PDENG PROJECTS 2021

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 - PDEng Projects 2021

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

Model-driven Software Engineering Meets Data Science

Software has proven to be a key technology for innovations in our society in the last decades. Software creates opportunities to improve the quality of our lives, communication, transportation, etc. Our society would have stopped functioning this COVID-19 era if we would not have been able to use all kinds of ICT based solutions. We are able to continue our "normal" way of working thanks to software solutions including video conferencing and data sharing. Innovations are primarily enabled by data and software, amongst other for the high-tech systems in the Eindhoven region, but also in the agriculture and food sector. By means of performing analysis of these vast amounts of data we are able to improve performance, predict maintenance, etc. It is remarkable that model-driven software engineering and data science evolve hand-in-hand and fit perfectly together. Based on results of machine learning existing models can be fine-tuned. We are on the very edge of deriving models from the data by means of advanced machine learning techniques. We can create digital twins of complex (engineered) systems and use the data of the systems to do accelerate design, maintenance, and optimize the production processes.

In recent years the importance of data and AI is accelerating, and this becomes also visible in the ST training program and final project assignments from industry. In this booklet you will find project examples of the current generation of PDEng trainees on how model-driven software development techniques are combined with data science and AI techniques. An example is the project of Siemens Leuven, where the trainee worked on the design and implementation of a plugin architecture, to support the cooperation of multiple engineers from different disciplines to perform design space exploration and the rating of architectures of systems-of-systems for aerospace, automotive, etc. The complexity of the software can be tackled by a good and well-structured architecture, of the projects at ASML, proposed a layered software architecture to integrate conventionally developed software with software generated from verified models. Another project done at ASML investigated the use of AI techniques to plan and optimize the exposure sequences in the wafer scanner. Cloud-based solutions are applied in a broad range of domains, a project in the area of food production for the AgrifoodTech organization on the topic of Internet of Food, and also in the area of healthcare with Philips Healthcare in order to deal with security and privacy aspects of data exchange. The continuous delivery of software is also high on the agenda of the high-tech companies. In cooperation with Thermo Fisher Scientific the automatic updating of software packages to customers was explored and a prototype was realised to allow for flexible software distribution. One of the two projects at Rabobank focused on the design and implementation of a testing and validation framework for AI and ML algorithms in the domain of credit scoring. These are just a few examples of the interesting projects, the projects cover a broad range of application areas, from high-tech systems to the financial domain, but there are also projects with a high societal impact on food, agriculture, and CO_2 reduction.

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 to come up with innovative solutions that are eagerly adapted by our industrial partners in a challenging time. Since mid-March 2020, all trainees were forced to work from home and had restricted access to the companies and the university. Nevertheless, all trainees have managed to finalize their training and graduation project in time and with good results as you can read in the rest of this booklet, which provides a source of inspiration and shows the challenges to advanced software development. We would like to congratulate our trainees with their results and wish them a bright and successful career.

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

PDEng Software Technology programme
Eindhoven University of Technology

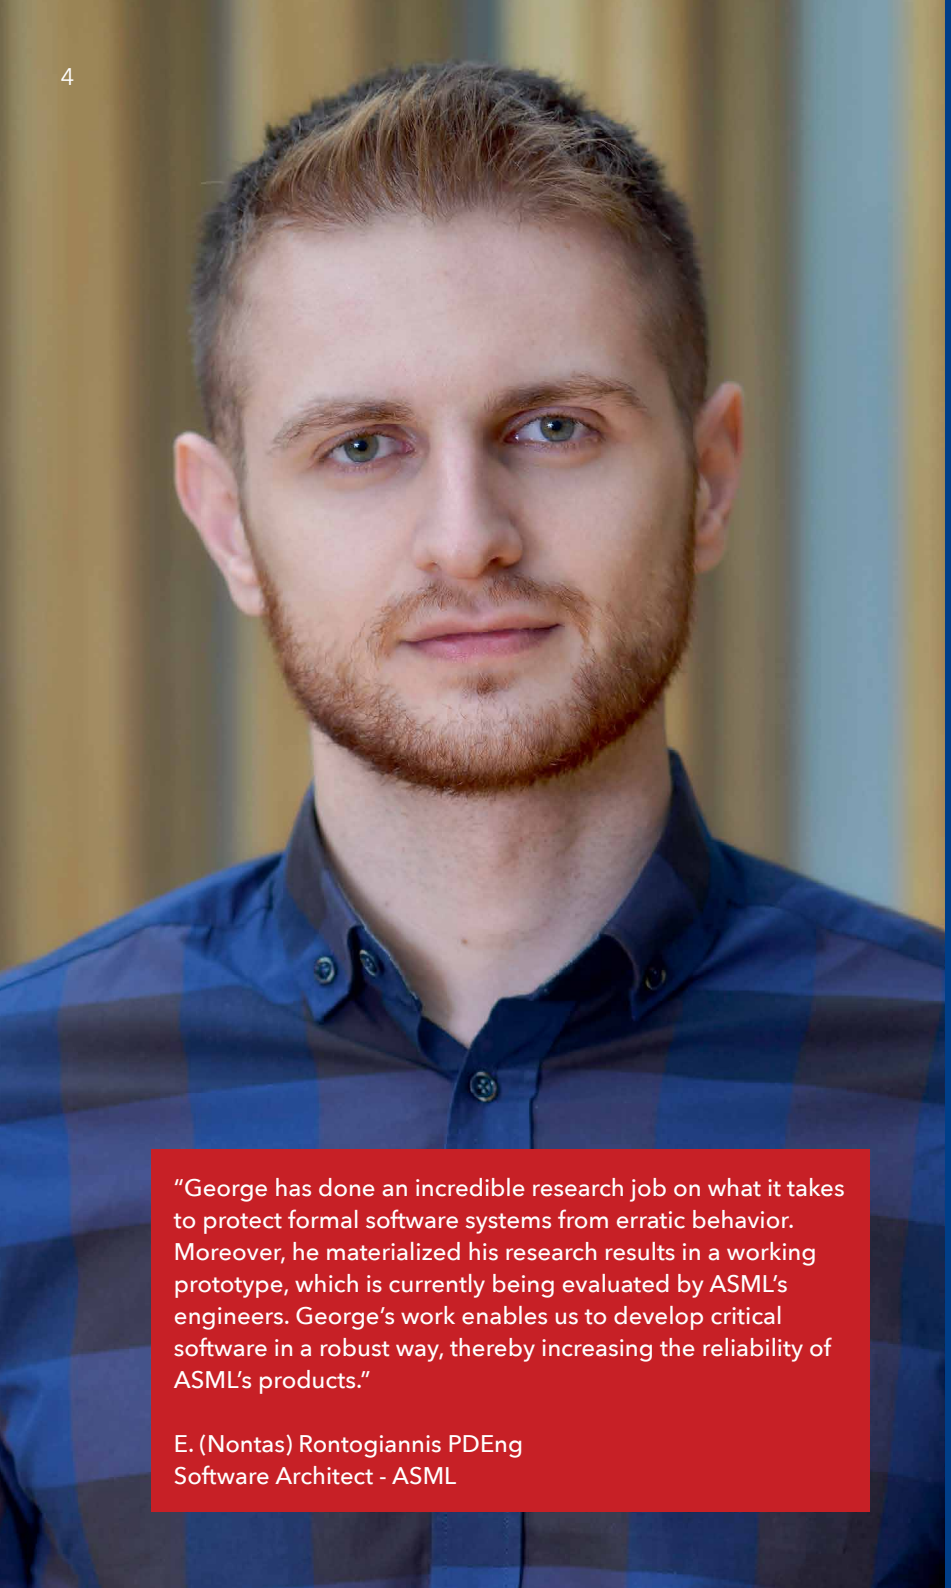

"George has done an incredible research job on what it takes to protect formal software systems from erratic behavior. Moreover, he materialized his research results in a working prototype, which is currently being evaluated by ASML's engineers. George's work enables us to develop critical software in a robust way, thereby increasing the reliability of ASML's products."

E. (Nontas) Rontogiannis PDEng
Software Architect - ASML

CHALLENGES

The main challenge of this project was to propose and assess robust integration strategies between conventional and formally implemented software. Furthermore, we had to design and implement a working prototype that facilitates the automatic application of the chosen integration strategy to the software environment of ASML. Additionally, we worked with development platforms that were under continuous development.

RESULTS

Firstly, we identified and documented guidelines to design formally implemented systems that can be safely integrated with conventional software. These guidelines aim to eliminate anti-patterns and corner cases that are hazardous to the formal system. Secondly, we provided a "correct by design" robust integration strategy between conventional and formally implemented software, assuming that the guidelines are followed. Thirdly, we designed and implemented a product that facilitates the automatic generation of the integration strategy. Finally, to validate the product, we used ASML production software.

BENEFITS

By following the proposed solution, we can safely integrate conventional and formally implemented software. In addition, the developed product:

- Automates the repetitive (and in several cases error-prone) task of integrating conventional and formally implemented software, by generating "correct by design" artifacts.
- Reduces the development time of Anti-corruption and Glue Code Layers from hours to seconds.

GEORGIOS AZIS PDEng

Armoring Formal Components against Foreign Behavior

Due to the increasing complexity of the software that runs on ASML's machines, ASML decided to adopt Formal Software Engineering (FSE). FSE is a robust methodology of developing software that allows engineers to explicitly define the behavior of software at all times. In addition, it uses automated mathematical analysis to prove the absence of certain defects, which are very difficult to discover with testing.

The integration of conventional with formally implemented software occurs based on an interaction protocol, which is modeled in a state machine. Formal verification ensures that the formally implemented software is guaranteed to adhere to the interaction protocol. Moreover, the formally implemented software is developed based on the assumption that its environment also adheres to this protocol. However, there are no guarantees that the conventional software actually adheres to this protocol. Hence, to avoid the execution of unexpected behavior (and even system crashes) a robust integration strategy is needed.

This project proposes a layered solution: The first layer (Anti-corruption Layer) protects the formally implemented software from unsupported behavior, whereas the second layer (Glue Code Layer) connects two technology spaces. Here, the formally implemented software can be safely used by conventional software, via its provided (top) legacy interface, and it can safely use conventional software, via its required (bottom) legacy interface. Finally, the project provides a product that generates "correct by design" Anti-corruption and Glue Code Layers, based on formal and legacy interfaces.

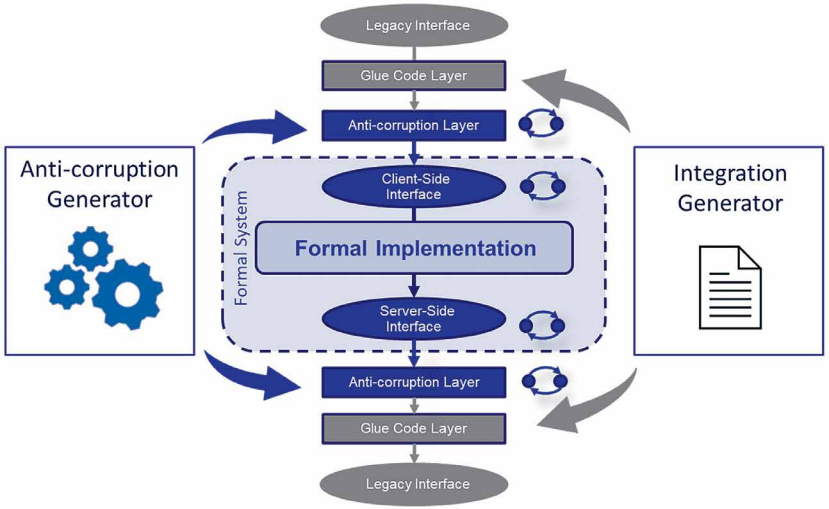

CHALLENGES

The challenge in this project was to understand and extract the main functionalities of a software component that has a long history and is not fully documented. The software component is also a complex component that holds many responsibilities and has a big size.

RESULTS

The result was a redesign of the software component in an object-oriented style that is more modular than the original design.
This design was created using popular object-oriented principles and proven design patterns.
A prototype implementation that maintains the main functionalities of the driver was provided.
In addition, a brief guideline reference that includes the steps and techniques used when migrating C component code to C++ was provided.

BENEFITS

The contribution of this project is to give ASML a good insight in the challenges one encounters when making a transition from a large legacy code in C to an object-oriented approach. It also provides valuable guidelines on dealing with commonly occurring patterns in ASML code in an object-oriented approach.

"Samsom made a quick start and steady progress in understanding the ASML legacy software and development environment, which has a steep learning curve. He showed us how some common coding patterns in ASML can be done much more simple and robust when using Object-Oriented techniques. Samsom's report gives us a good insight in the challenges one encounters when making the transition to C++ and provides valuable suggestions on how to deal with commonly occurring patterns in ASML code in an objected-oriented way."

Ir. Camiel Rouweler
Software Architect - ASML

SAMSOM TSEGAY BEYENE PDEng

Redesigning the Image Sensor Subsystem Driver using the Object-Oriented Paradigm

ASML is the leading provider of lithography systems in the world. These lithography systems are complex machines that are critical in the production of integrated circuits or microchips. Inside the ASML machines, there is an image sensor module that measures various optical system parameters such as lens aberrations, pupil illumination, and polarization state of the scanner. This module is controlled by a set of software components, with the camera driver being one of them. Over the years, the camera driver has become more complex as a result of incorporating new requirements. Most importantly, the driver is implemented in the C language, which is generally not efficient and not modular. In addition, as the software driver gets bigger and bigger, it becomes harder to maintain and modify.

A redesign of the camera driver was developed. In this redesign, different OOD techniques such as design patterns and SOLID principles were applied. Applying these techniques helps to divide and structure the responsibility of the camera driver into smaller and decoupled classes. This arrangement improves the modularity of the driver. For a proof of concept, a prototype implementation of this redesign in the C++ language was provided. This implementation was tested on an ATTEST, unit testing environment, the testing simulation environment (Devbench), and the testing environment involving real hardware (Testbench).

The redesign and prototype implementation show that the modularity of the driver is increased while maintaining its main functionalities. Hence, maintaining and modifying the driver component becomes easier.

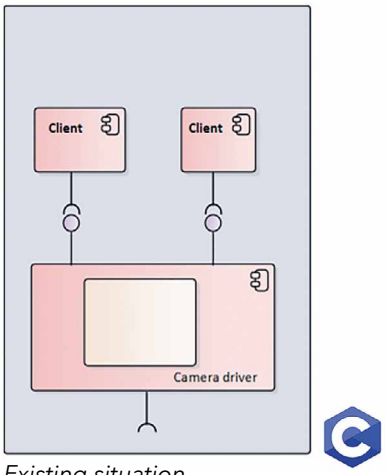

Existing situation

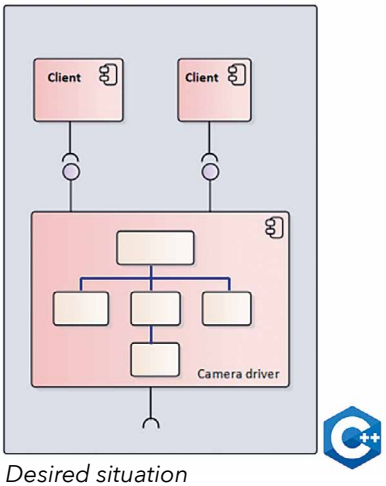

Desired situation

"During the assignment Jorge researched an interesting subject: the verification of a credit scoring analytic model through a monitoring engine is the outcome of the model in line with expectations. Or in other words, are the input data and the result connected. Rabobank is very happy with the results, and we hope that the monitoring engine will show its added value in the upcoming time."

Rik Bosman MSc
Rabobank

CHALLENGES

It is a common practice to monitor the quality predictions by using statistical metrics that require reference predictions for comparison. However, for some financial applications there are no reference predictions. This lack of reference prediction data was the main challenge of the project. To solve this problem, we designed monitoring methods that identify changes in the statistical properties of predictions without relying on reference prediction data.

RESULTS

For this project, we designed a solution architecture for monitoring analytic model predictions. In addition, this architecture describes a model deployment pipeline that ensures every deployed model can be executed and monitored in a production environment. We created a prototype to demonstrate how the proposed architecture can be implemented in the current codebase of the innovation project. This prototype shows how model monitoring and model deployment can be automated.

BENEFITS

Overall, the proposed solution architecture enables to monitor the quality of predictions in analytic models and simplifies the tasks required to development, deploy, and maintain analytic models. As analytic models and their requirements are expected to change over time, the architecture supports monitoring new properties of the analytic models. The automation of deployment tasks enables data scientists to focus more on designing models produce accurate predictions.

JORGE ALBERTO CORDERO CRUZ PDEng

Designing a Solution Architecture for Monitoring Credit Scoring Analytic Models

An innovation project at Rabobank aims at building bridges between smallholder farmers without financial records and credit providers. To achieve this goal, a credit scoring application leverages alternative (non-financial) data, such as historical weather conditions, field characteristics, and yield production, to calculate credit scores that credit providers can use for assessing the risk profile of their potential clients. These credit scores are predicted by analytic models created based on a set of variables identified by Rabobank data scientists together with scientists from the Wageningen University.

In accordance with Rabobank policies, analytic models must follow a model development process to ensure the quality of their predictions and minimize potential risks. Some of these potential risks can lead to reputation damage and fines due to wrong or biased predictions. The goal of this project was to design a solution architecture that enables monitoring the quality of analytic model predictions and helps data scientists during model development tasks.

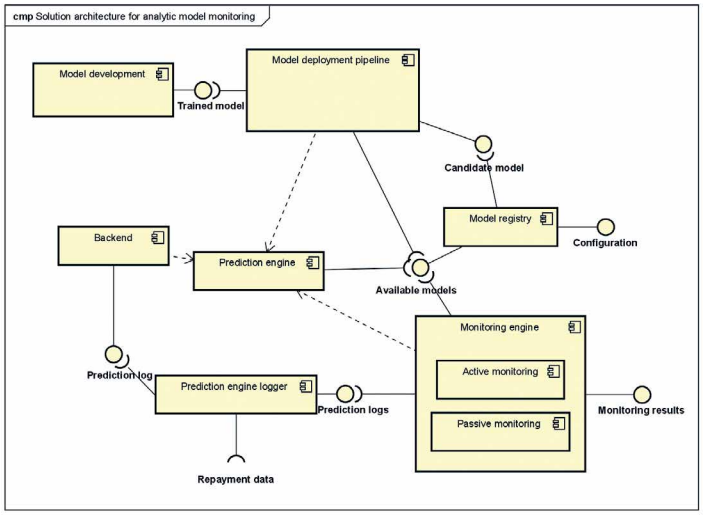

The proposed solution architecture describes methods for monitoring predictions with and without reference data. In addition, this architecture also explains how to implement a model deployment pipeline to automate testing and validation of analytic models. Overall, the solution architecture provides guidelines that data scientists can follow to develop, test, monitor, and maintain analytic models. Using a prototype, we demonstrated how this architecture can be implemented in the current codebase of the innovation project.

CHALLENGES

The challenge was to acquire the domain and the systems' knowledge of ASML in a very short period. Testing multiple product configurations at ASML leads to large amount of data in very large files stored in different formats. Understanding these files and the test process, developing an ETL solution, defining metrics and integrating various tools to create a streamlined workflow were challenges that were a part of the assignment.

RESULTS

The results are the development of a tool that automated and streamlined the test execution and reporting workflow. The tool enables ASML to reduce the time and manual effort consumed. Visualization of the data was achieved using Spotfire, a tool to easily create visualizations over data with minimum effort and no programming expertise. In order to achieve this, a concrete architecture was designed which is the backbone of the tool.

BENEFITS

The value of this project not only lies in reducing manual effort but also in enabling ASML to visualize test data easily. Test data is now stored in a central database. Using this database various visualizations can be easily created over the data using Spotfire. This enables ASML to safely guard test data and generate reports within a click of a button. This saves time and effort on a whole.

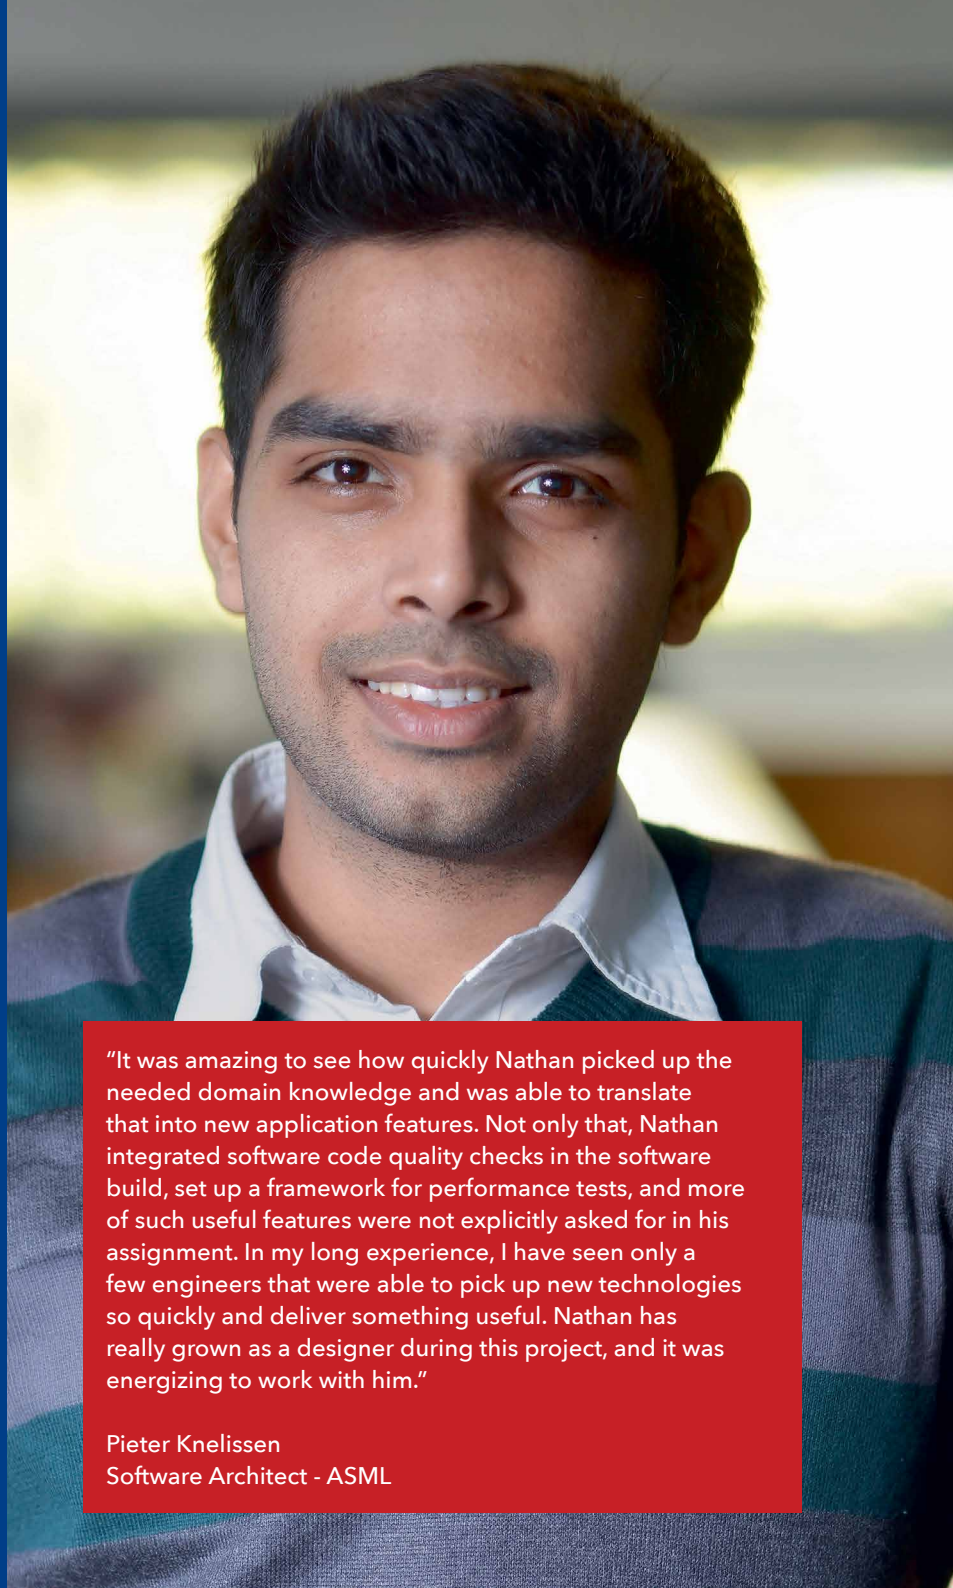

"It was amazing to see how quickly Nathan picked up the needed domain knowledge and was able to translate that into new application features. Not only that, Nathan integrated software code quality checks in the software build, set up a framework for performance tests, and more of such useful features were not explicitly asked for in his assignment. In my long experience, I have seen only a few engineers that were able to pick up new technologies so quickly and deliver something useful. Nathan has really grown as a designer during this project, and it was energizing to work with him."

Pieter Knelissen
Software Architect - ASML

NATHAN D'PENHA PDEng

A Tool to Improve Reporting on Test Execution and Test Coverage

ASML is the world's leading provider of lithography systems for the semi-conductor industry. ASML's lithography machines, along with other ASML products, form a holistic suite of products that are used by semiconductor manufacturers to produce chips that power the world's electronics industry. This suite of ASML products is complex, with many interactions between the products. ASML has implemented a novel test facility, a Digital Twin called Virtual Fab, in which a growing number of interoperability tests are being implemented. However, reporting regarding execution and coverage of tests is done manually using Excel, which comes with the following drawbacks: time-consuming, error-prone, utilizes expensive workforce, cannot accommodate complex analysis, reduces productivity, people dependent, and difficult to maintain.

The goal was to develop a tool that eliminates manual work needed to safeguard test data and generate the needed reports and populate dashboards by automating related tasks.

To achieve the goal, we designed and implemented a tool that automates the process of safeguarding test execution data and produces the needed reports in an automated method. To do this, the tool fetches test execution data from multiple sources and extracts data using a developed parser. The tool populates the database and dashboards. It also integrates with other tools (such as Jira, Confluence, and Artifactory) to reduce the manual effort consumed during a test process. Additionally, to easily create meaningful analysis, the tool integrates with Spotfire. Using Spotfire, simple and complex analysis over test execution data can be easily created.

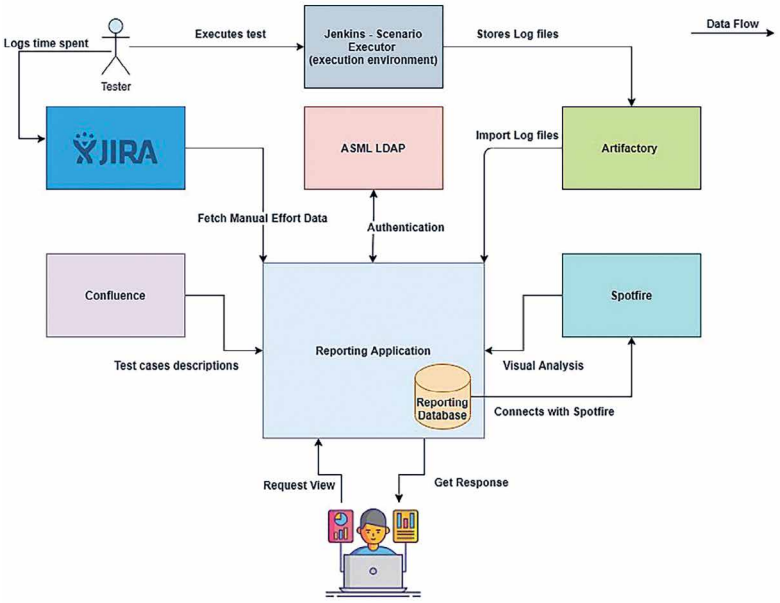

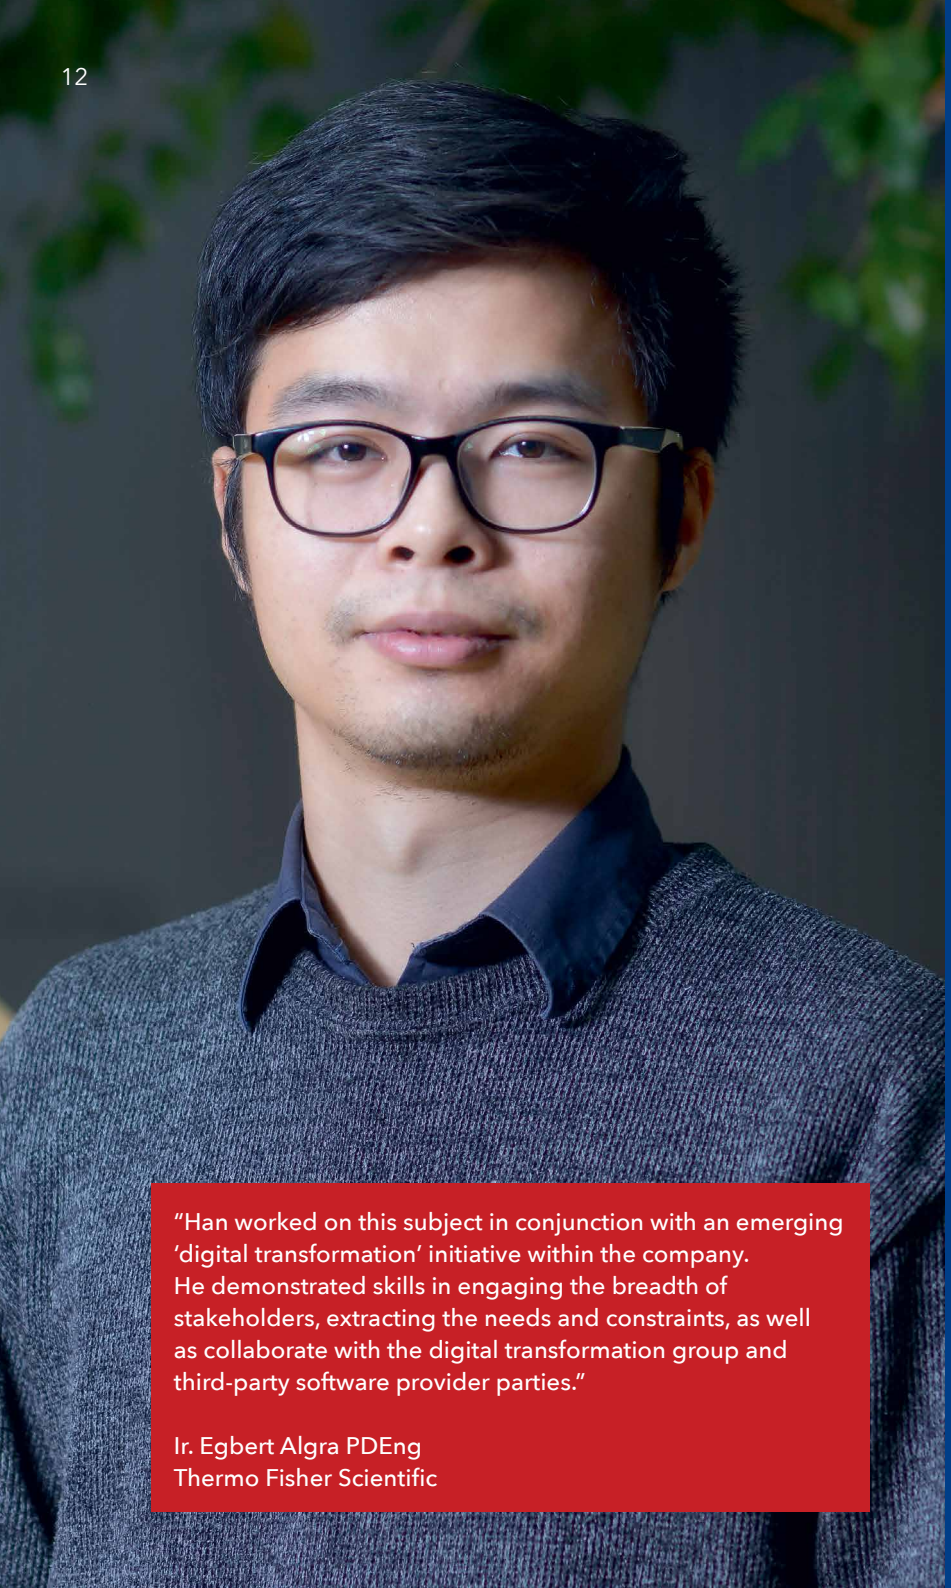

"Han worked on this subject in conjunction with an emerging 'digital transformation' initiative within the company. He demonstrated skills in engaging the breadth of stakeholders, extracting the needs and constraints, as well as collaborate with the digital transformation group and third-party software provider parties."

Ir. Egbert Algra PDEng
Thermo Fisher Scientific

CHALLENGES

Main challenge of this project was to integrate the domain knowledge and requirements acquired in fragments of different teams and individuals to identify and define reasonable project scope. Another challenge of this project was software repositories and consumption API choices that apply to the goal of the project.

RESULTS

The main result of this project is to use existing technology to realize the prototype of the pipeline consisting of software release and distribution. The prototype shows how to publish different format software packages with metadata to a publicly accessible repository master and distribute them to online and offline instruments through a repository replica and consumption APIs.

BENEFITS

The outcome of this project provides a reliable software distribution blueprint, which satisfies the breadth of Thermo Fisher Scientific software distribution requirements and applies to Thermo Fisher Scientific customers' various instrument scenarios with different network connections.

HAN GAO PDEng

Designing a Solution for Software Distribution towards Varying On-premise Deployments

The Materials and Structural Analysis Division (MSD) business of Thermo Fisher Scientific (formerly known as FEI company) has traditionally been a global leader in the innovation of Electron Microscopes. The software of the Electron Microscopes has played an increasing role in the delivery of innovations because of market needs.

The software release lifecycle is faster than that of hardware. The company saw an opportunity to deliver newly created software quickly and seamlessly to its customer independently instead of a box of hardware and software. Currently, the mechanisms that Thermo Fisher Scientific uses to deliver software to its customers are manual operations, which depend on different instrument types and customer scenarios. These human operations are inadequate in that they take a lot of time and they are costly processes. Therefore, the company needs a better solution to publish and distribute software to its customers.

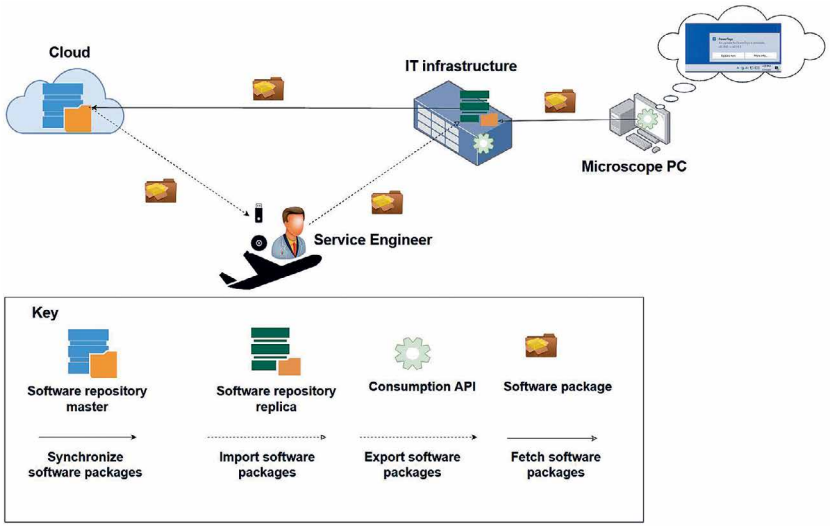

In this project, we proposed a software distribution solution towards varying Thermo Fisher Scientific customer on-premise deployments. The solution consists of three parts: a software repository master that centralizes all desired varying software packages; a software repository replica that synchronizes desired software packages to customer on-premise IT infrastructure; and consumption APIs that enable Thermo Fisher Scientific or its customer engineers to manage software. In addition, we realized a prototype which demonstrated the value of the software distribution solution.

CHALLENGES

The challenges in this project were to understand the project domain and creating a modular design that complies with the defined Reference Architecture. The project is executed in the healthcare domain, which uses different tools and standards, such as the HL7 FHIR Standard. These challenges were overcome using the extensive expertise from the Software Concepts team in Philips, who were very open to meetings and discussions.

RESULTS

A design and prototype for the HealthSuite Workflow Capability is presented. The prototype is based on the BPMN and DMN standards. The design is modular and allows for a Workflow Engine vendor replacement. Additionally, the design presented makes the system deployable on-site or in the cloud. The prototype shows the feasibility of the proposed design.

BENEFITS

The design presented provides the following value:

- A design with modularity of the key components in the system, ensuring no hard dependencies on a Workflow Engine vendor.
- A design that allows for deployability in the cloud.
- An interface definition between the Workflow Engine and the FHIR Store, which allows for Clinical Data to be used in Workflows.

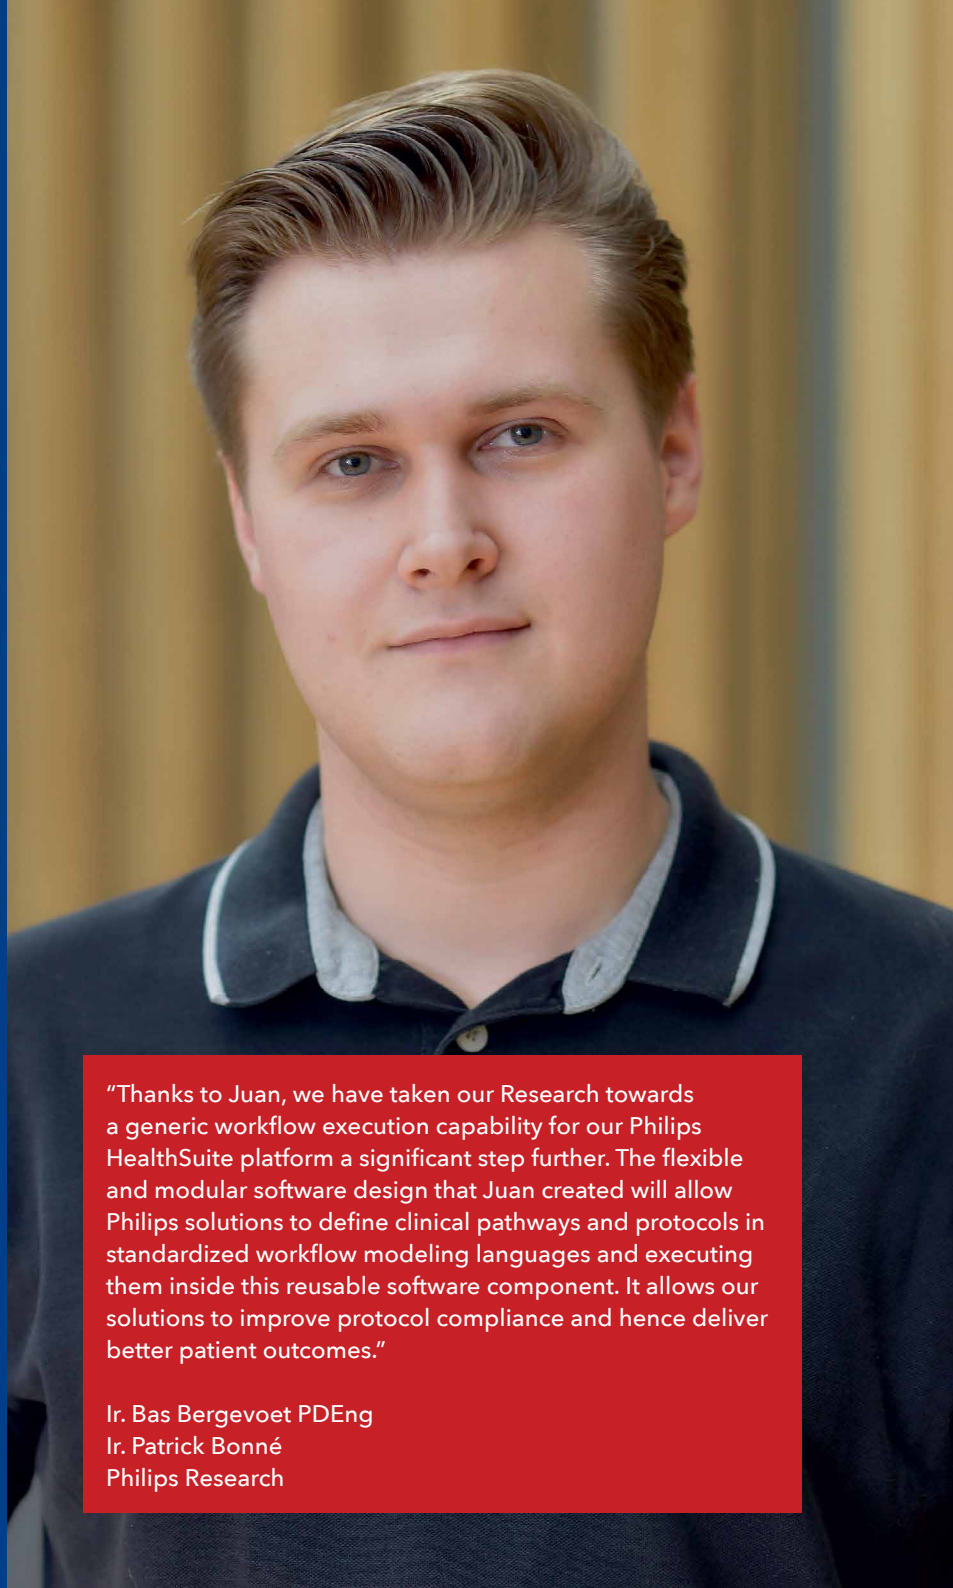

"Thanks to Juan, we have taken our Research towards a generic workflow execution capability for our Philips HealthSuite platform a significant step further. The flexible and modular software design that Juan created will allow Philips solutions to define clinical pathways and protocols in standardized workflow modeling languages and executing them inside this reusable software component. It allows our solutions to improve protocol compliance and hence deliver better patient outcomes."

Ir. Bas Bergevoet PDEng
Ir. Patrick Bonné
Philips Research

JUAN VAN DER HEIJDEN PDEng

Standardized Software Solution for Guidance of Clinical Workflows

During a patient's stay in the hospital, patients traverse through multiple clinical pathways, often having to comply with one or more defined clinical workflows. Information on how a clinical workflow must be executed is often described by large healthcare bodies. These clinical workflows are often locally adapted. Proper adherence to these workflows is often crucial for having the best clinical outcome for the patient.

Philips is researching the feasibility of creating a standardized software solution for guidance of clinical workflows, called the HealthSuite Workflow Capability. The HealthSuite Workflow Capability can help healthcare personnel with executing these clinical workflows more accurately. Accurate execution of clinical workflows helps healthcare institutions to improve the quality of their healthcare.

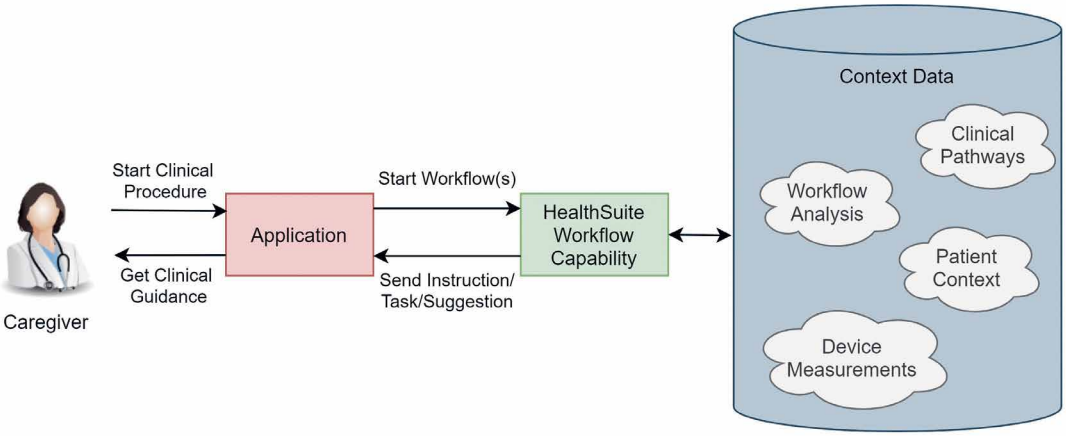

In this project, a detailed design and prototype for the HealthSuite Workflow Capability have been realized based on the BPMN and DMN standards. The proposed detailed design shows how a workflow engine vendor lock-in can be prevented. A modular design is presented that makes the system deployable either on-site or in the cloud. The design and prototype can be used as a baseline for future research.

"Jobaer was challenged to create a proof of concept solution to generate optimal expose sequence. His approach of modeling the expose image sequence in PDDL and using AI planning was successful. He managed to grasp the domain plus related technologies quickly. Moreover, he involved all his stakeholders and addressed their demands regularly. The prototype created by Jobaer gives us the desired confidence to start considering AI planning at certain ASML use cases and problems."

Dr. Damir Lotinac
Aristeidis Magklaras MSc
ASML

CHALLENGES

The main challenge of this project was understanding AI planning and turning the specification of the expose image subplan into a PDDL model. Another challenge was to find an off-the-shelf planner that could generate an optimal sequence for the expose image domain. Moreover, the requirement analysis was a challenge as this project was a feasibility study, the end goal was unclear.

RESULTS

The results demonstrate one possible approach of modeling the expose image sequence in PDDL to generate an optimal sequence using an off-the-shelf planner, proving the feasibility of using AI planning. In addition, a prototype Automated Sequence Planner (ASP) was developed to generate sequences, PDDL problem files, and C++ function signatures. The ASP shows the generated sequence in a Gantt chart that enables investigation with different specifications.

BENEFITS

Generating optimal sequences using ASP reduces the time for creating a sequence and improves system throughput. The generation of C++ function signatures from the ASP is a showcase for future implementation. This project provides a foundation for using AI planning to generate an optimal sequence for the expose image sequence and potentially other sequences of ASML.

JOBAER ISLAM KHAN PDEng

Automated Sequence Planner

A TOOL TO CREATE AN OPTIMAL SEQUENCE USING AI PLANNING

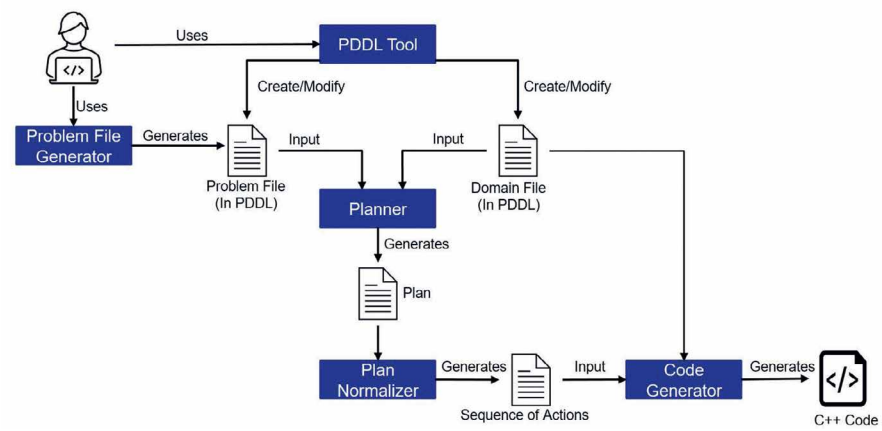

ASML is the global leader in providing photolithography systems (e.g. TWINSCAN) for the semi-conductors industry. In the TWINSCAN, for a wafer to be exposed, the wafer and the scanner must go through a series of actions to ensure successful image exposure. This exposure is done by following an expose image sequence that determines the actions and order of actions to expose images on the wafer. This sequence is created manually, which is time-consuming, not always optimal, and error-prone. After the sequence creation, code is written based on the sequence for the execution software. The code is then integrated into the TWINSCAN application.

This project was initiated to evaluate the feasibility of using Artificial intelligence (AI) planning to solve the aforementioned sequence creation problem. AI planning consists of a domain file, a problem file, and a plan. A planner takes a domain file and a problem file as input to generate a plan. The domain and problem files are defined in Planning Domain Definition Language (PDDL).

In this project, a modeling approach in PDDL for the expose image sequence was provided as a proof of concept to generate an optimal sequence. In addition, an Automated Sequence Planner (ASP) was developed by combining AI planner, problem file generator, and code generator. This tool provides the following functionalities:

- Create and modify PDDL files
- Generate optimal sequence
- Show sequence in a Gantt chart
- Generate problem files for a given domain file
- Generate C++ function signatures

CHALLENGES

The main challenge was to make the model training process reproducible and auditable. Because data scientists use different tools, it becomes hard to automate it end-to-end. More artifacts are to be managed beyond the code, such as data, parameters, hyper-parameters, ML library versions, and the resulted models. The second challenge is accomplishing inference in different environments and providing the same tool and approach for deploying to the cloud or on-premise devices.

RESULTS

The result was an ML pipeline that makes the process of making models transparent. The IntelLight+ infrastructure is ready to be used to help the data scientists manage and visualize the data, hyper-parameter tuning, and code resources used in experiments. It also covers deployment models in different target environments in an easy and convenient way.

BENEFITS

The emphasis on flexibility and modifiability of the IntelLight+ makes it possible to use different ML libraries for training models and compare them using a web-based UI. The IntelLight+ is also reproducible, it means it can deploy and use the trained models on different target environments for context recognition. Furthermore, the technologies like MLflow will allow researchers at Signify to move the IntelLight+ system towards a production-level environment without extensive reengineering.

"Developing a platform for experimentation of algorithms for different use cases is complex, especially when use cases and algorithms themselves are still expanding and developing. A new use case modifies inputs, outputs, and metrics. A new algorithm may also need to be trained with a different training strategy. Hossein has helped us in identifying the commonalities, and he has created a pipeline on which experiments are currently being executed. Results from these experiments will enable us to set the next step in delivering the best lighting experience."

Dr. Fetze Pijlman
Signify

HOSSEIN MAHDIAN PDEng

IntelLight+: Designing and Developing of Context-Aware Lighting Solution

As the world leader in providing lighting solutions, Signify delivers software-based solutions as a managed service to its customers. The IntelLight+ system plays an essential role in moving current solutions towards the next step, intelligent lighting. The goal of this project was to design and implement the infrastructure for data scientists to explore the ML approaches to detect the user's context.

IntelLight+ is designed and implemented to automize the process of ML by considering the extendability and flexibility in mind. By flexibility, it allows employing a wide range of training algorithms, libraries, and methods for context recognition. By extendability, it provides an API that can be implemented for adding new algorithms and sensors for context detection.

The training is an iterative process, which means data scientists launch several times the same run with slightly different parameters/data/preprocessing. If they rely on their memory to compare these runs, they will likely struggle to remember the best one. IntelLight+ visualizes the ML experiments with different parameters, data, and the code version that has been used to produce the model. Besides making the experiments reproducible, IntelLight+ allows data scientists to develop models in local machines and cloud then deploy the models on-premise or cloud by a shared storage. After deployment, it allows applying the received feedback from users. Besides the main parts of the ML pipeline, a software component was developed to enable giving feedback regarding the quality of the deployed model.

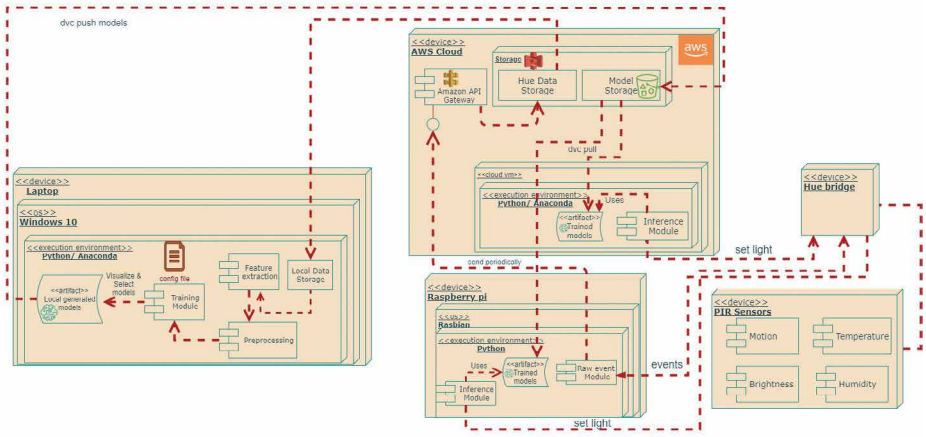

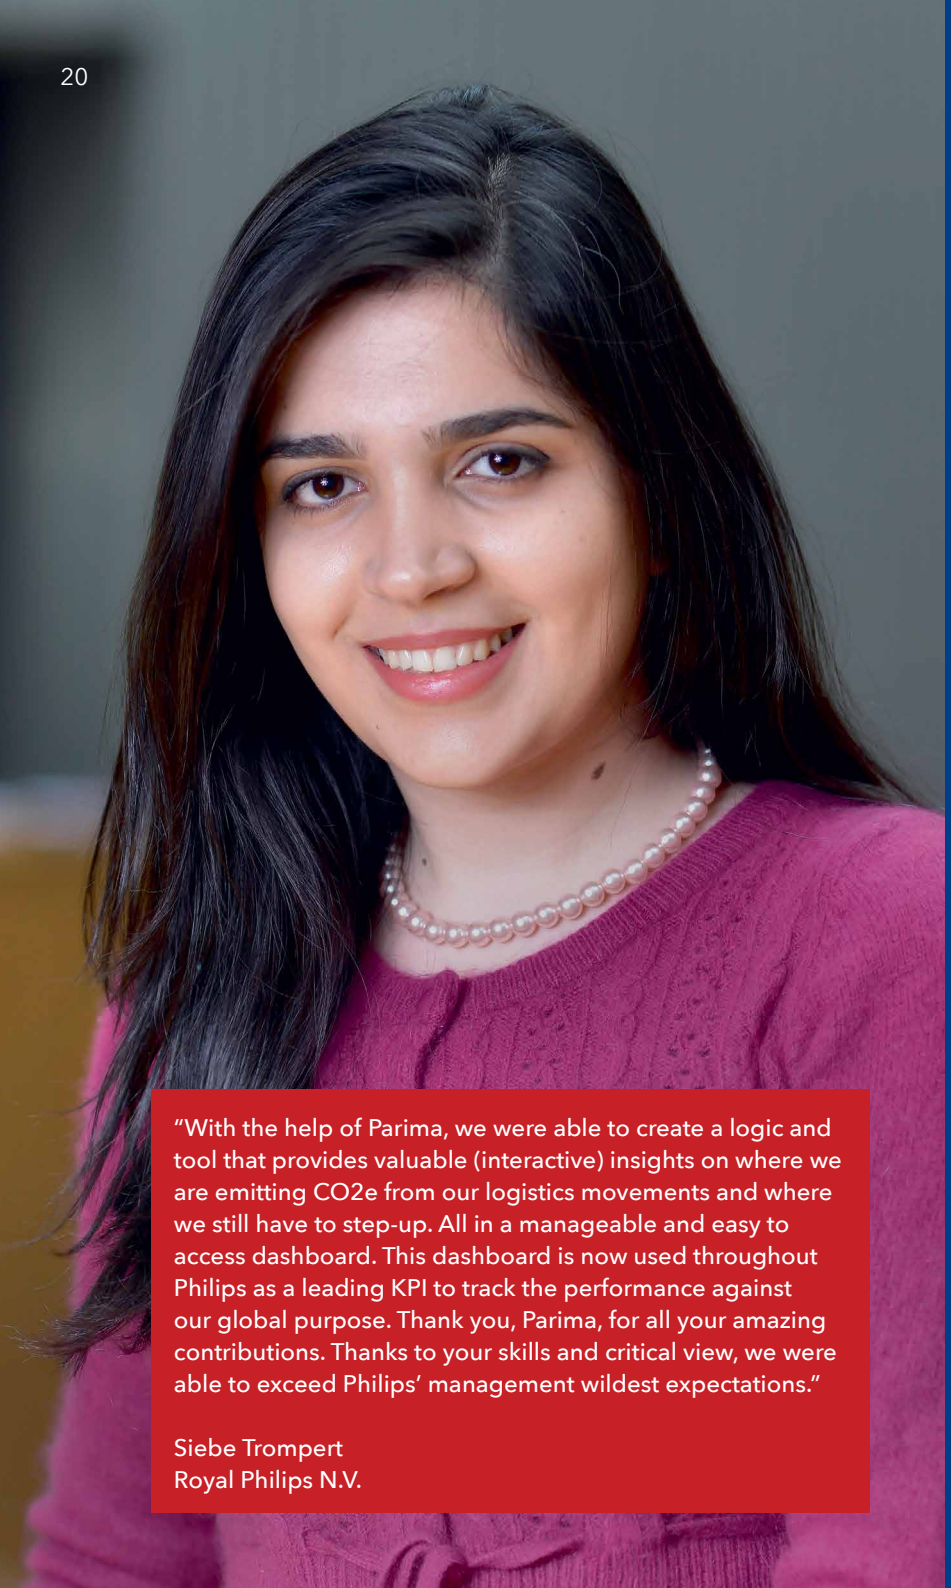

"With the help of Parima, we were able to create a logic and tool that provides valuable (interactive) insights on where we are emitting CO2e from our logistics movements and where we still have to step-up. All in a manageable and easy to access dashboard. This dashboard is now used throughout Philips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

Siebe Trompert
Royal Philips N.V.

CHALLENGES

The challenge in this project was gathering all the necessary information from different data providers and combining them. Merging the data was arduous as all the data sources had different data structures. Moreover, defining the required data cleaning and preparing steps, making them automated, was time-consuming.

RESULTS

The result was a dashboard representing the diffused CO_2 emissions based on three modes of transport; Ocean freight, Air freight, and Parcel shipment. Moreover, a distance calculator app was designed and implemented to calculate the inland and coastal distances. The output of this application will be used to enrich the Philips data lake.

BENEFITS

By providing this dashboard, managers have a holistic view of the amount of CO_2 emissions per the mode of transport. Thereby, the decision-making procedure is getting much easier for them. Furthermore, the application will help different markets and businesses in Royal Philips N.V. choose the more sustainable approach for shipping the products.

PARIMA MIRSHAFIEI PDEng

From Data to Insights to Drive Sustainable Change in Philips' Global Logistics Movements

Royal Philips N.V., as the pioneer in becoming Carbon-neutral in their procedure, started a new project to monitor the diffused CO_2 emissions in various sections. One of the most impactful sections is logistics. This project aimed to develop a comprehensive dashboard containing the Ocean freight, Parcel shipments and integrating the existing Air freight to have up-to-date data every day.

In order to implement the dashboard, four main procedures were considered. These procedures include identifying the primary data sources for each dashboard, specifying the requirements for each dataset and sending them to the data providers, cleaning the datasets and removing unnecessary information, and designing and implementing the dashboard based on the stakeholders' requirements.

Additionally, a distance calculator application was implemented to calculate the distances between different cities and ports. Therefore, the CO_2 emissions were computed more accurately and automatically. Furthermore, all the procedures from getting the data to cleaning and showing them on the dashboards are automatic, and there is no need for human intervention. Accordingly, all the managers can have access to the latest data and make sustainable decisions.

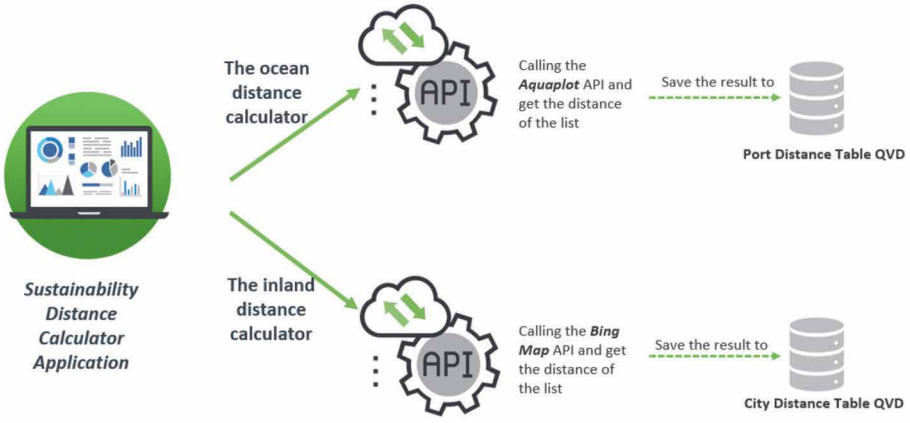

CHALLENGES

The biggest challenge was identifying the approach to integrating generative engineering and concurrent engineering to provide users with useful tooling. In innovative research projects, where the environments and requirements are rapidly changing, scoping is a non-trivial task. On the software level, providing synchronized data assets for user models incorporates challenges related to concurrency and consistency.

RESULTS

The result is an extensible, end-to-end plugin architecture pipeline, named COGENT. COGENT connects multiple generative users to each other for completing collaborative and simultaneous engineering tasks. Independent plugins were created to provide users with shared access to centralized storage, graph databases, experiment tracking software, visualization tools, and user-defined features. Users can extend the generative engineering capabilities of Siemens Simcenter Studio to the concurrent engineering domain.

BENEFITS

With COGENT, the optimal system architecture to meet mission and operational requirements is in the design space. Our pipeline architecture improves user collaboration, communication, and design decision traceability. Additionally, applying filters and constraints to the quantitative, qualitative, and architectural attributes allows subsystem designers to maximize the performance-specific criteria to meet their needs. Finally, system designers have traceability to formal standards and protocols, which ensures product safety and reliability.

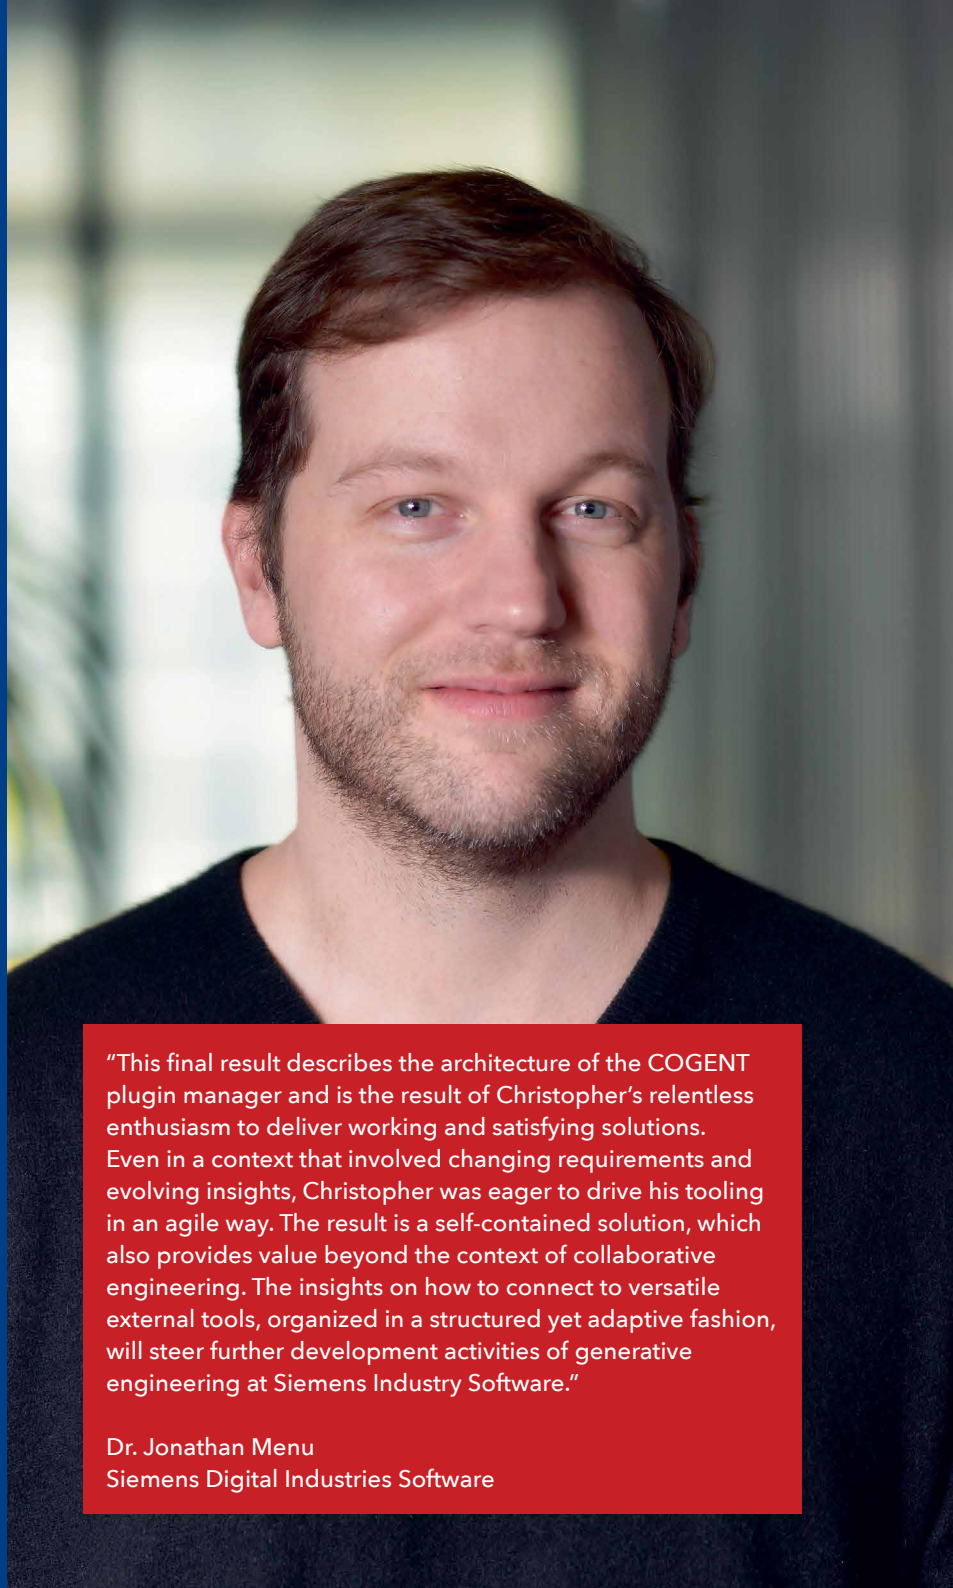

"This final result describes the architecture of the COGENT plugin manager and is the result of Christopher's relentless enthusiasm to deliver working and satisfying solutions. Even in a context that involved changing requirements and evolving insights, Christopher was eager to drive his tooling in an agile way. The result is a self-contained solution, which also provides value beyond the context of collaborative engineering. The insights on how to connect to versatile external tools, organized in a structured yet adaptive fashion, will steer further development activities of generative engineering at Siemens Industry Software."

Dr. Jonathan Menu
Siemens Digital Industries Software

CHRISTOPHER O'HARA PDEng

COGENT: Concurrent Generative Engineering Tooling

ENABLING CROSS-FUNCTIONAL TEAMS IN ARCHITECTURE DESIGN FOR SPACE SUBSYSTEMS

Siemens Digital Industries Software (SISW) is a global leader in industrial automation, lifecycle management, electrical/mechanical design, and digital innovation. Technology-oriented companies and institutions regularly employ Siemens' software within the automotive, aerospace, and intelligent manufacturing domains. Without exception, Siemens' technology solutions blur the boundaries between various engineering and technology domains, allowing for effortless, multidisciplinary product development. Siemens SISW is currently developing a new technology solution for automatic system architecture generation called Simcenter Studio (SCS).

At the start of this project, SCS was a single-user application. To enable cross-functional teams to engage in concurrent engineering, a methodology for simultaneous product development, SCS needed to be extended to a multi-user application. We decided to create a plugin architecture to allow users to concurrently interact with models, access centralized storage solutions, and connect to third-party technologies. This solution is known as the Concurrent Generative Engineering Tooling platform (COGENT). With COGENT system designers can focus on their primary concerns, goals, and constraints.

As a solution architecture, COGENT extends SCS from a single-user user application to a multi-user application with simultaneous tool access. Additionally, plugins were developed to allow users to utilize external or custom software including graph databases, experiment tracking, and user-defined modules. Furthermore, a centralized cloud solution was used to store objects, assets, metadata, and models in which end-users have concurrent and synchronous information. COGENT was designed as a high-level solution that can be extended to Industry 4.0 domains containing system-of-systems such as aerospace, automotive, industrial robotics, biomedical devices, and embedded Internet-of-Things.

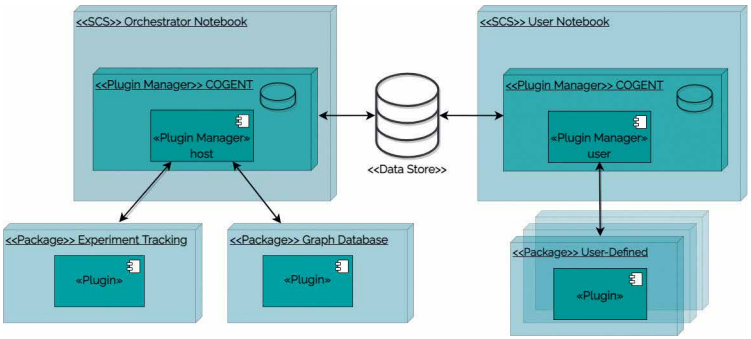

"The hard work of Tuvi resulted in an initial version of the next generation laser simulator software. It implements a sub-set of the complete functionality and already forms a solid basis for completion. The new object-oriented design provides the right levels of abstraction and drastically improves the maintainability, scalability and remote control capabilities of the software. She showed that a good software design in combination with a powerful and reliable platform and software stack solves the problems we currently have with the laser simulator.
I am very impressed by Tuvi's achievements and it was a true pleasure to supervise her!"

Sander van den Berg MSc
ASML

CHALLENGES

The challenge of this project was to reengineer the Windows OS specific legacy Application leveraging its performance in the interface. Additionally, the Application was demanded to be extensible and reliable. Moreover, the project was heavily dependent on the hardware, and the integration of software with the Dose Control Board (DCB) was quite challenging in the remote working situation.

RESULTS

The proposed design approach satisfies the following:

- significant reduction in the time to handle DC laser command as a request via a USB interface with the hardware, thereby improving the reliability of the laser simulator tool
- ability to monitor the laser simulator remotely

BENEFITS

The emphasis on a layered architecture of the laser simulator tool will help not to impact the hardware functionalities and the multithreading improvements help to increase performance in a USB transfer. Furthermore, the laser functionality as a core will allow developers to maintain and extend the software without extensive reengineering of code base.

TUVI PUREVSUREN PDEng

Redesigning a Laser Simulator Application

ASML manufactures lithography TWINSCAN NXT machines to print a circuit pattern on a wafer using a laser. Within these machines, the Dose Control (DC) subsystem controls the laser pulse by pulse to deliver the right amount of photons on the water at right moment in time. Due to the complexity of the laser technology, service pack qualification of the DC software uses a laser simulator tool that provides a realistic simulation behavior of a light source. With the years, the tool has proven that it responds to the DC requests late; therefore, service pack qualifications become cumbersome. Hence, the decision was made to replace the software of the tool. The goal of this project is to re-design, implement, and integrate this new software.

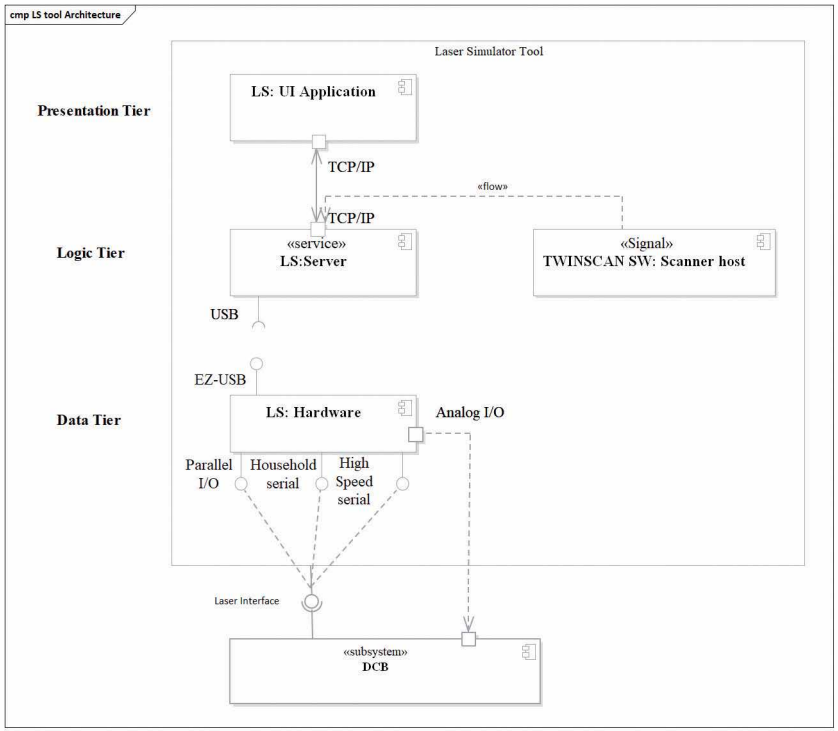

There were a few main requirements that shaped the design of the solutions: a need for an extensible and reliable laser simulator application for existing simulator hardware, thereby improving response timing of the tool and a need for an ability to monitor the laser simulator remotely.

In this project, a functional prototype was designed and developed for the laser simulator tool. The performance of the Laser Simulator (LS) new design was measured and compared to the existing application using a built-in test tool of the DC subsystem. The prototype provided a significant reduction in the timing of the hardware and software interface, thereby improving the reliability of the simulator tool.

CHALLENGES

This project was a feasibility study where the feasibility of the project was not clear at the beginning. The project was also greenfield, and it demanded creativity and proactive attitude in deciding the direction and scope of the project. Another challenge was that the domain was new to us, and we needed to understand the big overview of the Transmission Electron Microscopy (TEM) software architecture domain in a short time.

RESULTS

We designed and developed SoftWare execution Auto Tracing (SWAT). SWAT consists of one preprocessing module and multiple visualization modules. The preprocessing module processes the TEM log files into the uniformed format for the visualization modules. The visualization modules provide trace call visualizations based on the TEM log data. There are three visualization types: Graph Network using KBN Network, Time Series Pattern using Trace, and Common Execution Pattern using Process Mining.

BENEFITS

SWAT can speed up the execution path analysis by the TEM engineers. It can also help troubleshoot the software issues and assess the usage of the specific parts of the software stack. Additionally, the value of this project is not only related to the provided prototype but also to the possible improvements that can be done in the TEM logging environment.

"Yusril designed and created a software framework that helps us to reason about the execution architecture of the TEM software. Thanks to his efforts, we have tools that unify the format of the logs produced by different software components. Moreover, Yusril explored different ways of visualizing the execution paths and combining them with other information about the software running on the microscope.
His work makes it easier for us to analyze the software execution problems. The feasibility study conducted by Yusril is an inspiration to continue his work in the future."

Marcin Gramza MSc
Thermo Fisher Scientific

YUSRIL MAULIDAN RAJI PDEng

Conducting a Feasibility Study and Design for Software Execution Auto Tracing (SWAT)

The Transmission Electron Microscopy (TEM) software execution consists of a set of processes divided further into threads, forming multiple execution paths. The software execution paths of individual threads are recorded in log files that are separated logically according to the TEM software reference architecture. Roughly, the architecture is divided into several functional blocks, called building blocks. Each software building block has its own log format. In consequence, analysis of such logs to track calling sequences made by the software is difficult. In addition, there is a need to analyze the execution paths for solving the software issues (e.g., performance issues and integration issues) or assessing the usage of the specific parts of the software stack.

We designed and implemented a proof of concept system called Software execution Auto Tracing (SWAT) that allows TEM developers to easily visualize the calls between software components inside TEM.

SWAT system consists of one preprocessing module and multiple visualization modules. The preprocessing module is intended to clean and process the TEM log files into the uniformed format for the visualization modules. The visualization modules are intended to provide trace call visualizations based on the TEM log data and consist of three visualization types which are Graph Network visualization using KBN Network, Time Series Pattern visualization using Trace, and Common Execution Pattern visualization using Process Mining.

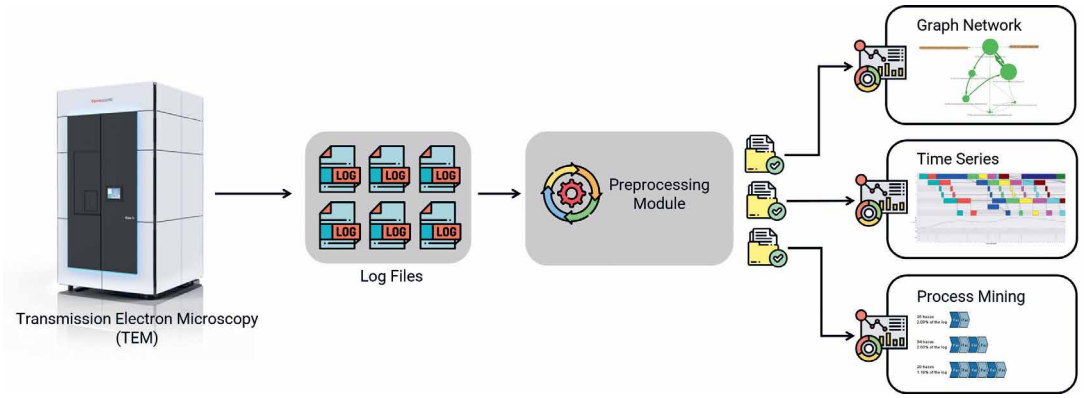

"The data pipelines developed in this project are an important first step to develop an AI-based system to automatically detect individual behavior of farm animals to improve animal welfare and identify which animals are genetically the most suitable for living peacefully in large groups."

Prof.dr. J. (Jakob) de Vlieg
Professor of Integrative Data Science
Chair AgriFoodTech@TU/e
Lead JADS AgriFood&Data program line

CHALLENGES

The challenge in this project was understanding the necessities of several non-technical key stakeholders and developing a software pipeline to help them in their research. Animal behavior analysis is a time-consuming qualitative task, to be performed on hundreds of individuals on the same farm. Therefore, computer vision was chosen to detect, track, and analyze individual animals belonging to huge groups. The output of these AI algorithms (i.e., text files) was too shallow, difficult to grasp for animal experts.

RESULTS

The result was a data pipeline, Insight, that allowed the researchers to analyze AI-generated data to obtain graphical results. This data pipeline focuses on simplicity to use and extend. It allows multiple configurations to obtain the analysis required by the situation. The current pipeline allows obtaining an analysis of trajectories, distances, velocities, proximity between animals, basic behavior analysis, and social networks. It can be extended to include new types of plots, analyses, or animals.

BENEFITS

Thanks to Insight, animal behavior researchers have a helpful tool to complement the use of AI to understand the social behavior of animals. This data pipeline can be further integrated into their interdisciplinary data platform needed to process the huge amount of heterogeneous data that AgriFood projects produce. The pipeline can be extended with new analytics and adapted for new kinds of animals or different projects.

LUIS ROMA BARGE PDEng

The Design of the Insight Pipeline for Behavioral Animal Science & Animal Breeding

Even in the best conditions, livestock animals live in crowded farms and may suffer stress. In most cases, this stress causes the animals to hurt fellow animals as an instinctive response. Additionally, this negative behavior is contagious; if one animal starts doing it, others will follow. These stress-caused reactions have a negative impact not only on the welfare of the animals but also on the economy of the farm.

Until now, the only way of monitoring behavior was manual (i.e., actively observing the animals). This task results in difficult and tedious work considering that the average farm counts hundreds (or thousands) of animals. Reacting on time to stop negative behavior and minimize losses has proven challenging to animal breeders.

The IMAGEN program aims to automatically study, analyze, and recognize animal behavior by using Computer Vision technology and, at a later stage, link this behavioral information to specific genes of individual animals. Cameras have been set up in multiple livestock farms. These cameras are used to collect animal data (videos). These videos contain animals living on farms and how they behave in a social group. These videos are used to produce AI models that automatically detect and track individual animals.

Insight is a data pipeline that analyzes the output of AI algorithms to help researchers conduct animal behavior research. By analyzing a specific time period of a video of animals, these pipelines produce the data plots that contain animal information such as trajectory and distance traveled, proximity encounters, basic animal state, and social networks based on activities.

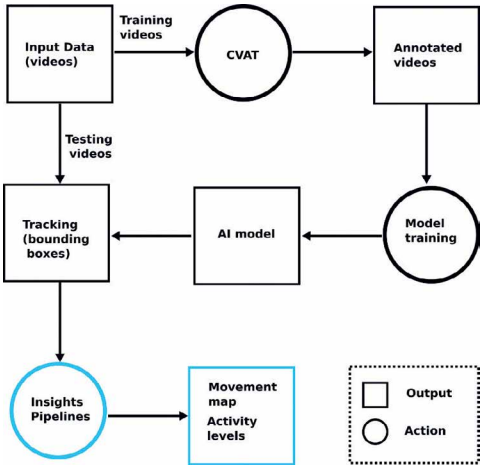

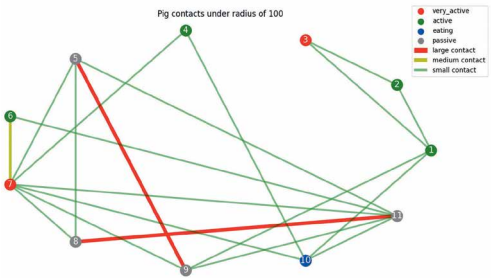

CHALLENGES

This project was aimed and developed to research and design the integration of Digital Working Capital (DWC) platform with an Operational Data Store (ODS) and to harmonize the raw data for subsequent operational reporting. The Global Client Data System (GCDS) integration with DWC was also critical since related functional data flowing into DWC was required in the first phase of the enterprise data strategy initiative.

RESULTS

As part of this project, we presented an analysis and proof of concept solution for data propagation, enabling DWC to generate operational reports and analysis. In general, the integrations presented in this project should be considered as the basis for the subsequent connection to the data lakes, which, in turn, provides ample opportunities for data analysis and forecasting for the entire department.

BENEFITS

Digitalization of value chain finance area in general and functional reporting for DWC system in particular (through the creation of automatic DWC-ODS integration) expands the range of products provided by W&R department to their clients with the same or reduced costs. GCDS-DWC integration is essential for full DWC-ODS pipeline and also helps to reduce number of errors in user-uploaded data and accelerates the document preparation process.

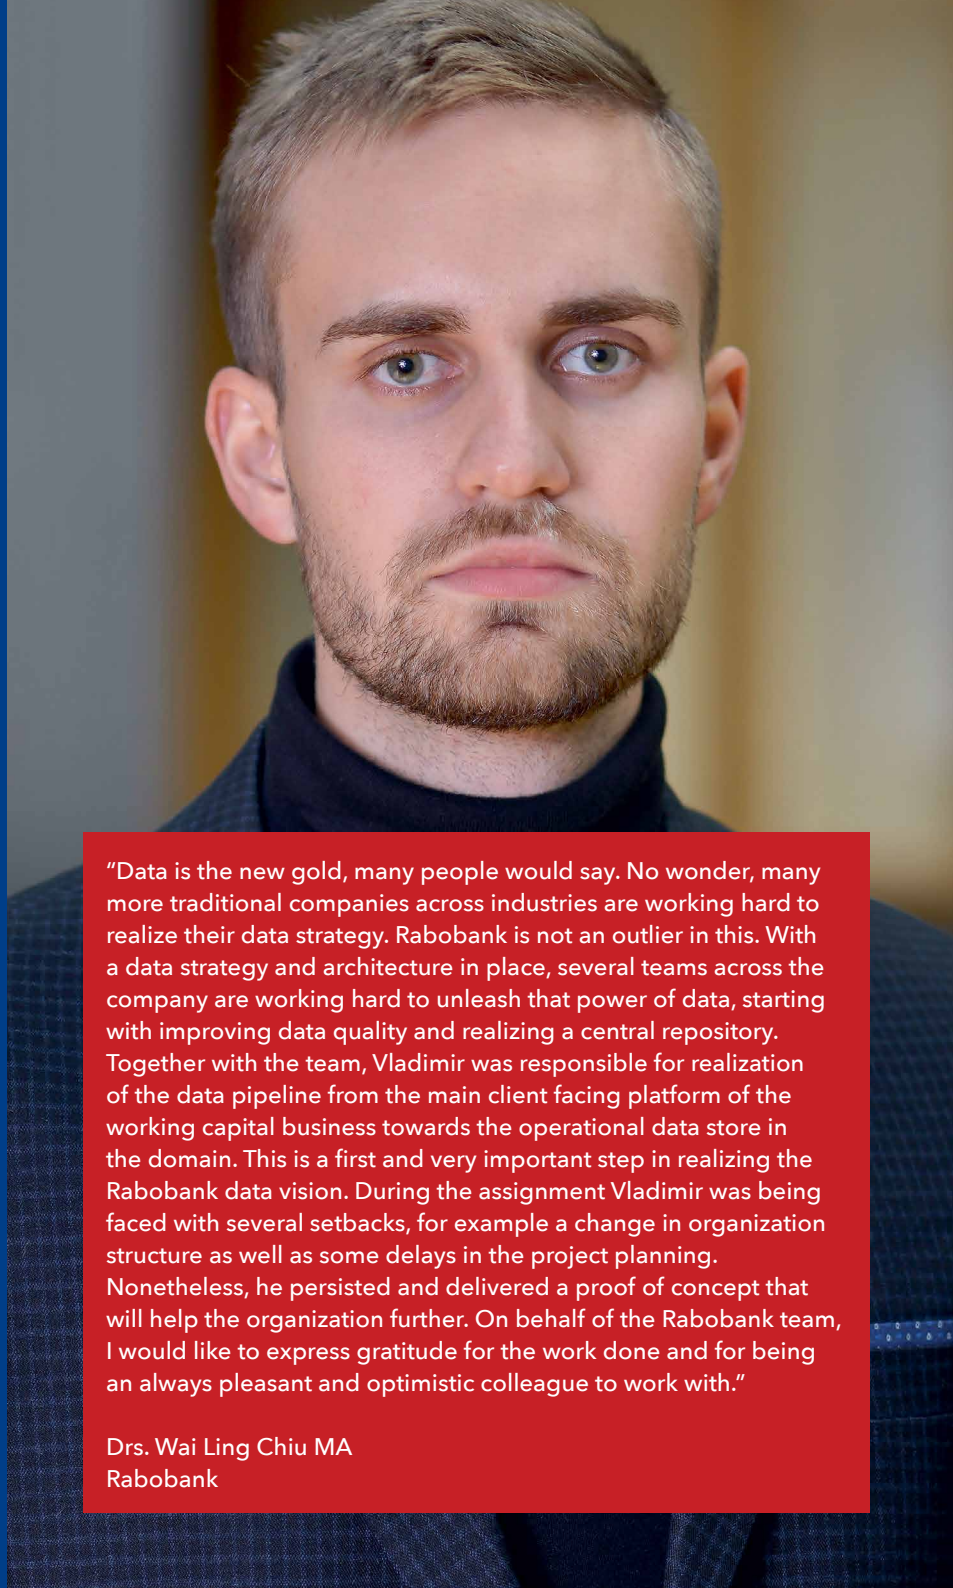

"Data is the new gold, many people would say. No wonder, many more traditional companies across industries are working hard to realize their data strategy. Rabobank is not an outlier in this. With a data strategy and architecture in place, several teams across the company are working hard to unleash that power of data, starting with improving data quality and realizing a central repository. Together with the team, Vladimir was responsible for realization of the data pipeline from the main client facing platform of the working capital business towards the operational data store in the domain. This is a first and very important step in realizing the Rabobank data vision. During the assignment Vladimir was being faced with several setbacks, for example a change in organization structure as well as some delays in the project planning. Nonetheless, he persisted and delivered a proof of concept that will help the organization further. On behalf of the Rabobank team, I would like to express gratitude for the work done and for being an always pleasant and optimistic colleague to work with."

Drs. Wai Ling Chiu MA
Rabobank

VLADIMIR ROMASHOV PDEng

Better Business by Integrating Heterogeneous Data from the Entire Value-chain

This project is affiliated with the Wholesale and Rural (W&R) department that is engaged but not limited to the creation of a single client journey for all the working capital needs of Rabobank's wholesale customers. Digital Working Capital (DWC) is a system that provides working capital solutions to our international wholesale clients and serves as a starting point for customers.

Rabobank is moving towards initiatives to develop and implement an enterprise data strategy that will integrate and complement local departmental initiatives. In order to be able to provide centralized reporting, Rabobank presented an Operational Data Store (ODS) that was designed to harmonize operational data for subsequent operational reporting and further distribution. Digitalization of functional reporting for DWC system particular through the creation of automatic DWC-ODS integration will help to expand the range of products provided by W&R department to their clients.

As part of this project, we presented an analysis and proof of concept solution for data propagation, enabling DWC to generate operational reports and analysis. In general, the integrations presented and developed in this project should be considered as the basis for the subsequent connection to the data lakes. Thus, in this project we considered both the principle data delivery alternatives and connectivity alternatives.

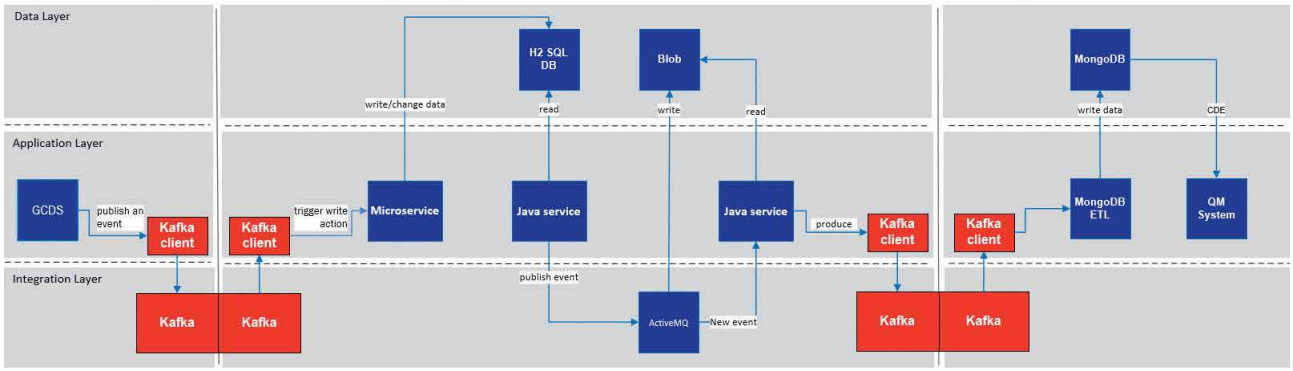

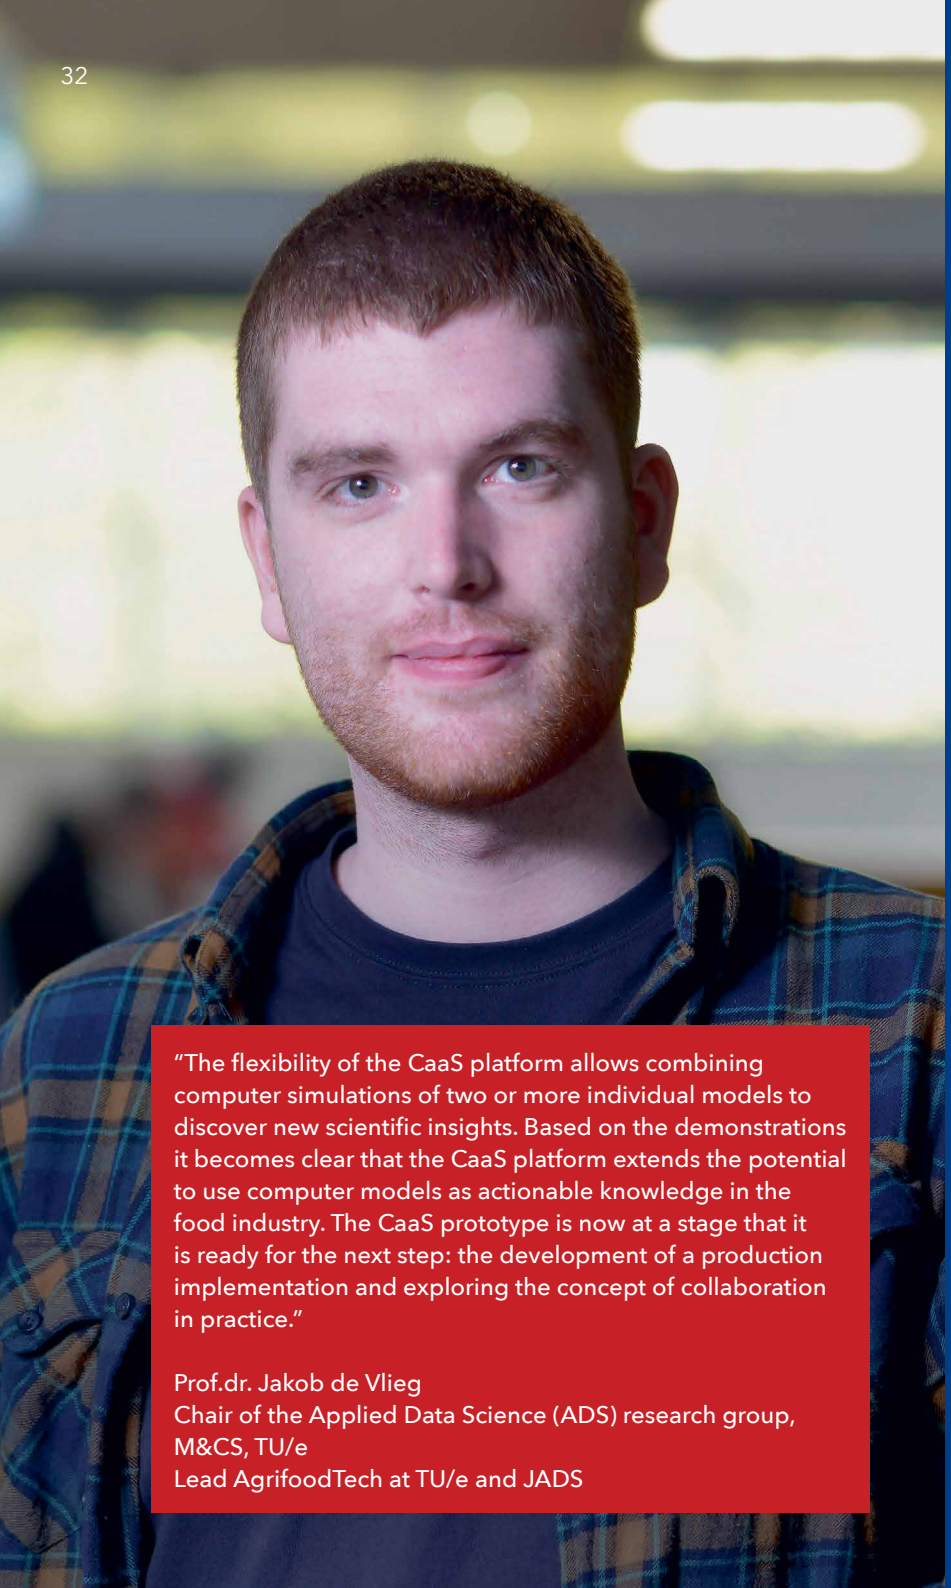

"The flexibility of the CaaS platform allows combining computer simulations of two or more individual models to discover new scientific insights. Based on the demonstrations it becomes clear that the CaaS platform extends the potential to use computer models as actionable knowledge in the food industry. The CaaS prototype is now at a stage that it is ready for the next step: the development of a production implementation and exploring the concept of collaboration in practice."

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

CHALLENGES

The main challenge in this project was to find a way to support the varying example use cases provided by industry partners. The models and data that are part of these use cases have very different structure and interface requirements. The platform solves this by putting a gateway between the data and model sources and the central sharing platform. The interface provided by the gateway exposes the underlying data source or model, while also supplying an ontology that further describes them.

RESULTS

The main result of the project is the Collaboration as a Service (CaaS) platform. It adds data gateways that allow organizations to share data from their data sources according to the FAIR principles. The model gateways and the model simulation module have also been updated to follow the data-focused approach. To show the capabilities of the new approach, three new use cases were implemented, in total involving five new model gateways and four data gateways. Modular data and model gateways were created to simplify the creation of all these gateways, as well as the integration of future use cases.

BENEFITS

Moving data and metadata storage away from the central platform gives organizations more control over what data they share in what way. It also makes the platform more flexible for adding new use cases: models and data sources can be integrated with the platform easily by adding a gateway. The use of ontologies also allows for additions to be made in the metadata by organizations. They can easily add their own metadata and relate it to the metadata that is required by the CaaS platform.

NIELS ROOD PDEng

Collaboration as a Service

EXTENDING THE CAPABILITIES OF THE MODEL AS A SERVICE PLATFORM FOR THE INTERNET OF FOOD

The Sustainable Food Initiative is a collaboration between companies in the food industry and academia that want to improve the sustainability of the food industry by exploring innovative technologies. One of the projects conducted by the Sustainable Food Initiative is the Internet of Food project. The Internet of Food project aims to find ways for organizations to collaborate in in-silico food product development, making food production more efficient and produce less waste.

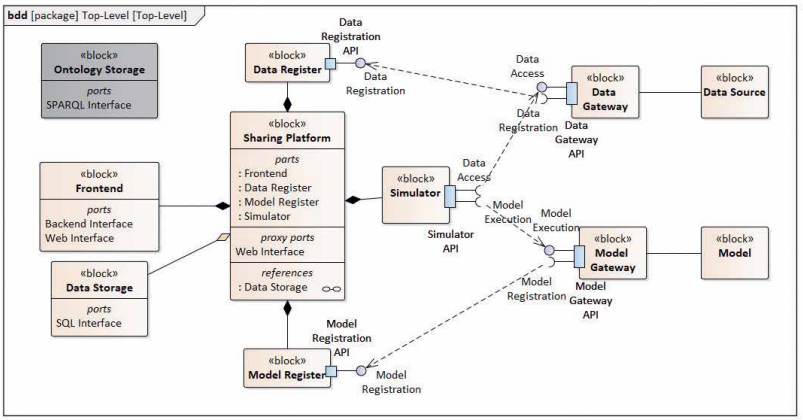

The Internet of Food project produced the Model as a Service (MaaS) prototype, which showed the possibilities for organizations to share their food-development-related models without giving them away. While this prototype was a good starting point, it was not flexible enough to support the example use cases provided by industry partners. It also stored all shared data on a central server, making it less attractive for organizations to share their data. The goal of the Collaboration as a Service project was to improve on the MaaS prototype to remove these shortcomings.

The Collaboration as a Service (CaaS) platform takes a data-focused approach. An ontology vocabulary was developed that allows organizations to describe the data provided by their data sources, and the data required and returned by their models. This approach allows for users to define on a per-column level what data should be provided to what model. Data storage is also moved to the owner of the data, bringing it at the same level as model sharing. To show the potential of this new approach, three new use cases were implemented.

CHALLENGES

The main challenge in this project was to find a proper and efficient pattern recognition solution with comparable accuracy, performance, and flexibility as its commercial counterpart. Finding a trade-off between accuracy and performance played a major role. A few more challenges faced are mainly related to the broad domain or way of working.

RESULTS

RECOG, which is the proposed pattern recognition solution, is designed and implemented in C# and it can be extended readily to address new concerns. ART is a pattern recognition toolbox with several features that facilitates recipe conversion, batch recognition, and finding the best set of parameters that yields the best pattern recognition result. Finally, RECOG was integrated into the YieldStar to demonstrate the feasibility of this integration and evaluate its performance and accuracy.

BENEFITS

RECOG, as a Cognex replacement, could significantly cut building costs at ASML. Each license costs about 3.5k euros, each machine needs a license, and there are more than 600 machines. RECOG doesn't require any hardware dongles, which reduce the system complexities from different aspects, especially for developers. As the proposed solution is transparent to developers, extending the solution or customizing it will be easy and straightforward.

"Raha was challenged to construct a replacement for Cognex (a pattern recognition tool used for wafer alignment) with the potential to: (1) reduce the cost, (2) enable further improvements. Raha worked hard and cleverly to find a very good replacement to Cognex, implemented the toolbox, and delivered the results to the YieldStar code archive. I am very happy to have had Raha to work on this project. I am very impressed by the result and the way she could work independently in this unique period where most of us are working remotely. I wish her all the best in ASML and we will work together again soon."

Dr. Zhifeng Sheng
ASML

RAHA SADEGHI PDEng

Feasibility and Prototype of Replacing Commercial Off-the-shelf Pattern Recognition Solution

ASML, as one of the leading lithography tool providers in the world, has delivered a proven holistic lithography solution to maximize the patterning process performance and control. In the ASML holistic lithography, YieldStar is a metrology tool that provides closed-loop feedback to scanners by measuring on-product errors such as overlay and focus. YieldStar utilizes image pattern recognition techniques to measure the position shifts of wafer alignment marks. It uses this position information to build up a high-order wafer model to guide measure target positioning. YieldStar uses a Commercial-Off-The-Shelf (COTS) software library (Cognex), which has expensive license costs and requires a hardware dongle, complicating machine build-up.

In this project, a new pattern recognition solution as a library with C# interface (RECOG) was proposed as a Cognex replacement in the context of YieldStar wafer alignment. RECOG was developed by taking advantage of free open-source libraries. A new GUI application based on RECOG was designed and implemented that enables pattern recognition. The figure depicts the high-level architecture of the YieldStar wafer alignment module, and how both libraries can be supported simultaneously. As extensibility was one of the design criteria, RECOG can support new algorithms with minimum effort.

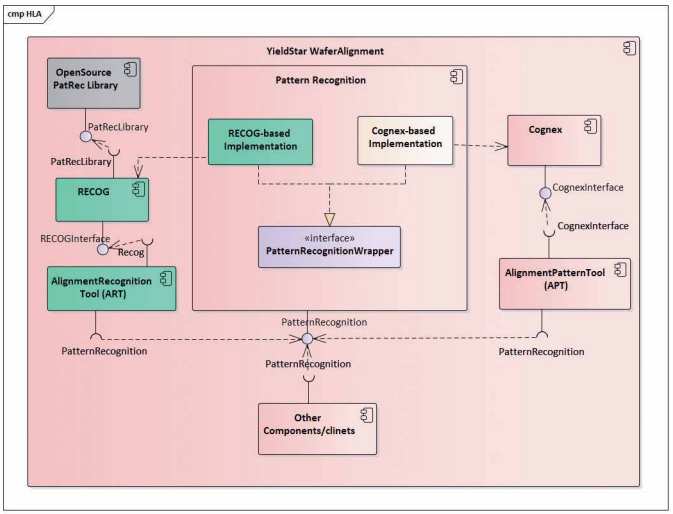

The results showed that RECOG could recognize patterns accurately with almost less than a pixel difference with Cognex. RECOG met the performance specifications. According to the results, RECOG is flexible enough to recognize various patterns even in noisy and low-contrast or dark images in more than 95% of cases and as good as Cognex.

"Akram's work in the MX Innovation Hub team was valuable, and she completed the goals of the project. One of Akram's biggest challenges at the start of the project was the huge amount of interest in the MaaS project. Akram had to pitch the MaaS project multiple times to different end-users cross ASML D&E, not only within MX. Some development teams already had ideas for hundreds or more metrics to be measured. She created a shortlist of requirements which she later focused on completing these requirements.
During the detail design phase, the product owner was already required to provide metrics for the current status of the whole MX cluster development process. The definition of these key indicators took a lot of discussion and alignment; however, gathering the data and creating the first dashboard was not a problem for Akram. The dashboard was shared with the MX cluster staff and set a benchmark, which showed the impact of the improvements made by the development teams and the MX Innovation Hub team."

Johan Krielen
ASML

CHALLENGES

The challenge in this project was designing a system to monitor metrics from all possible types of software development and deployment tools varying from on-premises to cloud-based. The other challenge was defining an extendable data schema for metrics that imposes a well-defined structure on data while allowing the addition of new metrics.

RESULTS

In this project, MaaS was designed and developed to monitor three software development tools by automatically collecting and visualizing a set of metrics from each of them. These three data sources were examples of how the system should be used. Several dashboards containing different graphs were created to visualize the metrics and provide insight into the efficiency and performance of the tools and their usages.

BENEFITS

The monitored data sources provided insight into the software development process for different teams. They help managers and developers observe the trends and current status of their development process by specifying the level of details or the time interval that the metrics are visualized using the provided filters on dashboards. Using these observations, teams can plan improvements and changes in their software development approaches to deliver high-quality software at a higher speed.

AKRAM SHOKRI PDEng

Monitoring as a Service (MaaS)

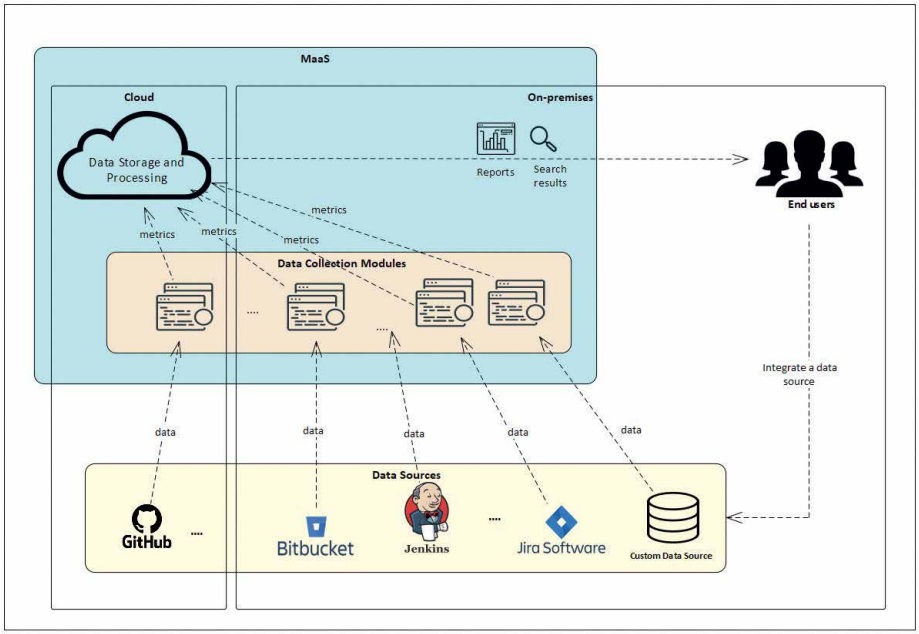

The Quality and Continuous Improve group of ASML, part of the Metrology and Machine Control (MX) cluster, is responsible for assessing the quality of development and deployment processes and constantly finding bottlenecks, and suggesting improvement steps to assist teams in delivering high-quality products at high speed.

The problem that they were addressing was that the software development speed was not fast enough, and scaling teams did not improve it. Therefore the software development processes should be improved. To evaluate and improve software development and deployment processes, the Innovation-Hub team in this cluster proposed the Monitoring as a Service (MaaS) project as a framework to enable developers and managers to have a better view of the performance of their process based on actual data, by gathering, analyzing, and visualizing metrics from various software development and deployment tools.

The goals of this PDEng project were to design MaaS, decide about the tools and technologies to be used for its realization, and develop the first version of it, which all were successfully achieved. The developed product consists of automatic metrics collection and monitoring for Jenkins, Patch Decision Form (PDF) tool, and T&I Database, as three different software development and deployment tools (data sources), and several dashboards to visualize the efficiency of these tools' usages. These dashboards provide filtering over the visualized metrics based on different conditions such as time, department, cluster, and delivery type. These filters help managers and developers to see different levels of details by combining them.

CHALLENGES

The primary design challenges for this project were to

- bidirectionally connect cloud applications with on-premises applications as well as applications cross-premises
- identify and address (globally deployed) applications
- ensure end-to-end secure communications between applications

Additionally, the solution had to be generic and acceptable such that it works for, and is accepted by, at least 90% of Philips' (global) customers.

RESULTS

The project results consist of a

- system design that solves all the design challenges
- proof of concept implementation that validates the design
- qualitative design validation study involving system experts, users, and customers and resulting in support for the design

BENEFITS

The benefits of this extension for HealthSuite, for both Philips and its customers, are that we have

- support for secure bidirectional hybrid cloud and cross-premises customer solutions
- an integration platform that is highly secured and controlled

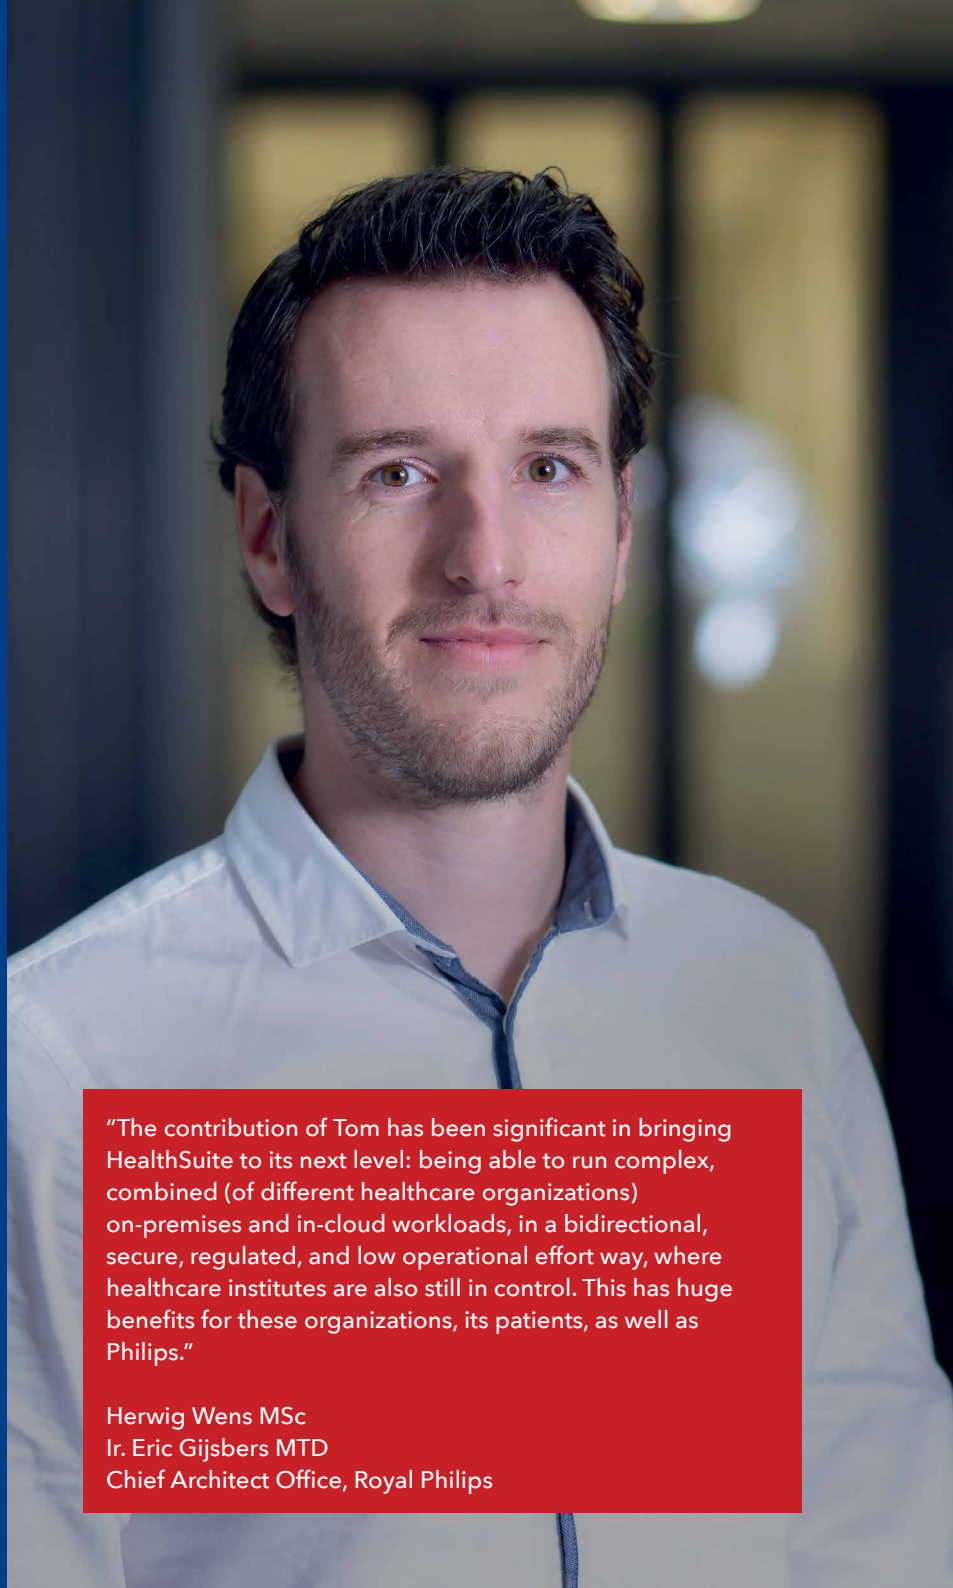

"The contribution of Tom has been significant in bringing HealthSuite to its next level: being able to run complex, combined (of different healthcare organizations) on-premises and in-cloud workloads, in a bidirectional, secure, regulated, and low operational effort way, where healthcare institutes are also still in control. This has huge benefits for these organizations, its patients, as well as Philips."

Herwig Wens MSc
Ir. Eric Gijsbers MTD
Chief Architect Office, Royal Philips

TOM VRANCKEN PDEng

Design and Development of a Secure Omnidirectional Communications Relay for Philips HealthSuite

Within Philips, we see that the Healthcare landscape is changing rapidly. Powerful drivers like excessive cost, staff scarcity and overload, and consumerization call us to rethink Healthcare at scale and require us to leverage digitization, data insights, and virtual care. Philips, as a leading health technology company, has the ambition to connect healthcare organizations to its certified hybrid cloud based healthcare solution called HealthSuite. HealthSuite allows these organizations to still run healthcare based workloads on-premises, but, at the same time, also benefit from HealthSuite Cloud. The goal of this project, was to design and develop an extension of HealthSuite that enables integration of applications in the Cloud with applications deployed on-premises, in a secure, unified, and standardized manner.

One of the key rationales behind this “hybrid” focus is a raised privacy awareness among people and organizations because health information is considered as extremely privacy sensitive. Increasingly, support for concepts like data locality and data sovereignty are being demanded. Additionally, traditional concepts like data confidentiality, integrity, and systems auditability must be facilitated.

We designed and developed a secure omnidirectional communications relay that enables end-to-end secure bidirectional integration of cloud applications with on-premises applications, and between applications cross-premises. The solution supports data locality constraints, data confidentiality and integrity protection, and at the same time offers the levels of transparency and auditability that are required by law. Our HealthSuite extension enables new and improved use cases for Philips and its customers and contributes to an improved healthcare ecosystem.

tue.nl/softwaretechnology

Follow us:

linkedin.com/company/tue_pdengst

facebook.com/TUePDEngST

twitter.com/tue_pdengst

CREDITS

Edited by: Dr. Y. Dajsuren PDEng and M.L.L. van Oorschot, Software Technology Programme, Eindhoven University of Technology

Text: Graduates of generation 2019 of the PDEng Software Technology Programme

Photography: Rien Meulman

Production: Software Technology Programme, Eindhoven University of Technology

Design: Grefo Prepress

Printing: De Digitale Drukker

Georgios Azis PDEng; Armoring Formal Components against Foreign Behavior ■ **Samsom Tsegay Beyene PDEng;** Redesigning the Image Sensor Subsystem Driver using the Object-Oriented Paradigm ■ **Jorge Alberto Cordero Cruz PDEng;** Designing a Solution Architecture for Monitoring Credit Scoring Analytic Models ■ **Nathan D'Penha PDEng;** A Tool to Improve Reporting on Test Execution and Test Coverage ■ **Han Gao PDEng;** Designing a Solution for Software Distribution towards Varying On-premise Deployments ■ **Juan van der Heijden PDEng;** Standardized Software Solution for Guidance of Clinical Workflows ■ **Jobaer Islam Khan PDEng;** Automated Sequence Planner – A Tool to Create an Optimal Sequence Using AI Planning ■ **Hossein Mahdian PDEng;** IntelLight+: Designing and Developing of Context-Aware Lighting Solution ■ **Parima Mirshafiei PDEng;** From Data to Insights to Drive Sustainable Change in Philips' Global Logistics Movements ■ **Christopher O'Hara PDEng;** COGENT: Concurrent Generative Engineering Tooling – Enabling Cross-Functional Teams in Architecture Design for Space Subsystems ■ **Tuvi Purevsuren PDEng;** Redesigning a Laser Simulator Application ■ **Yusril Maulidan Raji PDEng;** Conducting a Feasibility Study and Design for Software Execution Auto Tracing (SWAT) ■ **Luis Roma Barge PDEng;** The Design of the Insight Pipeline for Behavioral Animal Science & Animal Breeding ■ **Vladimir Romashov PDEng;** Better Business by Integrating Heterogeneous Data from the Entire Value-chain ■ **Niels Rood PDEng;** Collaboration as a Service – Extending the Capabilities of the Model as a Service Platform for the Internet of Food ■ **Raha Sadeghi PDEng;** Feasibility and Prototype of Replacing Commercial Off-the-shelf Pattern Recognition Solution ■ **Akram Shokri PDEng;** Monitoring as a Service (MaaS) ■ **Tom Vrancken PDEng;** Design and Development of a Secure Omnidirectional Communications Relay for Philips HealthSuite

The PDEng (Professional Doctorate in Engineering) 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