status change, processor overload alert) and operations (e.g. software upgrade, server reboot, link decommissioning). Based on various selection criteria, such as the coverage of our target knowledge domain and the potential enrichment of existing ontologies, we design the NORIA-O ontology so that it builds upon some of these existing semantic models as described in Section 3.

3. Methodology

In this section, we describe the knowledge engineering methodology we used to develop the NORIA-O ontology. First, we capture Competency Questions (CQs) from a panel of experts familiar with network operation issues and derive archetypes from these CQs for further analysis (Section 3.1). Second, we show how we designed the conceptual model (Section 3.2).

3.1. Competency Questions and Conceptualization

We gathered experts from several entities in the fields of engineering, operations, supervision, and incident management on networks and data centers, including teams from Network Operation Centers (NOCs) and Security Operation Centers (SOCs). This panel consists of 16 experts who collectively represent 150 operations team members. To effectively capture the needs and knowledge of the experts, we followed a user-centered design methodology combined with ontology engineering methods.

From our review of the literature, the Competency Question approach [19] turns out to be the most intuitive and straightforward with respect to how NOC and SOC teams use and talk about their tools. Indeed, this approach involves extracting the conceptual model of the knowledge domain by analyzing user queries expressed in natural language through a set of semantic patterns. During several iterations of knowledge capture meetings on a shared notebook, the experts could validate, invalidate, add and modify competency questions. At the end of this stage, the teams validated 26 CQs presented in Table 1. This includes questions on events, resources (e.g. server, router), applications (e.g. Domain Name System, Video-on-Demand service platform), log and alarms (e.g. login, CPU overload) and operation plan (e.g. SSL/TLS certificate renew, IS-IS interface re-prioritization).

From the set of CQs, we derived a conceptual model of the domain of discourse by applying the “Competency Question archetype mapping” approach [19, §4.3]. Table 2 provides an example of a specific competency question (CQ1) being mapped to an archetype. We also adhered to the guidelines of the Linked Open Terms (LOT) methodology [20], which notably include reusing or aligning with existing vocabularies.

Scrutinizing the conceptual model and candidate vocabularies, one may observe the two following characteristics. Firstly, we remark that concepts derived from the Table 1 can be referred to as “atomic concepts” or “primitive concepts” (e.g. application, alarm, resource). These terms represent concepts that are not defined by composition but are considered basic or indivisible in nature. Thus, we can expect to use simple relationships (potentially hierarchical) between concepts during the modeling and implementation phases. Secondly, we observe that the knowledge domain can be structured into four facets, which enables addressing and discussing each facet independently, namely: structural, functional, dynamic and procedural. Indeed, adjacent research fields break down their analysis of complex and dynamic systems according to similar and compatible principles. For instance, understanding the influence of network structure in diffusion processes [21] involves separate encoding of events and reactions. Similarly, the deduplication and coalescing of software code [22] rely on semantic analysis and function hierarchization for categorizing code clones into families such as “behavioral” (identical inputs and outputs), “conceptual” (identical design patterns), “structural” (identical sequence of steps), and “model” (identical classes and relations). Additionally, automated reasoning, drawing on cognitive science ideas, introduces the concept of “active inference” [23] where procedures are akin to calculus processes with pre/post-conditions. This reflects an operational knowledge that adapts to observed conditions. We further define these four facets in the next section.

3.2. Domain of Discourse and Modeling Strategy

Facets. Considering dynamic ICT systems with constrained and multi-level functional behavior, we define the four following facets for structuring the knowledge domain. An illustration of these facets is provided in Figure 1.

High-level concepts	CQs	Arch. ID	AT Eval.
agent, document, object	1) Which resource/application/site is concerned by a given incident?	1	OK
object	2) What assets are shared by a given asset chain?	6	OK
event, object	3) What logs and alarms are coming from a specified resource?	1	OK
object	4) Which metrics are coming from a specified resource?	1	OK
event	5) To which event family does this log belong and is this event normal or abnormal?	3	OK
event	6) What events are associated with a given event?	1	OK
agent, event, object	7) Which agent/event/resource caused the event under analysis?	1	OK
event	8) What do the various fields in the log refer to?	1, 3	OK
event	9) Is there any pattern in a given set of logs/alarms?	1, 6	AI
document, object	10) What interventions were carried out on this resource that could have caused the incident?	1, 6	OK
document, event	11) What was the root cause of the incident?	6	AI
document, event	12) Which sequence of events led to the incident?	6	OK
event, object	13) On which resource did this sequence of events take place and in which order?	1	OK
document, event	14) What past incidents are similar to a given incident?	6	AI
document, event, procedure	15) What operation plan (automation, operating procedures, etc.) could help us solve the incident?	1, 3	AI
document, event, procedure	16) What corrective actions have been carried out so far for a given incident?	1	OK
document, event, procedure	17) What is the list of actions taken that led to the resolution of the incident?	1	OK
document, event, procedure	18) Given all the corrective actions carried out so far for the incident, what assumptions covered the actions taken?	1, 4	AI
document, event, procedure	19) What has been the effect of the corrective actions taken so far for the incident?	1	OK
document, event, procedure	20) Given all the corrective actions carried out so far for the incident, what possible actions could we still take?	6	AI
document, event	21) What is the summary of this incident and its resolution?	1	OK
agent, document	22) Which agents were involved in the resolution of the incident?	1	OK
document	23) What is the financial cost of this incident if it occurs?	2	Extension
document, event	24) How long before this incident is resolved?	1	AI
document, object	25) What are the vulnerabilities and the associated risk levels of this infrastructure?	1, 2	AI
event, object	26) What is the most likely sequence of actions that would cause this infrastructure to fail?	6	AI

Table 1

CQs with their archetype and the authoring tests results.

CQ	<i>“Which resource/application/site is concerned by a given incident?”</i>
Pattern	Which [CE1] [OPE] [CE2] ? (Archetype ID: 1)
Components	CE1: Asset (resource/application/site), OPE: areContainedIn, CE2: Incident

Table 2

NORIA-O CQ break down on CQ#1 using the “Archetype 1” as defined in [19, Table 1].

- **The structural facet** describes the physical and logical elements of the network. It allows modeling the equipment classes, connections and compositions. This facet aims to support calculations on network objects and properties (direct or deduced) and calculations on the physical and logical structures (real or patterns).
- **The functional facet** describes services provided and diffusion areas. This facet makes it possible to meet the need for functional isomorphism (e.g. replacing one piece of equipment with another performing the same function). It allows modeling the service types, interactions between them, and compositions. This facet allows calculations on network domains and their properties (direct or inferred) and calculations on services and streams (e.g. “end-to-end” notion).
- **The dynamic facet** describes the sequence of events. It allows modeling the occurrence of an event on a given equipment or service as well as precedence relationships. This facet aims to support time calculations (absolute, relative, membership) and causality calculations (first order or probabilistic).
- **The procedural facet** describes how things work and should be interpreted. Automation principles (e.g. fail-over mechanisms of redundant systems) or operation principles (e.g. doubt removal procedures) are expected parts of this facet. Associated application goals are deductive/abductive reasoning over facts and reflection over knowledge for automated learning (discovery/recommendation) of procedures (e.g. evolutionary search over targeted goals, composition calculus over sequences of events).

Modeling strategy. Considering the domains of anomaly detection and incident management, we consider incidents as a central concept for i) computing and reasoning about anomaly signatures, and ii) linking trouble tickets

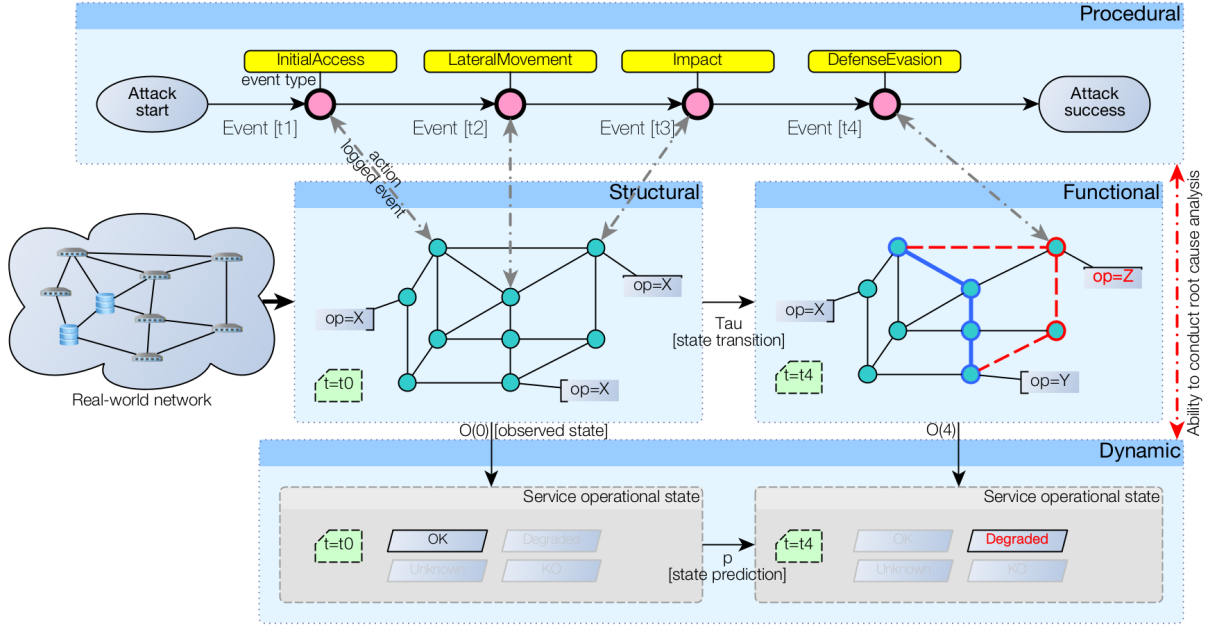


Figure 1. ICT system state transition model and relations to the NORIA-O facets. The representation of a real network can be divided into four facets: *structural* = technical architecture and parameters, *functional* = combination of structural elements into services (the blue path indicates an operational flow, the red path a faulty flow), *dynamic* = monitors the evolution of service states and technical indicators, *procedural* = network/user activity models (learned or asserted) for contextualizing event traces from the dynamic facet. The procedural example relates logged events of some course of action to cybersecurity attack tactics, where tactics' names are from the MITRE ATT&CK matrix [24]. *Tau* stands for state transition, $O(t)$ for observed state at time t , and p for state prediction.

to anomaly signatures for root cause analysis tasks. We introduce the `noria` namespace, which encompasses the entire set of NORIA-O concepts and relations in a unified manner. To align with risk management and business modeling practices (such as the incident logging and categorization steps in the ITIL's Incident Management Process model – IMP [25]), we adopt a top-to-bottom modeling strategy, starting from the process and extending to objects. Due to the “atomic concepts” characteristic of the knowledge domain (Section 3.1) and our top-to-bottom approach, we primarily focus on an axiomatization based on subsumption with RDFS [26], and with OWL [27] for cases that require additional logical structuring (e.g. class disjointness, property qualification)³.

Model re-use. Following the best practices in ontology development, we aim to re-use existing data models and vocabularies as a base and extend them to represent domain-specific classes and properties. From RDF-based ontologies, we interconnect and/or extend the following models: **BBO** [7] for describing activities from the business process modeling perspective in conformance to the Business Process Model Notation (BPMN); **BOT** [6] for describing resource locations and enabling geographical neighboring analysis for root cause analysis tasks; **DC-TERMS** for standard management of NORIA-O instances as parts of a catalog; **DevOpsInfra** [12] for enabling potential interactions of the NORIA-O model with the DevOps perspective; **FOAF** [29] for describing social organizations; **FOLIO** [3] for enabling Root Cause Analysis (RCA) tasks based on the Failure Mode and Effect Analysis (FMEA) approach; **ORG** [5] for describing stakeholders and related organizations; **SEAS & PEP** [1, 2] for describing technological systems, measures, commands, and results; **UCO** [4] for enabling cybersecurity risk assessment on instances of the NORIA-O model; **SLOGERT** [14] for describing system logs and enabling potential usage of the SLOGERT log interpretation framework.

From non RDF-based data models, we take advantage of the concept hierarchy and vocabulary definitions from the **TMForum Data Model**⁴ for enabling an interoperable definition of trouble tickets and change requests with

³During the implementation phase, we did not prioritize the use of cardinality restrictions as we did not consider them crucial at that time. Instead, we perceive cardinality restrictions as more beneficial in a post-implementation data quality task using SHACL [28].

⁴<https://github.com/tmforum-apis>

third-party Operations Support Systems (OSSs) and Decision Support Systems (DSSs), ITU-T [30, 31] for standard definitions of notifications and ways to handle them within the telecommunication industry, IETF for precise use of terminology in the context of a Request for Comments (RFC)⁵.

We propose implementing alignment with third-party data models and vocabularies on a class or property basis when relevant, using the dedicated OWL and RDF constructs such as `owl:equivalentClass`, `rdfs:subClassOf`, or `rdfs:isDefinedBy`. Similarly, we aim to provide guidelines for directly instantiating these vocabularies in cases where aligning a class or property would be redundant.

Modeling observations. Considering observables and their state change (e.g. the operational state of a network interface, the temperature measurements from a device sensor), we observe that modeling and logging observations can be done: (a) as a string, (b) as a concept from a controlled vocabulary, (c) as an instance, (d) as an instance with time property or time instance (e.g. using reification, or following the `sosa:Observation` model⁶). These four options are relevant for the NORIA-O application domain. The concern is not about choosing one option for all situations, but how we can mix them. Hence, we adopt the following selection criterion: 1) use (a) and (b) for invariant properties, 2) use (c) and (d) for time-dependent and/or specific use-case extensions to NORIA-O (i.e. additional observables are defined in a side vocabulary so the main ontology remains stable).

Controlled vocabularies. Because of potentially heterogeneous data incoming from varied ICT systems and incident situations to handle, we take notes of terms from datasets and other ontologies for building up a controlled vocabulary. This aims at efficient management of anomaly detection patterns, rules and methods by reducing the lexical range of possible situations to interpret. For this, we propose a set of domain-specific vocabularies (e.g. Incident Management Process, Application, Notification vocabularies) modeled as SKOS concepts within concept schemes (e.g. the milestones of the Incident Management Process). We add, whenever available, alternate definitions of the concepts for reconciliation of similar object attribute values through a single concept reference (e.g. communication devices may report the same status of network interfaces with varied terms such as “active”, “up” or “enabled”). We also use the concept scheme approach for enabling multiple interpretation of a similar concept. For example, an event may be categorized as an `integrityViolation` based on the analysis of the event text, which allows us to reason on the event type and infer a `SecurityAlarm` thanks to a dual membership of the `integrityViolation` concept definition. The implementation of the vocabulary reconciliation task (e.g. relating the observed network interface administrative status to the adequate concept reference with help of natural language processing – NLP) is out of the scope of this paper and is left to the NORIA-O user’s choice.

4. NORIA-O: Formalization and Implementation

We have formalized and implemented the NORIA-O conceptual model in RDFS/OWL-2. The NORIA-O ontology consists of 59 classes, 107 object properties, and 71 datatype properties. It is organized with the four facets presented in Section 3.2 and illustrated in Figure 2. Its expressivity is *ALCHOI(D)* as per Protégé 5.1⁷. In this section, we introduce some of the main concepts and properties.⁸

4.1. Resources, Network Interfaces, Network Links and Applications

Within computer science, a resource is some “part contributing to the functioning of an ICT system.” Similarly, as per the TMForum Data Model⁹, a resource is “an abstract entity that describes the common set of attributes shared by all concrete resources (e.g. TPE, EQUIPMENT) in the inventory”. Therefore, we define the `Resource` class for describing any physical or logical manageable entity composing the network at hand. Defining the type of

⁵<https://datatracker.ietf.org/>

⁶<https://www.w3.org/TR/vocab-ssn/#SOSAObservation>

⁷<https://protege.stanford.edu/>

⁸Data model directly available at <https://w3id.org/noria/ontology/> using an ontology browser, documentation at <https://w3id.org/noria/doc/>

⁹https://github.com/tmforum-apis/Open_Api_And_Data_Model

a resource is made possible through object properties such as `resourceType` (i.e. controlled-vocabulary concepts such as rack, server, router, virtual machine, etc.) and `resourceProductModel` (i.e. entity model instances). Additional properties allow for identifying the resources based on their logistic identifier, hostname, installation date, etc.

Locating and reasoning over a physical entity from a geographical standpoint is available with a chain of `bot:containsZone` and `bot:hasElement` properties, starting from a `bot:Site` with `bot:hasZeroPoint` property, down to a `Locus` concept for precise Resource location within a `Room` (i.e. a specialization of `bot:Space`). Locating a resource is also available through a dependency relationship with the `seas:subSystemOf` object property from the SEAS SystemOntology¹⁰. This allows for describing and reasoning with parts from various levels of organization (e.g. a virtual router instance in a router, a hard drive in a server, a server in a rack, a rack in a `bot:Site`, etc.).

Describing the network topology itself is defined with the `NetworkInterface` and `NetworkLink` classes. We align with the SEAS CommunicationOntology¹¹ model through object properties such as `networkInterfaceOf` and `networkLinkTerminationResource`. It should be noted that this approach is compatible with advanced networking features such as sub-interfaces, link aggregation, virtual channels, etc. Operational characteristics for interface and links are available with properties such as `networkInterfaceOperationalStatus` and `networkInterfaceRoutingPriorityMetric`.

The `Application` concept enables to define models of purpose (e.g. Internet access, network time, alarm monitoring) for sets of resources, and to categorize these with respect to their nature (i.e. controlled-vocabulary concepts such as infrastructure, service platform, etc.). An `ApplicationModule` is a concrete instance of a given model (e.g. national federated Internet access, corporate network time service, monitoring for in-production devices). This grouping level enables to relate specific technical skill centers, such as a named IP backbone engineering or support team (Section 4.4), to a given module for specific expertise (e.g. re-engineering, diagnosis and repair). Additional properties allow for finer grain resources and events management at the module level such as `applicationModuleSlaLevel` for prioritizing servicing teams, or `applicationModuleHotlineEnabled` for triggering night shift support teams.

We also define the `Service` concept in accordance to the TMForum TMF638 Service Inventory API¹² and the IETF SFC Architecture [33] for grouping instances of `ApplicationModule`, and thus enabling the data path and application composition perspectives of the functional facet (Section 3.2). The network topology related to a given service is inferred from the set of resources, network interfaces and network links included in each application that is part of the service. We observe that, although deterministic, the data path granularity calculus for some communication session (e.g. a time-bounded IP/http query with its response) depends on the specificity of the resources included in `ApplicationModule` instances. For example, the resulting granularity for a “national IP backbone infrastructure” application instance will correspond to the routing domain.

4.2. Logs and Alarms

As per the International Telecommunication Union (ITU), “the log is a repository for records” (ITU-T Rec. X.735) [30] and an event log record “represents the information stored in the log as a result of receiving notifications or incoming event reports” (ITU-T Rec. X.721) [31]. Based on this definition, we define the `EventRecord` class for storing any event coming from managed objects (e.g. `Resource`, `Application`) such as system logs [34], SNMP Traps [35] and application specific messages (e.g. user applications, operational support systems, processing platforms). Fundamental properties such as `loggingTime`, `logText`, `logOriginatingManagedObject` and `logOriginatingManagementSystem` allow for keeping track of the event origin and content. Details about the message meaning are managed with the `dcterms:type` property that refers to a controlled-vocabulary for event type tagging¹³ (e.g. state change, processing error alarm, integrity violation). The `alarmSeverity`

¹⁰<https://w3id.org/seas/SystemOntology>

¹¹<https://w3id.org/seas/CommunicationOntology>

¹²<https://github.com/tmforum-apis>

¹³Event type tagging can be carried-out at the data integration stage, or through a posteriori language processing of the “logText” property.

property complementarily provides an indication of how it is perceived that the capability of the managed object has been affected, or how serious are the service affecting conditions (including for security alarms). Additional properties related to alarm management and interpretation are available with alignment to the DCTERMS and PEP models, such as: `alarmMitigatedBy` and `dcterms:relation` for aggregating events and building event signatures; `dcterms:conformsTo` for root cause analysis and repair planning; `dcterms:mediator` for acknowledgment and responsibility follow up.

4.3. Trouble Tickets and Change Requests

We define the `TroubleTicket` concept accordingly to the TMForum DataModel where a trouble ticket is “a record of an issue that is created, tracked, and managed by a trouble ticket management system”¹⁴. It is not an event per se (Section 4.2), but a mean to efficiently manage targeted resource/service (e.g. `troubleTicket-RelatedResource` property) restoration operations through collaboration. Hence, we also consider trouble tickets as a product of the ITIL’s Incident Management process [25], and relate them to ITIL’s Problem Management process [36] and the BPMN by alignment to the `BBO:DataResource` class.

Corrective maintenance actions are logged as `TroubleTicketNote` and related to the parent `TroubleTicket` with the `dcterms:isPartOf` property. Actions’ accountability is implemented with the `dcterms:creator` in relation to the `foaf:Agent` class (Section 4.4). Correlating actions to the digital traces (i.e. `EventRecord`, Section 4.2) they produce at the structural and functional level (e.g. login, configuration change, upgrade) is available with the `dcterms:relation` property towards a `pep:ProcedureExecution-Container` entity.

We provide additional properties for improving the incident diagnosis stage efficiency (e.g. `dcterms:hasPart` for hierarchical grouping of tickets), and moving towards root cause analysis based on the notion of Known Error Database (KEDB) (e.g. `troubleTicketCategory` and `problemCategory` for a priori and a posteriori categorization, respectively) and primary/secondary anomaly (cause/effect) with alignment to the FOLIO model. With greater details, a trouble ticket is a document transitively referencing a set of corrective maintenance actions that can be abstracted into an issue remediation `OperationPlan` for solving the `AnomalyPattern` at hand. Reaching such abstraction from actions’ digital traces is enabled by considering the PEP model with `TroubleTicket` as a specialization of a `pep:ProcedureExecutionContainer`, actions as `pep:ProcedureExecution` and `OperationPlan` as `pep:Procedure`.

Similarly to trouble tickets, we define the `ChangeRequest` concept according to the TMForum DataModel¹⁵ for tracking scheduled change operations (as sets of `pep:ProcedureExecution` carried-out in correspondence to a given `OperationPlan`) with structural or functional impact, and computing (potential) causality for trouble tickets based on the set of correlated resources/applications and operations start/end time.

4.4. Agents, Teams and Organizations

From the incident management perspective, finding experts in short time is key for operational efficiency. A typical organization for this is to build teams based on technical expertise (e.g. routing and international backbone, servers and virtual machines, forensics and malware retro-engineering), assign teams to the management of a fleet of equipments or services, and rely on external support/engineering services for specific cases. We take advantage of the FOAF [29] and ORG [5] ontologies for interoperability with complementary knowledge bases (e.g. `foaf:Person`, `org:OrganizationalUnit`, `org:Organization`) and we model relationships with IT entities (e.g. `Resource`, `Application`) using properties such as `elementManagedBy` or `applicationModuleRelatedParty`. We also define the `CorporateUserIdentifier` class as a specialization of `foaf:OnlineAccount`, and provide a controlled vocabulary for finer grain role description for agents, teams (e.g. Technical Support Group) and organizations (e.g. Manufacturer). This notably enables cyber security out-of-policy approach (i.e. what is not defined is not allowed) for tracking non-legitimate operations (unless

¹⁴<http://datamodel.tmforum.org/en/master/Common/TroubleTicket/>

¹⁵<http://datamodel.tmforum.org/en/master/Common/ChangeRequest/>

facing an insider) by asserting access control groups as `org:OrganizationalUnit` and scrutinizing observed or declared user actions (e.g. `eventLogOriginatingAgent`, `dcterms:creator`). For this, we assume that companies' human-resource databases are reliable and accurate sources of truth.

5. Evaluation

We have evaluated the NORIA-O [model implementation](#) according to the ability of the model to answer the CQs that were collected in Section 3.1. The CQs have emerged from an iterative and collaborative process of capturing knowledge from domain experts. Therefore, we consider that translating these CQs into Authoring Tests (ATs) [19, 37] and obtaining a satisfactory answer to these SPARQL queries from the knowledge graph constitute a sound evaluation of NORIA-O. This evaluation aims to check that all the concepts and relations that are important for the experts' needs are included in the ontology. [The set of authoring tests, available at https://w3id.org/noria/evaluation](https://w3id.org/noria/evaluation), has been defined and tested on two knowledge graph instances structured by NORIA-O. The first one describes a fictitious case of supervision and is publicly available (Section 6). The second one has been generated from Orange internal data (10 data sources encompassing 128 features over 15 tables) using an in-house data pipeline [38]; the size of the resulting RDF dataset is approximately 4 million triples for 400K entities, including streamed events spanning over 111 days.¹⁶

After this evaluation, we distinguish three different situations depicted in column "AT Validation" of Table 1. First, a large number of CQs (16/26) can be answered using a single or several simple SPARQL queries and the ontology ("OK" in Table 1). 9/26 CQs ("AI" in Table 1) are partially satisfied using SPARQL queries, to which complementary AI techniques should be added to fully address the underlying objectives posed by these CQs (e.g. incident context capture [43], anomaly detection algorithms). For example, to answer CQ#11 ("What was the root cause of the incident?"), the explicit representation of alarms and logs associated with a given incident is not enough and needs to be enhanced with root cause analysis algorithms. Another example is CQ#25 ("What are the vulnerabilities and the associated risk levels of this infrastructure?") that can be answered only by looking for non-desirable network topology shapes or relations to third-party cybersecurity vulnerability entities based on structure and security scanners. Third, 1/26 CQs requires the introduction of new concepts or relations via an extension of the NORIA-O model ("Extension" in Table 1). The CQ #23 ("What is the financial cost of this incident if it occurs?") involves information about the cost of an incident. This will be the object of a future extension of NORIA-O leveraging the SEAS Failable System ontology [2].

6. Use Case: Modeling a Complex IT Infrastructure

To illustrate the usage of NORIA-O and the expressiveness of the model, this section proposes an example of instantiation of the semantic model on a fictitious case of supervision of a network infrastructure of a large company. [Figure 3 summarizes this use case with both the network topology and the NORIA-O model instances](#). The `term1` laptop tries to connect to the application `APP_PFS001` that is hosted on the server `srv1`. A trouble ticket `TT_TOY2022TT` is issued for tracking the diagnosis and restoration actions of the service. A set of events originating from the resources involved in the data path and hosting infrastructure is available for analysis and filtering-out primary from secondary events (cause vs effects): the connection fails because of a `communicationsSubsystemFailure` situation on the access switch `as2`, with fault spreading phenomenon leading to `LOG_srv1[...]` record for the network interface state change of `srv1`. A member of the support team for the `as2` switch applies the `PROC_NetCardReset` procedure for remediation. The corresponding command-line interface action is logged (`LOG_as2[...]` with `attributeValueChange` type), and related to the action report note at the trouble ticket level (`ldp:member` relationship) along with the resulting `equipmentAlarm` clearing effect (`alarmMitigatedBy` relationship).

The dataset corresponding to this scenario is available at <https://w3id.org/noria/dataset>. 660 triples are needed for representing the full scenario with additional resources, organization and root cause analysis details.

¹⁶Due to confidentiality, this dataset is not made public.

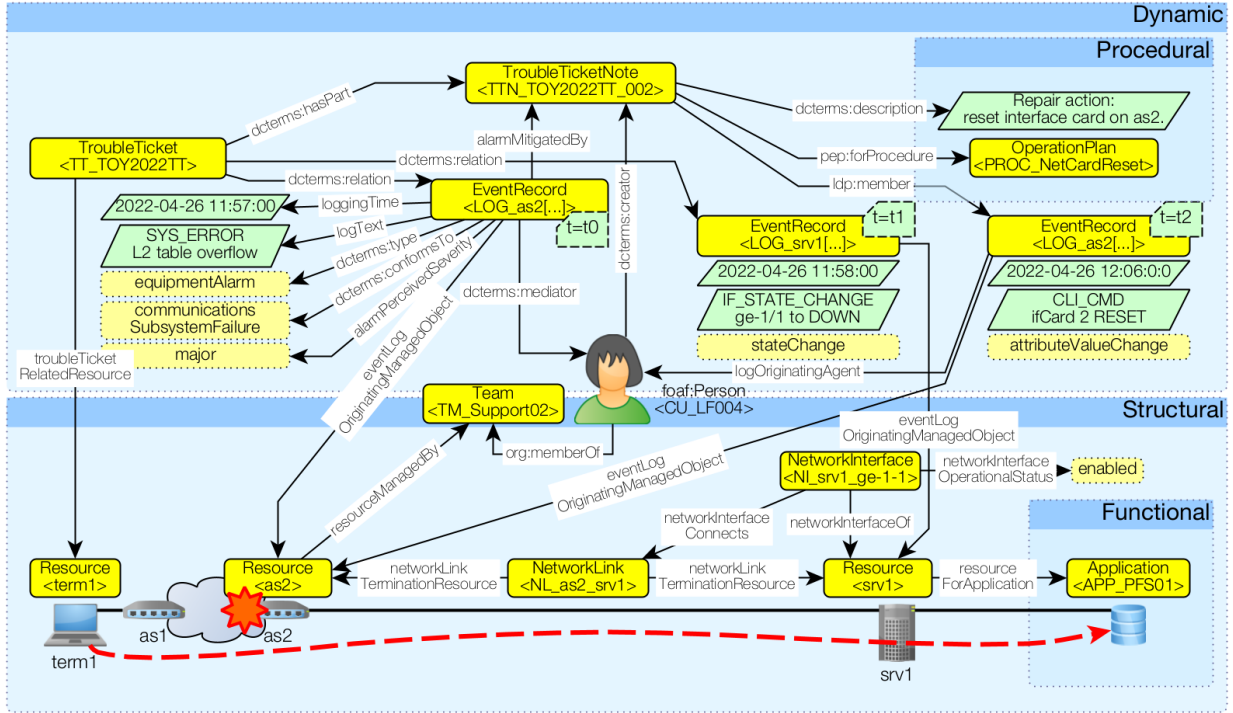


Figure 3. NORIA-O model instantiation example. Instantiation of the semantic model on a fictitious case of supervision of a network infrastructure. *norია* is the default namespace for classes and properties. The red star indicates the location where the fault occurs in the network. The entity names and property values correspond to the NORIA-O public dataset. Values for the “SKOS scheme” blocks correspond to the vocabulary provided in NORIA-O.

7. Conclusion and Future Work

In this paper, we have presented NORIA-O, an ontology for representing network infrastructures, incidents and maintenance operations on networks, that relies on and extends well-known semantic models such as BBO, BOT, FOAF, FOLIO, SEAS and UCO. The NORIA-O ontology is available at <https://w3id.org/noria> under a BSD-4 License, along with its documentation and a sample dataset. We conducted a thorough evaluation of the NORIA-O model implementation using the Competency Questions & Authoring Tests methodology [19]. This evaluation demonstrates the suitability of the model according to the expert needs.

With the Anomaly Detection and Incident Management applications in mind, we observe that computing and reasoning on *anomaly signatures*, and linking *trouble tickets* to *anomaly signatures* for root cause analysis tasks, are features that naturally go beyond logical inference, and thus require additional AI methods. First, event logs from heterogeneous data sources depicting an identical phenomenon need to be parsed and categorized in the same way. Typical options are to focus on log parsing [39, 40] and semantization [14], either before or after the data integration stage. We note the importance of NLP-related techniques such as Named Entity Recognition [41], Topic Modeling [42] and Vocabulary Reconciliation [15]. Careful consideration is required in regard to the semantic interpretation of log messages due to the current lack of accurate technical language models. **Second, learning and exploiting anomaly models**, such as discussed in [43] for the classification of trouble tickets with graph embeddings and a random forest model, requires to filter-out event logs and alarms on both trouble tickets’ timespan and impacted resources characteristics. Thereon, we argue that “link prediction with a confidence metric” is a pivotal task for inferring new relationships such as *relatedEvent*, *hasProbableCause* or *similarOperationPlan* (e.g. relating events with attack scenarios [44]). Recent research efforts on dynamic graphs with event streams show promising results by combining graph embedding, Top-K ranking and co-occurrence encoding techniques [45–47]. We observe that computing an average timespan on learned issue remediation operation plans enables predicting the mean-time-to-repair (MTTR) metric for trouble tickets.

Finally, we note that network resilience and cybersecurity application domains will benefit from extensions of a NORIA-O knowledge graph with third-party data collection tools. For example, network topology anti-patterns (e.g. with the Shapes Constraint Language toolset – SHACL [28]) and semantic interpretation of the ICT resources configuration [48] could be related to the network performance and issues. Similarly, integrating data from vulnerability scanners and Cyber Threat Intelligence tools (e.g. OpenCTI¹⁷) could enable cybersecurity risk evaluation (e.g. Common Vulnerability Scoring System – CVSS [49]) and minimization (e.g. optimizing countermeasure placement [50]).

References

- [1] Maxime Lefrançois, Planned ETSI SAREF Extensions Based on the W3C&OGC SOSA/SSN-compatible SEAS Ontology Patterns, in: *Workshop on Semantic Interoperability and Standardization in the IoT (SIS-IoT)*, 2017.
- [2] Maxime Lefrançois, Jarmo Kalaoja, Takoua Ghariani and Antoine Zimmermann, SEAS Knowledge Model, Deliverable, 2.2, ITEA2 12004 Smart Energy Aware Systems, 2016.
- [3] Bram Steenwinckel, Pieter Heyvaert, Dieter De Paepe, Olivier Janssens, Sander Vanden Haute, Anastasia Dimou, Filip De Turck, Sofie Van Hoecke and Femke Ongenaë, Towards Adaptive Anomaly Detection and Root Cause Analysis by Automated Extraction of Knowledge from Risk Analyses, in: *9th International Semantic Sensor Networks Workshop (SSN)*, 2018.
- [4] Zareen Syed, Ankur Padia, M. Lisa Mathews, Tim Finin and Anupam Joshi, UCO: A Unified Cybersecurity Ontology, in: *AAAI Workshop on Artificial Intelligence for Cyber Security*, 2016.
- [5] Dave Reynolds, The Organization Ontology, W3C Recommendation, W3C, 2014.
- [6] Mads Holten Rasmussen, Maxime Lefrançois, Georg Ferdinand Schneider and Pieter Pauwels, BOT: The Building Topology Ontology of the W3C Linked Building Data Group, *Semantic Web Journal* (2020). doi:10.3233/SW-200385.
- [7] Amina Annane, Nathalie Aussenac-Gilles and Mouna Kamel, BBO: BPMN 2.0 Based Ontology for Business Process Representation, in: *20th European Conference on Knowledge Management (ECKM)*, 2019.
- [8] Mattijs Ghijsen, Jeroen Van Der Ham, Paola Grosso, Cosmin Dumitru, Hao Zhu, Zhiming Zhao and Cees De Laat, A Semantic-Web Approach for Modeling Computing Infrastructures, *Computers & Electrical Engineering* (2013). doi:10.1016/j.compeleceng.2013.08.011.
- [9] Alessandro Oltramari, Loria Cranor, Robert Walls and Patrick McDaniel, Building an Ontology of Cyber Security, in: *9th Conference on Semantic Technologies for Intelligence, Defense, and Security (STIDS)*, 2014.
- [10] Qianru Zhou, Alasdair J. G. Gray and Stephen McLaughlin, ToCo: An Ontology for Representing Hybrid Telecommunication Networks, in: *16th European Semantic Web Conference (ESWC)*, 2019. doi:10.1007/978-3-030-21348-0_33.
- [11] A. Brazhuk, Threat Modeling of Cloud Systems with Ontological Security Pattern Catalog, *International Journal of Open Information Technologies* (2021).
- [12] Oscar Corcho, David Chaves-Fraga, Jhon Toledo, Julián Arenas-Guerrero, Carlos Badenes-Olmedo, Mingxue Wang, Hu Peng, Nicholas Burrett, José Mora and Puchao Zhang, A High-Level Ontology Network for ICT Infrastructures, in: *20th International Semantic Web Conference (ISWC)*, 2021. doi:10.1007/978-3-030-88361-4_26.
- [13] Anastasia Dimou, High Quality Linked Data Generation from Heterogeneous Data, PhD thesis, University of Antwerp, 2017.
- [14] Andreas Ekelhart, Fajar J. Ekaputra and Elmar Kiesling, The SLOGERT Framework for Automated Log Knowledge Graph Construction, in: *18th European Semantic Web Conference (ESWC)*, 2021. doi:10.1007/978-3-030-77385-4_38.
- [15] Pasquale Lisena, Konstantin Todorov, Cécile Cecconi, Françoise Leresche, Isabelle Canno, Frédéric Puyrenier, Martine Voisin, Thierry Le Meur and Raphaël Troncy, Controlled Vocabularies for Music Metadata, in: *19th International Society for Music Information Retrieval Conference (ISMIR)*, 2018.
- [16] Mathieu Bastian, Sebastien Heymann and Mathieu Jacomy, Gephi: An Open Source Software for Exploring and Manipulating Networks, in: *3rd International AAAI Conference on Weblogs and Social Media (ICWSM)*, 2009. doi:10.1609/icwsm.v3i1.13937.
- [17] Krzysztof Janowicz, Armin Haller, Simon Cox, Danh Phuoc and Maxime Lefrançois, SOSA: A Lightweight Ontology for Sensors, Observations, Samples, and Actuators, *SSRN Electronic Journal* (2018). doi:10.1016/j.websem.2018.06.003.
- [18] Noam Ben-Asher, A. Oltramari, R. Erbacher and Cleotilde González, Ontology-Based Adaptive Systems of Cyber Defense, in: *10th Conference on Semantic Technology for Intelligence, Defense, and Security (STIDS)*, 2015.
- [19] Yuan Ren, Artemis Parvizi, Chris Mellish, Jeff Z. Pan, Kees van Deemter and Robert Stevens, Towards Competency Question-Driven Ontology Authoring, in: *11th European Semantic Web Conference (ESWC)*, 2014. doi:10.1007/978-3-319-07443-6_50.
- [20] María Poveda-Villalón, Alba Fernández-Izquierdo, Mariano Fernández-López and Raúl García-Castro, LOT: An industrial oriented ontology engineering framework. Engineering Applications of Artificial Intelligence, *Engineering Applications of Artificial Intelligence* (2022).
- [21] Manish Thapa, Jose Espejo-Urbe and Evangelos Pournaras, Measuring Network Reliability and Repairability against Cascading Failures, *Journal of Intelligent Information Systems* (2019).
- [22] Harpreet Kaur and Raman Maini, Identification of Recurring Patterns of Code to Detect Structural Clones, in: *6th International Conference on Advanced Computing (IACC)*, 2016. doi:10.1109/IACC.2016.80.

¹⁷<https://www.opencti.io>

- [23] Ben Goertzel, Cassio Pennachin and Nil Geisweiller, *Engineering General Intelligence, Part 1: A Path to Advanced AGI via Embodied Learning and Cognitive Synergy*, Atlantis Press, 2014. ISBN 978-94-6239-026-3 978-94-6239-027-0.
- [24] Peter E. Kaloroumakis and Michael J. Smith, *Toward a Knowledge Graph of Cybersecurity Countermeasures*, Technical Report, The MITRE Corporation, 2021.
- [25] Stefan Kempter, *IT Process Maps – Incident Management*, 2007.
- [26] Dan Brickley and Ramanathan V. Guha, *RDF Schema*, W3C Recommendation, W3C, 2014.
- [27] Sean Bechhofer, Frank van Harmelen, Jim Hendler, Ian Horrocks, Deborah L. McGuinness, Peter F. Patel-Schneider and Lynn Andrea Stein, *Web Ontology Language (OWL)*, W3C Recommendation, W3C, 2004.
- [28] Holger Knublauch and Dimitris Kontokostas, *Shapes Constraint Language (SHACL)*, W3C Recommendation, W3C, 2017.
- [29] Dan Brickley and Libby Miller, *Friend of a Friend (FOAF) Vocabulary Specification*, 2004. <http://xmlns.com/foaf/spec/>.
- [30] CCITT, *ITU-T Rec. X.735 (09/92) Information Technology - Open Systems Interconnection - Systems Management: Log Control Function*, Recommendation, ITU, 1992.
- [31] CCITT, *ITU-T Rec. X.721 (02/92) Information Technology - Open Systems Interconnection - Structure of Management Information: Definition of Management Information*, Recommendation, ITU, 1992.
- [32] Silvio Peroni, *Graffoo: Graphical Framework for OWL Ontologies*, 2013.
- [33] J. Halpern and C. Pignataro, *Service Function Chaining (SFC) Architecture*, 2015.
- [34] R. Gerhards, *The Syslog Protocol*, 2009.
- [35] Mark Fedor, Martin Lee Schoffstall, James R. Davin and Dr. Jeff D. Case, *Simple Network Management Protocol (SNMP)*, 1990.
- [36] Stefan Kempter, *IT Process Maps – Problem Management*, 2007.
- [37] Jędrzej Potoniec, Dawid Wiśniewski, Agnieszka Ławrynowicz and C. Maria Keet, *Dataset of Ontology Competency Questions to SPARQL-OWL Queries Translations*, *Data in Brief* (2020).
- [38] Lionel Tailhardat, Yoan Chabot and Raphaël Troncy, *Designing NORIA: a Knowledge Graph-based Platform for Anomaly Detection and Incident Management in ICT Systems*, in: *4th International Workshop on Knowledge Graph Construction (KGCW)*, 2023.
- [39] Jieming Zhu, Shilin He, Jinyang Liu, Pinjia He, Qi Xie, Zibin Zheng and Michael R. Lyu, *Tools and Benchmarks for Automated Log Parsing*, in: *41st International Conference on Software Engineering: Software Engineering in Practice (ICSE-SEIP)*, 2019. doi:10.1109/ICSE-SEIP.2019.00021.
- [40] Shilin He, Pinjia He, Zhuangbin Chen, Tianyi Yang, Yuxin Su and Michael R. Lyu, *A Survey on Automated Log Analysis for Reliability Engineering*, *ACM Computing Surveys* (2021). doi:10.1145/3460345.
- [41] Aritran Piplai, Sudip Mittal, Anupam Joshi, Tim Finin, James Holt and Richard Zak, *Creating Cybersecurity Knowledge Graphs From Malware After Action Reports*, *IEEE Access* (2020). doi:10.1109/ACCESS.2020.3039234.
- [42] Ismail Harrando, Pasquale Lisena and Raphael Troncy, *Apples to Apples: A Systematic Evaluation of Topic Models*, in: *Recent Advances in Natural Language Processing (RANLP)*, 2021. doi:10.26615/978-954-452-072-4_055.
- [43] Lionel Tailhardat, Raphaël Troncy and Yoan Chabot, *Leveraging Knowledge Graphs For Classifying Incident Situations in ICT Systems*, in: *18th International Conference on Availability, Reliability and Security (ARES)*, 2023. doi:10.1145/3600160.3604991.
- [44] Aviad Elitzur, Rami Puzis and Polina Zilberman, *Attack Hypothesis Generation*, in: *European Intelligence and Security Informatics Conference (EISIC)*, 2019. doi:10.1109/EISIC49498.2019.9108886.
- [45] Tianxing Wu, Arijit Khan, Huan Gao and Cheng Li, *Efficiently Embedding Dynamic Knowledge Graphs*, *Knowledge-Based Systems* (2019).
- [46] Chengjin Xu, Mojtaba Nayyeri, Fouad Alkhoury, Hamed Shariat Yazdi and Jens Lehmann, *Temporal Knowledge Graph Embedding Model Based on Additive Time Series Decomposition*, in: *19th International Semantic Web Conference (ISWC)*, 2020.
- [47] Yan Li, Tingjian Ge and Cindy Chen, *Data Stream Event Prediction Based on Timing Knowledge and State Transitions*, *VDLB Endowment* (2020).
- [48] Wassim Sellil Atoui, *Toward Auto-configuration in Software Networks*, PhD thesis, Institut Polytechnique de Paris, 2020.
- [49] Peter Mell, Karen Scarfone and Sasha Romanosky, *Common Vulnerability Scoring System*, *IEEE Security Privacy* (2006).
- [50] Yassine Naghmouchi, Nancy Perrot, Nizar Kheir, Ali Ridha Mahjoub and Jean-Philippe Wary, *A New Risk Assessment Framework Using Graph Theory for Complex ICT Systems*, in: *8th ACM CCS International Workshop on Managing Insider Security Threats*, 2016. doi:10.1145/2995959.2995969.
- [51] A. Miles and S. Bechhofer, *SKOS Simple Knowledge Organization System Reference*, W3C Recommendation, W3C, 2009.
- [52] Bram Steenwinckel, *IBCNServices/Folio-Ontology*, 2019.
- [53] Giovanna Castellano, Anna M. Fanelli and Maria A. Torsello, *Web Usage Mining: Discovering Usage Patterns for Web Applications*, in: *Advanced Techniques in Web Intelligence-2: Web User Browsing Behaviour and Preference Analysis*, Springer, 2013.
- [54] Object Management Group (OMG), *Business Process Model and Notation (BPMN)*, Version 2.0, Technical Report, OMG, 2014.
- [55] Wil van der Aalst, *Process Mining: Discovery, Conformance and Enhancement of Business Processes*, Springer, 2011. ISBN 978-3-642-19344-6 978-3-642-19345-3.
- [56] A. Meroño-Peñuela and C. Martinez, *grlc: the git repository linked data API constructor*, 2021. <https://github.com/CLARIAH/grlc>.