Vision on System Qualities Reasoning in the High-Tech Equipment Industry

TNO Public ⟩ TNO 2026 R10422
February 18, 2026

ICT, Strategy & Policy
www.tno.nl
+31 88 866 50 00
info@tno.nl

TNO 2026 R10422 – February 18, 2026

Vision on System Qualities Reasoning in the High-Tech Equipment Industry

Author(s)	Bram van der Sanden, Jacques Verriet
Classification report	TNO Public
Title	TNO Public
Report text	TNO Public
Number of pages	22 (excl. front and back cover)
Number of appendices	0

Disclaimer
In the creation of this document, Microsoft 365 Copilot was used to improve the use of English. In all instances where Copilot was used, the authors thoroughly reviewed the results to ensure that the end result is accurate and free from errors, thus allowing the authors to take responsibility for the content.

Acknowledgement
The research is carried out under the responsibility of TNO-ESI. This research is supported by the Netherlands Organization for Applied Scientific Research TNO, the Netherlands Ministry of Economic Affairs and Climate.

Summary

This document presents TNO-ESI's Vision on System Qualities Reasoning for the Dutch high tech equipment industry. It is written for industry leaders, system architects, R&D strategists, and researchers who need to understand how system qualities such as productivity, energy efficiency, resource usage, and cost can be systematically engineered across increasingly complex, interconnected, and software-defined systems.

The document explains why system qualities must become a first-class engineering concern and outlines a long-term roadmap for research and innovation. It introduces key industry challenges in context of emerging technological and societal trends and defines strategic directions to realize the vision by developing methodologies that allow engineers to stay in control of system qualities through the entire lifecycle.

The research roadmap is structured around the three major industry needs, each with several strategic research directions:

1. Analysis, prediction, and optimization of system qualities over the full lifecycle
    - Reasoning with uncertainty
    - Quantitative modeling
    - Model extraction and calibration
    - Performance verification
    - Runtime monitoring and optimization
2. Systematic reasoning on multiple system qualities
    - Trade-off analysis using quantitative models
    - Multi-objective optimization
3. System-dependent analysis, prediction, and optimization of system qualities
    - Performance engineering for systems of systems
    - Performance engineering for product families
    - Performance engineering for software-defined systems
    - Performance engineering for AI-enabled systems

Addressing the needs together forms a connected strategy: from early design to runtime adaptation, from expert-driven to democratized engineering with quantitative reasoning, and from individual components and individual systems to ecosystems.

Contents

1 Introduction

This document gives an overview of TNO-ESI's vision on reasoning on system qualities, in short, our "System Qualities vision". It describes the global trends and drivers being faced by the Dutch high-tech equipment industry, the key industry needs that should be addressed regarding reasoning on system qualities, and strategic research topics that are key to address those needs.

The System Qualities vision focuses on *system performance* in a broad sense, referring to how effectively and efficiently a system operates to meet its intended goals. This includes qualities like productivity, effective usage of the available resources, reliability, responsiveness, and energy consumption. The vision targets the development of methodologies that enable system engineers to be in control of the system qualities that give value to their systems in the market. This includes reasoning about qualities in isolation and in their interplay, as well as providing methods to predict qualities in the early-design phase, perform design-space exploration, trade-off analyses and optimize design choices to name a few. We address reasoning on system qualities both at design time, for example balance total cost of the system versus productivity, and at runtime, for example to optimize resource usage for a given workload. The vision also targets the development of performance diagnostics methodologies to analyze the root-cause when a system performs below expectations in terms of its qualities.

System performance can be seen in scope of the broader definition of *system operational effectiveness* (see Figure 1.1) [1] that considers the effective performance of the system over the full lifecycle, including all activities to develop and maintain the system. Here, system performance is defined in terms of functions, capabilities, and priorities. Functions relate to the abilities that the systems should have in the operational environment. For example, an X-ray medical device that can generate images of parts of a patient's body. Capabilities are the desired performance attributes and measures of the system. Continuing the example, the X-ray device can perform imaging with a specific resolution and contrast, given a patient with certain characteristics, and under a specified X-ray skin dose. Priorities reflect the stakeholders' expectations and capture the relative weightings of the capabilities as applied to different scenarios.

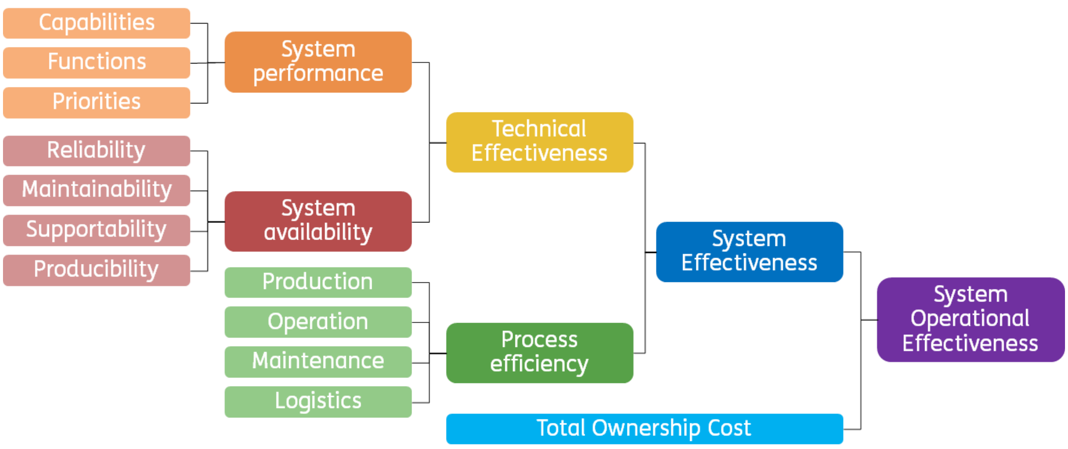

Figure 1.1: System Operational Effectiveness (adapted from [1])

The System Qualities vision connects closely with the field of *loss-driven systems engineering (LDSE)*, that focuses on addressing possible losses associated with system development, use, and maintenance. Systems engineering is typically “driven by capabilities, but not on integrated attention to potential losses associated with developing and using systems” [2, p. 191]. Qualities like performance, reliability, usability, and security are often considered as specialty areas in isolation and can greatly benefit from comprehensive approaches that leverage commonality and synergy in modeling and reasoning about them. We anticipate that our System Qualities vision will provide significant contributions to the field of LDSE. Our vision also connects to standards like ISO/IEC 25010 [3]. That standard defines a taxonomy of qualities, that can be used when evaluating a software product.

The remainder of this document outlines the trends shaping the Dutch high-tech equipment industry in Section 2, presents the core vision and its three key engineering needs in Section 3, and then details the strategic research directions that address these needs in Section 4. It concludes by positioning this vision within related national and international roadmaps in Section 5 and summarizing its implications for industry and research stakeholders in Section 6.

2 Strategic Context

2.1 Dutch High-Tech Equipment Industry

The Dutch high-tech equipment industry develops world-leading systems in a diverse set of market segments, including equipment for the semiconductor industry, healthcare imaging, production printing, electron microscopy, logistic process automation, and combat management systems. Many innovations are related to precision- and opto-mechatronics, whereas the added value today is increasingly around software.

The Dutch high-tech equipment companies share many similarities, despite their different market segments. They all target high-end markets, serving customers with systems tailored to their needs. They build complex systems in relatively low numbers, ranging from a few to hundreds per year. The high-tech systems produced by these high-tech equipment companies have a long lifespan, typically being operational for multiple decades. Next to the product similarities, another characteristic is that companies typically collaborate in the Dutch innovation ecosystem, learning from each other and leveraging the network effect to stay in a world-leading position. TNO-ESI plays an orchestrating role in this ecosystem, bringing industry partners together, driving applied research on common challenges, and helping companies adopt fundamental research results.

2.2 Trends

Societal, economic, and technological trends and drivers are impacting the Dutch high-tech equipment industry, providing new opportunities as well as leading to challenges that need to be addressed. This section describes the key trends and drivers that have been identified in TNO-ESI's vision on Systems Engineering & Systems Architecting [4] in the context of reasoning on system qualities.

2.2.1 T1: Increasing Dependence on Complex Systems

Society is increasingly relying on the correct operation of high-tech systems. These systems need to work in an increasingly uncertain, changing, and potentially hostile system context. For example, the luggage handling at airports is increasingly automated to deal with the increasing amount of luggage and should be robust against delayed flights. System performance needs to be guaranteed under varying system usage scenarios and in the interaction with other systems that the system depends on. Scenarios may not always be known at design time to further add to the potential complexity. System performance should degrade gracefully if performance can no longer be guaranteed under certain circumstances. The system should also be able to optimize its performance at runtime, depending on the usage and system context.

2.2.2 T2: Growing Diversity and Customization

Companies offer a large range of features and options to tailor systems to the usage needs of customers. This leads to more diversity in the product portfolio and to more variations of the system components. System performance should be guaranteed for all relevant system configurations, without requiring expensive physical prototypes of each configuration available for testing and validation.

2.2.3 T3: Growing Role of AI

The increased capabilities of artificial intelligence provide opportunities to improve the performance of high-tech systems by runtime optimization, but also by enhancing usability of the systems, and increasing their autonomy. Upon encountering performance issues, AI provides opportunities to automatically analyze the root cause of performance issues (e.g., using ML to analyze deviations from normal operation) and take corrective actions to resolve the issues (e.g., using reinforcement learning). AI also provides opportunities to improve efficiency and effectiveness of the system development process itself, by providing ways to automatically create models based on other artifacts (e.g., using agentic AI), or provide digital assistance to make better design decisions and predict system qualities. Next to these opportunities, AI may also introduce complexity and uncertainty to the system, for example to know in which situations the AI agent in a system can diagnose or optimize the system at runtime.

2.2.4 T4: Continuous Innovation and Updating on Long-Life Systems

High-tech systems typically have a lifespan of multiple decades. Customers expect that the systems will be kept up-to-date and stay performant, or even improve performance, over time. Existing systems are upgraded to improve and ensure performance in a continuously changing system context with new systems to interact with. System performance should be guaranteed in this setting after rolling out new software updates and applying system patches.

2.2.5 T5: Climbing the Value Chain

High-tech equipment manufacturers are climbing the value chain, providing not only the equipment but also additional services for their customers. Consequently, there is a shift from focusing on equipment performance to performance of the equipment in a specific workflow with other systems and users operating the systems. With this shift, aspects like availability, unscheduled downtime, and switching between usage scenarios become more prominent. System manufacturers need to be able to predict how the system will perform in the workflow of the customer, incorporating knowledge of the other systems in the workflow. This is especially important for high-tech systems that are the bottlenecks in the workflow, determining the overall workflow performance.

2.2.6 T6: Innovation in Ecosystems

In the Dutch high-tech equipment industry, OEMs typically collaborate with first- and second-tier suppliers, knowledge institutes, and engineering consultancy companies to develop and maintain their systems. Parts of the system are often outsourced, being developed by external companies, while the OEM still needs to guarantee performance of the assembled system. The industry needs to improve innovation in ecosystems and effectively work together in the setting of digital engineering, collaborating more with shared models, data, and knowledge.

2.2.7 T7: Growing Scarcity of Engineering Experts

The Netherlands is facing an aging population, and at the same time many of the experts that created the high-tech systems in the past decades are going into retirement. As the systems have become more complex over time, the learning curve to train new engineers has increased substantially. Currently, only a few experts still know the full system and can oversee the consequences of design decisions on system-level qualities. High-tech equipment companies are looking for ways to “democratize” their systems, making them easier to comprehend and to assess the system-level qualities from the design. At the same time, they also want to democratize the engineering process, have assistance available to predict system qualities, and determine the right budget decompositions to decompose budgets for aspects like timing latency or energy to the component level.

2.2.8 T8: Growing Emphasis on Sustainability

The European Union has set the goal to be climate-neutral,[1] addressing the threats of climate change and environmental degradation. The Dutch high-tech equipment industry must also move towards sustainable business models through net-zero emissions of greenhouse gasses by 2050 (according to the European Green Deal) and a circular production model. This move is also emphasized by the shift of the manufacturing industry to Industry 5.0, focusing on sustainability, human centricity, and resilience. Sustainability concerns will affect the design of next-generation high-tech systems. For example, lowering the energy consumption of systems may mean that systems go into a hibernation mode, and need a longer ramp-up time to resume operation; or the interplay between sustainability and the quality of the products produced with the equipment, where a higher-quality product could mean a higher usage of scarce materials being consumed; or reducing energy consumption by optimizing code, potentially decreasing system performance. Getting a clear grip on these relationships and having the means to make the right trade-offs is key. Balancing sustainability with other system qualities is part of TNO-ESI’s vision on Lifecycle Excellence in the High-Tech Equipment Industry [5] and therefore not elaborated upon in this vision.

[1] See for example https://ec.europa.eu/newsroom/know4pol/items/664852.

3 Vision

In a world where high-tech systems are increasingly complex, interconnected, and software-defined, the Dutch high-tech equipment industry must deliver systems that not only function but excel—across qualities like productivity, energy efficiency, cost, and robustness. The TNO-ESI System Qualities vision provides a roadmap to empower system engineers to stay in control of these qualities throughout the system lifecycle. The vision is based on three key industry needs:

1. Analysis, prediction, and optimization of system qualities over the full lifecycle
2. Systematic reasoning on multiple system qualities
3. System-dependent analysis, prediction, and optimization of system qualities

Addressing these three industry needs together form a connected strategy: from early design to runtime adaptation, from expert-driven to democratized engineering with quantitative reasoning, and from individual components and individual systems to ecosystems. We believe that this strategy will empower system engineers to stay in control of all key system qualities throughout the system lifecycle, delivering systems with excellent operational effectiveness.

This connected strategy transitions the industry from traditional, expert-driven engineering toward a democratized, model-based approach, where quantitative reasoning is accessible to a wider group of engineers. Ultimately, it empowers teams to maintain control over system qualities from initial design through to runtime adaptation, supporting operational excellence at every stage of the system lifecycle. In the remainder of this section, we introduce the three industry needs in more detail.

Analysis, Prediction, and Optimization of System Qualities over the Full Lifecycle
Systems should be designed for system qualities from the outset, not as an afterthought. System architects require the ability to understand the consequences of design choices on system qualities and emerging system behavior. Having this ability ensures that decisions result in a system that is fit for purpose. It mitigates uncertainty and risks in the design process and avoids costly rework during later system development, during system integration, or even at runtime. It is our ambition to enable early analysis, prediction, and optimization of system qualities, when the system is not yet realized or not operational in the field. Predictions should be made as early as possible, with a level of accuracy appropriate for the current engineering phase, using the assumptions, data, and knowledge that are available. We believe that digital engineering and model-based engineering are required to achieve early analysis and prediction, where quantitative models are used to predict how design decisions will impact the resulting system qualities. In our vision, these models can be used over the full system lifecycle, for example for performance verification and validation, and to optimize the system at runtime once it is operating in the field.

Systematic Reasoning on Multiple System Qualities
When predictions can be made for multiple individual system qualities, these can serve as the starting point to make informed trade-off decisions. Systematic reasoning about system qualities is needed to analyze the interplay between qualities. The qualitative and quantitative models and predictions that are used in this decision-making process should be understandable and trustworthy to the various stakeholders involved. Quantitative models can be used in semi-automated design-space exploration to find potential system designs quickly

based on the design decisions being considered. In such an exploration, predictions feed into multi-objective optimization, to accelerate the process of finding (near-)optimal design choices, balancing competing qualities like energy efficiency, cost, and throughput.

System-Dependent Analysis, Prediction, and Optimization of System Qualities
Prediction and optimization of the key system qualities are highly dependent on a system's type and context, as they determine what aspects are important to consider. We see specific needs for systems of systems and the systems that are part of systems of systems, product families, for software-defined systems, and for AI-enabled systems, incorporating AI components in the system.

Systems of systems Performance engineering of systems of systems considers the increasing scope, where systems operate within (customer) workflows involving multiple systems. New methods for workflow optimization are essential to guarantee end-to-end system-of-systems effectiveness. In this context, it is important to reason about how the qualities of the constituent systems contribute to those of the system of systems.

Product families Systems are generally developed for various purposes and contexts, resulting in multiple system variants. To effectively manage these variants, performance engineering for product families is needed. With product families, system configurations are derived from a common platform, enabling re-use of building blocks of the system. Strategies are needed to ensure that system performance can be optimized for the configurations derived from the platform.

Software-defined systems Software accounts for a significant portion of the value delivered by today's high-tech equipment. Understanding the software's influence on system qualities is key to dealing with the system-level complexity introduced by the software. Furthermore, software execution architectures and deployment strategies are needed such that the system can respond at runtime to changes in the system workflow and adapt the required resources dynamically.

AI-enabled systems AI components are increasingly being integrated into high-tech systems. Prediction and optimization techniques should be able to incorporate reasoning on these components, requiring performance engineering of AI-enabled systems to understand the AI components' probabilistic and dynamic behavior, whether the component is fit for the context in which it operates, and how the AI component influences the relevant system qualities.

4 Strategic Research Directions

This chapter outlines how we have translated the three key industry needs of our vision into a focused research agenda, detailing strategic research topics and corresponding research questions for each need. By systematically addressing these research questions, solutions can be developed to address the identified industry challenges and drive progress to realize the vision.

4.1 Analysis, Prediction, and Optimization of System Qualities over the Full Lifecycle

Reasoning with Uncertainty (RWU)

The first key research topic in enabling early analysis and prediction of system qualities is how to reason with uncertainty. This includes dealing with incomplete or imprecise input data and assessing the accuracy of resulting performance predictions. The uncertainty comes from the fact that the system design and exact usage profile of the system are not yet fully known. Tackling this topic is key to ensure that systems can work in an increasingly uncertain and changing system context, where at the same time society is relying on a correct system functioning (trend T1). Key questions include:

- RWU1 ▪ How to predict and validate key system qualities and their interrelations in the early-design phase?
- RWU2 ▪ How to obtain sufficiently accurate performance numbers for system components (e.g., based on aspects like CPU usage, energy consumption, or heat generation) in a non-intrusive way (e.g., low overhead)? (related to ED3 in [6])

Quantitative Modeling (QM)

Systematic reasoning is often done using conceptual models on a whiteboard or qualitative models. Qualitative models capture the relationships between system qualities that affect each other and help to determine the relationships between system design parameters and the resulting system quality values. Qualitative models are not executable to automatically compute the system qualities directly from the model. We see quantitative modeling as the next step, where aspects like energy usage, timing, and reliability are quantified using models that capture the relevant parts of the system design and are formalized to semi-automatically compute the resulting system quality values. Such models help to make tacit expert knowledge accessible within the organization, alleviating the issue of a growing scarcity of engineering experts (trend T7). The models can be used for systematic reasoning and trade-off analysis, which is the focus of the research topic Trade-Off Analysis using Quantitative Models (TAQM). Currently, there is limited automated support to model and quantify system qualities and reason about them in a systematic way. The models should show the underlying assumptions and accuracy of the computed quality values. Research questions in this direction include:

- QM1 ▪ What are key metrics that can be used to quantify different system qualities?
- QM2 ▪ How to determine the critical design parameters or component properties that influence one or multiple system qualities?

Model Extraction and Calibration (MEC)

Our vision is to leverage model-based analysis and simulation to predict system qualities during the early design phase using models that are semi-automatically created from existing requirements, operational data, and high-level specifications. For example, using AI-based techniques to learn and calibrate these models (trend T3). Building accurate models is typically time-consuming and requires a good understanding of the system, making it a challenge and bottleneck, particularly in established (brownfield) environments, and in organizations that see a growing scarcity of engineering experts (trend T7). We see model extraction and calibration as an important research topic, including the following questions:

- MEC1 ▪ How to combine static and dynamic data like traces, logs, documentation, existing code, and measurement data with expert knowledge to infer models that capture the structure and dynamics of the system?
- MEC2 ▪ How to infer quantities that can be added to the models as inferred in MEC1 with available performance data? For example, extracting the durations of actions in a performance model with timings observed in execution traces.
- MEC3 ▪ How to assess the model accuracy of models as inferred in MEC1 and MEC2?

Performance Verification (PV)

Verifying and validating system performance is typically done late in the system development process, after assembly of the system. Performance is difficult to verify in early stages, as hardware is typically needed that is not yet available or cannot be tested without the integration with the rest of the system. Testing using virtual prototypes, for example with hardware virtualization, often delivers results that have a limited accuracy, caused by abstractions and not knowing precisely how the system will be used. Besides the technical challenges, testing as a discipline is strongly driven by functional testing, and has much less focus on testing system qualities. Testing should be broadened to check for system qualities, besides focusing on testing functionality. We see performance verification and validation over the full system lifecycle as an important research topic to address the challenge of continuously innovating and updating systems (trend T4). Relevant questions include:

- PV1 ▪ How to verify performance in the early-design phase, for example using techniques like virtual prototyping or hardware-in-the-loop testing?
- PV2 ▪ How to determine a performance test plan that minimizes test effort using models created during design-time?
- PV3 ▪ How do we verify performance of rapidly evolving applications in a scalable manner?

Runtime Monitoring and Optimization (RMO)

System qualities need to be balanced at runtime, as it is often not feasible to optimize system performance during the system design phase for all possible runtime situations [6]. For example, when the system usage changes, external systems start interacting differently with the system, or operating conditions impacting the system performance change. Runtime monitoring is required to detect runtime changes. Monitoring can be done as part of the system, or by an external monitor. Runtime optimization is needed to swiftly react and ensure a good performance under changing conditions, such that users can trust the systems (trend T1). Optimization might be autonomous, where the system optimizes itself (e.g., using AI-based techniques (trend T3)), or the system asks for operator assistance, providing a detailed analysis of observed issues and possible root causes. The system should know within which boundaries it may optimize itself and when to escalate to a higher-level optimization loop. Important research questions on runtime monitoring and optimization are:

- RMO1 ▪ How to develop a framework for runtime system monitoring and optimization, e.g., based on well-known frameworks like MAPE-K [7] or the Observe-Orient-Decide-Act loop by Boyd [8]?
- RMO2 ▪ How to make trade-offs and optimize system performance at runtime, given the limited available time and computational resources to find a sufficiently good solution (e.g., using anytime algorithms [9])?
- RMO3 ▪ What are good strategies to design and dimension architectures for resilient, high-performing systems that can respond to changes in the system context?
- RMO4 ▪ How to compare the system with related systems in the field (e.g., using fleet analytics) to identify performance degradations or optimization opportunities?

4.2 Systematic Reasoning on Multiple System Qualities

Trade-Off Analysis using Quantitative Models

Using quantitative models, trade-off analysis on different system designs can be performed (semi-)automatically, showing how improving one system quality affects other qualities, and showing which designs are optimal with respect to the optimization criteria. Such techniques help to address the growing scarcity of engineering experts (trend T7). AI-based techniques may provide added value to reason under uncertainty (trend T3). Research questions on trade-off analysis using quantitative models include:

- TAQM1 ▪ How can quantitative models predict system qualities under uncertainty, e.g., due to varying workloads?
- TAQM2 ▪ How can quantitative reasoning help to make trade-offs between system qualities? For example, balancing qualities such as performance, cost, and energy usage.
- TAQM3 ▪ How to accurately capture complex, interdependent relationships among various system quality attributes?
- TAQM4 ▪ How to perform trade-off analysis when considering multiple envisioned system users?
- TAQM5 ▪ How can AI be integrated with traditional quantitative methods to enhance reasoning about system qualities?

Multi-Objective Optimization

There is a lack of techniques that can effectively deal with multiple optimization criteria [10]. Due to this lack of techniques, multi-objective analysis questions are often transformed into a single-objective analysis question by combining criteria in one high-level criterion using a weighting scheme or by introducing budget constraints [10]. Solving multi-objective optimization problems cannot be done by a single technique but rather requires the combination of multiple methods [10], e.g., to evaluate the different system qualities using analytical models, simulations, or based on system data. A key challenge is also to determine which design points to evaluate in a large design space, for example using heuristics or AI-based techniques (trend T3) like reinforcement learning. Research questions on multi-objective optimization include:

- MOO1 ▪ How to determine which exploration methods to select for multi-objective optimization (e.g., constraint programming or simulation-based optimization)? Often this decision is guided by the experience and preferences of the engineer, rather than by best practices.
- MOO2 ▪ How to explore a multi-objective optimization space by accounting for design- and trade-space structure, especially when certain key design or quality parameters heavily influence decisions, rather than relying solely on sampling?
- MOO3 ▪ How to synthesize a system design that is optimized for multiple objectives using robust optimization, to find a design that can handle variations and deviations?
- MOO4 ▪ How to effectively use surrogate modeling to reduce the time needed to analyze individual design points?

4.3 System-Dependent Analysis, Prediction, and Optimization of System Qualities

Performance Engineering for Systems of Systems (PESOS)

High-tech equipment systems often operate within a customer workflow involving multiple systems, where OEMs are looking for opportunities to provide additional services in such workflows (trend T5). In such a context, there are often many unknowns, including the lack of in-depth knowledge about third-party systems or outsourced systems or system parts (trend T6). Furthermore, systems in a system-of-systems setting should adapt when other systems are dynamically added or removed over time. There is also the risk that systems in a system of systems have conflicting optimization criteria, leading to sub-optimal performance of systems they interact with as well as of the overall system of systems. As the scope increases, there is also the challenge in scalability of the analysis. We see AI-based techniques as an opportunity to improve reasoning and optimization in the setting with partially unknown systems. Key research questions to guarantee overall systems of systems performance include:

- PESOS1 ▪ How to design system-level interfaces that are sufficiently rich to assert quality aspects of the integrated system and support very late binding?
- PESOS2 ▪ How can systems effectively adapt to guarantee or optimize performance at the system-of-systems level, where systems can be dynamically added or removed, and systems might have their own goals?
- PESOS3 ▪ How to learn performance models of partially unknown systems?
- PESOS4 ▪ How to efficiently predict system qualities for a specific user context?
- PESOS5 ▪ How does a change to an individual system influence the customer workflow's performance?

These research questions connect to the ESI vision on Systems of Systems in the High-Tech Equipment Industry [11]. The use of AI-based techniques for optimization in a systems-of-systems context is in line with the TNO study on leveraging AI for Systems of Systems [12].

Performance Engineering for Product Families (PEPF)

High-tech equipment manufacturers often deliver tailored systems to customers (trend T2), which are often part of a product family. Developing product families presents significant challenges, including the difficulty of tailoring designs to specific markets using reusable and generic components, understanding how design choices in shared elements impact various system variants and their performance, and managing the potential for reusable assets to become too costly for lower-end market segments while also supporting high-end needs.

Additionally, it is often complex to reason how a particular decomposition of components affects system performance across different designs [6], and some reusable components may deliver suboptimal results. Performance engineering for product families is an important research topic, focusing on questions like:

- PEPF1 ▪ How to design composable architectures, where qualities of the components can be characterized in isolation, and this characterization can be used to assess the resulting system qualities?
- PEPF2 ▪ How to reason about system qualities across system variants in a scalable and effective way, e.g., with composable analysis? How to analyze the performance impact of design choices across system variants? (VAR2 in [6])
- PEPF3 ▪ How can we develop generic, reusable components that perform effectively across diverse operating conditions? (adapted from DC3 in [6])

Performance Engineering for Software-Defined Systems (PESDS)

Software has an important influence on the system performance, as it plays an increasingly important role in realizing system behavior. Software execution architectures therefore play an increasingly important role. They should be scalable and predictable to guarantee overall system performance, especially considering the growing dependence of society on these systems (trend T1). Another important challenge is that software typically evolves much more frequently than hardware. Key research questions on performance engineering of software-defined systems include:

- PESDS1 ▪ How to predict how changes in software components or hardware dimensioning will influence overall system behavior and system qualities?
- PESDS2 ▪ How to design software execution architectures that have a predictable use of shared resources?
- PESDS3 ▪ How to determine the right execution platform design for a system that uses a combination of device, edge, and cloud compute nodes? The decision should consider the resulting system performance, but also other factors like cost, end-of-life risks of the hardware, and development efficiency.
- PESDS4 ▪ How to dynamically coordinate and schedule a heterogeneous set of resources and optimize performance? (e.g., including compute nodes like CPUs, GPUs, and FPGAs)

Performance Engineering for AI-Enabled Systems (PEAS)

AI-enabled systems (trend T3) typically imply the use of data-driven, resource-intensive AI models. It is challenging to reason and balance the accuracy of an AI model's results with the associated computational cost. Another important challenge is the inherent uncertainty associated with AI-enabled systems and guaranteeing system qualities that might be affected by the AI components making autonomous decisions. With AI inside the system, many properties that were under control and could be validated and verified are now potentially unpredictable, making the system more complex. Key research questions on AI-enabled systems include:

- PEAS1 ▪ What is the impact of having AI components in the system on existing methodologies that provide the means to reason on emerging system properties and qualities?
- PEAS2 ▪ What metrics can be used for fine-grained monitoring of AI components?

- PEAS3 ▪ How to ensure predictable performance of AI components, that typically have probabilistic or dynamic non-linear behavior and might evolve by self-adaptation?

5 Related Visions

The System Qualities vision has connections to various other visions and roadmaps that focus on (model-based) systems engineering. In this chapter, we elaborate on the connections to these roadmaps.

5.1 TNO-ESI

This vision connects to various related visions of TNO-ESI. The TNO-ESI vision on Lifecycle Excellence in the High-Tech Equipment Industry [5] links to the system lifecycle focus of the first industry need of this vision. In this vision, the focus is on reasoning on system qualities that give value to their systems in the market. Typically, these are qualities that you can find on the “spec sheet” of the system. In [5], the focus is on operational excellence, targeting qualities like dependability, serviceability, and sustainability. The TNO-ESI vision on Systems Architecting/Systems Engineering [4] connects to our vision to use a model-based and data-based methodologies for systems engineering of high-tech equipment systems. The strategic context of our vision is in line with the trends and drivers in [4]. The TNO-ESI vision on Systems of Systems in the High-Tech Equipment Industry [11] matches with our vision to focus on system-dependent reasoning, including dedicated methodologies for systems of systems.

5.2 Software Engineering Institute

In 2021, the Software Engineering Institute of the Carnegie Mellon University developed a vision [13] on software engineering research & development for the United States for the next 10 to 15 years. This vision contains various related research topics to improve reasoning on qualities. We highlight the most important ones linked to our System Qualities vision:

- A *“theory of composability for model-integrated computing and quality attributes”*, that contains composition rules to create the desired behaviors in terms of both functionality as well as quality attributes, and model-driven engineering technologies to guarantee compositional correctness.
- *“New forms of evidence of quality”* that consider the potential gap between a system’s intended design and the qualities designed into the system. Focus is on developing the means to compare as-implemented designs to as-intended designs as part of continuous integration pipelines.
- *“New quality attributes based on human behavior at scale”* that highlight the trend of traditional qualities like reliability and performance to new quality attributes such as transparency in the context of societal-scale software systems.

5.3 INCOSE 2035 Vision

The International Council on Systems Engineering (INCOSE) defined a systems engineering vision for 2035 [14]. The digital transformation will play a pivotal role, leading to systems engineering being predominantly model-based.

> ***INCOSE 2035 vision on model-based engineering:***
> *"Modeling, simulation, analysis, and visualization of system designs and end-to-end solutions enabled by high fidelity digital representations will dominate the practices of all engineering disciplines."* [14, p. 6].

We share the same outlook in our vision. The INCOSE vision also highlights the increased focus on autonomous systems, and higher expectations on delivering system qualities, which is in line with our vision as well:

> ***INCOSE 2035 vision on increasing expectations in autonomy and delivering qualities:***
> *"As society benefits from advancements in system capabilities, consumers and users continue to expect more from these systems. This includes expectations that systems are more capable, dependable, sustainable, and affordable. They expect systems to be more socially acceptable by considering their impact on society and the environment. Users also expect systems to be more autonomous, enabling them to seamlessly interact, and understand and respond to their requests."* [14, p. v]

5.4 Model-Driven System Performance Engineering Vision

TNO-ESI, together with Eindhoven University of Technology, defined a vision on model-based system performance engineering [6]. In this vision, the focus is on timing-related aspects, where system performance refers to the amount of useful work done by a system, which is measured in production speed of products of a predefined quality. In our System Qualities vision, we take a broader definition of system performance, also covering other system qualities. Next to that, our vision also considers a system-of-systems setting and incorporates the potential value of AI to design and optimize systems. Many of the industrial challenges articulated in [6] link to research topics described in our vision, where [6] also presents various solution directions to address these challenges.

5.5 HTSM Roadmap

The ESI strategic research agenda is well-aligned with the Systems Engineering roadmap of the Dutch Top Sector High Tech Systems and Materials (HTSM SE roadmap). Like our System Qualities vision, the HTSM SE roadmap is also focusing on the Dutch high-tech equipment industry, and as such has a lot of common ground. Some of their key methodological challenges (in italic) relate closely to our System Qualities vision:

- *Virtual, model-based development (models, tooling, higher abstraction levels, languages, …).*
- *Systems thinking, including bridging multiple engineering disciplines and scalable modeling details.*
- *Architectures for data-intensive and/or AI-intensive systems (of systems), including IoT systems.*
- *Design for system (of systems) qualities, including safety, security, performance, install-ability, diagnosability, sustainability, re-use.*
- *Autonomous, self-organizing, self-learning (system of) systems.*
- *Design for customer adaptations, incl. adaptability, configurability, flexibility.*

Various scientific challenges (in italics) identified by the HTSM SE roadmap also relate closely to research topics (in bold) addressed in our vision:

- **Quantitative modeling and reasoning:**
  - *Systematic analysis of and design for key system values (e.g. performance, reliability, robustness, security, serviceability, energy).*
  - *Design space exploration (e.g. for power consumption, performance, reliability, adaptability) and quantified cross-domain design optimization.*
- **Runtime monitoring and optimization:**
  - *Adaptivity and run-time optimization.*
- **Performance engineering for software-defined systems:**
  - *Resource planning and scheduling (including multi-criticality, heterogeneous platforms, multicore, software portability).*
  - *Exploiting hybrid compute platforms, including efficient software portability.*

5.6 HTSC Research Propositions 2025

The Research Propositions 2025 vision [15] of the High-Tech Systems Center (HTSC), part of Eindhoven University of Technology, has several research topics that relate to our vision:

- “Hyper-automation for embedded and cyber-physical systems” links to our research topic on **quantitative modeling and reasoning**.
- “Extreme value budgeting: modeling, sampling, and design” links to our research topic of **reasoning with uncertainty**, focusing on using stochastic modeling to reason on the link between design parameters and error probabilities.
- “Explainable AI for certified AI” links to our research topic on **performance engineering for AI-enabled systems,** focusing on designing explainable AI frameworks that consider stakeholder requirements and can adapt to limited computational resources being available.
- “AI-assisted optimization of high-tech production systems” links to our research topic on **multi-objective optimization,** focusing on using AI to compute optimal production system designs.
- “Compositional model order reduction for interconnected systems” links to our research topic of **performance engineering for product families**, focusing on how to reason on composable architectures.

6 Conclusion

In this report, we have introduced our vision on how system engineers can stay in control of system qualities, designing systems that bring value to the market. To build our vision, we first considered the trends and drivers that are shaping the world of today. Based on these developments, we identified three key needs that should be addressed for the Dutch high-tech equipment industry to stay competitive:

1. Analysis, prediction, and optimization of system qualities over the full lifecycle
2. Systematic reasoning on multiple system qualities
3. System-dependent analysis, prediction, and optimization of system qualities

For each need, we identified strategic research directions that elaborate on the research questions that need to be answered. Building solutions in these directions is crucial to ensure that the Dutch high-tech equipment industry can keep delivering systems with excellent operational effectiveness.

This vision document complements the other vision documents of the TNO-ESI program. For readers from industry, it serves as starting point to discuss their needs and see how fundamental or applied research might contribute to solving those needs. For readers in academia and applied research institutes, it serves as inspiration for relevant research directions that play a key role in the high-tech domain. As TNO-ESI, we invite you to discuss your challenges, needs, research results, and ideas for methodological improvement, such that we can keep bringing industry partners together, driving applied research on common challenges, and helping companies adopt fundamental research results.

7 References

[1] D. Verma and B. Galloi, "System Design and Operational Effectiveness (SDOE): Interface Between System Developers and Users," in *International Conference on Engineering Design (ICED)*, Glasgow, 2001.

[2] INCOSE, Systems Engineering Handbook - A Guide for System Life Cycle Processes and Activities, San Diego: INCOSE, 2023.

[3] ISO/IEC, " ISO/IEC 25002:2024 Systems and software engineering — Systems and software Quality Requirements and Evaluation (SQuaRE) — Quality model overview and usage," ISO, 2024.

[4] S. Acur and T. Hendriks, "TNO report 2024 R10542 - PMC SA/SE Vision and Outlook for Systems Architecting and Systems Engineering in the High-Tech Equipment Industry," TNO, Eindhoven, 2024.

[5] L. Barbini, D. Bera and J. Marincic, "TNO report 2025 R13007 - Vision and Outlook for Lifecycle Excellence in the High-Tech Equipment Industry," TNO, 2025.

[6] B. van der Sanden, Y. Li, J. van den Aker, B. Akesson, T. Bijlsma, M. Hendriks, K. Triantafyllidis, J. Verriet, J. Voeten and T. Basten, "Model-Driven System-Performance Engineering for Cyber-Physical Systems," in *Proceedings of the 2021 International Conference on Embedded Software*, 2021.

[7] J. Kephart and D. Chess, "The vision of autonomic computing," *Computer,* vol. 36, no. 1, pp. 41-50, 2003.

[8] F. Osinga, Science, Strategy and War - The Stategic Theory of John Boyd, London: Taylor & Francis Group, 2007.

[9] S. Zilberstein, "Using anytime algorithms in intelligent systems," *AI magazine,* vol. 17, no. 3, 1996.

[10] J. Verriet, "TNO report 2024 P11832 - Multi-objective analysis of manufacturing systems," TNO, Eindhoven, 2024.

[11] J. Tretmans and J. Lukkien, "TNO report 2025 R11685 - Vision and Outlook on Systems-of-Systems in the High-Tech Equipment Industry," TNO, Eindhoven, 2025.

[12] S. Shirzadehhajimahmood and J. Tretmans, "TNO report 2025 R10133 AI in Systems of Systems," TNO, 2025.

[13] A. Carleton, M. Klein, J. Robert, E. Harper, R. Cunningham, D. de Niz, J. Foreman, J. Goodenough, J. Herbsleb, I. Ozkaya, D. Schmidt and F. Shull, Architecting the Future of Software Engineering: A National Agenda for Software Engineering Research & Development, Pittsburgh, PA: Software Engineering Institute, 2021.

[14] INCOSE, "Systems Engineering Vision 2035 - Engineering Solutions for a Better World," INCOSE, 2022.

[15] High-Tech Systems Center, "Research Propositions 2025: Pushing the limits of high-tech equipment," Eindhoven University of Technology, 2025.

ICT, Strategy & Policy

High Tech Campus 25
5656 AE Eindhoven
www.tno.nl

innovation
for life