ENGD PROJECTS 2022

Software Technology

TU/e

EINDHOVEN UNIVERSITY OF TECHNOLOGY

4TU.School for Technological Design, Stan Ackermans Institute offers 2-year postgraduate technological designer programmes and tracks. This institute is a joint initiative of the four technological universities of the Netherlands: Delft University of Technology, Eindhoven University of Technology, University of Twente and University of Wageningen. For more information please visit: 4tu.nl/en/

Software Technology - EngD Projects 2022

PROF.DR. M.G.J. VAN DEN BRAND
DR. Y. DAJSUREN ENGD

Data Science Strengthens Model-Driven Software Engineering

Software has proven and still proves to be a key technology for innovations in our society in the last decades. Software creates opportunities to improve the quality of our lives, including areas such as communication and transportation. Our society would have stopped functioning during the COVID-19 era if we would not have been able to use all kinds of ICT- based solutions, for example, to maintain social connections, enable virtual education, and work remotely. We were able to continue our "normal" way of living and working thanks to software solutions including video conferencing and data sharing. Current innovations are primarily enabled by data and software, not only for the high-tech systems in the Eindhoven region, but also in the agriculture and food sector. Modern systems produce vast amounts of data. By analysing the generated data, we can improve performance, predict maintenance, and adapt these systems for future needs. Model-driven software engineering and data science evolve hand-in-hand and fit perfectly together. Based on the results of machine learning, existing models can be fine-tuned, or new models can be derived. By creating digital twins of complex (engineered) systems and using the data of the systems, we accelerate design, maintenance, and optimize the production processes.

In recent years the importance of data and AI has been accelerating, and this has also become visible in the ST training program. In addition to the traditional software engineering-oriented workshops, the program offers a workshop on machine learning. We have already observed for a few years an increase in final project assignments from industry on data and AI. Another trend we have observed is the Cloud-based solutions and the architectural consequences.

In this booklet, you will find project examples of the current generation of EngD trainees on how model-driven software development techniques are combined with data science and AI techniques. An example is the application of machine learning models to replace conventional approaches in the metrology tool of ASML. The proposed solution offers outlier detection capabilities by separating the effect of focus on the printed wafer from the effect of other possible configurations and factors such as the type of the tool. Another example is from the life-science domain, where a data pipeline solution was developed to handle the data processing tasks, which includes ingesting vast quantities of data from distinct sources, identifying and resolving any data-related problems, configuring data storage, automating the entire data transfer flow, and facilitating access to the High Performance Computing (HPC) Services for AI model development. Philips Research is interested in the validation of automated Neural Architecture Search (NAS) techniques for their machine learning research and development projects. NAS approaches were studied and evaluated on research datasets as well as on a real-world problem, namely automatic fetal anatomy detection using ultrasound images. These solutions show how software technology facilitates AI-based solutions.

The development of Cloud-based solutions was also a recurring theme this year. In cooperation with Philips Healthcare, a suite of Cloud applications was designed and developed that enabled the clinical study to be conducted seamlessly. The results of this project show many improvements to the system currently in place and a potential way for Philips to carry on in future development of tools concerning clinical research. The goal of a project with Signify was to build a Cloud-based infrastructure with push notification capabilities. Signify needed a new system with a future-proof architecture that can process event streams, filter events based on defined rules, store them in an event log, and send them as push notifications.

Finally, you will find projects with a more "traditional" flavour. For instance, another project with Signify dealt with a tool that was developed to allow the creation of light scripts for light designers based on specific songs. Another more traditional project was related to finding technical debt and refactoring opportunities in the TWINSCAN software of ASML. It is impossible to summarize all final projects in one paragraph; therefore, we invite the reader of this booklet to go through the project descriptions.

Before finalizing this foreword, we would like to express our admiration to all our trainees who have once again shown that they are able to tackle tough design problems and are able come up with innovative solutions that are eagerly adapted by our industrial partners. Again, all trainees have managed to finalize their projects, satisfying customer needs and with good results. The project summaries in the rest of this booklet provide a source of inspiration and show the challenges to advanced software development. We would like to congratulate our trainees on their results and wish them a bright and successful career.

Mark van den Brand, Scientific Director
Yanja Dajsuren, Program Director

EngD Software Technology program
Eindhoven University of Technology

"We met you as a very motivated person that is extremely driven to achieve the best possible results. Before starting this project, you indicated you wanted to challenge yourself in the unknown area of music light scripting and writing C++ code. You were not afraid to jump into the deep, and that I really admired. I knew from the start that given these unknowns, it would not be an easy assignment for you. And although we had our challenges, you managed to solve them all and deliver a great product that will help us in validating the music light script with our end users."

Rob Polman
Signify N.V.

CHALLENGES

The challenge in this project was to design and develop an intensive UI-interactive tool using the Qt framework to innovatively build a tool that can generate a music-tailored light script while considering the entertainment setup of the Hue user. It was even more challenging to render the light effects on the screen while streaming them into the lights considering the subjectivity of the criteria for having a good light script.

RESULTS

MLST was the project result in which the light designer can upload a song and visualize the audio wave to indicate the audio events and translate them to the light effects compatible with Entertainment Development Kit (EDK) library. EDK made it possible to create a pulse effect, modify it, and stream it to the Hue lights. The current workspace can be saved in a JSON file for further development.

BENEFITS

The project result applies the user's creativity to generate a light script that better matches the song. MLST offers several functionalities, including audio wave visualization, timeline concepts with zooming possibility, Bridge connection, and streaming the light effects to the lamps. With the help of these functionalities, the users can design their light script for their ideal entertainment setup created in their Hue App.

SEDIGHEH ARASTEH EngD

Philips Hue Entertainment Music Light Scripting Tool

Signify launched the Hue+Spotify integration to make lights react to Spotify music. The existing implementation is a fully automated algorithm that makes the Hue lights react based on the Spotify-provided metadata. However, there was a need to add manually authored light scripts that matched better with the music to create a more immersive experience.

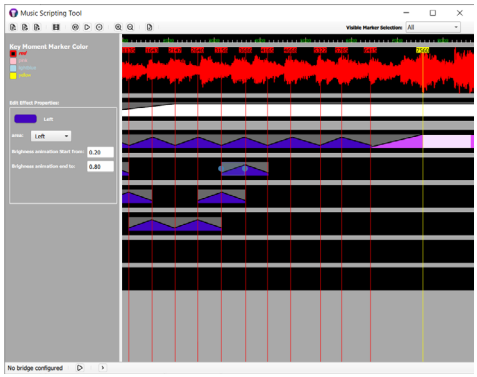

In this project, I developed a tool to create curated light scripts for music manually. This Music Light Scripting Tool (MLST) covers the automatic algorithm drawbacks such as metadata reliability, the lack of human creativity, and the user setup. It allows the light designers to manually generate a light script tailored to each specific song by using the predefined effect libraries. MLST makes the process of creating light scripts easy, employing the power of human creativity by creating high-quality, unique light experiences. In addition to MLST, I designed several scripts for different types of songs that illustrate the power of the MLST and that can be directly deployed to the Hue lights.

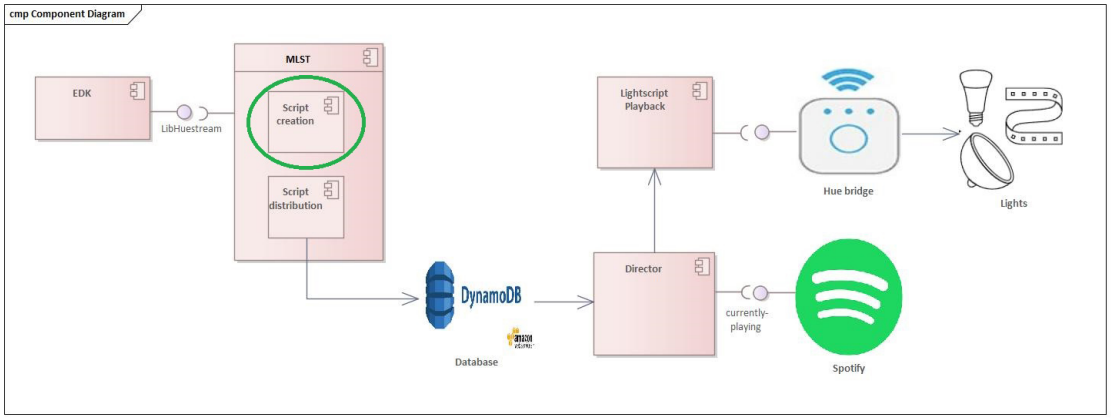

CHALLENGES

The primary design challenges for this project were:

- Develop a solution that can handle large time-based data volumes
- Develop a solution that can batch orchestrate jobs on the cloud for running algorithms
- Develop a solution that caters to clinical workflow

Furthermore, the solution needed to follow the FHIR HL7 standards for exchanging information and data elements between different systems. Additionally, the solution should also allow customers of Philips to evaluate technology before onboarding.

RESULTS

The primary result was the design and prototype of the entire system which caters to clinical workflow while allowing the customers of Philips to have an evaluation environment. The system also provides a new concept for the future called, Analytics as a Service, through the means of running algorithms on the cloud that abstracts the implementation from its users. The system is also capable of handling large time-based data volumes without compromising on performance.

BENEFITS

The benefits of the developed solution are:

- A modern way of conducting clinical research without hassles such as managing patients, devices, personal files.
- An evaluation environment for the customers of Philips, which enables the users to test Philip's technology while still abstracting the implementation of the algorithms with the help of Analytics as a Service.

"Akash has greatly contributed to the development of the user interface portal of Philips' remote monitoring kit, as well as to the architectural design of the system. The kit will significantly facilitate clinical studies and the evaluation of remote monitoring technology, eventually increasing the quality of life for patients. Akash's hard work has substantially shortened the time path to when the monitoring kit will be ready for use."

Ir. Paul Dillen
Royal Philips N.V.

AKASH ARORA EngD

Outpatient Monitoring Kit for Clinical Research

Philips is a global leader in personal health, connected care, and diagnosis and treatment. Philips' healthcare devices work with other software and hardware products to create a complete set of products. Many hospitals worldwide use Philips products to provide the best healthcare for patients. A lot of companies apply Philips' technology in their products as well. During the COVID-19 pandemic, a lot of people were asked to stay home and quarantine without direct access to the hospital care. A healthcare provider could not determine whether the patient needed to be hospitalized or not. Philips is trying to solve the problem with their Outpatient Monitoring Kit, and for the first steps wants to apply it for clinical research. The goal of this project was to design and develop a solution that enables easy management of clinical studies while also providing an evaluation environment for the customers of Philips.

An outpatient monitoring kit consists of wearable devices, a communications hub, and a cloud infrastructure. The cloud infrastructure includes a GUI dashboard and multiple backend microservices. The backend microservices are for different purposes, such as storing time-based data, storing patient data, running algorithms on the cloud, and monitoring and logging for all microservices.

We designed and developed a suite of cloud applications that can enable the clinical study to be conducted seamlessly. The results of this project show many improvements to the system currently in place and a potential way for Philips to carry on in future development of tools concerning clinical research. For example, the Barista system took 15+ seconds to render a single visualization for a metric but the developed system takes 4-5 seconds to render about 5 visualization metric graphs.

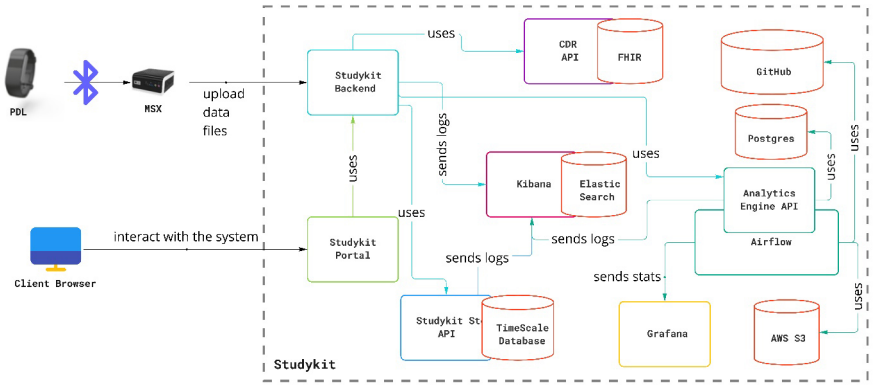

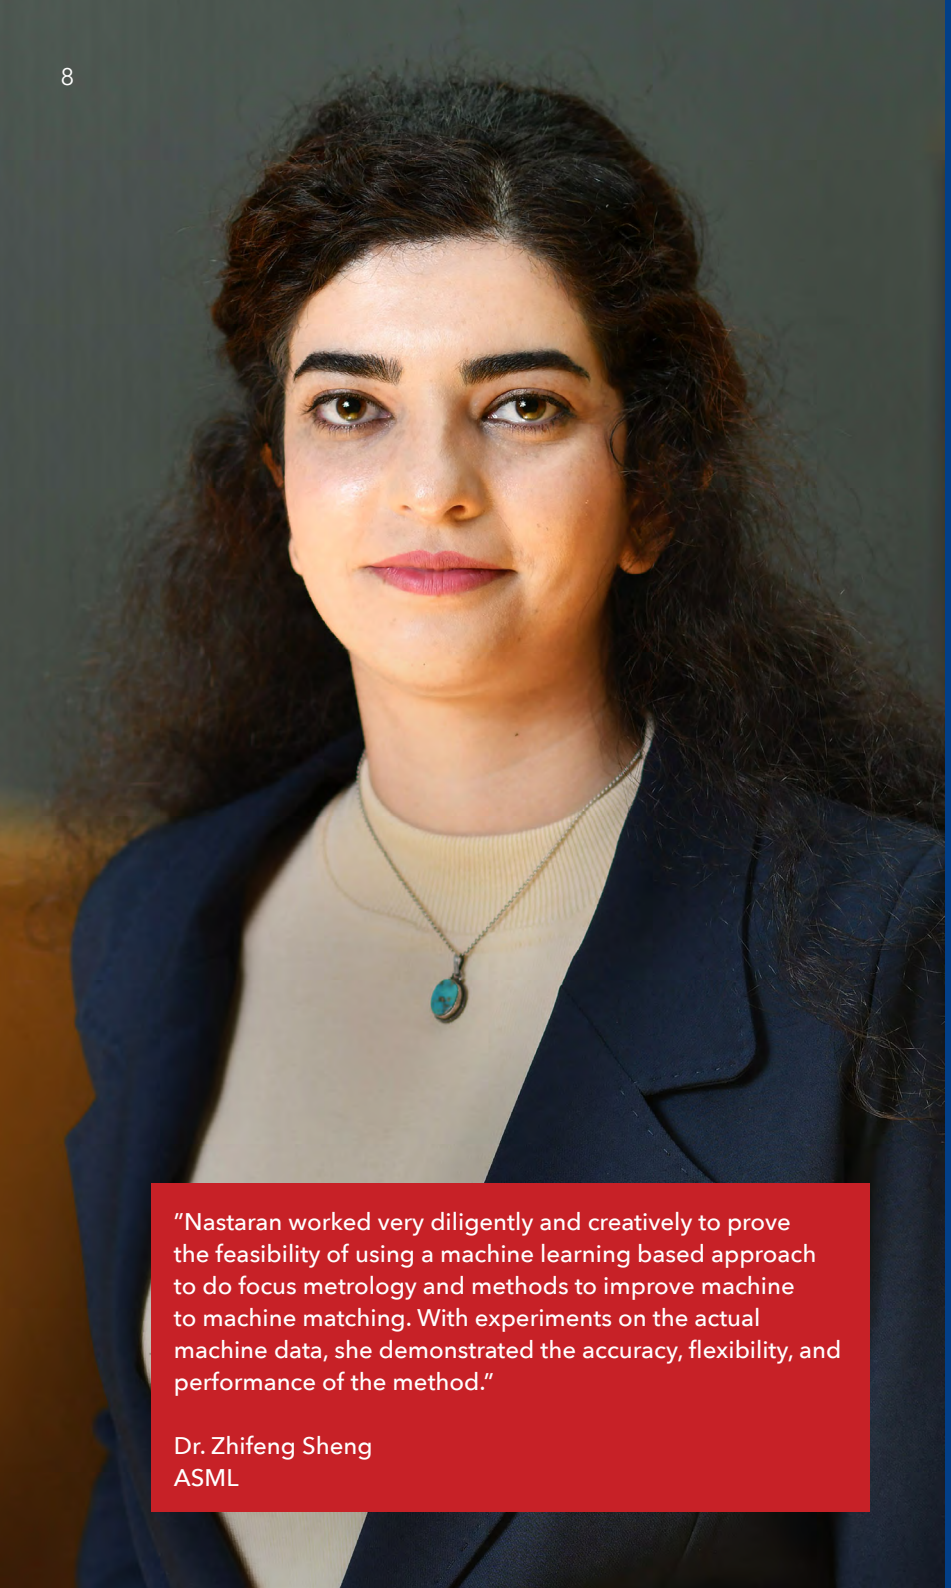

"Nastaran worked very diligently and creatively to prove the feasibility of using a machine learning based approach to do focus metrology and methods to improve machine to machine matching. With experiments on the actual machine data, she demonstrated the accuracy, flexibility, and performance of the method."

Dr. Zhifeng Sheng
ASML

CHALLENGES

The main challenge was designing a generalizable solution. For a tool to be validated, YS requires this tool to produce the same results on measurements from other tools. This challenge involved devising a way to separate the influence of different tools on the measurements and only keep the influence of the desired quality metric. Another challenge was designing the tool so that it can be trained using either images or numerical information.

RESULTS

The result of this project is a prototype that contains training, retraining, and inferencing functionalities for the designers. The recommended machine learning models are integrated within this tool. They enable designers to tune parameters and choose different architectures for analyzing the possibility of different models for each type of wafer or combination of wavelengths.

BENEFITS

The recommended machine learning model can identify process variations in datasets, while extracting quality metrics with comparable results to conventional methods. Furthermore, users can directly perform calculations on wafers from different lithography process configurations. The tool could retrain the pre-trained models to make them more powerful over time with the arrival of new datasets.

NASTARAM BAJALAN EngD

Feasibility and Prototype of Data-Driven Focus Extraction

ASML ensures reliable lithography with the help of its metrology machines. YieldStar (YS) is a metrology tool that provides closed-loop feedback to scanners by measuring on-product errors such as overlay and focus. Focus is one of the metrics that YS calculates to determine how well patterns are printed on the wafers. Currently, YS uses conventional regression approaches to extract focus from printed wafers. Within this project, the feasibility of replacing the current tool with Machine Learning (ML) models was investigated. Although replacing the current tool requires a great amount of validation and investigation, the accuracy, flexibility, and generalization properties of the current tool were optimized using ML models. The conventional approach cannot find process variations in form of outliers in the datasets. However, the proposed solution offers outlier detection capabilities by separating the effect of focus on the printed wafer from the effect of other possible configurations and factors such as the type of tool.

Furthermore, YS has established a set of rules to ensure the validity of the approaches on different tools. Therefore, the proposed approaches were also validated against these so-called matching rules.

This project investigated different ML-based regression models. These models could generate comparable KPI and matching results to the conventional tool. In addition to the KPIs, outlier-detection characteristics were added to the models using Autoencoders.

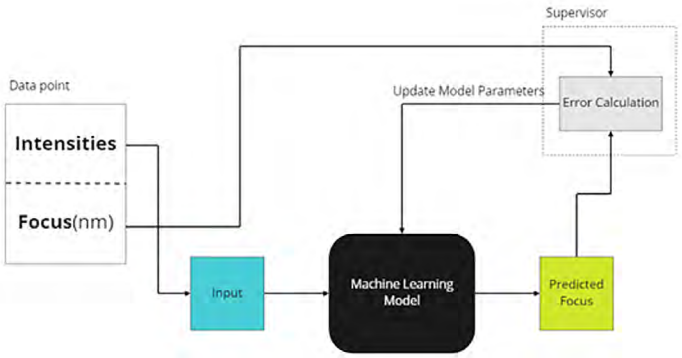

CHALLENGES

As a commercial software company, TomTom has little experience in safety-by-design. The project team had to slowly discover and understand how ISO26262 can be integrated in the TomTom software development lifecycle. One major challenge was the configuration of static code analysis tools for the proof-of-concept and company code, which took far longer than expected, due to the company being new to safety coding and configuration practices.

RESULTS

The development work done on the proof-of-concept helped estimate the cost of safety-by-design at 20-25% of the development process. This cost was reached after the required static code analysis and testing configuration was done and the safety processes were adopted and optimized. The advantages outweigh the costs, as safety provides access to the automotive market, shields the company from financial and judicial risk, and provides better customer experience.

BENEFITS

TomTom can now experiment with the safety-by-design process provided. Safety-by-design helps produce a more reliable and higher quality system. It improves the design and implementation skills of engineers, as they learn robust design and coding principles. It contributed to the step of entering and thriving in the automotive market, with customers and car manufacturers confident in TomTom solutions.

"It was a pleasure to supervise Dan during his traineeship. He quickly picked up on the complex subject of safety; the concepts of the ISO 26262 standard, how to apply it, and building a prototype. Along the way he got familiar with the TomTom way of making software and its complexities. Hurdles were taken or worked around. He was flexible in adapting his plan and he continuously communicated with relevant people within the organization. His work brought the TomTom organization one step closer to informed decision-making regarding safety."

Ir. Guus Holshuijsen
Solution Architect - TomTom

DAN-CRISTIAN CHIRASCU EngD

Safety-by-Design in Architecture of Automotive Software Systems

With the advent of "software-defined vehicles," Advanced Driver Assistance System (ADAS) functions are being integrated with In-Vehicle Infotainment (IVI), to provide situational awareness coupled with great user experience. TomTom is a leading geolocation and driving assistance company seeking to add lane-level guidance to its range of services, as part of its revolutionary IndiGo digital cockpit platform.

The automotive space has strong safety requirements for mitigating or preventing hazardous events caused by malfunctions. For example, a blind-spot highlighting failure may lead to the IVI system not reporting a vehicle in proximity. The driver may then dangerously steer towards it. Safety-by-design is required for such lane-level guidance applications to function reliably and warn the user of any failures. TomTom explored what it takes to enter the automotive space in terms of safety, and how safety can be integrated in software development.

In this project, we prototyped a blind-spot highlighting system, analyzed possible failures, and implemented the corresponding safety mechanism. We analyzed the code to spot any vulnerabilities. The results obtained show that it is relatively easy to meet the safety requirements for TomTom software products, given the many advantages safety-by-design provides. Technically, safety practices help produce a more reliable and higher quality system. The proposed system improves the design and implementation skills of engineers, as they learn robust design and coding principles.

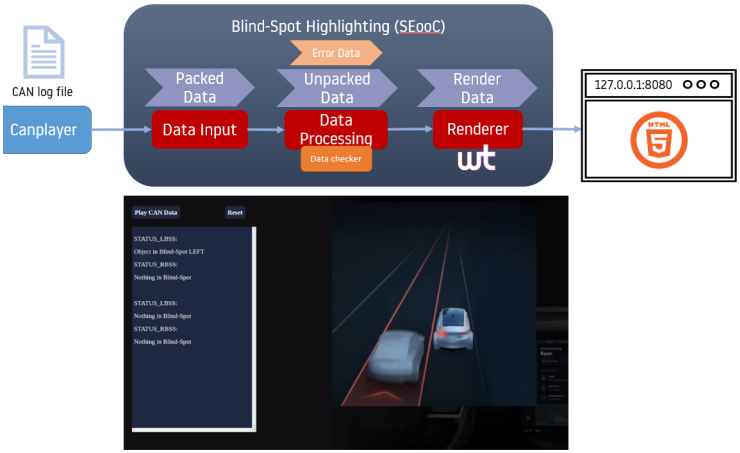

"In this project, Janice has taken a deep dive into the multidisciplinary ink handling domain and its associated design process. The resulting domain-specific modeling tool enables us to specify the ink handling topology concisely. Capturing the ink handling topology in models is an essential step towards easily creating Digital Twins of our ink handling component. We are looking forward to continue the path that Janice has started."

Ir. Eugen Schindler EngD
Dr.ir. Joost van Pinxten
Canon Production Printing Netherlands

CHALLENGES

The first challenge was to find the right level of abstraction and presentation of the components in the topology, so that they can both be understood by the experts, as well as deriving the simulations. The information has to be both concise and complete to function as an effective model.
The second challenge was to derive the model automatically and integrate it into the existing embedded software.

RESULTS

The result is a tool that allows one source of truth for the ink handling topology that generates code for different simulation tools. The tool allows automatic derivation from the designers' topology model to the designers' and embedded software engineers' simulation models.

BENEFITS

The tool developed in this project facilitates the collaboration between experts, reducing the need for iteration cycles. It can be expanded to other aspects that require information from the ink handling functionality, such as specifying diagnostic models. The tool decreases the time spent on manually creating models for simulation and visualizing the results. It is an essential step towards creating a full-fledged digital twin for ink handling.

JANICE CONCUET EngD

Designing a Digital Twin Factory for Ink Handling in Advanced Production Printers

Canon Production Printing is an international leader in high-tech production printing solutions, an important part of which are print engines. The print engines have at their core print processes, which include the multidisciplinary ink handling component. Canon Production Printing wants to easily create Digital Twins of their components. This project contributes to that vision by capturing part of the ink handling functionality in a domain-specific modeling tool.

The domain-specific models for the topology of the ink handling (including pipes, pumps, and print heads) replace models that were previously created in general-purpose tools such as Visio. In this project, a prototype graphical modeling tool was created in JetBrains MPS that allows the designers to specify the ink handling topology. The prototype can automatically derive Matlab Simscape models for exploring design alternatives, and Software-In-the-Loop models for evaluating the ink handling behavior in the embedded software. The automatic derivation ensures that the derived models are consistent and up to date with the specified ink handling topology.

By specifying the model in the specific jargon of ink handling, it also improves the communication between the domain and implementation experts. The captured topology and its derived simulation models will enable Canon Production Printing to also model the ink handling behavior, its failure modes, and its recovery scenarios, effectively creating a Digital Twin Factory for the ink handling component.

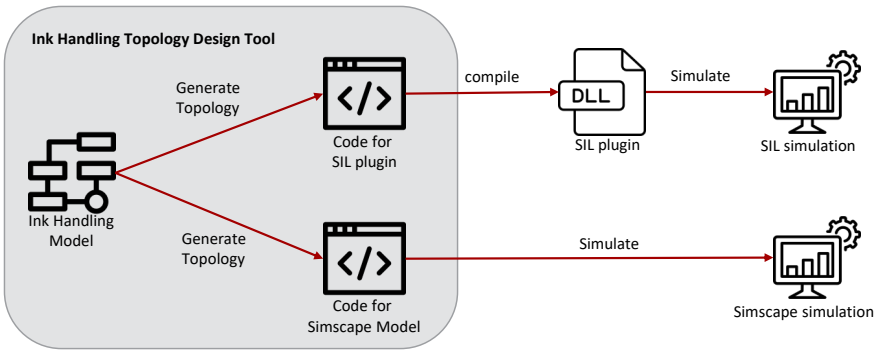

CHALLENGES

The main technical challenge was to design a system that can fulfill the specified performance requirements. The system should be able to process a large number of events per second, while the actual number was expected to differ drastically between different components of the system. Additionally, there were many challenges in managing the project as it feeds directly into Signify's product.

RESULTS

We successfully delivered a future-proof event-driven microservice architecture that fulfills our performance requirements. Using microservices, we can scale different components of the system up and down independently depending on their individual needs.

BENEFITS

The event log and push notification system and its flexible and scalable architecture will allow Signify to easily create new features in the future. These features can make use of push notifications to communicate important information to users about their Philips Hue system or new Signify offerings.

"It was amazing to see how Christian managed to adapt to our way of working, quickly learning about the Philips Hue system and Cloud infrastructure, while maintaining extraordinarily high standards in the work he delivered for the organization. Overall, we were all very happy to have had Christian onboard, helping us shape the future of the Philips Hue platform."

Nuno Ferreira
Signify

CHRISTIAN DEGOTT EngD

Hue Push Notifications

EVENT LOGS AND PUSH NOTIFICATIONS FOR PHILIPS HUE

Signify's Philips Hue brand is the leading smart home platform for lighting. It provides smart lighting to millions of customers world-wide. Customers can create an ambience that fits with their activities, use light recipes that provide health benefits, and control their lighting through local tangible controls and mobile applications. So far, Signify's own mobile application, the Philips Hue app, has not had any push notification capabilities.

The goal of this project was to change that and to build the necessary Cloud infrastructure supporting that. To achieve this goal, Signify needed a new system with a future-proof architecture that can process event streams, filter events based on defined rules, store them in an event log, and send them as push notifications.

Our event-driven microservice architecture is scalable and able to process event streams with large numbers of events per second. It is extendable and will allow Signify to support various use cases that send push notifications to communicate information to users.

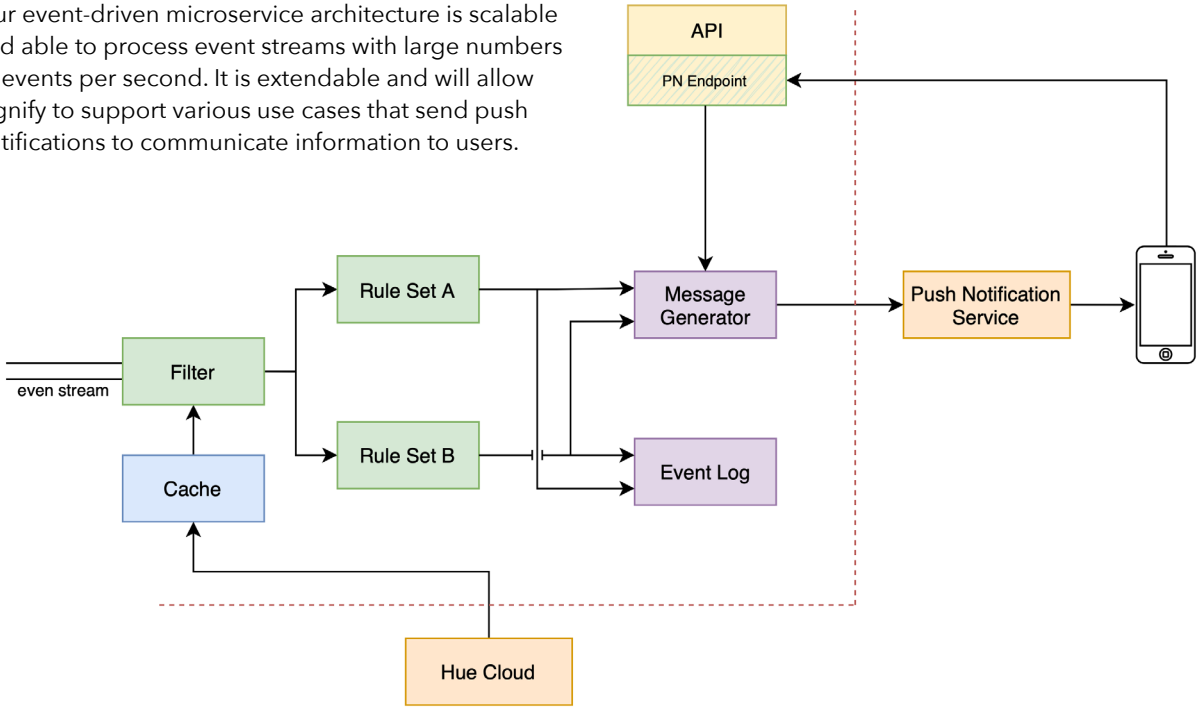

"Mohammad Ibrahim worked together with Hossain Muctadir on the development of model management tooling. He developed a prototype for transforming Simulink and SysML models into graphs in order to be able to detect (in)consistencies. It was a challenging endeavour for Mohammad. He enthusiastically worked on the assignment and delivered interesting and usable results."

Prof.Dr. M.G.J. van den Brand,
Software Engineering & Technology Cluster, TU/e

CHALLENGES

To remain relevant in a highly competitive marketplace, modern high-tech systems are becoming increasingly complex. Model-Based Systems Engineering (MBSE) has enabled engineers to design and develop such complex systems more efficiently and effectively. MBSE is being used to model and design systems across a variety of industries from aerospace to healthcare, making it a critical tool for modern engineering. However, managing these models and understanding the dependencies among them can be a challenging task, particularly when dealing with multiple tools and platforms.

RESULTS

We designed and developed a model management tool that extracts relationships and dependencies among multi-tool models and stores them using a graph database. Moreover, we designed a common graph data model that presents models and their relationships. This tool can be used to query more insights and identify inconsistencies among dependent models.

BENEFITS

This model management tool demonstrates the possibilities of a graph-based approach and how it can be used in managing consistencies among complex models. Also, it helps design engineers to get more insights about the system they are working with by querying information from the graph database.

MOHAMMAD IBRAHIM EngD

A Graph Database Design for Multi-Domain Model Management

Model-Based Systems Engineering (MBSE) has become increasingly popular in recent years as a way to improve the efficiency and effectiveness of system design and development. With its ability to provide a more comprehensive view of a system, MBSE helps design engineers to design and develop complex systems more efficiently, while also improving system quality and reducing costs. However, ensuring consistency among models is a complex task in the context of model management. It becomes even more complicated in a multi-tool and multidisciplinary setting.

The goal of this project was to develop a model management tool to ensure consistency among the models. To address the project's goal, we firstly investigated how to extract models' data that are developed in different modeling tools, in our case Simulink and Rhapsody SysML models. Secondly, we developed a parser for retrieving model information and identifiing essential elements and their properties. Thirdly, we designed a graph database for storing model information with a graph data model (meta ontology). Fourthly, we developed a data loader to store extracted data in the graph database. Finally, we developed a dashboard to visualize data from the graph database.

This tool can help design engineers to better manage and analyze complex models for detecting inconsistencies in the early stage of the development.

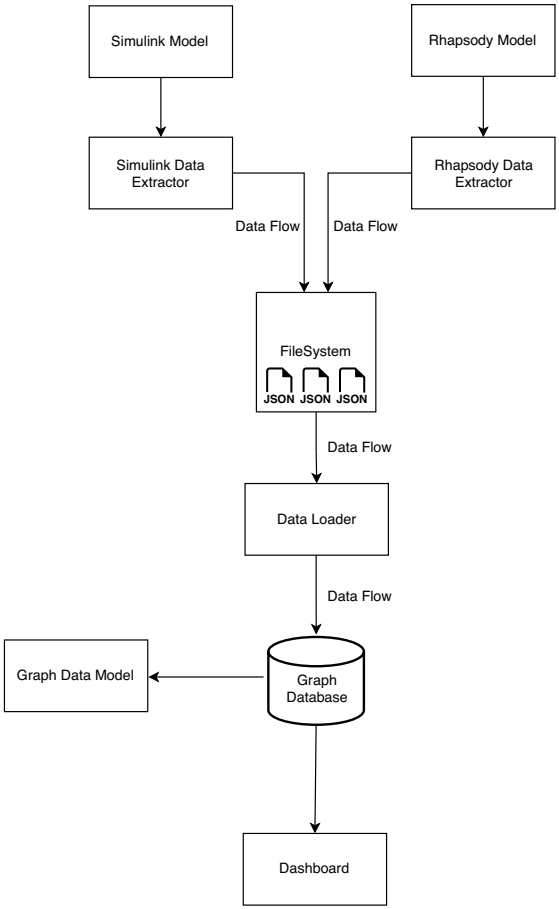

CHALLENGES

The primary obstacle in the project involved the ingestion of vast volumes of data from multiple sources. Additionally, we needed to employ preprocessing techniques to identify and resolve any data anomalies. Once the cleaned data was ready, it needed to be migrated to a permanent storage system to enable access to the High Performance Computing (HPC) Service for developing AI models.

RESULTS

We developed a Data Pipeline solution to handle the data processing tasks, which include ingesting vast quantities of data from three distinct sources, identifying and resolving any data-related problems, configuring data storage, automating the entire data transfer flow, and facilitating access to the HPC Services for AI model development.

BENEFITS

Through the implementation of the suggested solution, we can automatically transfer and preprocess the data and store it in a permanent storage location, making it easily accessible to the HPC Services for developing AI models.

"The integrative data pipelines developed in this project are an important step to translate interdisciplinary research activities into real-world data-centric A.I. solutions by linking data and minds."

Prof.dr. Jakob de Vlieg
Chair of the Applied Data Science (ADS)
research group, M&CS, TU/e
Lead AgrifoodTech at TU/e and JADS

HASAN KAPLAN EngD

The Design and Implementation of an Integrated Data Pipeline Solution for Behavioral Animal Science and Animal Breeding

This project was carried out under the IMAGEN and SmartTurkeys projects which aim to investigate scientific approaches and technological advancements for the large-scale automated detection of animal behavior. Specifically, the SmartTurkeys project focused on analyzing turkey behavior, while the IMAGEN program was responsible for analyzing pig and laying hen behavior.

In these projects, the data was gathered from livestock farms and stored using ad hoc techniques. The objective of this EngD project was to develop the architecture of the Data Pipeline and create three separate deployment instances for different sites. This Data Pipeline is responsible for automating the preprocessing and storage of the data, serving as a link between the data source and storage. After the data was stored in a scalable data storage system, it could be effortlessly exported to HPC Services for the development of machine learning models.

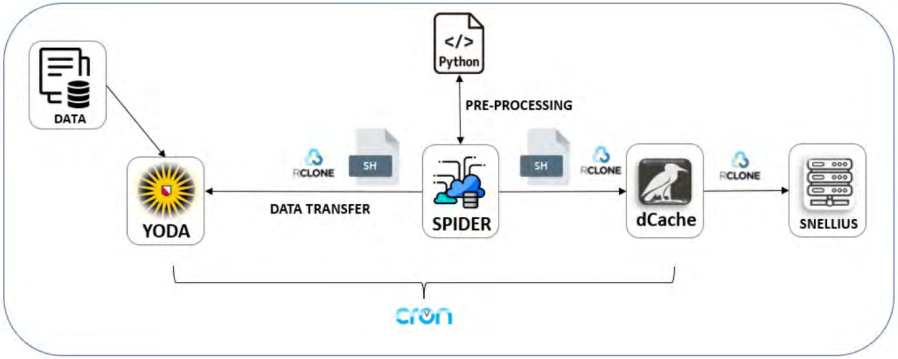

The proposed solution architecture outlines the essential components required for constructing the Data Pipeline. Additionally, this architecture details how to establish a model deployment pipeline. In summary, the solution architecture furnishes researchers with a set of guidelines to create, test, monitor, and deploy the Data Pipeline.

"In this project Lkham Nyambuu, has studied approaches that seek to automate the design and deployment of neural networks. In particular, she focused on a subproblem within Automated-Machine Learning (AutoML) that concerns automated exploration for a high-performance network topology conforming to the task. For this, Lkham had to understand the current state-of-the-art in Neural Architecture Search (NAS), short-list and implement the promising strategies and perform evaluation on research datasets and a real-world use case of fetal head detection in ultrasound images. Lkham did this in a very nice, independent, and well-structured way, for which I want to thank her warmly."

Rameez Ismail EngD
Philips Research

CHALLENGES

The main challenge of this project was to implement Neural Architecture Search (NAS), which automatically creates the optimal neural network architecture for a given dataset. The main use case was detecting the fetal anatomy position from ultrasound images. Moreover, we automatically created the optimal architecture for the CIFAR-10 dataset. Since NAS is an AI that creates smaller AIs, it is a very complex technique. Due to the complexity of the NAS, it was challenging to understand and implement it.

RESULTS

The main result of the project was a framework that is suitable for both image segmentation and classification tasks. The input of the framework is a dataset, and the output is an optimal neural network architecture for the input dataset. When we trained the optimal architecture, we managed to get 91% mIoU on the fetus anatomy detection dataset and 95% accuracy on the CIFAR-10 dataset.

BENEFITS

This project created a framework that automatically creates a novel optimal neural network architecture for a certain dataset using NAS. It has the following values:

- Automation of designing neural network architecture
- Exploration of novel architecture
- Reduction of experimentation time

The main use case, fetal anatomy detection, will help pregnant women to avoid complications during labor in under-resourced communities.

LKHAM NYAMBUU EngD

Automated Neural Architecture Search for Fetal Head Detection

Designing neural network architecture for Deep Learning (DL) is always challenging because there are countless numbers of possibilities for neural architecture. Neural Architecture Search (NAS) automates the manual design process of an artificial neural network and strives to find the best architecture. Therefore, Philips Research is interested in validation of such automation techniques for their machine learning research and development projects. In the scope of this project, we studied NAS approaches and evaluated them on research datasets as well as on a real-world problem: automatic fetal anatomy detection using ultrasound images.

To evaluate the NAS, we used two main datasets. First, the fetal anatomy detection dataset, which was provided by Philips Research, was used to evaluate the functionality of NAS in a Philips use case. More specifically, using this dataset, we aimed to detect the anatomy of a 23-week fetus. Second, the CIFAR-10 dataset, which is an open-source dataset for classification tasks, was used to evaluate the viability of the NAS.

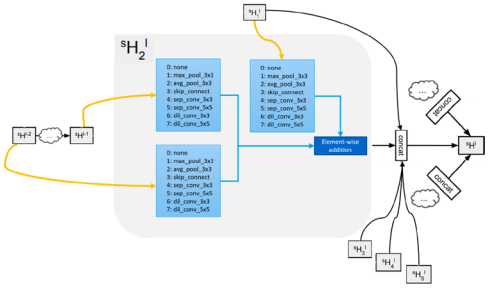

We implemented the NAS framework using the Auto-DeepLab technique, which searches for neural architecture in both cell and network levels. The framework consists of two main stages. First, it searches for the optimal neural architecture for a given dataset. Second, it retrains the optimal architecture, found by the first stage, on the same dataset to get the final result, which represents the accuracy of the model. Overall, our results show that the Auto-DeepLab is able to work on both image segmentation and classification tasks. It could reduce human labor for producing a new architecture for a custom dataset.

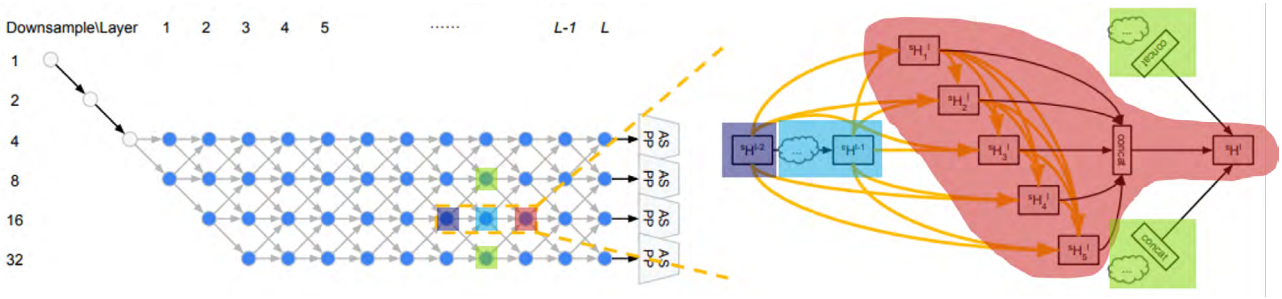

CHALLENGES

ThermoFisher Scientific provides a highly-configurable Infrastructure-as-a-Service to enable the advanced feature of the Electron Microscope (EM). The main challenge was to create a bridge between a highly-configurable infrastructure with an easy-to-use solution that must guide the user to configure the infrastructure. Another challenge was to align the implementation of the solution with the infrastructure that was under continuous development and involved multiple technologies.

RESULTS

We designed and implemented a solution for configuring the managed-service highly-configurable infrastructure. The solution transforms the existing installation script that was written in the procedural approach into an object-oriented style with an easy-to-use Graphical User Interface (GUI) that guides the user throughout the configuration process. This brought an improvement in terms of the usability, maintainability, and extensibility of the infrastructure configurator.

BENEFITS

The solution of the project opens the opportunity for the lesser experienced Service Engineers to execute the infrastructure configuration. By simplifying the infrastructure installation process, the task can be moved to a regular service worker, so the IT specialist can focus on higher level services when supporting ThermoFisher Scientific's customers.

"Respa has done a great job in filling this gap. He mastered the current state of the SDP configuration mechanism quickly and familiarized himself with the functions and limitations of the current system. Respa's solid structured way of working in all parts of the project (requirements definition, architecture and design choices, implementation, verification, and validation) contributed to the very usable result delivered in a timely manner, whilst making friends in the organization along the way."

Ir. Egbert Algra, EngD
ThermoFisher Scientific

RESPA PUTRA EngD

A Solution for Configuring an Infrastructure-as-a-Service

ThermoFisher Scientific is one of the largest suppliers of scientific instruments, reagents and consumables, and software services. The products are widely used in various market segments, such as the Electron Microscope (EM) that helps scientists in various research projects in the Life Science, Material Science, and Semiconductor domains.

The scientists employ the EM to magnify nanoscopic samples into a 3D high-resolution image. During this process, the EMs generate a large amount of data that requires a huge storage capacity and powerful computing resources to process the data. Addressing this concern, ThermoFisher Scientific provides an Infrastructure-as-a-Service (IaaS) called Software Delivery Platform (SDP) that enables the advanced feature of the EM. The SDP is configured by the Service Engineer from the Service Organization on each customer site. When it comes to the SDP installation process, it is challenging due to the limitation of the existing CLI-based configurator that requires an IT skillset to interact with. Therefore, an easy-to-use, intuitive and foolproof solution is needed.

In this project, we designed and implemented a solution that provides an easy-to-use guided configuration management with a Graphical User Interface (GUI). The solution aims to fill the gap between the highly configurable infrastructure and a configurator application that provides an overview, transparency, and guided configuration management. It also opened the opportunity for the lesser experienced engineers to interact with SDP configuration.

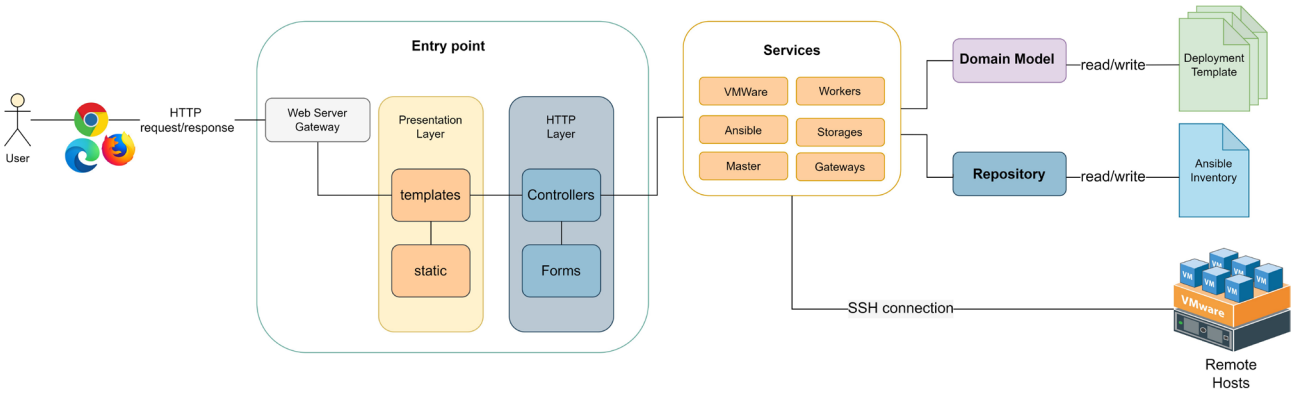

"Eager, friendly, flexible, stakeholder oriented, I started really liking the quality time meetings with Shubham."

Bob Peeters
Royal Philips N.V.

CHALLENGES

Radiology reporting is varied with no universally accepted standards, and hospitals adopt their own conventions. The reporting templates authored for capturing radiology reports are similarly varied. The primary challenge was the design of a data model that accurately captures the source model while supporting consumption by next-generation health applications. The secondary challenge was the development of a methodology to parse and codify templates with low clinical structure.

RESULTS

The output of this project is a tool that lowers the threshold for clinicians and hospitals to migrate existing templates to the HL7® FHIR® standard with minimal effect on the workflow. The output of the tool was verified using industry standard FHIR compliance tools. The models were validated using real-world reporting templates from a medical institution. To show added value, the output of one model was integrated with an in-house health application.

BENEFITS

The benefits of the project include liberation of existing reporting templates into a modern standard that can be consumed by next-generation health applications. The threshold for migration to the standard is lowered. Further, the tooling support to code output models from standard ontologies promotes capture of semantical data.

SHUBHAM RAWAL EngD

Automated Transformation of Picture Archiving and Communication System Reporting Templates to HL7® FHIR® Models

The radiology workflow involves medical professionals utilizing multiple informatics systems with data captured in different standards with low clinical interoperability. The reporting data is often captured as free text, which is not computer-interpretable and has low potential for reuse by other applications. Structured and semantically labelled data facilitates advancement of clinical workflows in next-generation health applications. Reporting templates form a vital part of radiological reporting. Structuring the templates to an open and computer-interpretable standard enables the capture of structured data at the source.

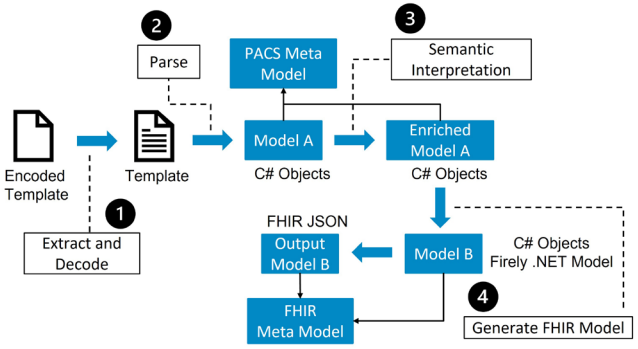

The aim of this project was the automated transformation of reporting templates from a Picture Archiving and Communication System (PACS) into the HL7 FHIR standard. Transformation of existing templates into an open, next-generation standard like FHIR can liberate the data for use in other applications such as clinical dashboards. Further, the tool provides a low-threshold path for migration of existing models to FHIR. To enable reusability, the FHIR models were enriched by automated addition of codes from standard medical ontologies.

The non-functional requirements for the system are support for the conversion of varied reporting templates and standardization. The tool supports reporting templates with standard PACS syntax. The first requirement is addressed by the design of an extensible tool to account for variability in reporting standards across medical institutions. The second requirement is addressed by validation using industry standard FHIR compliance tools.

CHALLENGES

This project's main challenge was to find suitable metrics to show the technical debts in terms of refactoring priority for the TWINSCAN software. Another challenge was to get available data from ClearCase to generate appropriate metrics. Furthermore, we had to design and implement a prototype that calculates the chosen metrics to the software environment of ASML.

RESULTS

For this project, we designed a solution to detect the TWINSCAN software's technical debt. We developed a prototype to see how the TWINSCAN software's code is evolving using the selected metrics.

BENEFITS

The contribution of this project is to give ASML good insight into the refactoring targets of the system. Overall, it provides some valuable information about the coupling of the TWINSCAN software.

"As a project manager, I joined the weekly project meetings in which Lamisha was well-prepared, so we could efficiently do the sessions remotely. It was good to see how she approached modularity, starting from literature, selecting metrics, and visualizing them in a tool. The assignment gave ASML good insights into how to approach and depict modularity."

Ir. William van Houtum
ASML

LAMISHA RAWSHAN EngD

Value/Cost Analysis of Modularity Improvements

ASML is the global leader in providing photolithography systems for the semi-conductor industry to produce integrated circuits. The TWINSCAN software controls the lithography machine. The software has a large code base with millions of lines written in several programming languages. As the software grows, the architects encounter many unwanted dependencies, reducing the code's easy maintainability and extendibility. Hence, the architects think the current codebase's modularity needs to be improved. Technical debt needs to be identified to improve modularity.

The project was initiated to identify the technical debt at the architectural level. In addition, there needed to be a rank to prioritize the debt based on maintenance costs. Therefore, tool support is needed to visualize and prioritize the debt of the TWINSCAN software to help the architects.

In this project, we researched the existing modularity tool and selected four metrics (hotspot, change coupling, complexity trends, and hotspot rank) based on the data available. Applying these four metrics helps show the system's technical debt using metadata of the source code (such as change frequencies, lines of code, and release data) from the version control system. We developed a tool, MoVACA, to present these four metrics in graphical and tabular format for a user-specified scope.

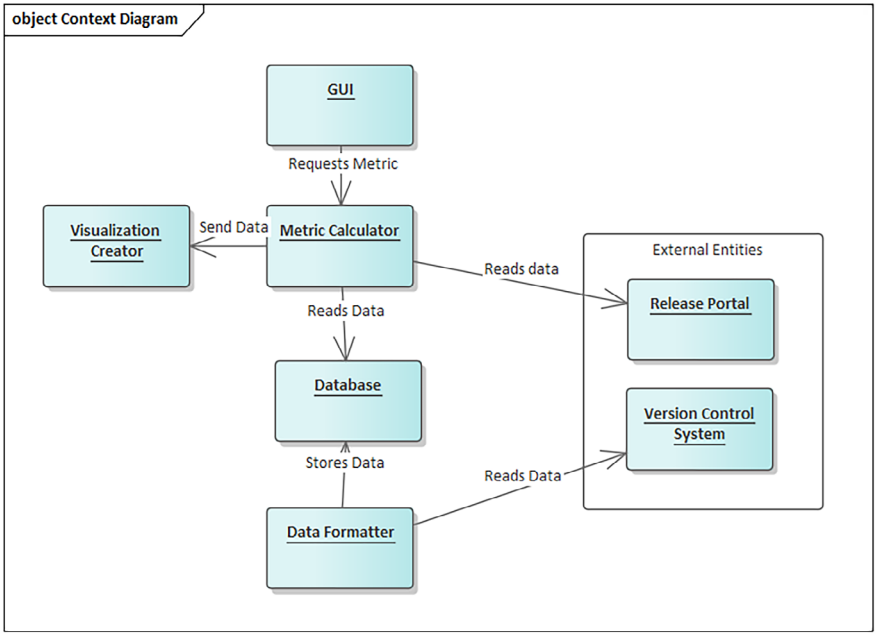

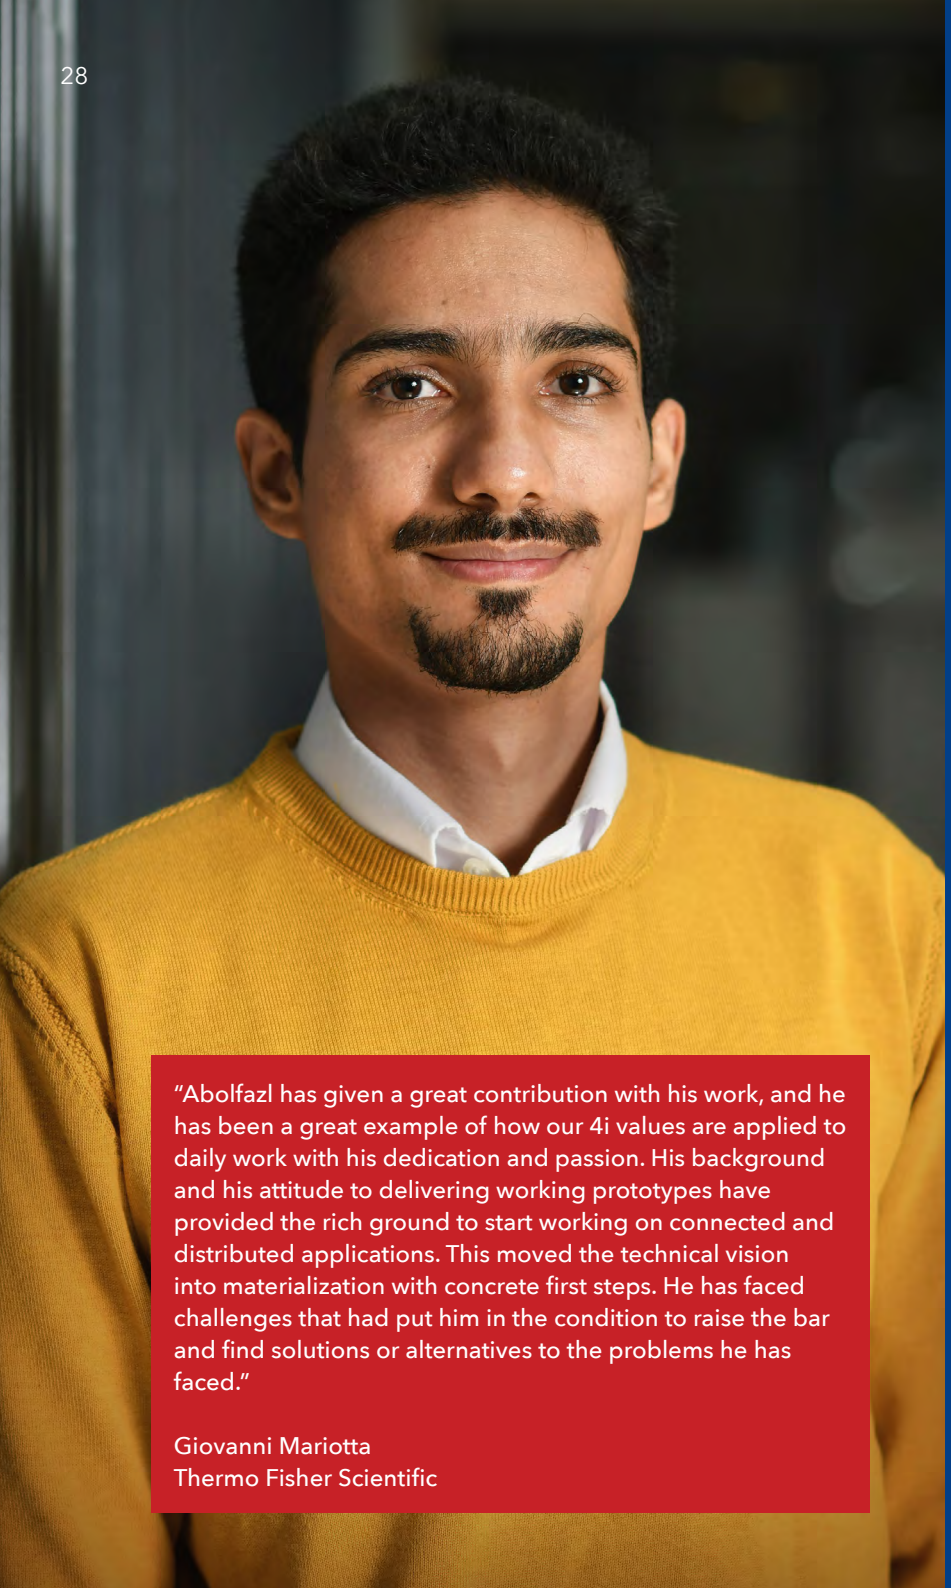

"Abolfazl has given a great contribution with his work, and he has been a great example of how our 4i values are applied to daily work with his dedication and passion. His background and his attitude to delivering working prototypes have provided the rich ground to start working on connected and distributed applications. This moved the technical vision into materialization with concrete first steps. He has faced challenges that had put him in the condition to raise the bar and find solutions or alternatives to the problems he has faced."

Giovanni Mariotta
Thermo Fisher Scientific

CHALLENGES

The main challenge of this project was to find suitable technologies satisfying project needs. After that, this remote layer had to be in place with the current system design with minimal changes to the code of any other layer. Finally, we had to test and build the remote layer alongside the entire project and make sure everything was aligned.

RESULTS

The result of this project was a prototype that showed how a remote layer gRPC can fit inside an MPC to provide remote access to the microscope.

BENEFITS

The remote prototype of the microscope brough insight into how this type of communication can be adapted to the current system. Additionally, it shed light on how a gRPC client-agnostic way of design can help Thermo Fisher Scientific achieve a faster, easier, and cost-effective update approach to its code base. All of these efforts contribute to Thermo Fisher's mission to make the world healthier, cleaner, and safer.

ABOLFAZL SARAVANI EngD

Remote Electron Microscope

Thermo Fisher Scientific supports science by creating lab equipments. One of them is an electron microscope which produces images of samples using electrons. This microscope is supported by an API called TemApp.Omp and a variety of applications that provide access to different parts of the microscope. These applications require direct and physical connection to the microscope. To have this direct link, all these applications run on a Microscope PC (MPC.)

The goal of this project was to investigate the possibility of breaking the boundaries of this local system using a remote layer design. The main requirements of the project were creating a remote server agnostic interface that supports any client, separating the interfaces of clients and server, and creating one code base for development.

Therefore, we created a solution based on a remote communication using Google Remote Procedure Call (gRPC.) The architecture consists of a remote server that connects directly to the microscope interfaces via MPC. Then, clients can access the microscope via a remote connection hidden behind a proxy class.

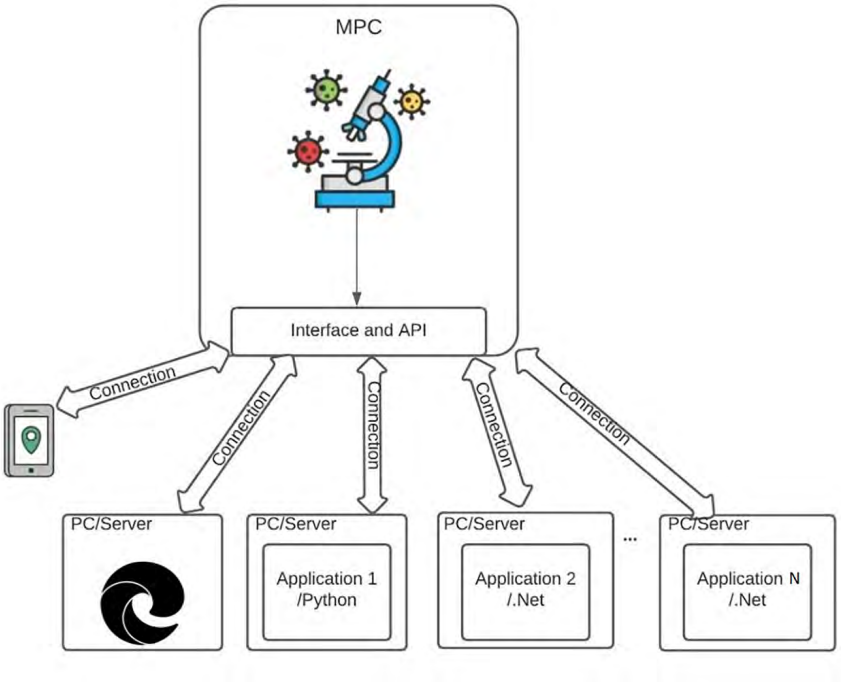

CHALLENGES

To provide an insightful analysis of Philips'road freight shipments, a structured process for cleaning raw data and compensating the missing value was needed. Although there was a system operating for other modes of transport, the existing design was low in efficiency and in level of modularization. The challenges lie in pinpointing the key shortcomings of the current system, identifying common data flaws, and evaluating a new architecture that solves the problems.

RESULTS

The result is a renovated system that integrated the emission calculations for all modes of shipments in Philips' logistics process. Emissions from two of the transport modes, the road transportation and ocean transportation, were calculated using the new system. The KPI numbers, together with emission trend lines and impact route maps, were displayed on a dashboard that allowed the user to manipulate the dashboard freely.

BENEFITS

The architecture provides a solid solution to the establishment of a reliable road freight emission calculation application. Moreover, the brand new design can be applied to other modes of transport. The flexibility and the level of abstraction are beneficial for further improvements.

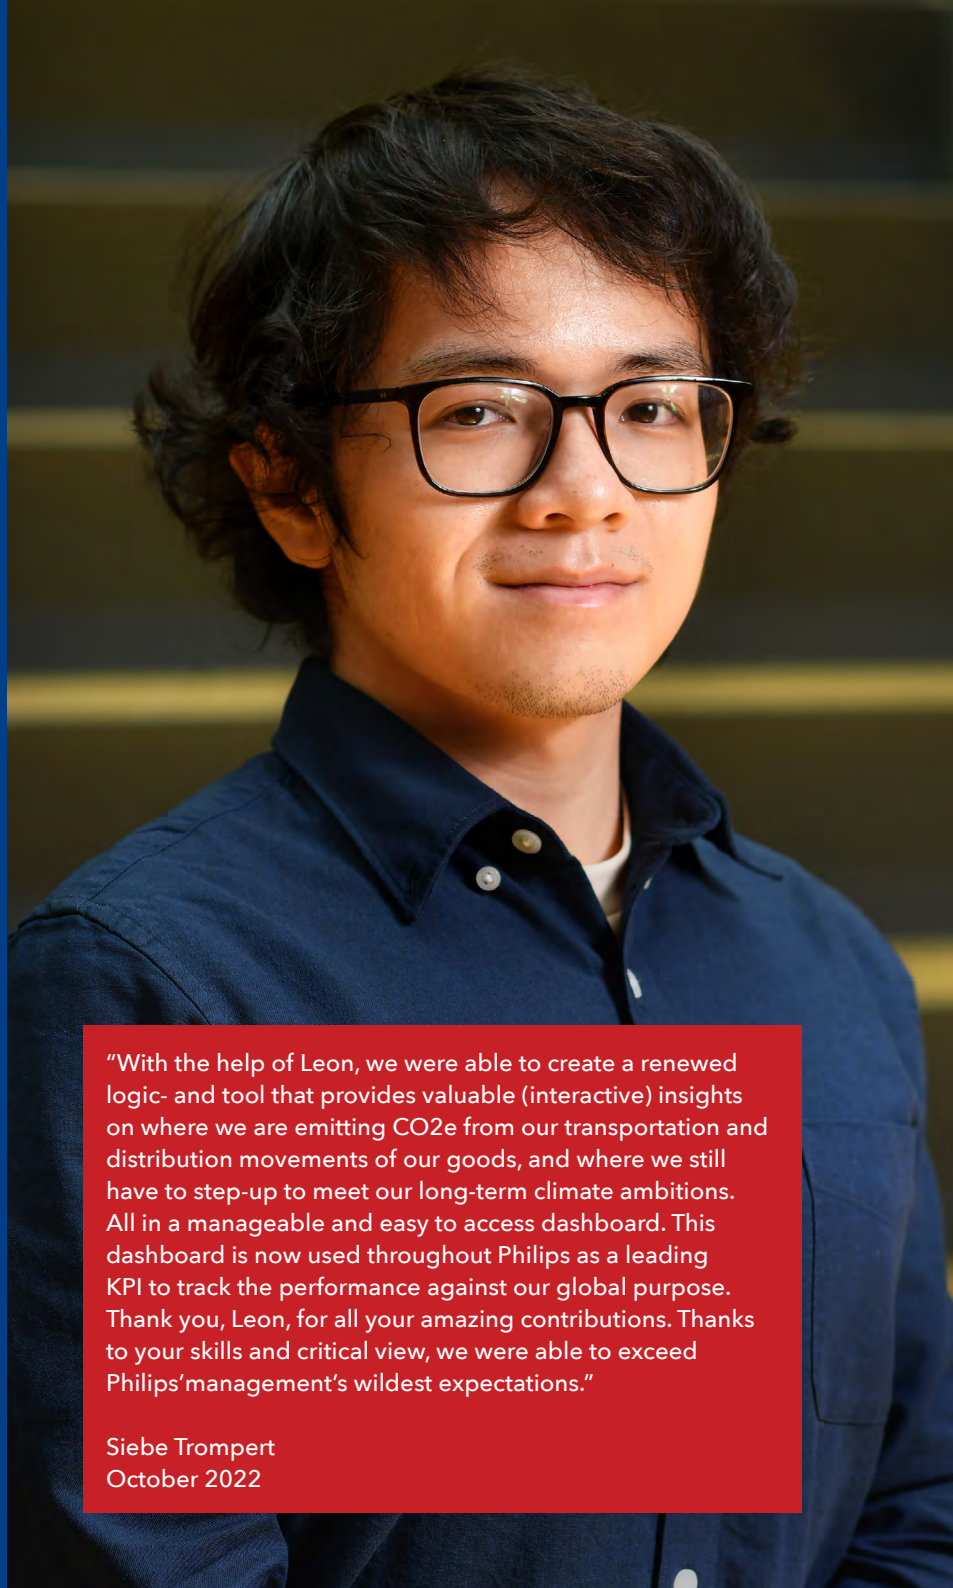

"With the help of Leon, we were able to create a renewed logic- and tool that provides valuable (interactive) insights on where we are emitting CO2e from our transportation and distribution movements of our goods, and where we still have to step-up to meet our long-term climate ambitions. All in a manageable and easy to access dashboard. This dashboard is now used throughout Philips as a leading KPI to track the performance against our global purpose. Thank you, Leon, for all your amazing contributions. Thanks to your skills and critical view, we were able to exceed Philips'management's wildest expectations."

Siebe Trompert
October 2022

LI-YANG WANG EngD

From Data to Insights to Drive Sustainable Change in Philips' Global Road Freight CO2e Emissions

As a company that owns businesses worldwide, Philips is dedicated to operating sustainably and environmentally friendly. The company set a target to reduce its carbon footprint across the value chain in line with a 1.5°C global warming scenario. To achieve such a goal, one of Philips' focusses is reviewing and improving the logistics processes. Following this strategy, a dashboard that analyses the logistics carbon emissions was developed in Philips. However, the analysis for road transportation logistics was still missing from the dashboard.

The outcome of the project successfully filled in the gap for a complete picture of Philips'logistics carbon emissions. The new applications process and clean the raw input data, calculate emission using shipment records, and display the analysis results. These applications were integrated into the existing system and adopted an identical reload schedule. Finally, a new sheet, along with an additional section on the landing page, was created in the user interface, the Logistics Freight Dashboard, to present the road freight carbon emission statistics.

Judging from the performance of the latest design in this project, we applied the architecture to ocean-freight-related applications as refinement and achieved success. We recommend accepting the structure depicted in this document as a new standard to improve the efficiency and maintainability of the Philips'logistics carbon emission analysis system.

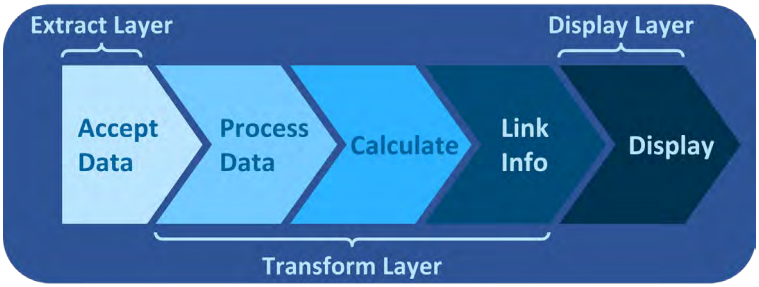

tue.nl/softwaretechnology

Follow us:


linkedin.com/company/tue_pdengst


facebook.com/TUePDEngST


twitter.com/tue_pdengst

CREDITS

Edited by: Dr. Y. Dajsuren EngD and K.A.H. Majoor, Software Technology Programme, Eindhoven University of Technology

Text: Graduates of generation 2020 of the EngD Software Technology Programme

Photography: Rien Meulman

Production: Software Technology Programme, Eindhoven University of Technology

Design: Grefo Prepress

Printing: De Digitale Drukker

Sedigheh Arasteh EngD; Philips Hue Entertainment Music Light Scripting Tool ■ **Akash Arora EngD;** Outpatient Monitoring Kit for Clinical Research ■ **Nastaram Bajalan EngD;** Feasibility and Prototype of Data-Driven Focus Extraction ■ **Dan-Cristian Chirascu EngD;** Safety-by-Design in Architecture of Automotive Software Systems ■ **Janice Concuet EngD;** Designing a Digital Twin Factory for Ink Handling in Advanced Production Printers ■ **Christian Degott EngD;** Hue Push Notifications - Event Logs and Push Notifications for Philips Hue ■ **Mohammad Ibrahim EngD;** A Graph Database Design for Multi-Domain Model Management ■ **Hasan Kaplan EngD;** The Design and Implementation of an Integrated Data Pipeline Solution for Behavioral Animal Science and Animal Breeding ■ **Lkham Nyambuu EngD;** Automated Neural Architecture Search for Fetal Head Detection ■ **Respa Putra EngD;** A Solution for Configuring an Infrastructure-as-a-Service ■ **Shubham Rawal EngD;** Automated Transformation of Picture Archiving and Communication System Reporting Templates to HL7® FHIR® Models ■ **Lamisha Rawshan EngD;** Value/Cost Analysis of Modularity Improvements ■ **Abolfazl Saravani EngD;** Remote Electron Microscope ■ **Li-Yang Wang EngD;** From Data to Insights to Drive Sustainable Change in Philips' Global Road Freight CO2e Emissions

The EngD (Engineering Doctorate) Software Technology programme is an accredited and prestigious two-year doctorate-level engineering degree programme. During this programme trainees focus on strengthening their technical and non-technical competencies related to the effective and efficient design as well as development of software for resource constrained and intelligent software intensive systems in multiple application domains for the High Tech Industry.

The programme is provided by the Department of Mathematics and Computer Science of Eindhoven University of Technology in the context of the 4TU.School for Technological Design, Stan Ackermans Institute.

For more information, visit our website:

tue.nl/softwaretechnology