I/O magazine

ICT RESEARCH PLATFORM NEDERLAND

VOLUME 20 APRIL 2023 NO 1

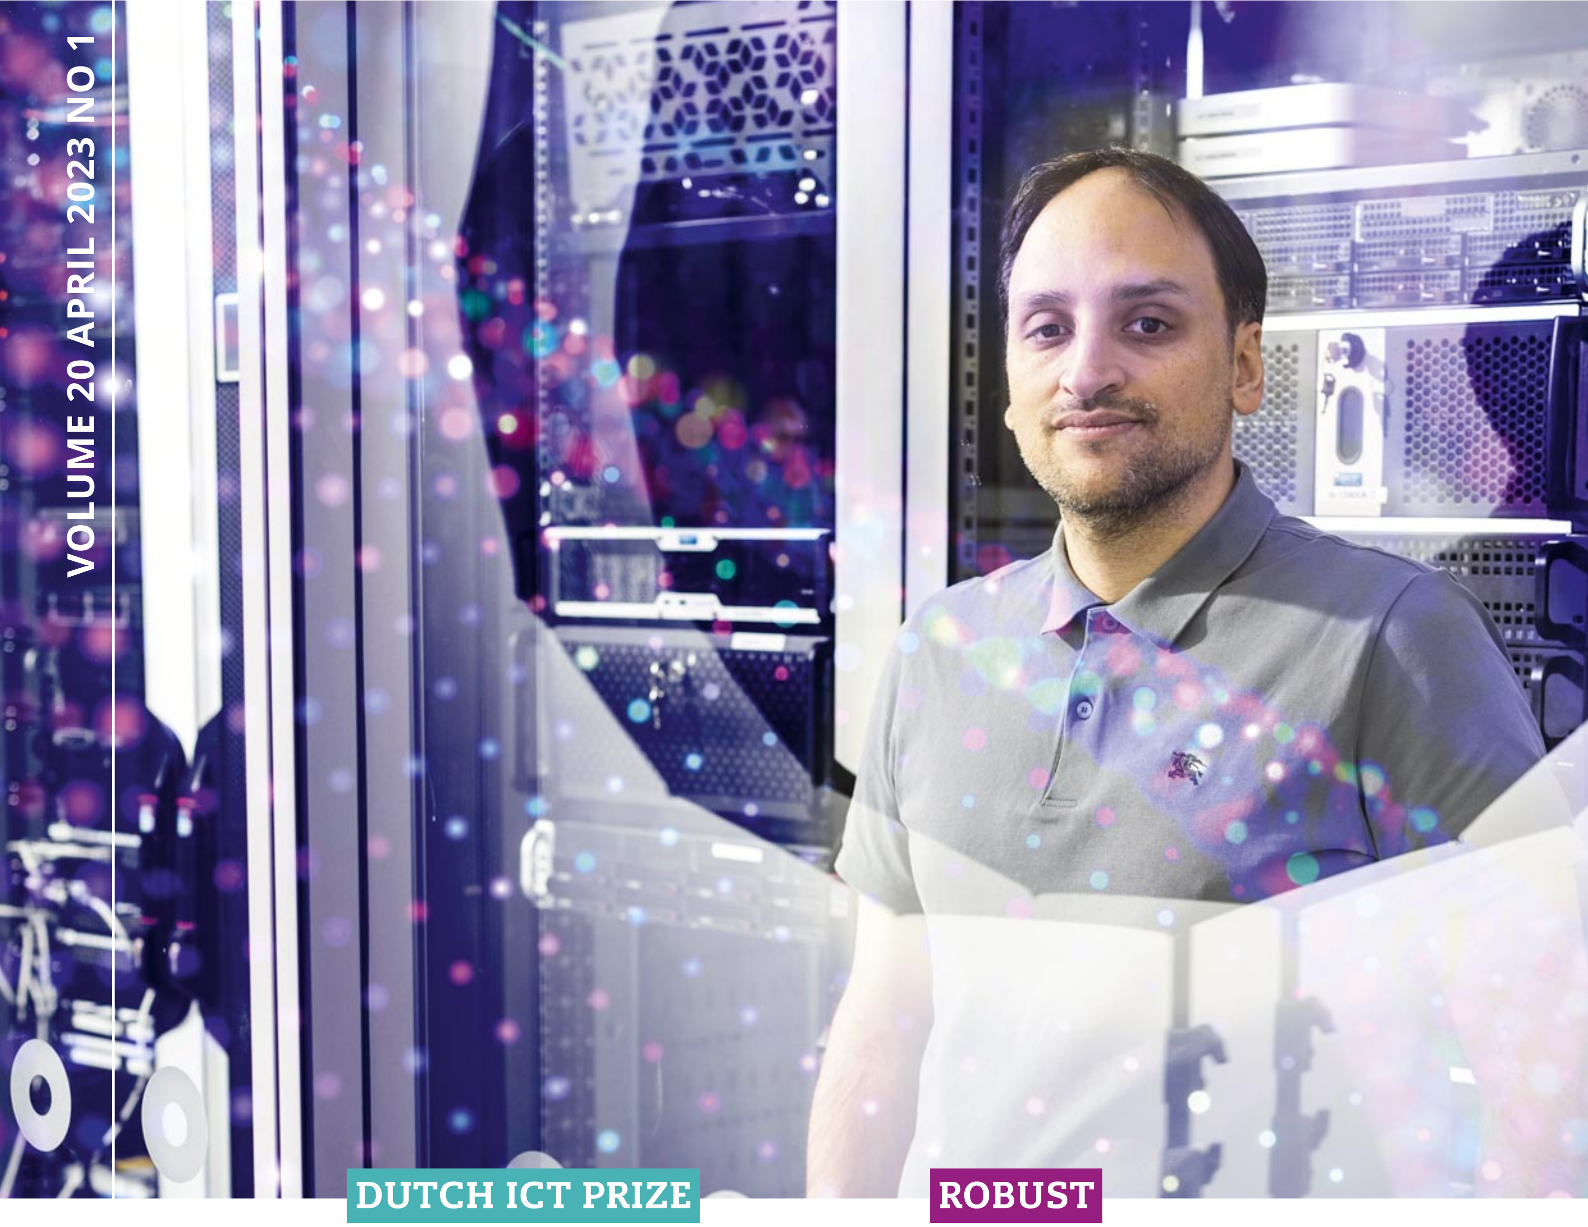

COLOFON

I/O Magazine is a publication of the ICT Research Platform Nederland (IPN) and is sent free of charge to ICT researchers and relations of IPN four times a year. IPN consists of the ICT research schools ASCI, IPA, and SIKS; the ICT-related themes of NWO domains Science (ENW) and Applied and Engineering Sciences (AES); the institutes of the technical universities, united in NIRICT; the institutes of the general universities; SURF; e-Science Center; CWI; Dutch Platform for Mathematics; Data Science Platform Netherlands; Dutch Tech Center for Life Sciences; VERSEN; TNO and COMMIT.

IPN (ICT Research Platform Nederland) unites all Dutch academic research groups that have ICT science as their core, and as such acts as a single point of contact for all matters relating to ICT innovation and its importance for our current and future society. IPN is supported by the NWO Domain Science.

EDITORS Margit Bleijs, Eveline van den Bos, Juul Brouwers, Sonja Knols, Melanie Lemmen, Femke Stephan **FINAL EDITING** Sonja Knols **TEXT CORRECTION** Dave Thomas, NST Science **CONTRIBUTORS** Marysa van den Berg, Leendert van der Ent, Sonja Knols, Bennie Mols **DESIGN** WAT ontwerpers, Utrecht - Leon Hulst **PHOTOGRAPHY** Ivar Pel (cover, p. 12, p. 16), ASML (p. 4), TU/e – Vincent van den Hoogen (p. 6), Canon Production Printing (p. 7), University of Twente (p. 7), iStock (p. 9, p. 10, p. 18), WAT ontwerpers (p. 19), Sjoerd van der Hucht (p. 24)

EDITORIAL ADDRESS Secretariat IPN, c/o NWO Science, PO Box 93460, 2509 AL The Hague, the Netherlands, +31 70 344 0772, ipn@nwo.nl, www.ict-research.nl

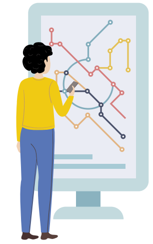

NWO

IPN ICT-ONDERZOEK PLATFORM NEDERLAND

PAUL KLINT

Quantification and gamification

Quantification and gamification dominate our lives, but do they lead to desirable results?

A lot has changed since the 1890s when Frederick Taylor was measuring the productivity of factory workers, stopwatch in hand. Taylorism (or scientific management as it is officially called) has led to productivity increase and prosperity. Today nearly everything is based on measurement: the news we read, the music we listen to, and the movies we watch.

Quantification has a deep societal impact too. The Dutch government assessed the behaviour of citizens based on metrics and this resulted in a massive disaster: the so-called "Toeslagenaffaire" where profiling was used to supposedly detect fraud with childcare benefits. Also, it is becoming harder and harder to get a mortgage if your metrics do not match the very specific ranges financial institutions expect. The list of examples is endless.

A close relative of quantification is gamification: based on metrics, a system can give an incentive to its users in the form of awards (points, levels, banners, privileges, ...). Cases in point are Genius levels for frequent users of Booking.com and levels and badges in Stackoverflow.com.

When there are rules and awards, people will always try to "game the game". This is true in society, organisations and ICT systems. The rules and awards become a proxy for reality, which sometimes leads to highly undesirable effects.

A personal anecdote may illustrate this. Over the past year, I have been using DuoLingo to learn some Italian. Nice app, good content and strongly gamified. At a certain stage, I realised that my behaviour was completely determined by my desire to reach the highest league and stay there. Instead of learning Italian in the best possible way, I was focused on actions that scored the most points. I have now turned off all gamification features: my pace is slower, but my focus is completely on language learning.

Quantification and gamification have a lot of merits. However, we should complement them with qualitative mechanisms to reduce their domination and negative impact.

BEATING WORLD RECORDS WITH NEXT-GENERATION HIGH-TECH MACHINES

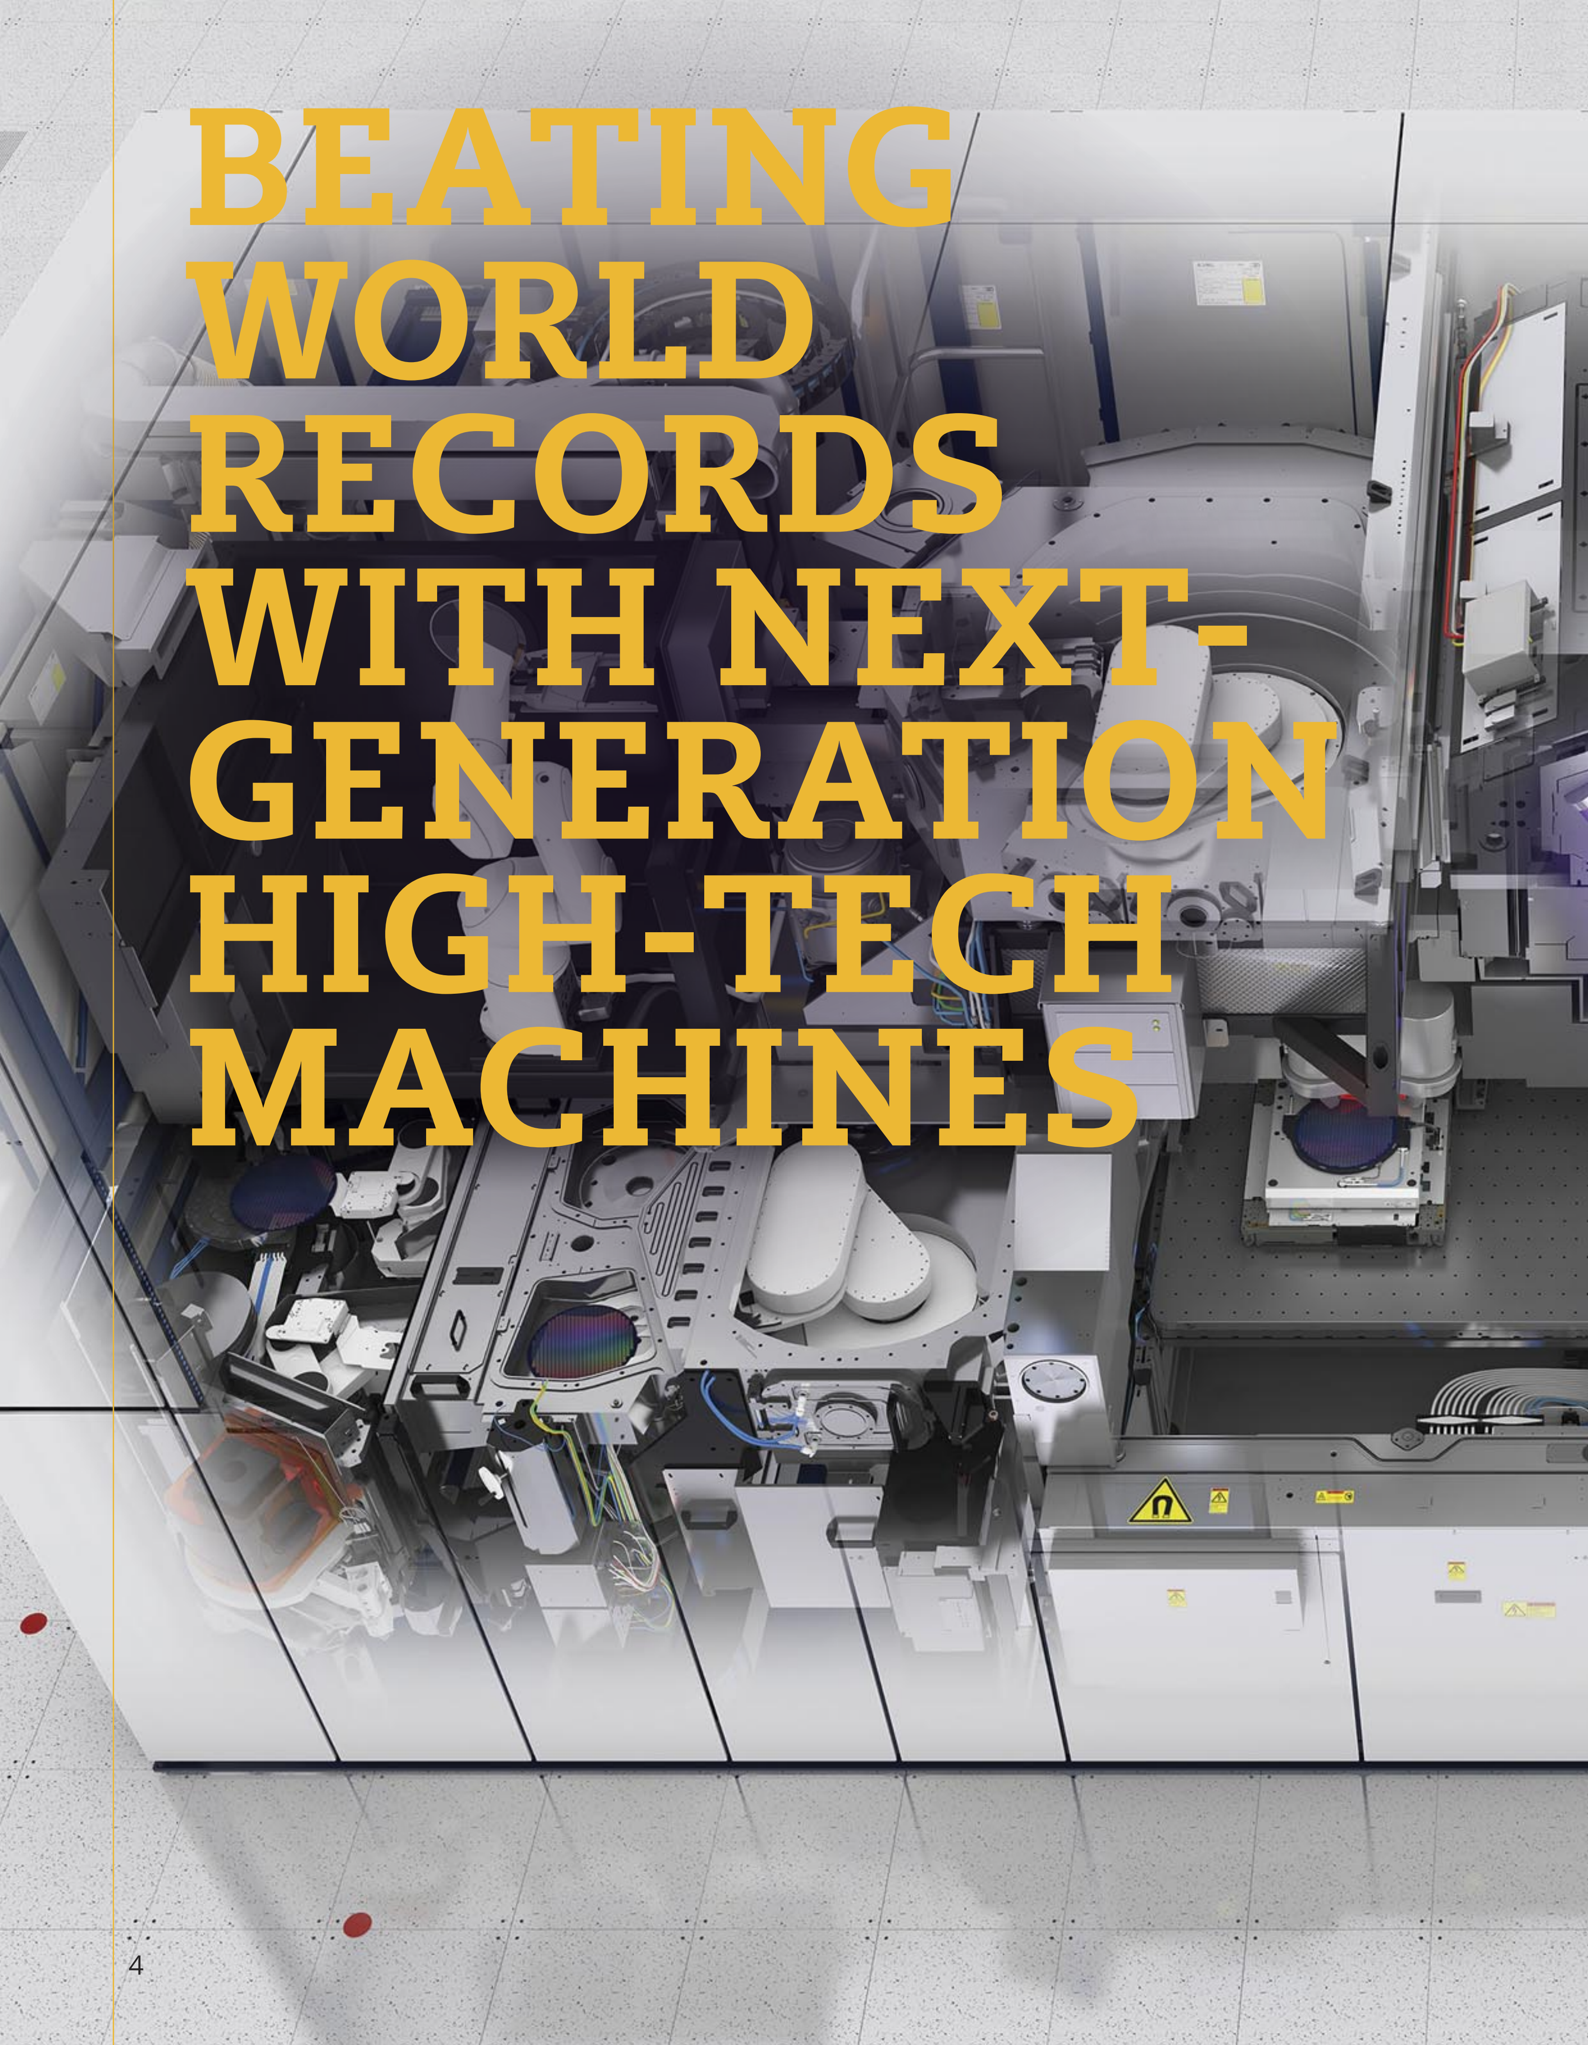

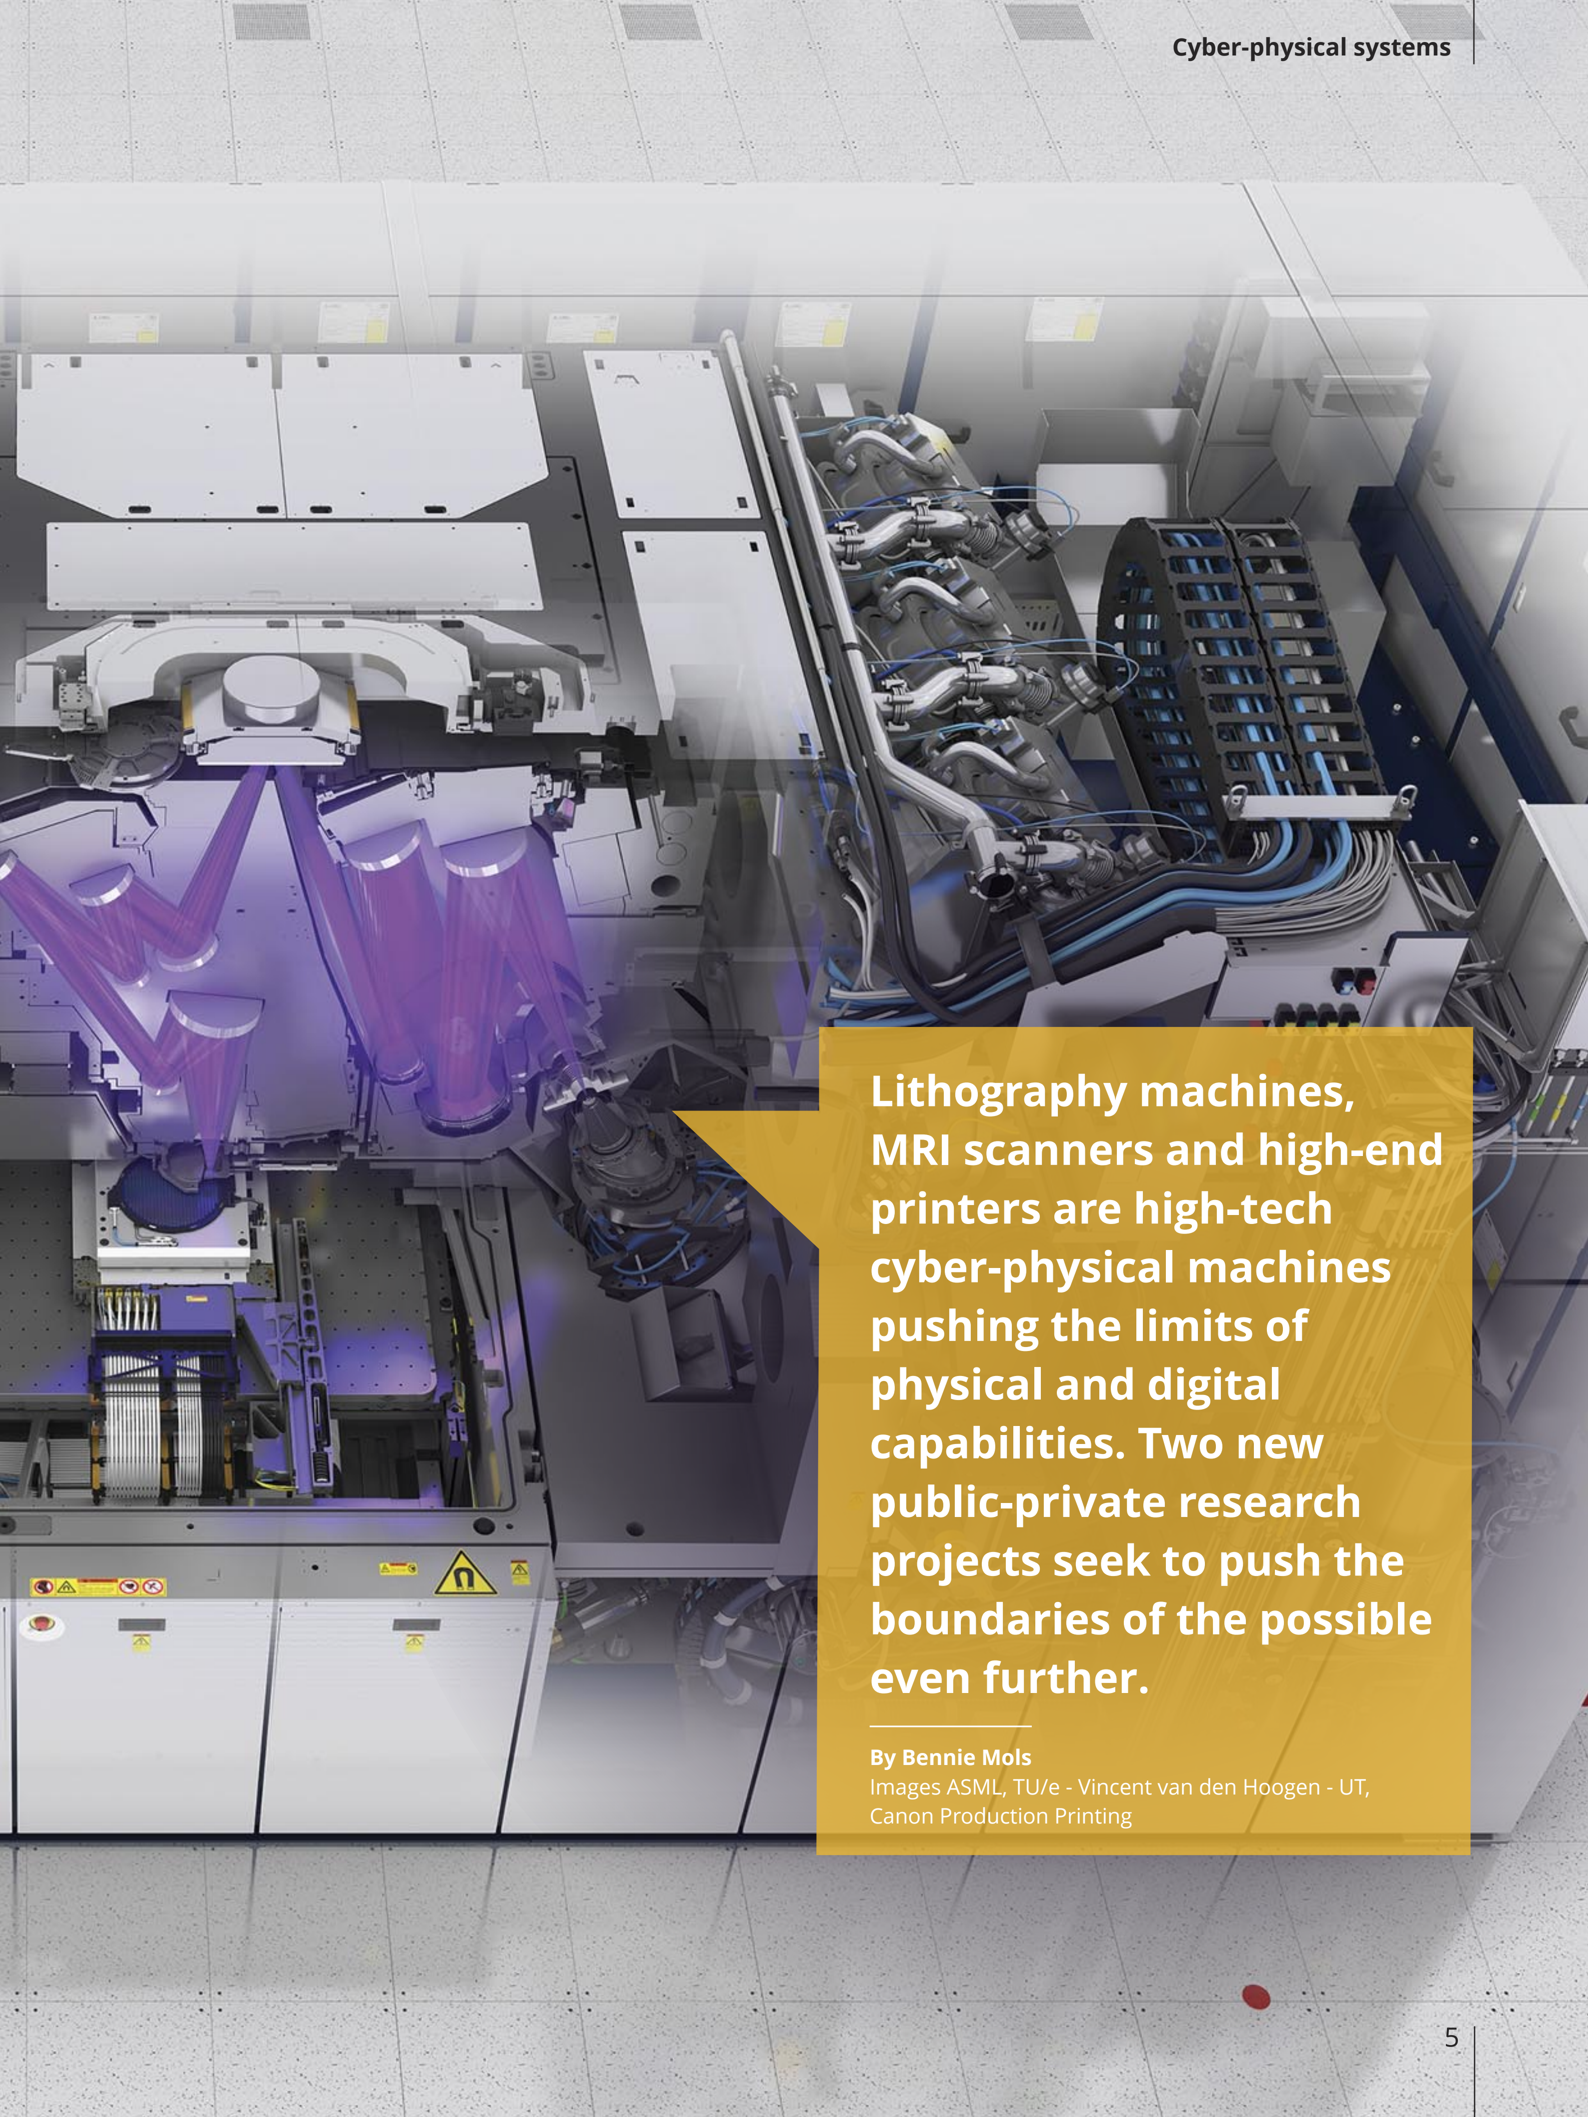

Lithography machines, MRI scanners and high-end printers are high-tech cyber-physical machines pushing the limits of physical and digital capabilities. Two new public-private research projects seek to push the boundaries of the possible even further.

By Bennie Mols

Images ASML, TU/e - Vincent van den Hoogen - UT, Canon Production Printing

An ASML lithography machine is a technological marvel, a masterpiece of engineering and precision. Like a painter's brush, the machine delicately etches patterns onto a silicon wafer, until complex computer chips are born. The muscles of such a lithography machine are the so-called actuators. They are responsible for moving and positioning with nanometre accuracy the silicon wafer, on which the computer chip is printed, and the reticle, a kind of photomask that holds the image that needs to be printed on the wafer.

Over the years, the positioning devices or actuators used for the wafer scanning process have developed as the demands on accuracy and acceleration have been continuously increasing. Unfortunately, all the improvements in acceleration and precision of the actuators that engineers have made possible over the years have also led to increased mechanical, electromagnetic and thermal stresses. This has increased the likelihood that the actuators fail. Problems like delamination of sub-components, other structural defects and partial discharge issues are already known to happen and can have potentially catastrophic effects for the lithography machine.

Elena Lomonova

'We specialise in the muscles of the machine'

Physics of failure

'The aim of our project *Reliable Next Generation Actuation Systems* is to understand the physics of the failure mechanisms in the actuators', states Elena Lomonova from Eindhoven University of Technology, who is leading the three million euro public-private research project. 'At the end of the project, we will give our tools and results to the people who are responsible for the cyber part of the machine, the part that contains the control algorithms. They program the brain of the machine, whereas we specialise in the muscles of the machine.'

Reliable Next Generation Actuation Systems is one of the two projects awarded within the KIC call 'NextGeneration High Tech Equipment'. The project will officially start in July 2023 and is a collaboration between researchers from Eindhoven University of Technology, Centrum Wiskunde en Informatica (CWI), TU Delft and two Dutch companies: ASML, world leader in the production of lithography machines, and Tecnotion, specialist in motor technologies, like the actuators used in ASML lithography machines. 'Once it is fully up and running, six PhD students and two postdocs will be working on the project', says Lomonova. 'There are a lot of connections between their work and they will do part of the tests and experiments at ASML and Tecnotion.'

An important part of the research into the failure mechanisms of the actuators is to study physical mechanisms at the microscopic level – what is happening on the scale of individual atoms – to the macroscopic level – what is happing on the scale of the whole actuator. 'Therefore, our consortium contains a multidisciplinary group of researchers', says Lomonova. 'They will do experiments on the various failure mechanisms and complement the experiments with modelling on multiple scales: atomic modelling to study interfacial dynamics and fracture, thermo-mechanical and electrostatic modelling to study strains and stresses on material interfaces, electric discharge modelling to study the role of partial discharges, and multi-physics modelling to study the combined effects of failure mechanisms.'

On the one hand, the project will address fundamental physical questions on the failures of actuators, but on the other, Lomonova also wants to develop practical solutions: 'We want to improve the reliability, lifetime and performance of actuators for our partners ASML and Tecnotion. Therefore, at the end of the project, we plan to build a demonstrator in which we want to integrate all the knowledge that we have acquired into new design principles for more reliable actuators.'

Toward zero downtime

Mariëlle Stoelinga from the University of Twente (UT) is leading a second project that was awarded within the KIC call: *Engineering for Zero Downtime in Cyber-Physical Systems via Intelligent Diagnostics* (ZORRO). Whereas Lomonova's project focuses on the physical part of high-tech cyber-physical machines, Stoelinga's project focuses on the cyber part.

'Our dot on the horizon is zero downtime for high-tech machinery like lithography machines, MRI scanners and high-end printers', says Stoelinga. 'We want to achieve this by developing intelligent diagnostics. Diagnostic algorithms are well-studied on the level of individual components. However, we need intelligent diagnostics that work on the level of the whole machine: that is where decisions on maintenance, spare management, or redesigns are made. Ideally, I envision a diagnostic system that explains what, when and why a particular part needs to be replaced because it might break down soon.'

In the ZORRO project, researchers from University of Twente, VU Amsterdam, TNO-ESI and Saxion University of Applied Sciences will collaborate with five industrial partners: ASML, Philips, Canon Production Printing, Thermo Fisher Scientific and ITEC. Five PhD students and three postdocs will be hired to do the bulk of the scientific research.

The researchers will work on four major scientific challenges, three on the cyber part and one on the physical part. ‘On the physical part, we want to know how to do better measurements’, says Stoelinga. ‘What do we need to measure? Where can we measure? How do we have to measure? In complex machines, you cannot place sensors in random places. The question is which sensors are absolutely necessary. We want to develop a reliable and resource-efficient monitoring system.’

Mariëlle Stoelinga

‘Our dot on the horizon is zero downtime for high-tech machinery’

The three challenges on the cyber part are typical computer science challenges. Incorporating knowledge into the diagnostic workflow is one of them. While it is tempting to unleash machine learning techniques on the data that high-tech machines already collect, in practice, it turns out that a lot of data is needed and that the machine learning methods lack transparency. ‘We want to combine a data-driven approach with a knowledge-driven approach’, says Stoelinga. ‘There is a lot of knowledge in the minds of technicians. We want to make that knowledge explicit in a graphical model. Such knowledge graphs have proven their value in medical diagnostics and information retrieval, for example.’

To these knowledge graphs, Stoelinga also wants to add techniques from reliability engineering in which failures on the system level can be broken down into failures in subsystems and sub-sub-systems. These techniques will help to understand and predict the causes of failures. Another computer science challenge is developing efficient and scalable diagnostic algorithms. Stoelinga: ‘There are many diagnostic systems at the component level but not at the systems level. How can we scale up component-based diagnostic methods to system-level diagnostics? That is the big challenge.’

Design for diagnostics

The final computer science challenge is the development of model-based system engineering methods that support design for diagnostics. ‘We want to integrate diagnostics into the normal operation of high-tech machines’, says Stoelinga. ‘Diagnostics should not be something that you put into your machine afterwards, but instead something you integrate into the machine already at the design stage. Part of this challenge is ensuring that sensors communicate correctly with each other and that they assign the same meaning to the measurements.’

Rather than first building a theory and then validating it in practice, the ZORRO project develops innovative solutions together with stakeholders. Stoelinga: ‘We apply this paradigm by working toward use cases of the industrial partners in the project. This allows us to practically shape and evaluate our methods, and to push fundamental science forward and develop concrete solutions for the challenges that our industrial partners are facing.’

KIC CALL ‘NEXTGENERATION HIGH TECH EQUIPMENT’

In January 2023, NWO awarded two projects that aim to improve the performance of high-tech cyber-physical systems. Both projects have a duration of six years, will start during the course of 2023 and involve several public and private partners.

The two projects came out of the Knowledge and Innovation Covenant (KIC) call ‘Key Enabling Technologies: NextGeneration High Tech Equipment’. This call has a total budget of 5.5 million euros and is aimed at supporting the Netherlands as an international leader in ultra-precise high-tech equipment, like the lithography machines of ASML. The programme focuses on integrating the key physical and digital components of a cyber-physical system. Digital technologies include artificial intelligence, big data and data analytics, and physical technologies include sensors, opto-mechatronics and robotics.

More information: **www.nwo.nl/cyberphysical**

Since 1 January 2023, Marieke Huisman has joined the Board of IPN. She succeeds Patricia Lago, whose second term as IPN Board member ended. Huisman: 'I want to make sure that everyone in the Dutch Computer Science community is involved in and feels represented by IPN.'

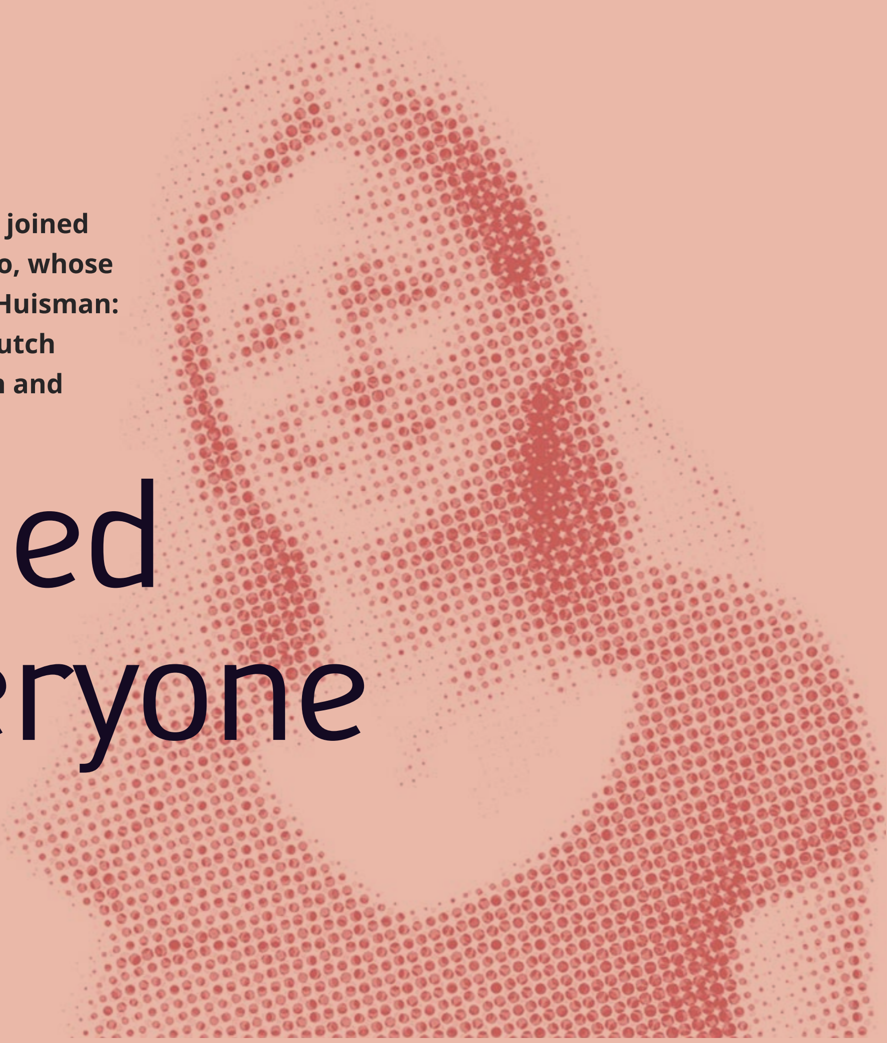

Determined to get everyone on board

By Sonja Knols

Marieke Huisman is no stranger to IPN, having acted as the representative for the University of Twente and the Special Interest Group for software engineering for many years. During those years, she witnessed the development of IPN into the broad platform it is today. 'Where the field of Computer Science in the Netherlands used to be somewhat divided, with people lobbying primarily for their own subdiscipline or institution, IPN has managed to unite the field over the past years. And that has paid off, for example in the sector plans and the resulting investments in our field.'

'I want to fire up ICT-Next Generation'

As far as Huisman is concerned, uniting the field and acting as a one-stop-shop for other stakeholders is the primary task of IPN. 'We act as the single interface to stakeholders, such as industry and government, and are able to speak for the entire community regardless of the subdiscipline or institution. We can identify shared concerns, challenges and strengths, and advocate the need to invest in Computer Science research. In our field, student-staff ratios are going through the roof. At the same time, society realises that there is a need for some counterweight to big tech companies that are now directing the future of digitalisation. As IPN, we need to make clear to policymakers why basic research in computer science is needed and how it can be used to counterbalance big tech, what role Dutch researchers can and should play in this, and what is needed for our country to keep pace with the fast technological developments in this field.'

MAKE VOICES HEARD

If IPN is to act as the voice of the Dutch Computer Science community, it is imperative that everyone is involved and has the feeling they are being heard, Huisman stresses. That means specific attention is needed to include the younger generation, people with different personalities, sexual orientations, and gender, and people from different cultural and ethnic backgrounds and socioeconomic environments. 'One of the priorities I have defined for myself is to fire up ICT-Next Generation. Together with its coordinator Cynthia Liem, I want to make optimal use of the possibilities and capabilities of this network of assistant and associate professors.' In addition to that, Huisman will further build on her work as a member of the IPN working group on Equity, Diversity and Inclusion, which, among others, issued a report on how to attract and retain more women to Dutch Computer Science.

If anything, Huisman wants to send out the message that IPN, first and foremost, is there to serve the interests of the field as a whole. 'So, if you have any ideas on how we could do better, if you want to raise topics we should discuss on a national level, or if you want to be involved in drafting new research agendas, please join one of our special interest groups or drop us an email, and make yourself heard. Together we can take a stand and advance our field.'

More information on IPN and the SIGs can be found at ict-research.nl

SECTOR PORTRAIT COMPUTER SCIENCE PUBLISHED

The government is structurally investing 200 million euros in scientific education and research through sector plans, using the resources of the Coalition Agreement 'Looking at each other' (2021-2025). To guide the process of allocating funds to the field of Computer Science in a very short period of time, the Sector Portrait Computer Science has been realised, coordinated by IPN and created in cooperation with all computer science faculties of Dutch universities.

The Portrait is complementary to other Dutch Sector Plans. It is partly covered by the Science Sector Plan (on behalf of the general universities), the Technology Sector Plan (on behalf of the technical universities) and the Social Sciences and Humanities Sector Plan (on behalf of Tilburg University).

The Sector Portrait Computer Science (in Dutch) can be found on the website of IPN, **ict-research.nl**.

TEACHER DELVES INTO CONVERSATIONAL AGENTS

Rhied Al-Othmani from the University of Applied Sciences Utrecht received a Doctoral Grant for Teachers to conduct PhD research at Utrecht University into automated Conversational Agents (e.g. chatbots and voice assistants) and their role in user services in the public domain. In the latest round, 23 laureates have received a Doctoral Grant for Teachers. With the grant, the teachers can further develop themselves and strengthen the link between universities and schools. In this twentieth funding round, 4.6 million euros was awarded.

ICT.OPEN2023

On Wednesday 19 and Thursday 20 April, the annual conference ICT.OPEN will be held in the Jaarbeurs Utrecht. ICT.OPEN2023 will showcase the best and most exciting developments in ICT research. More than 500 ICT researchers from academia and industry will come together to learn, share ideas and network.

Keynote speakers are Gail Murphy from the University of British Columbia, Harry Buhrman from QuSoft, Jos de Groot from the Ministry of Economic Affairs & Climate Policy and Omar Niamut from TNO. In 13 tracks, virtually all aspects of Computer Science will be covered. In addition, Cristiano Giuffrida from VU Amsterdam will receive the Dutch Prize for ICT Research 2023, worth € 50,000.

SOCIETAL IMPACT OF DATA2PERSON

The board of NWO Domain Science has awarded funding to six proposals within the call Data2Person: top-up for societal impact. The focus is on the way in which Data Science can contribute to a personalised offer and opportunities for self- and joint management of the individual health situation.

The Data2Person: top-up for societal impact was a closed call in which only project leaders of the ten current Data2Person projects could submit a proposal. The Data2Person call was part of the research programme Commit2Data.

The six granted projects cover topics like serious games for diabetics, self-monitoring based management of multiple sclerosis, personalised data for cancer patients, a clinical decision support system for medication-related fall risk, and voicebots for healthcare.

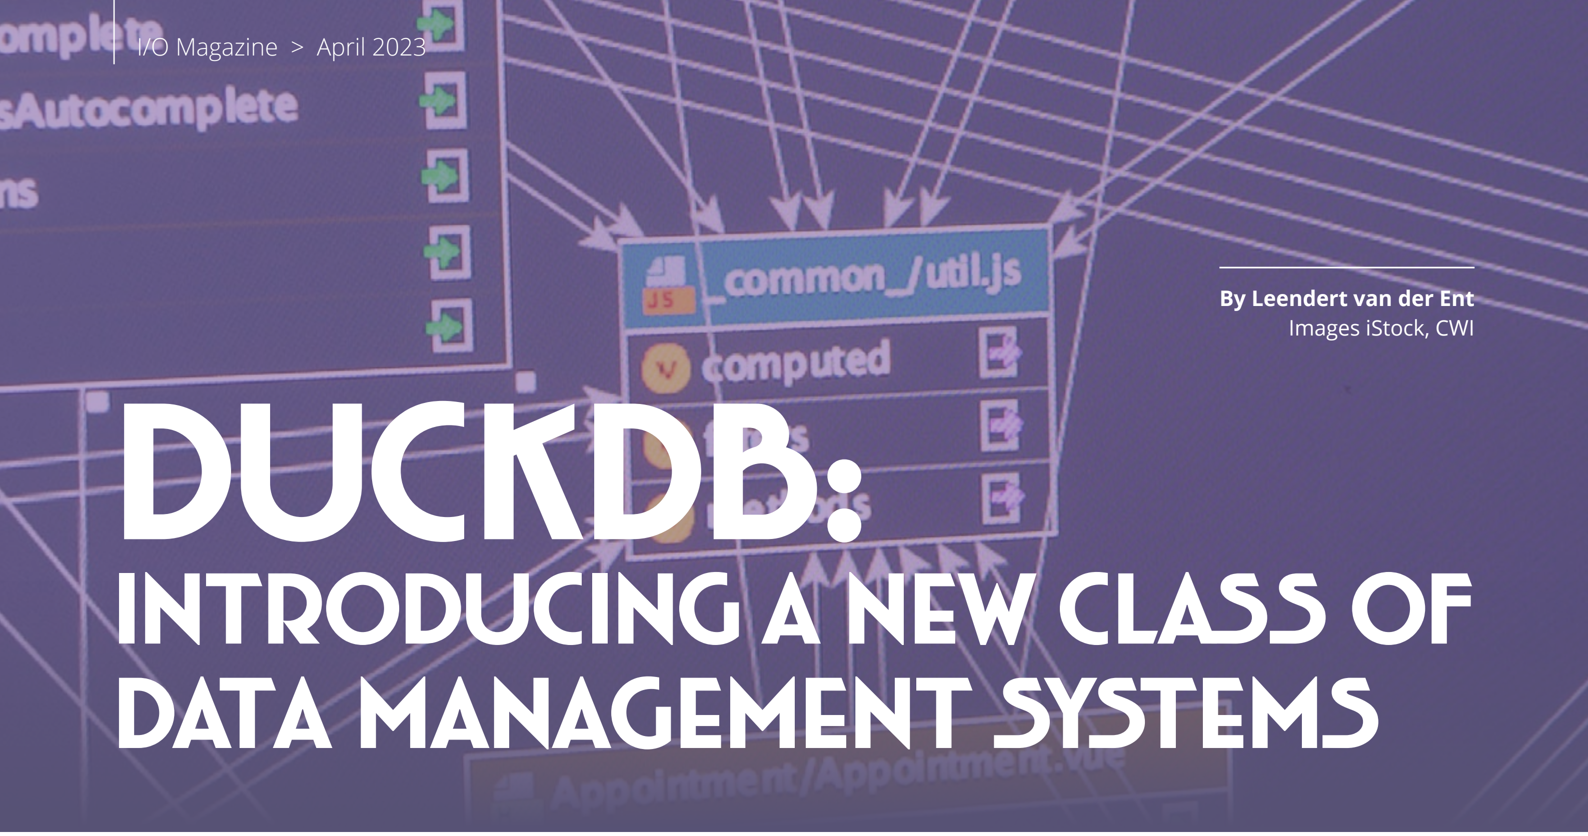

By Leendert van der Ent
Images iStock, CWI

DUCKDB: INTRODUCING A NEW CLASS OF DATA MANAGEMENT SYSTEMS

As CWI database architecture researchers, Hannes Mühleisen and Mark Raasveldt invented a new, much more efficient database technology for analysis. After its introduction in 2021 it now reaches two million downloads per month.

'The name? DuckDB comes from my late pet duck Wilbur', CEO and founder Hannes Mühleisen reveals. What DuckDB is, however, requires a more elaborate explanation. Mühleisen: 'As database architecture researchers at CWI, it struck co-founder Mark Raasveldt and I that out of four possible directions of database technologies, only three types existed.'

There are analytical and transactional workloads. Online Analytical Processing (OLAP) is optimised for queries and reports retrieved from large amounts of data. Online Transactional Processing (OLTP), on the other hand, supports the execution of a large number of real-time transactions. Another division in databases are the 'client/server' and the 'in-process' types. Put in a matrix, the quadrant for the combination of OLAP and in-process remained empty.

'That made it clear to us there was an opportunity for this new type of database technology', says Mühleisen. 'So we set out to develop a prototype in 2018. It didn't exist until then because it is complex to make. But Mark and I had a pretty good notion of how to tackle that. In 2021, we were ready to start a spin-off from CWI – which is what we did.'

OPEN SOURCE

As they already assumed, there is a lot of demand for in-process database analysis technology. Mühleisen: 'Especially as a component built into an application, it comes in very handy. Within two years, DuckDB managed to reach two million downloads a month worldwide.' More than one-third of the visitors to the website come from the USA; other users are located in Germany, Canada, France, the Netherlands, the UK and China. DuckDB is widely used in sciences which use huge datasets, such as genetics and astronomy. And it is even used in satellites.

The reason for the popularity: the state-of-the-art data engine. Its efficiency saves resources and energy. In practice, it enables analyses on a single laptop, which previously required dozens or even hundreds of computers. Jackpot! Well, not quite.

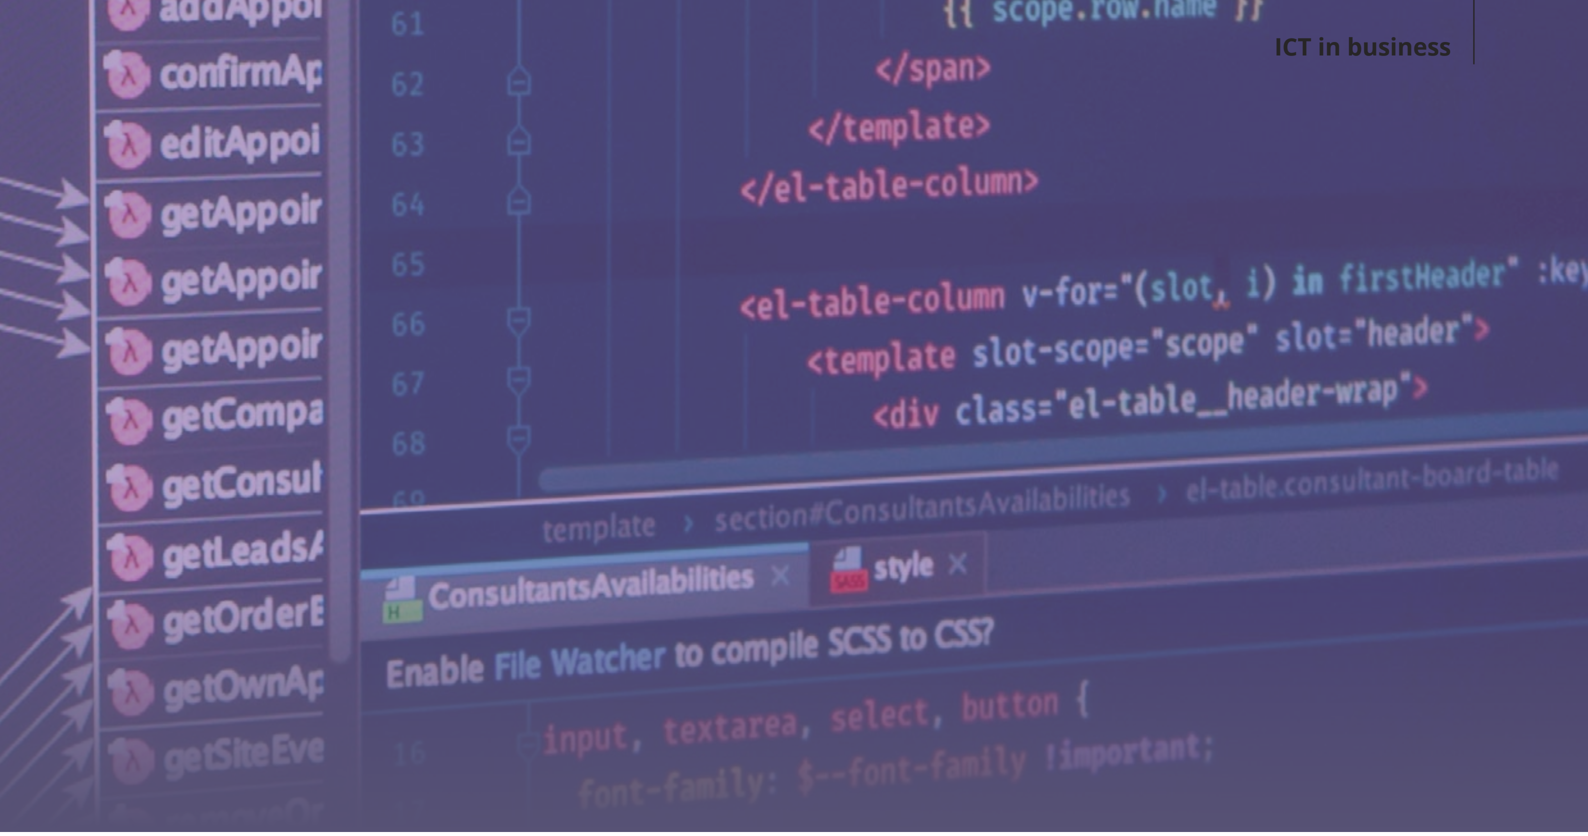

```
              {{ scope.row.name }}
            </span>
          </template>
        </el-table-column>

        <el-table-column v-for="(slot, i) in firstHeader" :key
          <template slot-scope="scope" slot="header">
            <div class="el-table__header-wrap">
```

template › section#ConsultantsAvailabilities › el-table.consultant-board-table

Enable File Watcher to compile SCSS to CSS?

```
input, textarea, select, button {
  font-family: $--font-family !important;
```

The founders chose to launch DuckDB as Open Source software. Mühleisen: 'We find it unethical to make proprietary software based on taxpayer-funded research. Researchers should always remember who pays their bill in the end.' The software project is managed by a non-profit foundation. If you donate to it, you get to provide input to the development roadmap.

COMMERCIAL SERVICES

The Open Source set-up does not mean there is no commercial activity though. The company DuckDB Labs offers paid services based on the DuckDB Open Source platform. Its customers include Google. Mühleisen explains: 'There is only one version of DuckDB. Our roadmap describes which features we'd like to add and when. Customers pay for us to develop extra features with priority or to develop features that we didn't envision so far but make sense to add. All these features become available to all other users.'

An interesting feature that DuckDB Labs developed is to work around computation and storage limitations as much as possible. Mühleisen: 'Part of our success comes from listening closely to people who work with DuckDB, which is quite rare in fundamental research. We learned that it is important for our users – mostly data analysts and not programmers – that they can always finish their query, regardless of hardware limitations. So we saw to it that when a query uses up all memory capacity, it doesn't abort but automatically switches to using other available options, such as using disk space.'

Mühleisen tends to think that DuckDB represents the future of data analysis. 'I'm biased, but I'm far from the only one to think that', he says. That counts for something, because as a senior researcher at CWI, professor of Data Engineering at Radboud University and employer of database specialists at DuckDB Labs, he is pretty well acquainted with the latest developments within the database research community. 'I can assure you that during the next couple of years, we will have enough work on our hands.'

Hannes Mühleisen

'PART OF OUR SUCCESS COMES FROM LISTENING CLOSELY TO USERS, WHICH IS QUITE RARE IN FUNDAMENTAL RESEARCH'

Data science for the common good

By Bennie Mols Images Ivar Pel

The Data Science group at Radboud University develops theory and methods for machine learning and information retrieval with a strong focus on social responsibility.

Tom Heskes

Franka Buytenhuijs

With the surge in the production of digital data and the explosion of machine learning applications over the past decade, it is no wonder that the Data Science group at Radboud University has grown significantly to some forty researchers. One of the group's key characteristics is its strong focus on social responsibility in general and a strong connection with applications in the health domain in particular, the latter via close cooperation with the Radboudumc hospital.

'Despite the growth of the Data Science group in recent years, we have decided to stick to our three core themes', says group leader Tom Heskes. 'We focus on causal reasoning for machine learning, biomedical applications of machine learning, and information retrieval and recommender systems. With the growth of the group, we considered whether we should split into more subgroups, but we decided not to do so precisely because there is a strong social cohesion running through the subgroups. Even though people are doing different things in terms of content, we do a lot of social activities together.'

One of the weaknesses of most machine learning applications developed and applied in recent years is that they are poor at reasoning about cause and effect, so-called causal reasoning. A machine learning application may conclude from data on smoking and lung cancer that the two are correlated but does not automatically understand that smoking can cause lung cancer. 'Machine learning techniques are based on making associations', says Heskes, 'but because they are bad at causal reasoning, they don't know what happens if you do an intervention, like banning smoking in public spaces.' In the field of causal reasoning for machine learning, the Data Science group is one of the largest research groups in the Netherlands.

Heskes and his colleagues try to extract more information from the data to reason about cause and effect. Heskes: 'One of the basic ideas that we use is that a model that goes from cause to effect is likely less complex than a model that goes from effect to cause. This essentially goes back to the philosophical thesis of Ockham's razor, which states that the simplest explanations are usually the best ones. For example, we have applied this idea to data about attention deficit hyperactivity disorder, ADHD. The data show that the attention deficit causes the hyperactivity and not the other way around. When machines get better at causal reasoning, it makes them much smarter and more robust in many applications.'

Artificial immune system

PhD candidate Franka Buytenhuijs has worked since 2020 in the second big theme of the Data Science group: biomedical applications of machine learning. She is part of the computational immunology subgroup. 'I work on a project called Artificial immune systems', she states. 'Just like the brain, the immune system is also a learning system. The brain was the inspiration for the development of neural networks. We are now looking for a system to describe the immune system's behaviour. How does the immune system learn? How does it remember? How does it forget? For example, we want to use the insights to study how the immune system determines which cells are harmful and which ones are harmless. My research focuses on a specific type of immune cell called T cells. From experimental data from Canadian colleagues, I am trying to find features that determine how strong T cells bind to viruses and the body's own cells.'

Before starting her PhD research in the Data Science group, Buytenhuijs had completed her Master's project in the same group. 'I have a background in AI, but I like to apply this knowledge in the medical domain', she says. 'Furthermore, I enjoy the broad diversity of topics in the Data Science group, the ease with which everybody can be approached and the minimum of hierarchy. And despite the diversity of topics, most group members use some form of AI technique, so we can still learn from each other. We get a chance to share our results in our bi-weekly seminar.'

Harrie Oosterhuis

Learning from clicks

The bi-weekly seminar is co-organised by assistant professor Harrie Oosterhuis whose research focuses on optimising ranking systems for search engines and recommender systems. His work is part of the third theme of the Data Science group: information retrieval and recommender systems. 'We develop statistical methods that learn from the click behaviour of users', says Oosterhuis. 'One of the applications is that search or recommender results that are displayed lower in the results list, but that people click on frequently, get an extra push to the top.' In recent years, the work of Oosterhuis has won three best paper awards at the top conferences in the information field.

People sometimes ask Oosterhuis if improving search and recommendation systems is not already solved by companies like Google or Microsoft. 'Of course, they are working on that as well', replies Oosterhuis, 'but companies like to solve it best for themselves and don't like to share their results. In our group, we think it is important that what we develop is freely available and open for investigation and improvement.' For example, group members Arjen de Vries and Djoerd Hiemstra are working on a European search engine that does not depend on large American tech companies. 'We do not focus solely on publishing papers but also aim to contribute to such socially responsible initiatives', says Oosterhuis.

GROUP PASSPORT

RESEARCH FIELD

Causal reasoning for machine learning,
Biomedical applications of machine learning,
Information retrieval and recommender systems.

INSTITUTION

The Data Science group is a section of the Institute for Computing and Information Sciences at Radboud University.

LABS

ICAI Lab 'Radboud AI for Health'
AI for Precision Health, Nutrition & Behaviour
AI for Energy Grids
AI for Parkinson

EMPLOYEES (as of March 2023)

8 professors (1 by special appointment, some with part-time appointments, in total 5 fte), 2 associate professors, 7 assistant professors, 4 postdocs, 43 PhD students (including 17 external PhD students), 2 support staff

WEBSITES

Data Science group: **www.ru.nl/datascience**
Radboud AI (campus-wide initiative connecting all activities on Artificial Intelligence and Data Science within Radboud University and Radboudumc): **www.ru.nl/ai**

Cristiano Giuffrida is an associate professor in the Computer Science Department at VU Amsterdam, where he received his PhD cum laude in 2014. He was awarded the Roger Needham Award and the Dennis M. Ritchie Award for the best PhD dissertation in Computer Systems in 2015 (Europe and worldwide), a Veni grant in 2017, a VMware Early Career Faculty Award in 2020, and a Jochen Liedtke Young Researcher Award in 2022.

Dutch Prize for ICT Research

The Dutch Prize for ICT Research is awarded annually to a scientific researcher who has carried out innovative research within 15 years of earning their PhD. New this year is that the prize is funded by members of IPN with a grant from COMMIT\, through the Royal Holland Society of Sciences and Humanities (KHMW).

FOCUS ON THE VULNERABILITY OF HARDWARE

‘Software is better protected than ever, whereas age-old vulnerabilities in hardware remained overlooked.’ That is the pivotal insight of Dutch ICT Prize winner Cristiano Giuffrida. So he started working on weaponizing hardware.

By Leendert van der Ent

Image Ivar Pel

What's the current status of software security?

‘Eighty percent of attacks exploit known vulnerabilities. Attackers focus on users who haven’t updated their software in time. Therefore, updates currently come as light updates installed on the fly – an essential procedure for data centres and power plants that can never reboot. Moreover, nowadays, users are more or less forced to update before they can continue working.’

When did you decide to focus your research on hardware?

‘Patches often used to contain errors, causing the system to falter. But the bar in software has been raised significantly, to the point where attackers target another weakest link: hardware vulnerabilities. I started looking into chip design in 2017 and found lots of problems. Criminals can use software-based attack models that exploit hardware vulnerabilities.’

Where do these vulnerabilities originate from?

‘In the 1990s, to speed up CPUs, chip designers chose to abandon an in-order instruction execution model and went for an out-of-order speculative execution model instead. This means running dozens of instructions simultaneously and preparing for assumptions that will become true in most cases, such as which instructions should be executed next. If an assumption proves wrong, a process will roll back to the old state, discarding many speculatively executed instructions and their user-visible side effects. While rollbacks do cause some delay, overall, this execution model works much faster. There is, however, a big but. As only user-visible side effects are rolled back, microarchitectural components deep inside the CPU such as CPU caches still preserve sensitive data traces of the incorrectly executed instructions. And attackers can exploit side-channel analysis techniques to leak such sensitive data even after rollback. Moreover, since attackers can control speculative execution, they can also lure the CPU into speculatively accessing very specific sensitive data they intend to leak, such as credit card numbers and passwords. This has been optimised to the point that a master’s student, after attending my two-month hardware security course on the matter, can retrieve a root password on a borrowed laptop in 300 microseconds.’

What can be done about the vulnerabilities?

‘In the long term, thorough CPU redesign is needed. For the medium term, it involves more principled software redesign, for instance, limiting the locations where data is placed in memory to secure domains. And for the short term, stop-gap solutions in software are often possible, such as selectively disabling speculation in sensitive locations. I have worked with all the major hardware (such as Intel and AMD) and software (such as Google, Microsoft, and Amazon) vendors to ensure such remedies are readily deployed in production.’

A SOLID THEORY FOR POST-QUANTUM CRYPTOGRAPHY

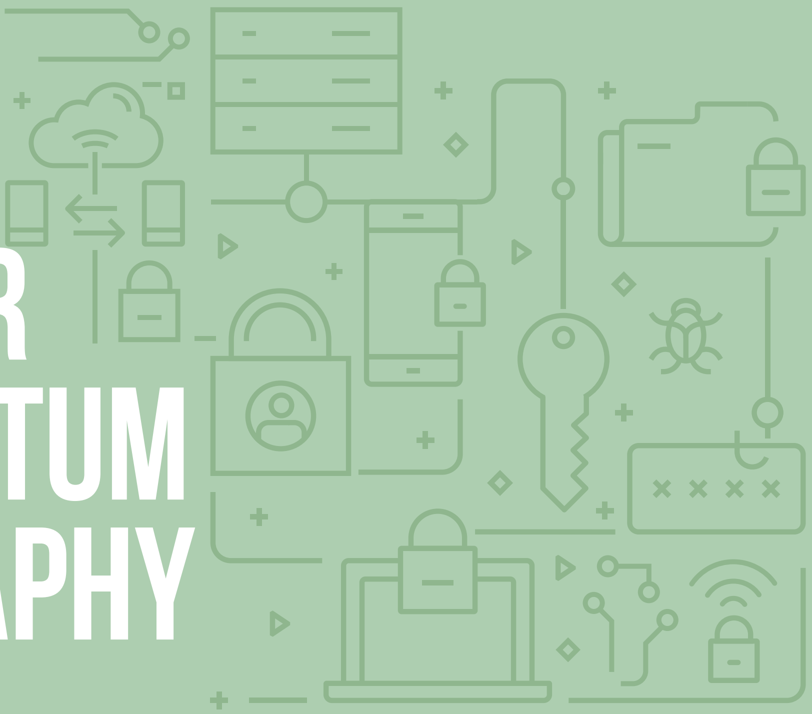

By Bennie Mols Image iStock

The advent of the quantum computer requires new cryptographic ways to secure our digital information and services. In his Vidi project, Andreas Hülsing develops novel methods to create secure cryptographic schemes and protocols.

After obtaining his master's in computer science, Andreas Hülsing worked as a security consultant for Fraunhofer SIT in Germany for a couple of years. 'In those days there was a discussion about electronic identity', Hülsing states. 'What struck me was how often handwaving or sometimes even political arguments were used to take security critical decisions. I found that very unsatisfying and preferred to base these decisions on mathematical arguments instead.'

Hülsing decided to pursue a PhD in the field of cryptography. Now he works as an associate professor at Eindhoven University of Technology, leading the group for applied and provable security at the Department of Mathematics and Computer Science. In 2021, he started his five-year Vidi research project *A solid theory for post-quantum cryptography*. While a practical quantum computer is still under development, it is already known that it will allow malicious actors to break today's public key cryptography. To protect digital information and services against attacks with a quantum computer, a new type of cryptography is needed, the so-called post-quantum cryptography.

New standards

In July 2022, the US National Institute of Standards and Technology (NIST) selected new, post-quantum crypto-standards for public-key encryption and digital signatures, all based on mathematical problems that are supposed to be extremely hard to break, even for quantum computers. All of these schemes are backed up by security proofs, some of them by Hülsing and his collaborators. One of his own cryptographic schemes called SPHINCS+ was selected as a new standard.

Hülsing: 'Over the next few years, the mathematical schemes chosen by NIST have to be transformed into protocols that are going to be used in practice. In my Vidi project, I want to make sure that we get secure connections. I want to bridge the gap between the mathematical building blocks and the technical protocols that are going to be used in practice. During my Vidi project, I think that we will be able to formulate the first generation of protocols.'

'I WANT TO MAKE SURE THAT WE GET SECURE CONNECTIONS'

Besides Hülsing, two PhD students and a postdoc researcher are also working on the Vidi project. One of the PhD students is studying the use of the computer as a mathematical proof assistant. Hülsing: 'In the NIST competition to look for new cryptographic standards, several mistakes have been found over the years. Often, a mistake can be fixed, but recently an ISO standard was shown to be entirely breakable. The proof came from one of the big names in the field, which shows how hard it has become to come up with a solid proof. They often cover dozens of pages. And cryptographic proofs get a lot more complicated if you take quantum attackers into account. Therefore I am interested in using the computer to check these cryptographic proofs.'

Since Hülsing tries to combine a theoretical and an applied perspective on post-quantum security, he is working closely together with industry and standardisation bodies. 'Post-quantum cryptography comes with additional costs for implementation. However, for information that needs to be secure for many decades, like some governmental or medical data, it will be obliged by legislation. The development of secure and practical post-quantum standards is urgently needed. Not only in the future but also today.'

By Rianne Lindhout
Images WAT ontwerpers

Robust algorithms with societal value

Universities from Groningen to Curaçao are involved, as are companies ranging from NS to RTL. In 17 labs, the ten-year research programme ROBUST will use artificial intelligence to provide sustainable solutions to societally relevant issues in areas such as healthcare, logistics, media, food and energy. 'Our approach with labs like this one has already proven itself.'

'The approach with labs like this has already proven itself'

NS wants a data-driven, dynamic system to make the best use of the increasingly crowded track and available seats. RTL wants to have the technology to offer personalised summaries of football matches, and does not want to be overtaken left and right by companies like Google and Meta. Grid operators want to make optimal use of the power grid by regulating the grid load with smart charging batteries.

Besides intelligent systems, companies also want something else. They want to become more visible to technical talent, people who often do do not know they can do interesting work at Siemens or ProRail. By collaborating with universities and letting in master's and PhD students for research, companies hope to increase the chances that students will later apply for a job with them.

If it were up to the companies, the programme leader of ROBUST Maarten de Rijke could have put together more than the current 17 labs. The professor of Artificial Intelligence and Information Retrieval at the University of Amsterdam explains how the ROBUST programme, which NWO gave the green light in January, will make important strides in research and applications of artificial intelligence.

PhD students in action

'The approach with labs like this has already proven itself', says De Rijke. 'We have been working in this way within ICAI since 2018, with Qualcomm, Bosch and Ahold Delhaize, among others.' In this Innovation Centre for Artificial Intelligence, knowledge institutions, companies, governmental organisations, and societal organisations work together. The ROBUST programme emerged from this approach.

The duration of a lab is always five years, slightly longer than a PhD track. ROBUST will have two rounds of labs, 17 at the start, and another 17 at the mid-term point, with each lab employing five PhD students. 'PhD students work one or two days a week at the company. They learn from the company's problems and context, share the expertise they build, and often co-create solutions.'

This makes it seem like every lab is a silo, in which the same wheel might be invented several times. That is not what De Rijke expects. 'All labs consider socio-technical problems and the challenges of trustworthy systems in a specific context. The contexts differ substantially. In selecting the right ad for a media consumer a mistake might not be that bad, but in medical applications, the algorithm must be particularly resilient against changes in the world and in user preferences.' The exchange of knowledge and best practices is something De Rijke keeps an eye on together with Bob Huisman and others. Huisman is research and development manager at NS, and in that capacity, a participant in one of the 17 ROBUST labs: the RAIL lab. He is also involved at the programme level as a principal investigator and chair of the overall users' committee, which promotes cooperation between universities and companies.

Huisman: 'At NS, we have been working with several universities for decades on systems for logistic planning that can find needles in a haystack. At the front, you see our fixed timetable, but at the back, it is a new puzzle every day. Due to track maintenance, breakdowns or events, for instance. That puzzle becomes increasingly complex as it gets busier on the tracks. With AI, you can solve those puzzles better and faster.'

Social sciences and humanities

An important aspect of the ROBUST labs, Huisman thinks, is that one out of every five PhD students comes from the social sciences and humanities. 'They study, for example, the aspects a human being considers in accepting advice from an AI system. What exactly does the user expect? The system needs to align with their preferences and somehow support their decision-making processes.'

That's the goal, but how do you get there? De Rijke outlines: 'You can't plan projects like this entirely. You put a dot on the horizon, a challenging transformation that you would like to achieve with your partners. You might not know how to get

‘With AI, you can solve logistic puzzles better and faster’

there, at least not in detail. But you do know the first steps. Once those are taken, you re-assess. Is this still what we want? What have other colleagues done? Naturally, we continuously monitor progress along the axes of trustworthy AI that we want to advance scientifically.’

Hallucinate

With the term axes of trustworthiness, De Rijke is referring to explainability, accuracy, reliability, repeatability, resilience, and safety. ‘We avoid launching algorithms that are not fit for purpose. For instance, an algorithm that is supposed to present factual information to users, should not hallucinate. But if it is only put to work in a creative context, it might help artists and designers to come up with new ideas.’

Small and medium-sized enterprises, which do not have the resources to fund a ROBUST lab, have not been forgotten in the programme. ‘Besides the 17 labs in which we conduct research and learn a lot by doing, we organise clinics for SMEs’, says De Rijke. ‘Companies can go there with questions about data and algorithms. Building on the extensive experience of JADS in Den Bosch, we have set up an intake process. Once a company has been admitted, master students get to work with the company for eight to twelve weeks on a paid basis, with guidance from our labs.’

ROBUST AT A GLANCE

BUDGET

87 million euros, of which NWO funds 25 million euros through the so-called long-term programmes (LTPs) of the Knowledge and Innovation Covenant (KIC) innovation programme, and 7.5 million euros is funded by the Ministry of Economic Affairs and Climate Change.

DURATION

10 years.

PARTNERS INVOLVED

More than 50, including universities, universities of applied sciences, 19 companies and 15 civil society organisations. The project is led by the University of Amsterdam.

IMPROVING COMPUTER SCIENCE EDUCATION

By Sonja Knols
Images Utrecht University

Johan Jeuring

Head Department of Information and Computing Sciences,
Utrecht University 2017-2023

'The whole process of developing a strategic document like a sector plan in itself can be a rather useful exercise. In 2017, when I started as head of the department, we wrote a strategic plan, identifying three main areas of interest for our research. At that time, Utrecht had one of the largest computer science programmes in the Netherlands. In the discussions with our external advisory board, we identified an important theme to focus on as a department with many students: computer science education.

For the first sector plan, we identified seven themes at a national level that the Netherlands should invest in. We looked at where we experienced the biggest pressure from education and tried to relieve the burden there. After talking to external partners and taking stock of who was working on which topics and providing what kind of education, we defined priorities for the positions to be filled in. At Utrecht University, that resulted in seven new junior positions on our, then three (now four, having added AI & Data Science, ed) themes Algorithms, Interaction, and Intelligent software systems, one of which we decided to fill in jointly with Eindhoven University of Technology.

As a head of department, you are constantly balancing budgets. What plan or ambition are we going to fund from which budget? Although for the Computer Sciences part of Utrecht University, the sector plan budgets are rather modest in size when compared to the overall departmental budget, the nice thing about them is the freedom to operate that they bring. For example, these budgets allow us to jointly appoint people together with other institutions, which is not that easy with other types of funding. And they offer possibilities to solve some of the bottlenecks we are confronted with, or to invest in upcoming areas of interest.

One of these areas is research into computer science education. What is the best way to teach someone to code or to develop new algorithms? Our Software Technology for Learning and Teaching group is one of the largest in this field in the Netherlands, with a clear and visible impact on both science and society. New hires need to bring in fresh perspectives. And that is exactly what we got with Hieke, not in the least because of her background and extensive practical experience in teaching.'

The two subsequent sector plans have resulted in a significant number of new hires at various Dutch universities. In this diptych, former Head of the Department of Information and computing sciences at Utrecht University Johan Jeuring reflects on how they decided on where to invest, and assistant professor Hieke Keuning explains what her appointment brings to the field.

Hieke Keuning

Assistant professor in the Software Technology for Learning and Teaching group since September 2020

'Soon after obtaining my Bachelor in Informatics from Hanze University of Applied Sciences in Groningen, I started working as a lecturer in software engineering at Windesheim University of Applied Sciences. As I wanted to expand on my expertise then in addition to my daytime job as a lecturer, I studied for my Master in Computer Science at the Open University. My master's thesis was about tutoring systems for learning programming. That's when I got interested in the subject.

When I heard about the NWO Doctoral Grant for Teachers, I applied, and, besides keeping up my job as a teacher for three days a week, I became a PhD researcher. Johan was one of my thesis supervisors, and during my PhD, I visited his group as a guest researcher on several occasions. So when this vacancy came up, briefly before I finished my PhD, I did not hesitate to apply for it.

My research focuses on the question of how to use software techniques such as automatic code analysis and automatic feedback systems to optimise the learning process of computer science students. What kinds of problems do students encounter when they learn, and what tools could help them overcome these? What is the best way to give feedback to them and what should this feedback consist of? How do students use these types of tools, and what problems do they encounter while using them?

My main aim is to help students learn to write good-quality code. Code should not only do what it is supposed to do but should also be understandable for and maintainable by others. In my research, I mostly focus on the student's perspective. So what do they think is good code? And how can we use tools to help them see the flaws in their work and improve on them? In addition, from a teacher's perspective, we look into what is perceived as valuable and helpful feedback and at what moment in time that should be given.

We already know a lot about how to teach courses related to computer science, and how to improve learning experiences. However, we could use that knowledge more in practice. Especially in computer science, we need to keep as many people on board as possible, ranging from the first-year undergraduates who have never written a line of code to those who started coding from a very young age.'

Joost Visser is Professor of Large Scale Software and Data Science at Leiden University, Programme Manager of the Master in ICT in Business and the Public Sector, and Head of the LIACS (Leiden Institute of Advanced Computer Science) Software Lab. Previously, he fulfilled R&D leadership positions at the Software Improvement Group (SIG).

A CRITICAL YET CONSTRUCTIVE ROLE

By Marysa van den Berg Image Sjoerd van der Hucht

As of January this year, Joost Visser of Leiden University became a member of the Advisory Council on IT Assessment. He states that the council has come a long way and is still improving.

'The primary role of the Advisory Council on IT Assessment (in Dutch: Adviescollege ICT-toetsing) is to critically assess the chances of success of large IT-projects – with budgets exceeding 5 million euros – of the national government and to provide recommendations for improvement. The council has been doing this successfully over the last eight years. One of its accomplishments is that the great majority of its recommendations are accepted and implemented by the project organisers.

Governmental IT projects are as large and as complex as IT projects get. And they are often surprisingly innovative. The amount of data, users, requirements and connections with other systems can be overwhelming. As a council, we try to be critical but also fair and constructive. It is easy to point out apparent problems, such as budget overruns or failed deadlines. Uncovering the root causes and indicating effective steps for correction and control form a bigger challenge.'

CRITICAL ASSESSMENT

'I do like the challenge of the job. In the past, I have been assessing IT projects in the private and public sector in a different role with the Software Improvement Group. Assessing IT projects and systems is interesting and rewarding, as it requires a lot of critical thinking, careful analysis, but also situational awareness, and clarity of presentation. Things I enjoy as a scientist as well. And I think the role of independent assessor of complex IT in general, and especially the council's role within the governmental IT landscape, is essential, as it ultimately helps to successfully develop our digital society.

In addition, the council must also critically reflect on its own activities and welcome feedback. In fact, several years ago, I wrote an opinion article with a colleague where we, as outsiders, formulated several suggestions for improvement for the council, which was then still called Bureau ICT Assessment (BIT). One of these was to further standardise its assessment procedures and to increase transparency. Now I have joined the council, I am happy to see that it has grown and developed in these and other aspects over the years. I feel the responsibility to help continue that movement.'