Ph.D. thesis booklet

Qualitative models in resilience assurance

Imre Kocsis

Topic advisor: Dr. András Pataricza, DSc

Budapest University of Technology and Economics
Department of Measurement and Information Systems

Budapest, 2019

MŰEGYETEM 1782

Motivation and challenges

Critical systems have to comply with requirements imposed on their extra-functional properties; including dependability, performance, and resilience. The activities that contribute to assuring compliance with extra-functional requirements span the whole system lifecycle and are supported by a wide variety of modeling and analysis approaches. At the highest level of abstraction, where domain experts and engineers specify and interpret the "meaning" and impact of requirements, behavior and assurance mechanisms, qualitative models (QMs) are regularly used.

QMs are *discrete state space structural and behavioral models that carry direct interpretation in a given extra-functional domain* – in this dissertation, dependability, performance, and resilience.

QMs either complement *quantitative* models as a semantic cover model of behavior, or form part of them: for instance, as the logical structure of discrete quantitative models, as the state space of discrete *modes* of hybrid system models, or as an abstraction of continuous behavior. QMs are by nature parametric, and thus potentially reusable and portable, due to abstracting continuous system behavior into discrete state spaces through discretization.

Resilience

Resilience is a system aspect that, in contrast to dependability and performance, is a relatively new requirement. Resilience is the"*persistence of service delivery that can justifiably be trusted, when facing changes*" [3]. "Change" is to be understood here in a quite general sense: it includes every environmental change that the originally assumed design envelope does not cover, such as shifting work-, fault- and attack loads. A resilient system should either tolerate these changes or respond to them with graceful degradation.

Resilience can be characterized with the same core primitives as dependability and performance: fault and error states in terms of deviations from the originally intended operational parameters, and failure modes in terms of the trustworthiness of keeping to the desired runtime quality attributes. *Resilience state spaces* [4][5] capture the resilience behavior of a system through discrete operational and service states.

Figure 1 presents a resilience state space example with two example services. Service 1 suddenly crashes when load exceeds its capacity and needs a restart and a warmup phase to bring it back to operation. In contrast, service 2 protects itself against sudden workload spikes by beginning to queue requests; this gains enough time to scale out the service to more nodes. As a result, while the operational state of the system is temporarily degraded, the service state remains acceptable. Intuitively, service 2 is "more resilient".

While the exact definition of resilience is still somewhat disputed, it is clear that it is not equal to dependability in general, and fault-tolerant operation in particular. Instead, while the mechanisms of fault-tolerant operation *are* key to resilient operation, design for dependability addresses a specific design envelope and emphasizes the planned and deterministic avoidance of failures; resilience, in contrast, targets graceful degradation and low-impact transients when the original assumptions change.

From a practical point of view, resilience is *not* the capability to tolerate fault and error

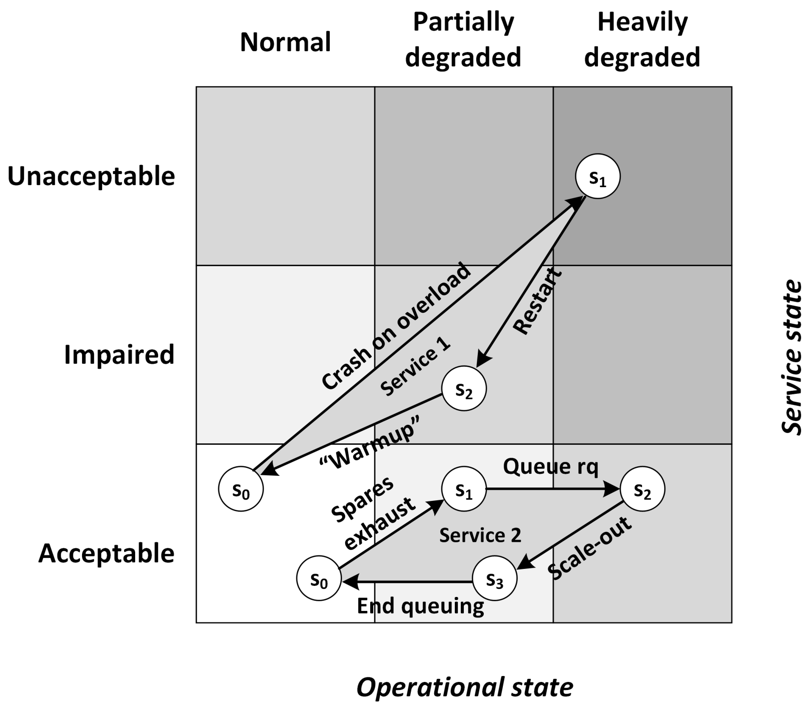

Figure 1: Qualitative modeling of resilience behavior with state spaces

modes that were originally entirely ruled out; instead, the capability to be able to imperfectly, but cope with changing rates and points of manifestation.

From the particular point of view of qualitative *modeling and analysis*, the three extra-functional domains of dependability, performance, and resilience pose similar challenges; thus, most results of the dissertation are applicable in all three domains.

Qualitative models and life cycle

Generally speaking, in the dependability, performance and resilience assurance of software-based systems, qualitative models can appear at all stages of the life cycle; either in an explicit way or as the structure of quantitative models.

During requirement definition and its related activities, QMs can describe notions of requirement compliance and the expected system behavior for various aspects, building on existing domain concepts and their ontologies. Using lightweight ontology processing methods, even non-semantically captured, but structured and syntactically standards-compliant requirement descriptions can be transformed into ontologies that can serve as a conceptual basis for QMs (see [Pat+17]). Models at this stage do not specify internal behavior – that hasn't been *designed* yet; and also may be *parametric* in the sense that the *landmark values*[1] are not set. An example is the workload capacity of a specific software system; while specific deployments determine its numeric value, the behavior in case of overloads is a requirement which has to be captured.

The design processes of critical systems involve qualitative and quantitative modeling and analysis of dependability, performance and resilience [6]. Qualitative EPA has a key role

[1]Numerical boundary values of qualitative discretization.

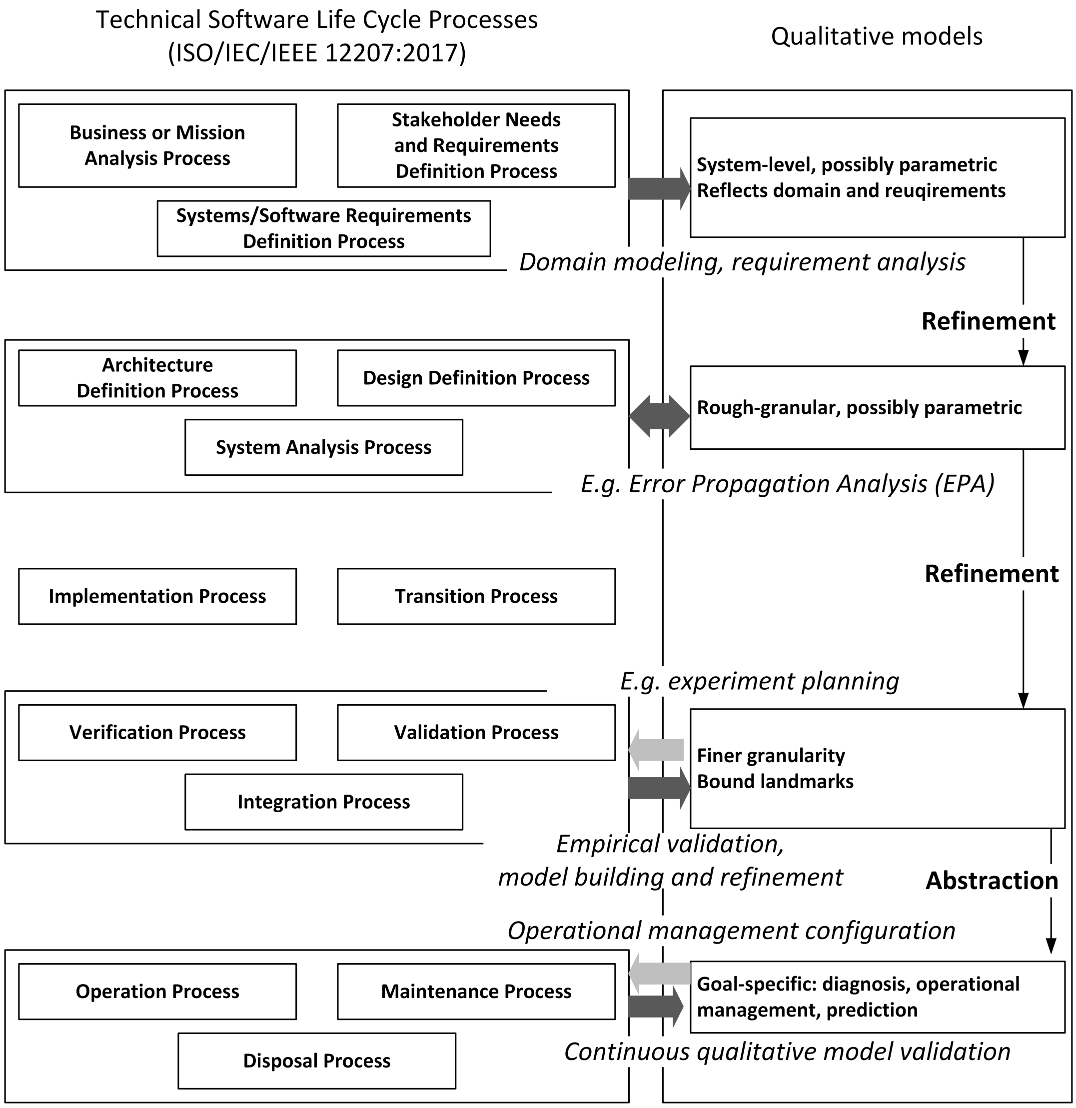

Figure 2: Typical roles of qualitative models in the software lifecycle

here. These processes consume as well as produce QMs; ideally, as explicit refinements of requirement-level models.

Verification, validation, and integration processes check the validity of the qualitative assumptions that design-related analysis made, reduce ambiguity and sharpen behavior, and establish the value of unbound parameters. In addition to incorporating evidence from the application of formal methods that focus on eliminating systematic design and implementation faults, *empirical* methods: testing and benchmarking are of paramount importance in these processes.

During the operational phase, QMs can form the high and typically human-interpretable level of operational management logic, from diagnosis to prediction. Sophisticated models – fine-grained discrete, hybrid or fully continuous ones – may be run in parallel to the qualitative ones; in such cases, it can be still good practice to also track state in QMs that are a) human-interpretable and b) check and supervise the sophisticated logic in easily understandable and verifiable terms.

Figure 2 demonstrates the general relationship between the potential use of QMs and the major life cycle phases, based on the ISO/IEC/IEEE 12207:2017 process standard [7] (a very generic process model which is compatible with a wide range of domain-specific ones). Conceptually, the depicted qualitative models form a refinement-abstraction hierarchy. The presence, style, and trustworthiness of QMs during the life cycle does differ significantly between different application domain and criticality classes of systems; leading to the main motivation of the thesis.

Resilience of soft critical systems and qualitative models

Service criticality fundamentally influences the strictness of the extra-functional requirements a system has to adhere to; and consequently, the magnitude of the effort that design and operation time activities can and should rationally devote towards assuring requirement compliance. *Safety critical* functions, where service "loss would present an unacceptable safety hazard during the transition to reduced capacity operations" ([8], p16), are modeled, verified and validated extensively according to standardized processes if they are present in a system – usually at considerable costs.

The quality, design sophistication and trustworthiness of the solutions intended to ensure extra-functional guarantees decrease together with criticality. At the other end of the spectrum – beyond mission- and business-critical systems – lie general-purpose computing systems, with rough-granular runtime monitoring and predominantly manual problem resolution. In this thesis, systems that have a definite attached criticality, but not an extremely high one, will be collectively referred to as *soft critical* in contraposition to *hard critical* ones. Classic criticality level models attach numeric values to their categories; in contrast, this distinction is qualitative and aims to denote systems where designed resilience is becoming increasingly necessary, but the knowledge- and effort-intensive tools and processes of design for dependability are still not viable options.

Today, extra-functional requirements, as well as the complexity of meeting them, is undergoing a significant shift in soft critical systems. As computing permeates all aspects of everyday life, notably in the form of Cyber-Physical Systems (CPSs), resilience is emerging as a requirement. At the same time, these systems rely more and more on component and service integration: software that is functionally capable, but essentially a black-box from the resilience characteristics point of view, or various cloud and network services with lax guarantees. On the other hand, the practice of regular dynamic reconfigurations of the ingested services and externally sourced runtime platforms – e.g., elastic scaling in Infrastructure as a Service (IaaS) – means that the ability for resilience assurance through automated *runtime adaptation* is also present.

The design, verification, and validation of hard critical systems rely on capturing the conceptual understanding of the system and its environment in QMs and reusing models throughout the development process. Similarly, methodically planned and assessed resilience assurance for soft critical systems *would* require QMs. But in this aspect, soft critical systems lag behind hard critical ones, even though portable and parameterizable, state-based supervisory models were proposed by the Autonomic Computing (AC) initiative as a general mechanism more than a decade ago [9].

Qualitative models in policy-based management

At least since AC, it has been widely accepted that deriving operational management actions in complex IT systems should be a) based on an explicit model that's continuously updated with ongoing system monitoring and b) intelligently automated to the largest sensible – and

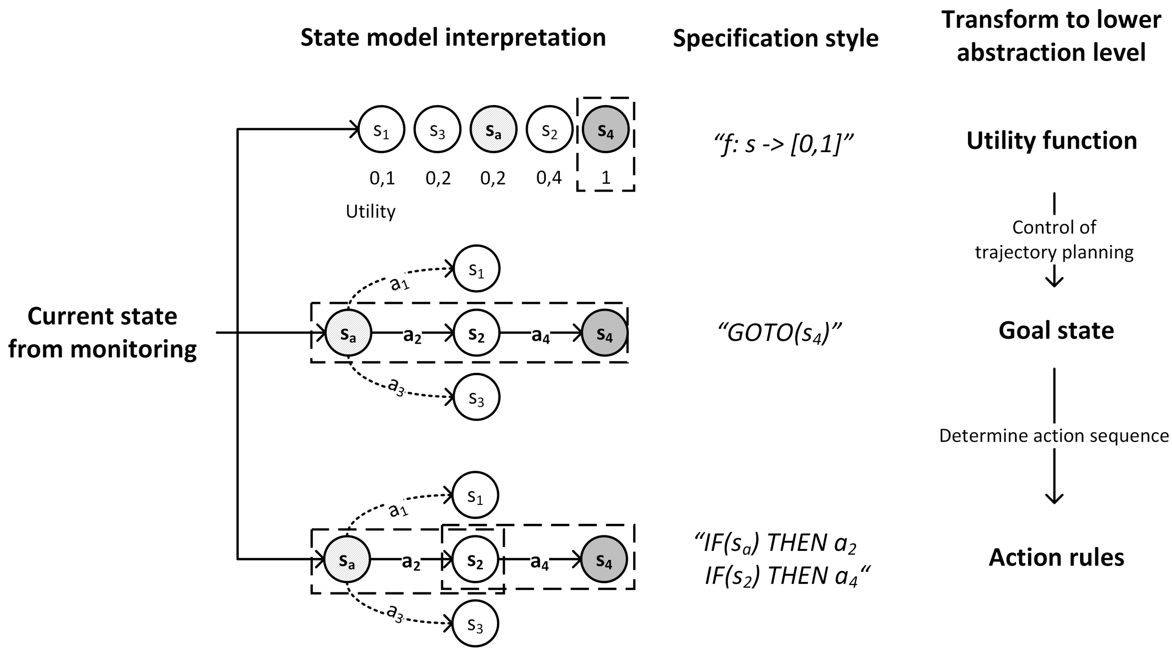

Figure 3: Autonomic computing policy styles and their relationships

safe – extent. For non-critical systems, the core incentive for this approach was increasing operational efficiency, especially for dynamically composed and highly reconfigurable systems. For soft critical ones, “autonomic” assurance mechanisms implemented as autonomic management functions have to complement structural defenses to meet extra-functional, as well as efficiency requirements at the same time.

AC puts an emphasis on supporting control through sophisticated reasoning. Its methods range from artificial intelligence [9] to control theory [10]. AC championed the systematic application of QMs in the management of general purpose IT systems, too, through emphasizing *policy-based management.*

Policy-based management assumes a system state model where the “current state” (or set of possible current states) is estimated by ongoing monitoring, and over this state model, a hierarchy of management policy declaration styles is defined (see, e.g., [9]). While *action policies* are simple event-condition-action rule sets, *goal policies* declare only “the state to reach”. At the most abstract level, *utility function policies* only express the utility of each state. For goal policies, it falls on the management framework to determine the trajectory to take; and for utility function policies, also the state the system should transition to (see Figure 3). For such policies to be human-interpretable and human-verifiable, the state space not only has to be discrete and finite but also qualitative in the semantic sense.

AC-style policies constitute a fundamental tool of the resilience assurance of contemporary soft-critical systems, which on the one hand can, and on the other hand have to rely on dynamic reconfigurations at the functional composition as well as the (typically cloud) deployment level. However, to systematically design and validate policies for a system, the *protected* system and the risks to be mitigated have to be also known at least at the QM level. Currently, this is extremely challenging. The creation of QMs during application development is not a typical goal in the soft-critial domain; and externally sourced components and runtime platforms largely come without “resilience specifications” that meaningfully enable resilience design.

Challenges and goals

The thesis addresses the following specific challenges and research goals.

Challenge 1: externally sourced components and services do not provide reusable and composable models of their extra-functional behavior; descriptive statistics as Mean Time Between Failures (MTBF), if at all available, are of insufficient resolution for planning resilience mechanisms. Observation-driven, *data analysis* based, typically quantitative empirical behavior characterization plays only a limited role in the practice, due to:

- insufficient model reusability – even across dynamic reconfigurations and redeployments of the same system;
- insufficient observational coverage of the qualitative system states, and
- the limited ability of general purpose data analysis methods to faithfully and effectively characterize rare events, as failures and unanticipated changes.

Goal 1: to provide practical, repeatable data analysis methods for extracting reusable and portable QMs from observations, in order to support the model-based resilience assurance of soft critical systems throughout the lifecycle.

Challenge 2: the EPA methods of the hard critical domains could be applicable for assessing and planning dependability, performance, and resilience in soft critical systems; however, they cannot efficiently cope with the high levels of QM nondeterminism involved in the partially speculative modeling of soft critical systems.

Goal 2: to extend modern EPA towards gaining the capability to effectively explore and assess sets of competing error propagation hypotheses that arise due to partial model nondeterminism.

Challenge 3: runtime diagnosis and action planning based on the runtime monitoring of qualitative states and transitions is a key application of QMs during the operational phase of the system lifecycle. To this end, an appropriate subset of a large number of potentially observable system metrics has to be chosen that is capable to determine robustly, and in the optimal case also to predict the runtime qualitative state. Additionally, modern soft critical systems pose the specific challenges that

- their increasingly typical runtime environment, IaaS cloud platforms can exhibit rare, high-impact, directly unobservable quality disturbances;
- due to the necessary use of observation- and expert knowledge based covering models, and to rectify the imperfections of automatic metric selections, methodical expert assessment of monitoring configurations is necessary, which is seriously hindered by the wide range of possible platform-level quality disturbances.

Goal 3: to provide practical monitoring metric set selection and assessment methodologies for qualitative monitoring in soft critical systems, taking into account the specific properties of the domain.

Application domains

The contributions of the thesis are dominantly results and generalizations of research efforts that targeted specific types of modern soft-critical systems. These can be categorized the following way, as the main target application domains of the dissertation with validation results.

- Cloud, and in particular, IaaS platforms themselves, the qualitative extra-functional characterization of which is usually a prerequisite to designing service level resilience.

- QoS-sensitive, engineered systems that have to undergo "cloudification". These are typical, e.g., in the financial and the telco sector[2], moreover, are already created with extrafunctional guarantees, provided that configuration, environmental and workload prerequisites are met. The key challenge, in this case, is that the classic "sizing guides" that numerically specify extra-functional properties for various dedicated hardware-software configurations simply cannot be created for the new environment, due to the platform behavior variability and runtime scaling that virtualization brings with itself. Additionally, cloud platforms can introduce new external fault modes, with an unknown associated impact on resilience behavior.

- Existing cloud-native components and applications that are sometimes developed and shipped even without regimented extra-functional integration testing. The ad-hoc configured elastic scaling and virtual machine (VM) restart strategies of these solutions have to be replaced with methodical resilience assurance, when the service becomes "critical enough" due to changes in its usage.

Newly created, soft critical CPS applications which heavily rely on the integration of external (unreliable) services, utilize cloud platforms and interact with the physical world (in a non-safety critical manner) are also expected to be an immediate application area of the results of the dissertation.

Last but not least, the intended application areas of the EPA contributions are much broader. They are relevant to the (usually prescribed) EPA activities in the development of hard critical systems – and, as a matter of fact, mostly targeted classic hard-critical challenges originally. On the other hand, in hard critical CPSs, as various autonomous vehicles and certain smart city functions, there is a very powerful tension between fulfilling safety, reliability and availability requirements in a trustworthy manner and the innovation potential of using essentially "untrustworthy", but highly intelligent remote services and cloud platforms. In [Koc18a], I propose error propagation space exploration from the second contribution group as a tool that supports effective pessimistic reasoning with all the ways these services can fail, and provide a heavily simplified smart railroad crossing demonstrational example.

[2] In telecommunications, the term Network Function Virtualization (NFV) covers the activity of transforming dedicated systems to Virtual Network Functions (VNFs), that can be deployed onto cloud platforms. NFV has been an ongoing strategic push in the whole sector for years.

Contribution 1: qualitative model extraction from empirical data

The first contribution addresses the fundamental mismatch between the knowledge required for compositional reasoning in the resilience assurance of soft critical systems, and the typical lack thereof.

Twin models in the resilience assurance of soft-critical systems

Knowledge reuse in the form of QMs across the lifecycle phases and instances of soft-critical systems is currenlty heavily underutilized. Borrowing the *twin model* concept from the Digital Twin [14] paradigm, I propose the systematic use, management and maintenance of QMs throughout the lifecycle. I identify the series of twin model elements that should ideally be introduced at each stage of the development process of soft critical systems (see Figure 4) and argue that the special significance of twin models in the specific problem domain is connected to more and more (re)design activities shifting to the operational phase.

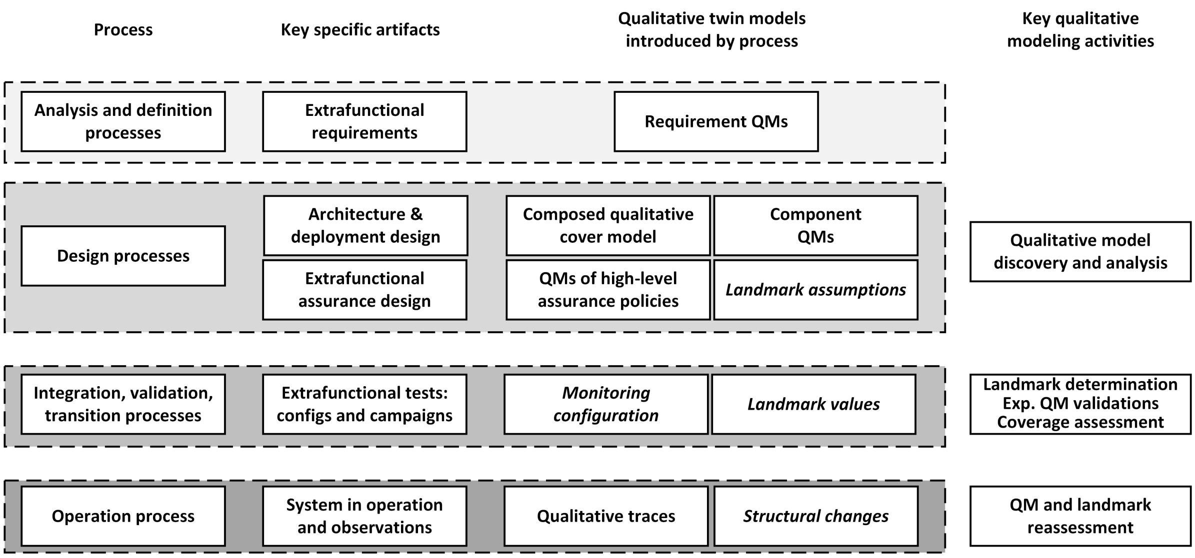

Figure 4: Process-wise introduction of twin model elements for QM-based resilience assurance. (Italics: artifacts directly supporting QMs)

In addition to the idealized development process pattern, based on project experiences, I describe the empirical properties of *post-design resilience increase* and *"cloudification"* projects and their fit to the process model.

Qualitative model extraction with Exploratory Data Analysis

Highly visual, interactive Exploratory Data Analysis (EDA) [11][12] strongly outperforms purely algorithmic data analysis in extracting QMs of dependability, performance and resilience behavior of components and moderate-size systems. While EDA as a "school of statistics" has rich traditions and application history, its use in this context is novel.

The proposed methodology relies on the monitoring of platform, application and service metrics, gathered either during operation or targeted experiments as performance and dependability benchmarks [13]. The interactive, visual exploration of these observations reveals visual artifacts that are used to create and refine QM fragments (see Table 1). The search for visual artifacts and their interpretation during incrementally building a QM can follow many workflow patterns; the thesis specifically elaborates a backward inference approach which, following the logic of sequential fault diagnosis, constructs a QM that phenomenologically "explains" service-level qualitative states and state changes by a hierarchy of lower level ones.

Data feature	Key plot types	Visual artifacts
continuous observations	2d/3d scatterplots parallel coordinates scatterplot matrix	clusters function images outliers
continuous summaries	histograms (1... n dim) heatmaps parallel boxplot	clusters superposed distrib. outliers
time series	line graphs summary-series, e.g., boxplots	trends/"seasonalities" lagged correlations local outliers
categorical summaries	barcharts mosaic plot variants	existing and not existing combinations

Table 1: Visual artifacts used for constructing and refining qualitative states and state transitions

Process embedding: ASUM-DM

EDA is an exploratory activity in general purpose data analysis; its counterpart, Confirmatory Data Analysis (CDA) serves to confirm the validity of the hypotheses that EDA delivers. In general purpose data analysis, in particular, data mining, practice-proven process templates and mature tooling support are available for planning and conducting the data analysis activity. I have proposed the specialization of the well-known ASUM-DM [15] process template for the qualitative EDA of dependability, performance and resilience experiments, through the particular application of exploratory analysis on fault injection (FI) experiments (Figure 5).

Specific contributions

C1.1: Qualitative twin models in the resilience assurance of soft critical systems. I have proposed the application of the twin model concept for the handling of qualitative models in the resilience assurance of soft critical systems. Qualitative twin models enable the systematic reuse of resilience-related knowledge about the system, its deployment targets and its runtime assurance policies during design; support iterative refinements through runtime observations; and facilitate shifting resilience assurance (re)design activities to the system operation life cycle phase.

Specific results: A template for the points-of-creation of possible qualitative twin models

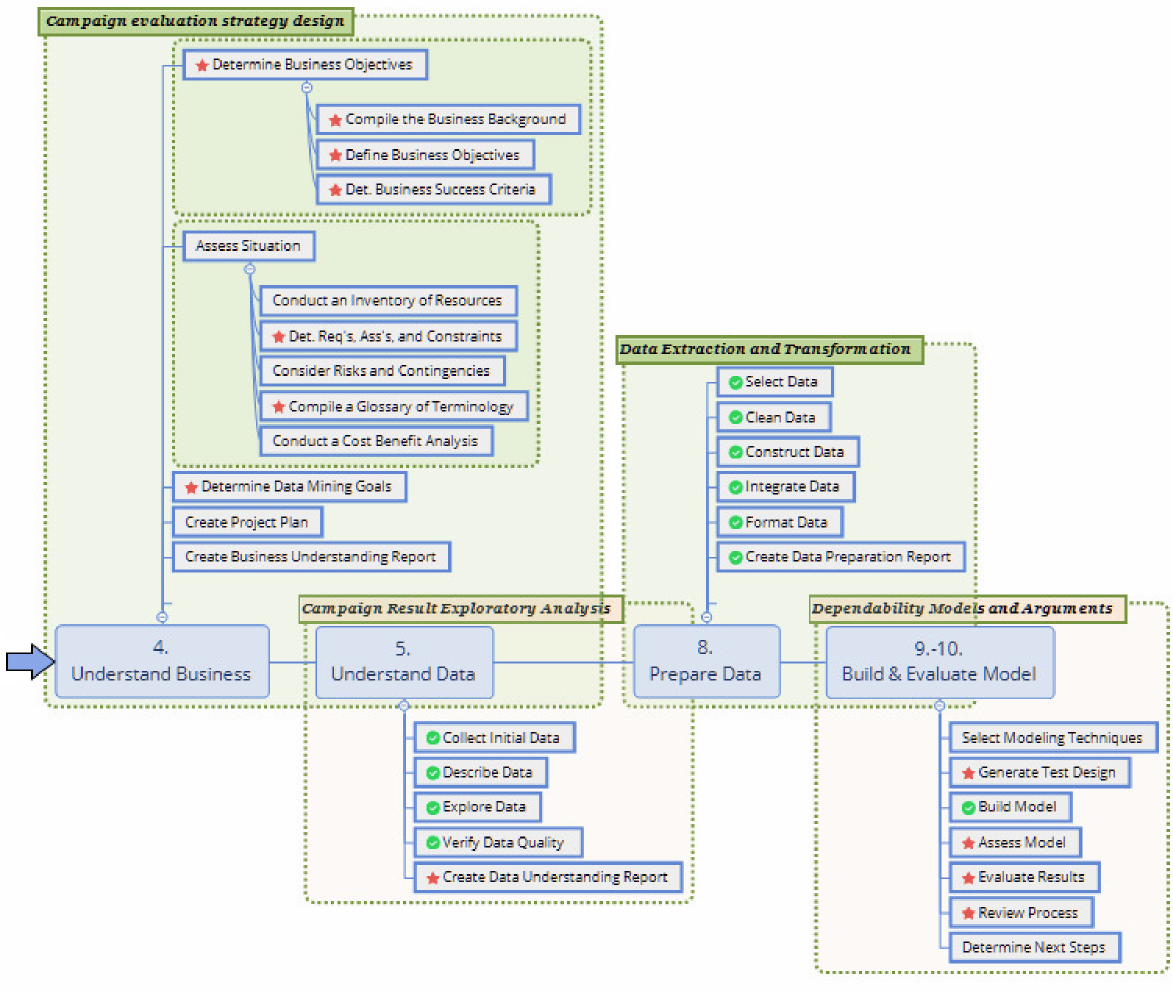

Figure 5: Key ASUM-DM activities in FI analysis. Red star: specialized step; green mark: directly reusable guidance and tooling

and associated artifacts in the life cycle; application characteristics for the "cloudification" of soft-critical engineered systems and the resilience assurance of cloud-native, non-critical ones; as a specific example, supervisory environment and configuration reuse (with the abstractions embedded in the configurations) in dependability benchmarking.

Supporting publications: [KSP16], [Koc+08b]

C1.2: Qualitative model extraction from empirical data. I have proposed a methodology to discover qualitative resilience-related models in measurement data, using Exploratory Data Analysis. The key innovation is the harmonization of fundamentally qualitative dependability processes – such as sequential fault diagnosis – and interactive, visual Exploratory Data Analysis steps in data analysis workflows.

Specific results: a mapping between visual artifacts and potential qualitative model features; a sequential fault diagnosis-inspired, iterative refinement based workflow template for qualitative model discovery through iterative, interactive exploratory data analysis with specific examples; the identification of experiment quality control considerations specific to qualitative EDA.

Supporting publications: [KSP16], [Pat+13]. [Koc18b] has been submitted to and is under consideration at the AARMS Journal

C1.3: Process embedding and specialization. I have specialized the ASUM-DM workflow pattern for the application of qualitative EDA in fault injection experiments, which is a representative case of a large number of resilience, dependability, and performance data analysis tasks. The embedding enables the reuse of general-purpose methodological as well as direct tooling support from general-purpose data analysis.

Specific results: identification of ASUM-DM tasks that have to be specialized for the domain and specific specializations; identification of tasks with potential for direct technical and process tooling reuse.

Supporting publications: [Cer+18]

Highlighted applications

EDA of system and service measurement data is an ongoing focus of research at the research group where I performed my research. The methods of the current contribution have been applied for the performance and performability analysis of consortial blockchain networks [Koc+18] as part of an ongoing series of research efforts.

Elements of the contributions are being incorporated into the application-centric education of EDA [Koc14]. The visual exploratory techniques formed a key part of the 2014 SERENE conference tutorial "Measurement-Driven Resilience Design of Cloud-Based Cyber-Physical Systems" and the FAS* 2016 conference tutorial "Model-based Cloudification of Critical Applications", both held by the author.

In industrial R&D, EDA and the proposed approaches were used to analyze the performance and dependability of the virtualized desktop infrastructure of a major multinational financial company. The methods formed part of a cloud performance and benchmarking environment development project for a major telecommunications equipment provider and applied on cell tower performance data for a telco operator. (The specifics of these projects are under NDA.)

As part of scholarships funded by Ericsson Hungary, the qualitative EDA methods were validated in a network function virtualization context and presented at the local 2016 Ericsson University Day. Initial concept exploratory and validation work for integrating "Algorithm as a Service" cloud offerings – specifically, the IBM Watson Analytics SaaS product – into qualitative EDA has been performed and presented at the 2016 International IBM Cloud Academy Conference [Pat+16].

Contribution 2: qualitative model analysis for resilience assessment

In [16], Prof. Pataricza formulated a mathematically precise, abstraction-based theory of Error Propagation Analysis (EPA), which forms the foundation of the second group of contributions.

EPA extensions

I extend the fundamental EPA approach and some later extensions of it (described in the thesis, but not proposed as individual contributions) with a standard XML language based problem representation scheme, a systematic approach towards error propagation solution "sharpening" via qualitative refinement, and a simulation-based application of EPA for software FMEA at early system design phases. `XCSP3` is a standard language to describe Constraint Satisfaction Problems (CSP). The significance of my easily traceable engineering model mapping scheme to it is that it enables the lightweight assembly of EPA problems from a wide range of engineering modeling languages that have natural "component" concepts, while remaining human-interpretable and portable across a wide range of solvers. Error propagation solution sharpening through refinement is fully worked out for the special EPA application of requirement change management. Early software FMEA is based on an executable subset of UML that can be used to specify behavior already during early design; and it follows the same idea that the presented EPA framework follows, i.e., composing all aspects into a single, analyzable description. Only in this case, the executable specification is transformed into an *executable error propagation specification*, which can be simulated using the same existing – not EPA specific – tooling.

The three EPA extensions were originally elaborated under the aegis of the CErtification of CRItical Systems (CECRIS) EU FP7 project. CECRIS targeted the development, V&V and certification processes of hard critical systems; I was involved in activities that targeted enabling behavior-driven EPA in the early design phase and the semantic management of requirements. However, my specific contributions are also directly relevant to the resilience inference of soft-critical systems, the overarching theme of this dissertation.

Application of FCA for hypothesis set exploration

Under nondeterministic assumptions that typically stem from limited knowledge, e.g., regarding the error propagation behavior of individual system components, EPA problems have multiple solutions that are all consistent with the problem formulation: sets of *error propagation hypotheses*. Easy human interpretation and exploration of these is an open challenge, for which I propose the use of Formal Concept Analysis (FCA). FCA takes a set of objects that are equipped with attributes and creates the maximal subsets of objects sharing the most attributes. These so-called formal concepts have a lattice superconcept-subconcept structure, that can be visualized and interactively refined/abstracted using Hasse diagrams. As a simple example, see Figure 6. Each node in such a diagram is a formal concept; regarding interpretation, nodes collect attributes "upwards" and objects "downwards". For instance, on the diagram `Subtle` + `Data error` and `Sign bug` + `Sign bug + overload` + `Too fast + sign bug` form the formal concept belonging to the node labeled `Subtle`.

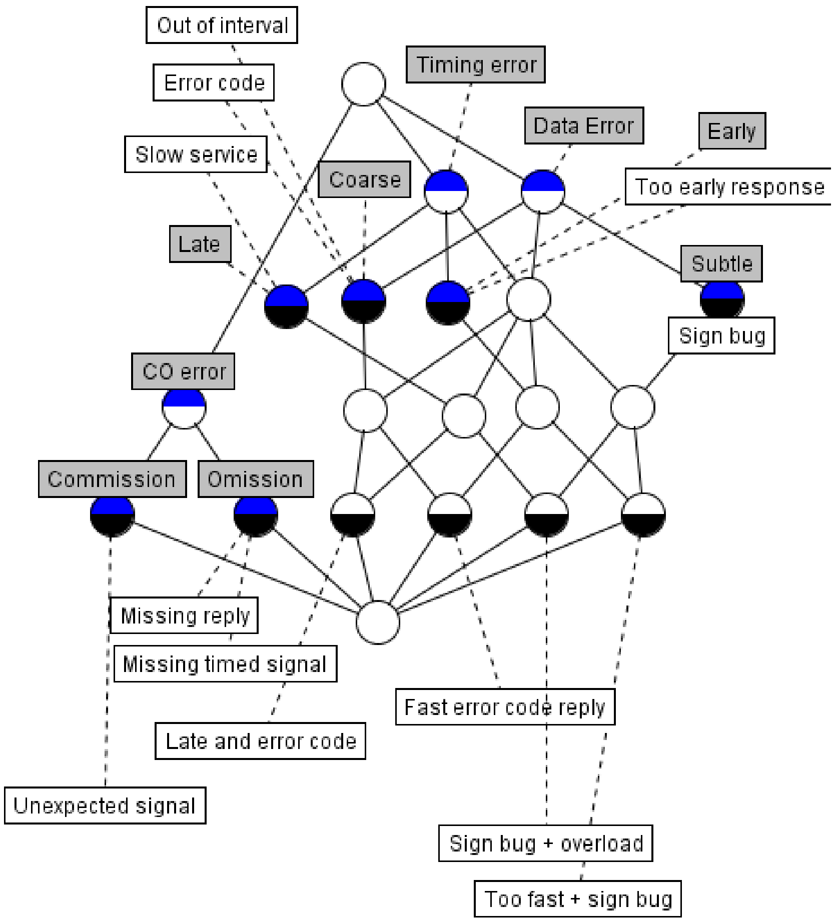

Figure 6: Line diagram for an example lattice of failure modes and attributes

When applied for fault activations as objects and error/failure manifestations as properties, FCA [17] and its visualization techniques enable the human assessment of error propagation hypothesis sets – through the series of effects that gradually partition the set of fault activations into equivalence partitions. Many aspects of FCA have direct diagnostic interpretations; these parallels are elaborated by the thesis.

Error propagation space analysis

True error propagation space exploration is expected to be a process that step by step assesses and modifies full error propagation hypothesis sets. Using a simple, relational algebra based axiomatic framework, I formalized the central concept of hypothesis set assessment: error propagation covers, and defined the key refinement and filtering operators of the process.

Specific contributions

C2.1: EPA extensions supporting reasoning for resilience. I extend the fundamental EPA approach of Prof. Pataricza and its existing application extensions with a portable problem representation scheme, a systematic approach towards error propagation cover reduction via qualitative refinement, and simulation-based application of EPA at early system design phases.

Specific results: A traceable, directly solvable, but solver-independent EPA problem representation scheme in `XCSP3`; patterns of qualitative EPA refinement, with application for requirement change impact estimation in safety critical design; the definition of an EPA simulation approach through executing composed "error propagation programs" in Alf.

Supporting publications: [Bon+17], [Pat+17], [Koc+08a]

C2.2: Structured interpretation of EPA hypotheses. I introduce FCA as a tool for deriving and visualizing the internal abstraction hierarchy inherent to solution hypothesis sets in EPA. I also show that FCA of hypothesis sets has direct interpretations with constructive applications in qualitative fault isolation.

Specific results: diagnostic interpretation guide for FCA additive line diagrams, diagnostic application of FCA algorithms, application for a business process EPA example.

Supporting publications: submitted to PPEECS: [KP18], [Pat+17]

C2.3: Error propagation space exploration. I have defined the key operators needed for error propagation space exploration over error propagation hypothesis sets – importantly, a qualitative variant of error propagation covers and conditioning-like filter operators for stepwise sharpening fault assumptions as well as failure and error requirements. This way, error propagation under uncertainty can be gradually constrained in to the required specificity by strengthening propagation assumptions, ruling out cases and taking decisions on the deployed error propagation mechanisms.

Supporting publications: [Koc18a]

Applications

I applied EPA in the DESEREC FP6 project in the context of analyzing error propagation in self-reconfiguring systems [Koc+08a] and in the CECRIS project [Bon+17].

The original application for business processes [Urb+14] was followed by an application for data quality error propagation in data analysis processes [18] (faculty student conference submission co-advised by the author).

The experiences of repeated applications with various prototype implementations lead to the consolidated, practical set of engineering application patterns described in the engineering modeling chapter and the first sub-contribution. The proof of concept implementation underlying [Koc18a] and [KP18] uses the documented techniques (except for fUML-based error propagation simulation); a dedicated tool paper will be justified when the integration of a key set of modeling languages (BPMN 2.0 and ArchiMate) has been finished. Integration of MDD-based constraint solving to directly implement EPA and error propagation space exploration with hypothesis sets is an ongoing, but unfinished, experiment.

Contribution 3: qualitative model based monitoring

The third group of contributions starts with the assumption that dependability, performance or resilience QMs of a soft critical system have been created during the design activities and high-level operational management policies rely on these for qualitative runtime diagnosis. The remaining challenge at this stage is ensuring that platform, application and service monitoring is appropriately configured and mapped into the qualitative state spaces. As specified in the description of the challenge, the specific contributions of the group target three specific subproblems.

mRMR based feature selection for qualitative monitoring

The first element of the contribution group is an experiment-based qualitative feature selection scheme for soft critical systems with a single, quantized output Quality of Service (QoS) variable. The methodology relies on creating roughly balanced sets of observations for each qualitative target state and performing wide coverage monitoring during the experiments. Observations are labeled with qualitative states and the minimum Redundancy, Maximum Relevance (mRMR) [19] algorithm is used for selecting the variables that best distinguish the qualitative states with minimum (information-theoretic) redundancy. The underlying intuition is that mRMR can handle highly nonlinear changes in the statistical dependence between input variables – in this case, the changes are associated with qualitative state transitions. The regression power of mRMR-based feature sets is evaluated and a qualitative state forecasting scheme is demonstrated and evaluated.

Mystery shoppers for IaaS

The second element of the contribution group introduces so-called mystery shoppers for soft-critical services deployed on cloud platforms. In Infrastructure as a Service (IaaS) cloud services, rare and short platform service disturbances, such as performance interferences in available CPU capacities can lead to major and long disruptions in the deployed applications. As cloud platforms do not necessarily provide support for monitoring these disturbances, or provide guarantees for their absence, it falls on the platform user to assess their prevalence before deployment and to monitor their occurrence thereafter. To this end, I proposed using modified benchmark applications that in resource usage mimic the serving of workloads, have configurable resource requirements and deliver their benchmark results as a fine-granular time series of measurements. This way, they can be used for platform assessment by "standing in" for real applications and workloads; and they can provide a low-impact monitoring method of spare platform capabilities, when deployed alongside real applications.

I provide an initial design methodology for such mystery shoppers in the IaaS context; the concept was validated for the open source `BigBlueButton` teleconference server, where, after discovering the extreme CPU interference sensitivity of the application, I successfully created a simple mystery shopper that can detect VM-internal synthetic interferences the same way as noisy neighbor related ones (see Figure 7).

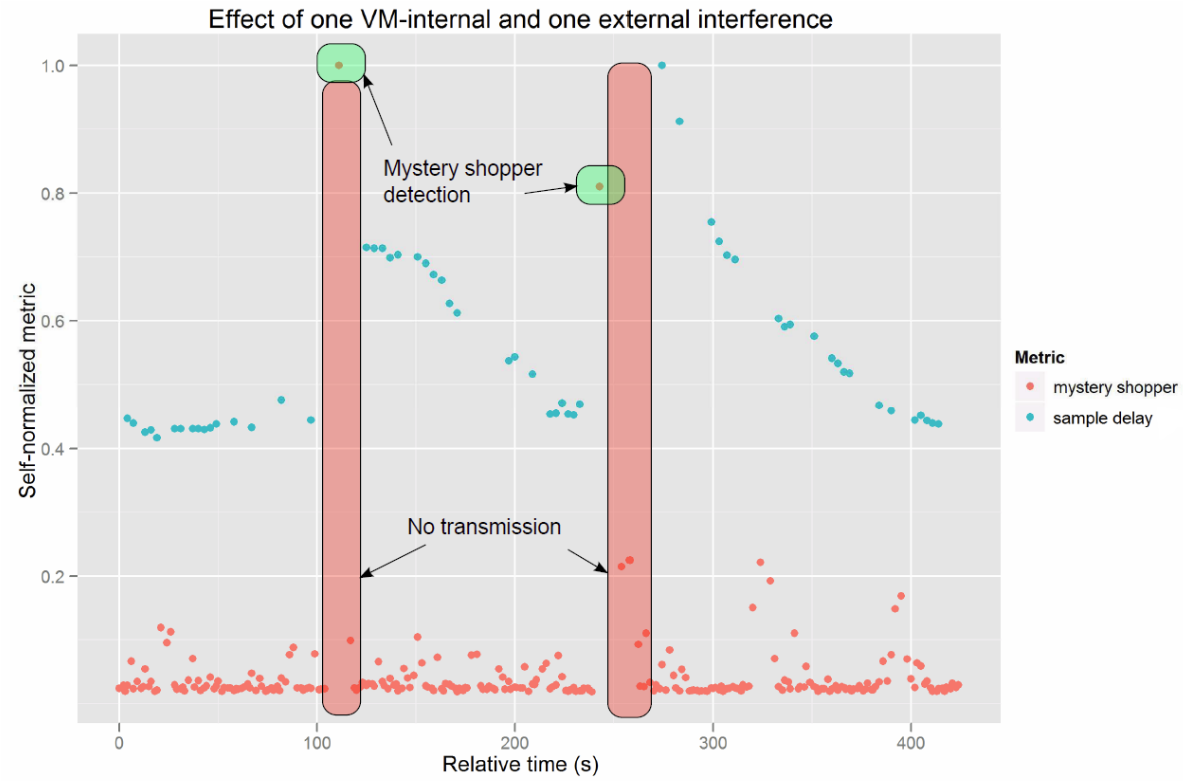

Figure 7: One VM-internal and one VM-external CPU interference injection: effect on conference call delay and mystery shopper metric. For visual analysis, voice transmission delay (blue) and mystery shopper metric (red) were brought to a common scale. The red bars denote a full outage in voice transmission.

Taxonomy-based monitoring evaluation

The third element of the contribution group addresses two problems that historically arose from the first two elements. Automated feature selection can be very helpful due to the large number of monitoring variable candidates; however, domain experts will want to review the set of variables selected, e.g., to assure that no overly sensitive variable is present, or that every necessary runtime aspect is covered.

To support methodical assessment of monitoring variable sets, I created a multi-aspect IaaS Key Quality Indicator (KQI) taxonomy (see Figure 8). The structure of the hierarchical taxonomy expresses multiple aspects that are *combined* necessary to define specific KQIs: resource type, resource service type and quality type as primary attributes, and temporal scope, population scope and statistical aspect of quality (e.g., central tendency) as secondary ones. Based on this taxonomy, I described how an expert can methodically select the important platform KQIs for a deployed component using the well-known Goal-Question-Metric (GQM) [20] reasoning approach. Then, the appropriateness of the selected metrics to reflect the KQIs can be individually assessed. Under expert uncertainty, as well as for applications-platform combinations where many sensitivities exist, the KQIs can be rank-ordered using well-known protocols, e.g., the Analytic Hierarchy Process (AHP) [21].

Specific contributions

C3.1: Feature selection and qualitative monitoring. For qualitative monitoring and qualitative service state prediction based directly on continuous variables without direct load and QoS observation, I proposed the use of the mRMR feature selection algorithm.

Specific results: using a TPC-W testbed and load profiles that transition between qualitative magnitudes, an experimental evaluation of current-QoS regression with mRMR in comparison to forwards selection, in general, and in specifics for the different qualitative states; experimental evaluation of qualitative forecasting with classifiers built on mRMR-selected feature sets.

Supporting publications: [Pal+10a], [Pal+10b], submitted to PPEECS: [KP18], submitted to AARMS: [Koc18b]

C3.2: Benchmark-like monitoring of services with mystery shoppers. I elaborated a data-driven methodology to discover, characterize and monitor high-impact, relatively infrequent platform service quality failures of IaaS services. Long-running benchmarks with configurable service capacity utilization serve to discover the effects themselves and accompany the workload during operation for runtime detection. This concept is called "mystery shopper" and provides a generic, low impact way to monitor disturbances for which the workload to be protected is sensitive.

Specific results: Definition of the mystery shopper concept; methodological guidance on creating mystery shoppers; experimental validation for IaaS CPU interferences.

Supporting publications: [Koc+13]

C3.3: Methodical assessment of monitoring variable sets. Based on existing taxonomies and their shortcomings, I have designed a multi-aspect IaaS KQI taxonomy. I proposed using standard methods of structured metric selection and preference-based rank ordering for methodically performing the task of qualitative monitoring variable set assessment for applications using IaaS.

Specific results: IaaS KQI taxonomy; assessment process template.

Supporting publications: [KSP16]

Applications

The contributions described in this chapter were applied in the Metrics working group of the DESEREC FP6 project, as well as in multiple private and public cloud performance assessment R&D projects. The results also formed part of industrial education activities on cloud performance.

The contributions were directly applied in the IBM Faculty Award "Measurement-based Qualitative Performance Modeling of Blockchain-based Distributed Ledgers" project (awarded to Prof. András Pataricza).

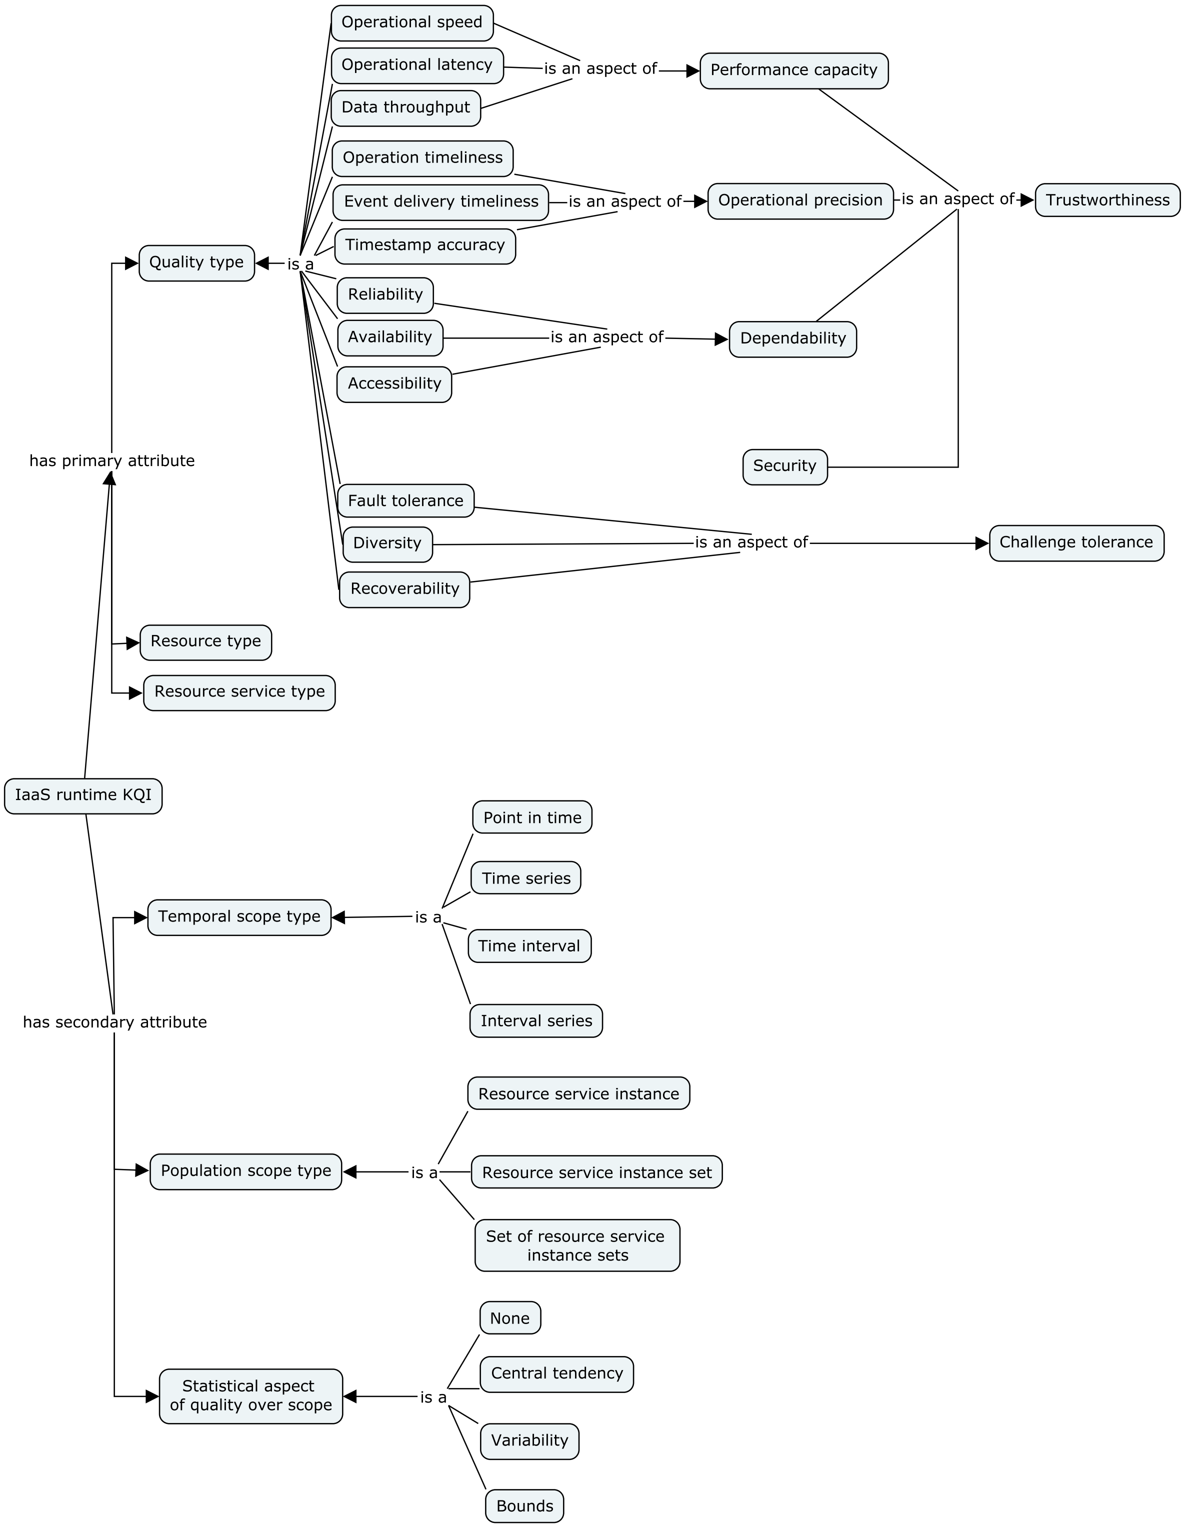

Figure 8: The proposed IaaS KQI taxonomy.

Future work

The contributions of the thesis lead to a number of interesting further challenges, that the author hopes to be able to address in further research.

Contribution 1 left open the so-called ”cognostics” (computer-aided diagnostics) aspect of EDA; i.e., the computer should rank order visualizations for the analyst by their potential for ”being interesting”. As cognostics in general began to apply even AI and ML techniques lately, there's a significant body of already existing ideas that should be evaluated for applicability and to assess whether fundamentally new ones are needed in the ”qualitative cognostics” context. Another, low complexity, but potentially large impact possibility is the integration of ”Algorithm as a Service” solutions.

On the process side, including the qualitative twin model concept, further validation is certainly necessary.

Contribution 2 states, but does not solve, the Multi-valued Decision Diagram (MDD) based operationalization of the operators defined for error propagation space exploration. In addition to the engineering challenges, the appropriate heuristics to use (e.g., whether variable ordering should simply reflect component network topology or other orderings are more appropriate) and mapping out the performance envelope of the approach are key open problems. At the same time, AND-OR graphs may prove to be an even choice than MDDs. Additionally, utilizing FCA to semantically connect requirements and domain knowledge with analysis (and, as a matter of fact, observations) is an intriguing possibility.

Contribution 3 proposes the use of specific feature selection and regression/classification methods; this avenue of research I do not plan to further pursue – in contrast to incorporating existing knowledge on causality [22] into feature selection in the target domains of the thesis.

List of publications

The thesis-supporting publications of the author map to the stated contributions as presented in the table below. The list of publications also includes the further publication activity of the author in a clearly separated way.

A number of these papers is also relevant to the thesis and its arguments, but due to redundancy in the message, a minor role of the author, or the connection to the thesis being only tangential these papers are not presented as foundations of the contributions.

Publication	1.1	1.2	1.3	2.1	2.2	2.3	3.1	3.2	3.3
[Bon+17]				X					
[Pat+17]				X	X				
[KSP16]	X	X							X
[Koc18b]*							X		
[KP18]*		X			X		X		
[Koc+13]								X	
[Cer+18]			X						
[Koc18a]						X			
[Pat+13]		X							
[Pal+10a]							X		
[Pal+10b]							X		
[Koc+08a]				X					
[Koc+08b]	X								

Mapping of publications to contributions. *: not accepted yet

Book chapters

[Bon+17] Valentina Bonfiglio, Francesco Brancati, Francesco Rossi, Andrea Bondavalli, Leonardo Montecchi, András Pataricza, Imre Kocsis, and Vince Molnár. “Composable Framework Support for Software-FMEA through Model Execution”. In: *Certifications of Critical Systems - The CECRIS Experience*. Ed. by Andrea Bondavalli and Brancati, Francesco. River Publishers Series in Information Science and Technology. River Publishers, 2017. URL: `http://www.riverpublishers.com/book_details.php?book_id=450`.

[Pat+17] András Pataricza, Imre Kocsis, Francesco Brancati, Lorenzo Vinerbi, and Andrea Bondavalli. “Lightweight Formal Analysis of Requirements”. In: *Certifications of Critical Systems - The CECRIS Experience*. Ed. by Andrea Bondavalli and Brancati, Francesco. River Publishers Series in Information Science and Technology. River Publishers, 2017. URL: `http://www.riverpublishers.com/book_details.php?book_id=450`.

[KSP16] Imre Kocsis, Ágnes Salánki, and András Pataricza. "Measurement-Based Identification of Infrastructures for TCPS". In: *Trustworthy Cyber-Physical Systems Engineering*. Ed. by Alexander Romanovsky and Fuyuki Ishikawa. Chapman & Hall/CRC Computer and Information Science Series. Chapman and Hall/CRC, Aug. 21, 2016.

Journal articles

[Koc18b] Imre Kocsis. "Measurement-Based Performance Modeling". In: *Submitted for consideration to Academic and Applied Research in Military Science (Journal on)* (2018).

[KP18] Imre Kocsis and András Pataricza. "Semantic Data Management in IT Service Performance Assurance". In: *Submitted for consideration to Periodica Polytechnica Electrical Engineering and Computer Science* (2018).

[Koc+13] Imre Kocsis, András Pataricza, Zoltán Micskei, András Kövi, and Zsolt Kocsis. "Analytics of Resource Transients in Cloud-Based Applications". In: *International Journal of Cloud Computing* 2.2-3 (Jan. 1, 2013), pp. 191–212. DOI: `10.1504/IJCC.2013.055267`.

Conference and workshop papers

[Cer+18] Frederico Cerveira, Imre Kocsis, Raul Barbosa, Henrique Madeira, and András Pataricza. "Exploratory Data Analysis of Fault Injection Campaigns". In: *18th International Conference on Software Quality, Reliability, and Security*. Lisbon, Portugal: IEEE, July 16–20, 2018. DOI: `10.1109/QRS.2018.00033`.

[Koc18a] I. Kocsis. "Design for Dependability Through Error Propagation Space Exploration". In: *2018 48th Annual IEEE/IFIP International Conference on Dependable Systems and Networks Workshops (DSN-W)*. 2018 48th Annual IEEE/IFIP International Conference on Dependable Systems and Networks Workshops (DSN-W). June 2018, pp. 172–178. DOI: `10.1109/DSN-W.2018.00059`.

[Pat+13] András Pataricza, Imre Kocsis, Ágnes Salánki, and László Gönczy. "Empirical Assessment of Resilience". In: *Software Engineering for Resilient Systems*. International Workshop on Software Engineering for Resilient Systems. Lecture Notes in Computer Science. Springer, Berlin, Heidelberg, Oct. 3, 2013, pp. 1–16. DOI: `10.1007/978-3-642-40894-6_1`.

[Pal+10a] G. J. Paljak, Z. Égel, D. Tóth, I. Kocsis, T. Kovácsházy, and A. Pataricza. "Qualitative Performance Control in Supervised IT Infrastructures". In: *2010 International Conference on Dependable Systems and Networks Workshops (DSN-W)*. 2010 International Conference on Dependable Systems and Networks Workshops (DSN-W). June 2010, pp. 59–65. DOI: `10.1109/DSNW.2010.5542618`.

[Pal+10b] Gergely János Paljak, Imre Kocsis, Zoltán Égel, Dániel Tóth, and András Pataricza. "Sensor Selection for IT Infrastructure Monitoring". In: *Autonomic Computing and Communications Systems. AUTONOMICS 2009*. Vol. 23. Lecture Notes of the Institute for Computer Sciences, Social Informatics and Telecommunications Engineering. Springer, Berlin, Heidelberg, 2010, pp. 130–143. URL: `https://link.springer.com/chapter/10.1007/978-3-642-11482-3_9` (visited on 07/24/2017).

[Koc+08a] I. Kocsis, G. Csertán, P. L. Pásztor, and A. Pataricza. "Dependability and Security Metrics in Controlling Infrastructure". In: *2008 Second International Conference on Emerging Security Information, Systems and Technologies*. 2008 Second International Conference on Emerging Security Information, Systems and Technologies. Aug. 2008, pp. 368–374. DOI: 10.1109/SECURWARE.2008.62.

[Koc+08b] Imre Kocsis, Zsolt Kocsis, Imre Malatinszki, and András Pataricza. "Dependability Benchmarks and System Supervision". In: *Workshop on Sharing Field Data and Experiment Measurements on Resilience of Distributed Computing Systems at IEEE SRDS 2008*. Naples, Italy, 2008.

Indirectly supporting and unrelated publications

[Gar+18] P. Garamvölgyi, I. Kocsis, B. Gehl, and A. Klenik. "Towards Model-Driven Engineering of Smart Contracts for Cyber-Physical Systems". In: *2018 48th Annual IEEE/IFIP International Conference on Dependable Systems and Networks Workshops (DSN-W)*. 2018 48th Annual IEEE/IFIP International Conference on Dependable Systems and Networks Workshops (DSN-W). June 2018, pp. 134–139. DOI: 10.1109/DSN-W.2018.00052.

[Koc+18] Imre Kocsis, Attila Klenik, András Pataricza, Miklós Telek, Flórián Deé, and Dávid Cseh. "Systematic Performance Evaluation Using Component-in-the-Loop Approach". In: *International Journal of Cloud Computing* 7.3-4 (2018), pp. 336–357. DOI: https://doi.org/10.1504/IJCC.2018.095401.

[KMZ17] Imre Kocsis, Zoltán Ádám Mann, and Dávid Zilahi. "Optimized Deployment of Critical Applications in Infrastructure-as-a-Service Clouds". In: *International Journal of Cloud Computing* 4.6 (2017), pp. 342–362.

[Sal+17] Ágnes Salánki, Gergő Kincses, László Gönczy, and Imre Kocsis. "Data Analysis-Based Capacity Planning of VCL Clouds". In: *International Journal of Cloud Computing* 4.6 (2017), pp. 370–383.

[Koc+16] Imre Kocsis, Áron Tóth, Zoltán Szatmári, Tamás Dabóczi, András Pataricza, and Gábor Guta. "Towards Cyber-Physical System Technologies over Apache VCL". In: *International Journal of Cloud Computing* (2016).

[Pat+16] András Pataricza, Imre Kocsis, László Gönczy, Katalin Tasi, Gergő Kincses, Dávid Cseh, and Flórián Deé. "Cloud Application Performance Analysis with IBM Watson Analytics (Extended Abstract)". In: *4th International IBM Cloud Academy Conference*. Edmonton, Canada, June 2–3, 2016.

[Koc14] Imre Kocsis. "Vizuális Analízis". In: Péter Antal et al. *Intelligens Adatelemzés*. Budapest: Typotex Kiadó, 2014, pp. 58–69.

[Koc+14] Imre Kocsis, András Pataricza, Gábor Huszerl, Benedek Izsó, Zoltán Szatmári, Áron Tóth, Dániel Varró, and András Vörös. "Informatikaoktatás Felhőben: Egy Új Oktatási Modell Bevezetése". In: *BeleSTEM - Felsőoktatási Jó Gyakorlatok a Tudomány, a Technológia, a Műszaki Tudományok És a Matematika Szolgálatában*. Ed. by Nóra Révai. Budapest: Tempus Közalapítvány, 2014, pp. 36–41.

[Urb+14] Gábor Urbanics, László Gönczy, Balázs Urbán, János Hartwig, and Imre Kocsis. "Combined Error Propagation Analysis and Runtime Event Detection in Process-Driven Systems". In: *Software Engineering for Resilient Systems.* International Workshop on Software Engineering for Resilient Systems. Lecture Notes in Computer Science. Springer, Cham, Oct. 15, 2014, pp. 169–183. DOI: 10.1007/978-3-319-12241-0_13.

[Sal+11] Ágnes Salánki, Imre Kocsis, András Pataricza, and Zsolt Kocsis. "Qualitative Characterization of Quality of Service Interference between Virtual Machines". In: *Konferenz: ARCS 2011 - 24th International Conference on Architecture of Computing Systems.* 24th International Conference on Architecture of Computing Systems. Como, Italy: VDE Verlag, Feb. 22, 2011, p. 8.

[Pal+10c] Gergely Paljak, István Szombath, Imre Kocsis, Tamás Kovácsházy, and András Pataricza. "Model-Based Control of IT Infrastructures". In: *International Carpathian Control Conference - ICCC 2010.* Eger, Hungary, May 26–29, 2010.

[Gus+09] M. Gusat, C. Decusatis, C. Minkenberg, L. Mckenna, K. Bhardwaj, G. J. Paljak, A. Pataricza, and I. Kocsis. "Benchmarking the Ethernet-Federated Datacenter". In: *First Workshop on Data Center-Converged and Virtual Ethernet Switching.* Issy-les-Moulineaux, Sept. 14, 2009.

[Koc08] Imre Kocsis. "Model Driven Design for Structural Reconfiguration Based Dependability". In: *Proceedings of the 15th PhD Mini-Symposium.* Ed. by Béla Pataki. Budapest, Hungary: Budapest University of Technology and Economics, Department of Measurement and Information Systems, 2008, pp. 24–27.

[Pin+08] Gergely Pintér, Zoltán Micskei, András Kövi, Zoltán Égel, Imre Kocsis, Gábor Huszerl, and András Pataricza. "Model-Based Approaches for Dependability in Ad-Hoc Mobile Networks and Services". In: *Architecting Dependable Systems V.* Ed. by Rogério de Lemos, Felicita Di Giandomenico, Henry Muccini, and Cristina Gacek. Lecture Notes in Computer Science. Springer, Berlin, Heidelberg, 2008, pp. 150–174. DOI: 10.1007/978-3-540-85571-2_7. URL: https://link.springer.com/chapter/10.1007/978-3-540-85571-2_7.

[Koc07] Imre Kocsis. "Towards the Model Based Design of Supervisory Systems". In: *Proceedings of the 14th PhD Mini-Symposium.* Ed. by Béla Pataki. Budapest, Hungary: Budapest University of Technology and Economics, Department of Measurement and Information Systems, 2007, pp. 112–113.

[PPK07] András Pataricza, Péter László Pásztor, and Imre Kocsis. "Bridging Model-Based Design for Dependability and Empirical Observations". In: *Workshop on Reliability Analysis of System Failure Data.* Cambridge, United Kingdom, 2007.

[KFN02] I. Kocsis, L. Farkas, and L. Nagy. "3G Base Station Positioning Using Simulated Annealing". In: *The 13th IEEE International Symposium on Personal, Indoor and Mobile Radio Communications.* The 13th IEEE International Symposium on Personal, Indoor and Mobile Radio Communications. Vol. 1. Sept. 2002, 330–334 vol.1. DOI: 10.1109/PIMRC.2002.1046716.

References

[1] A. Avizienis, J.-C. Laprie, B. Randell, and C. Landwehr. "Basic Concepts and Taxonomy of Dependable and Secure Computing". In: *IEEE Transactions on Dependable and Secure Computing* 1.1 (Jan. 2004), pp. 11–33. DOI: `10.1109/TDSC.2004.2`.

[2] Robert Hanmer. *Patterns for Fault Tolerant Software*. John Wiley & Sons, July 12, 2013. 329 pp.

[3] Jean-Claude Laprie. "From Dependability to Resilience". In: *38th IEEE/IFIP Int. Conf. On Dependable Systems and Networks*. 2008.

[4] James P. G. Sterbenz, Egemen K. Çetinkaya, Mahmood A. Hameed, Abdul Jabbar, Shi Qian, and Justin P. Rohrer. "Evaluation of Network Resilience, Survivability, and Disruption Tolerance: Analysis, Topology Generation, Simulation, and Experimentation". In: *Telecommunication Systems* 52.2 (Dec. 7, 2011), pp. 705–736. DOI: `10.1007/s11235-011-9573-6`.

[5] James P. G. Sterbenz, David Hutchison, Egemen K. Çetinkaya, Abdul Jabbar, Justin P. Rohrer, Marcus Schöller, and Paul Smith. "Resilience and Survivability in Communication Networks: Strategies, Principles, and Survey of Disciplines". In: *Computer Networks*. Resilient and Survivable Networks 54.8 (June 1, 2010), pp. 1245–1265. DOI: `10.1016/j.comnet.2010.03.005`.

[6] Kishor S. Trivedi and Andrea Bobbio. *Reliability and Availability Engineering: Modeling, Analysis, and Applications*. Cambridge University Press, Aug. 3, 2017. 729 pp.

[7] ISO/IEC/IEEE. "ISO/IEC/IEEE International Standard - Systems and Software Engineering – Software Life Cycle Processes". In: *ISO/IEC/IEEE 12207:2017(E) First edition 2017-11* (Nov. 2017), pp. 1–157. DOI: `10.1109/IEEESTD.2017.8100771`.

[8] Eric Bauer and Randee Adams. *Service Quality of Cloud-Based Applications*. John Wiley & Sons, Nov. 14, 2013. 302 pp.

[9] J. O. Kephart and W. E. Walsh. "An Artificial Intelligence Perspective on Autonomic Computing Policies". In: *Proceedings. Fifth IEEE International Workshop on Policies for Distributed Systems and Networks, 2004. POLICY 2004.* Proceedings. Fifth IEEE International Workshop on Policies for Distributed Systems and Networks, 2004. POLICY 2004. June 2004, pp. 3–12. DOI: `10.1109/POLICY.2004.1309145`.

[10] Joseph L. Hellerstein, Yixin Diao, Sujay Parekh, and Dawn M. Tilbury. *Feedback Control of Computing Systems*. John Wiley & Sons, Sept. 21, 2004. 452 pp.

[11] Stephan Morgenthaler. "Exploratory Data Analysis". In: *Wiley Interdisciplinary Reviews: Computational Statistics* 1.1 (2009), pp. 33–44. DOI: `10.1002/wics.2`.

[12] Martin Theus and Simon Urbanek. *Interactive Graphics for Data Analysis: Principles and Examples*. CRC Press, Oct. 24, 2008. 295 pp.

[13] Karama Kanoun and Lisa Spainhower. *Dependability Benchmarking for Computer Systems*. John Wiley & Sons, Oct. 3, 2008. 382 pp.

[14] Michael Grieves and John Vickers. “Digital Twin: Mitigating Unpredictable, Undesirable Emergent Behavior in Complex Systems”. In: *Transdisciplinary Perspectives on Complex Systems: New Findings and Approaches.* Ed. by Franz-Josef Kahlen, Shannon Flumerfelt, and Anabela Alves. Cham: Springer International Publishing, 2017, pp. 85–113. DOI: 10.1007/978-3-319-38756-7_4. URL: https://doi.org/10.1007/978-3-319-38756-7_4 (visited on 11/24/2018).

[15] IBM. *Analytics Solutions Unified Method - Implementations with Agile Principles.* datasheet. IBM, 2016. URL: ftp://ftp.software.ibm.com/software/data/sw-library/services/ASUM.pdf (visited on 07/22/2018).

[16] András Pataricza. “Model-Based Dependability Analysis”. DSc thesis. Hungarian Academy of Sciences, 2008.

[17] Rudolf Wille. “Formal Concept Analysis as Mathematical Theory of Concepts and Concept Hierarchies”. In: *Formal Concept Analysis: Foundations and Applications.* Ed. by Bernhard Ganter, Gerd Stumme, and Rudolf Wille. Lecture Notes in Computer Science. Berlin, Heidelberg: Springer Berlin Heidelberg, 2005, pp. 1–33. DOI: 10.1007/11528784_1. URL: https://doi.org/10.1007/11528784_1 (visited on 09/16/2018).

[18] Szilvásy Noémi and Urbán Balázs. *Adatelemzési Folyamatok Diagnosztikája És Adatminőség Érzékenység Elemzése.* 2015. URL: https://tdk.bme.hu/VIK/Intelligens1/Adatelemzesi-folyamatok-diagnosztikaja-es (visited on 12/08/2018).

[19] Hanchuan Peng, Fuhui Long, and C. Ding. “Feature Selection Based on Mutual Information Criteria of Max-Dependency, Max-Relevance, and Min-Redundancy”. In: *IEEE Transactions on Pattern Analysis and Machine Intelligence* 27.8 (Aug. 2005), pp. 1226–1238. DOI: 10.1109/TPAMI.2005.159.

[20] D. M. (Rini) van Solingen and Egon W. Berghout. *The Goal/Question/Metric Method: A Practical Guide for Quality Improvement of Software Development.* McGraw-Hill, 1999. URL: https://research.tue.nl/en/publications/the-goalquestionmetric-method-a-practical-guide-for-quality-impro (visited on 12/15/2018).

[21] Thomas L. Saaty. “How to Make a Decision: The Analytic Hierarchy Process”. In: *European Journal of Operational Research.* Desicion Making by the Analytic Hierarchy Process: Theory and Applications 48.1 (Sept. 5, 1990), pp. 9–26. DOI: 10.1016/0377-2217(90)90057-I.

[22] Judea Pearl. *Causality.* Cambridge University Press, Sept. 14, 2009. 486 pp.